

CHRISTOS PAPADIMITRIOU
SHUZHONG ZHANG (Eds.)

LNCS 5385

INTERNET AND NETWORK ECONOMICS

4TH INTERNATIONAL WORKSHOP, WINE 2008
SHANGHAI, CHINA, DECEMBER 2008
PROCEEDINGS

 Springer

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Christos Papadimitriou Shuzhong Zhang (Eds.)

Internet and Network Economics

4th International Workshop, WINE 2008
Shanghai, China, December 17-20, 2008
Proceedings

Volume Editors

Christos Papadimitriou
Computer Science Division
University of California at Berkeley
EECS Department
Berkeley, CA, USA
E-mail: christos@cs.berkeley.edu

Shuzhong Zhang
Department of Systems Engineering
and Engineering Management
The Chinese University of Hong Kong
Shatin, N.T., Hong Kong
E-mail: zhang@se.cuhk.edu.hk

Library of Congress Control Number: 2008940843

CR Subject Classification (1998): H.4, K.4.4, J.1, H.3, H.2, G.1.2

LNCS Sublibrary: SL 3 – Information Systems and Application, incl. Internet/Web and HCI

ISSN 0302-9743
ISBN-10 3-540-92184-2 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-92184-4 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2008
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12582938 06/3180 5 4 3 2 1 0

Preface

This volume contains the papers presented at the International Workshop on Internet and Network Economics held during December 17–20, 2008, in Shanghai, China, for its fourth edition. WINE 2008 provided a forum for researchers from different disciplines to communicate with each other and exchange their researching findings in this emerging field.

WINE 2008 had ten invited speakers: Fan Chung Graham, Matthew Jackson, Lawrence Lau, Tom Luo, Eric Maskin, Paul Milgrom, Christos Papadimitriou, Herbert Scarf, Hal Varian and Yinyu Ye. There were 126 submissions. Each submission was reviewed on average by 2.5 Programme Committee members. The Committee decided to accept 68 papers. The programme also included 10 invited talks.

This final program contained papers covering topics including equilibrium, information markets, sponsored auction, network economics, mechanism design, social networks, advertisement pricing, computational equilibrium, network games, algorithms and complexity for games.

December 2008

Christos Papadimitriou
Shuzhong Zhang

Organization

Programme Chairs

Conference Chair	Herbert E. Scarf (Yale University)
Program Co-chair	Christos Papadimitriou (UC Berkeley)
Program Co-chair	Shuzhong Zhang (Chinese University of Hong Kong)

Local Organizing Committee Chairs

Committee Chair	Yifan Xu (Fudan University)
Committee Co-chair	Duan Li (Chinese University of Hong Kong)
Committee Co-chair	Shouyang Wang (Chinese Academy of Sciences)
Committee Co-chair	Xiaoping Zhao (SSE INFONET Ltd.)

Invited Speakers

Special Session Speakers

Lawrence Lau	Chinese University of Hong Kong
Eric Maskin	Princeton University

Plenary Speakers

Matthew Jackson	Stanford University
Paul Milgrom	Stanford University
Hal Varian	University of California at Berkeley
Yinyu Ye	Stanford University

Tutorial Speakers

Fan Chung Graham	UC San Diego
Tom Luo	University of Minnesota
Christos Papadimitriou	University of California at Berkeley
Herbert Scarf	Yale University

Programme Committee

Gagan Aggarwal	Google Inc.
Dirk Bergemann	Yale University
Jesus-Mario Bilbao	University of Seville
Tian-Ming Bu	East China Normal University
Yiling Chen	Harvard University

Xi Chen	Institute for Advanced Study, Princeton University
Chuangyin Dang	City University of Hong Kong
Gabrielle Demange	Paris-Jourdan Sciences Economiques
Qizhi Fang	Ocean University of China
Joan Feigenbaum	Yale University
Paul Goldberg	University of Liverpool
Rica Gonen	Yahoo! Research
Ming-Yang Kao	Northwestern University
Marek Karpinski	University of Bonn
Spyros Kontogiannis	University of Ioannina
Elias Koutsoupas	University of Athens
Stefano Leonardi	Università di Roma, La Sapienza
Mohammad Mahdian	Yahoo! Research
Alberto Marchetti-Spaccamela	Università di Roma, La Sapienza
Evangelos Markakis	CWI, Amsterdam
Marios Mavronicolas	University of Cyprus
Burkhard Monien	University of Paderborn
Rajeev Motwani	Stanford University
Kazuo Murota	The University of Tokyo
Paul Spirakis	Research Academic Computer Tech Institute (CTI)
Bernhard von Stengel	London School of Economics and Political Science
Éva Tardos	Cornell University
Shanghai Teng	Boston University, Akamai Technologies Inc. Tsinghua University, Microsoft Research
Vijay Vazirani	Georgia Institute of Technology
Makoto Yokoo	Kyushu University
Kiho Yoon	Korea University
Yunhong Zhou	HP Labs

Local Organization

WINE 2008 was organized by the Department of Management Science, Fudan University, Shanghai, China.

External Reviewers

Saeed Alaei	Vincenzo Bonifaci
Esteban Arcaute	Konstantin Busch
Bahman Bahmani	Rahul Deb
Luca Becchetti	Gabrielle Demange
Dirk Bergemann	Dominic Dumrauf
Vittorio Bilò	Robert Elsässer

VIII Organization

Itay Fainmesser
Angelo Fanelli
Rainer Feldmann
Dimitris Fotakis
Martin Gairing
Donato Gerardi
Arpita Ghosh
Gagan Goel
Ashish Goel
Rica Gonen
Mingyu Guo
Nicole Immorlica
Atsushi Iwasaki
Alexis Kaporis
Chinmay Karande
Fuhito Kojima
Lucas Maestri
Aranyak Mehta
Henning Meyerhenke
Luca Moscardelli

Rong Pan
Panagiota Panagopoulou
Giuseppe Persiano
Georgios Piliouras
Tim Roughgarden
Matthias Ruhl
Amin Saberi
Yuko Sakurai
Florian Schoppmann
Ulf-Peter Schroeder
Aneesh Sharma
Paul Spirakis
Mukund Sundarajan
Tobias Tscheuschner
Angelina Vidali
Berthold Voecking
Kiho Yoon
Martin Zinkevich

Table of Contents

Invited Talks 1: Special Session

Mechanism Design Theory: How to Implement Social Goals	1
<i>Eric Maskin</i>	
Thirty Years of Chinese Economic Reform: Reasons for Its Success and Future Directions	2
<i>Lawrence J. Lau</i>	

Invited Talks 2: Plenary Session

Average Distance, Diameter, and Clustering in Social Networks with Homophily	4
<i>Matthew O. Jackson</i>	
Assignment Exchanges	12
<i>Paul Milgrom</i>	
Search Engine Ad Auctions	13
<i>Hal R. Varian</i>	
Computational Economy Equilibrium and Application	14
<i>Yinyu Ye</i>	

Invited Talks 3: Tutorial Session

Four Graph Partitioning Algorithms	15
<i>Fan Chung Graham</i>	
Dynamic Spectrum Management: Optimization and Game Theoretic Formulations	16
<i>Zhi-Quan (Tom) Luo</i>	
Some Recent Results in Algorithmic Game Theory	17
<i>Christos Papadimitriou</i>	
The Elements of General Equilibrium Theory	18
<i>Herbert E. Scarf</i>	

Session A.1: Market Equilibrium

A Fast and Simple Algorithm for Computing Market Equilibria	19
<i>Lisa Fleischer, Rahul Garg, Sanjiv Kapoor, Rohit Khandekar, and Amin Saberi</i>	

A FPTAS for Computing a Symmetric Leontief Competitive Economy Equilibrium	31
<i>Zhisu Zhu, Chuangyin Dang, and Yinyu Ye</i>	
Online and Offline Selling in Limit Order Markets	41
<i>Kevin L. Chang and Aaron Johnson</i>	
Predictive Pricing and Revenue Sharing	53
<i>Bobji Mungamuru and Hector Garcia-Molina</i>	
Dual Payoffs, Core and a Collaboration Mechanism Based on Capacity Exchange Prices in Multicommodity Flow Games	61
<i>Luyi Gui and Özlem Ergun</i>	
Session B.1: Congestion Games	
Graphical Congestion Games	70
<i>Vittorio Bilò, Angelo Fanelli, Michele Flammini, and Luca Moscardelli</i>	
How Hard Is It to Find Extreme Nash Equilibria in Network Congestion Games?	82
<i>Elisabeth Gassner, Johannes Hatzl, Sven O. Krumke, Heike Sperber, and Gerhard J. Woeginger</i>	
On the Road to \mathcal{PLS} -Completeness: 8 Agents in a Singleton Congestion Game	94
<i>Dominic Dumrauf and Burkhard Monien</i>	
Conflicting Congestion Effects in Resource Allocation Games	109
<i>Michal Feldman and Tami Tamir</i>	
The Price of Malice in Linear Congestion Games	118
<i>Aaron Roth</i>	
Session C.1: Information Markets	
Parimutuel Betting on Permutations	126
<i>Shipra Agrawal, Zizhuo Wang, and Yinyu Ye</i>	
Strategies in Dynamic Pari-Mutual Markets	138
<i>Tian-Ming Bu, Xiaotie Deng, Qianya Lin, and Qi Qi</i>	
Truthful Surveys	154
<i>Nicolas Lambert and Yoav Shoham</i>	
Correlated Equilibrium of Bertrand Competition	166
<i>John Wu</i>	

Diffusion of Innovations on Random Networks: Understanding the Chasm	178
<i>Marc Lelarge</i>	

Session A.2: Nash Equilibrium I

An Efficient PTAS for Two-Strategy Anonymous Games	186
<i>Constantinos Daskalakis</i>	
Equilibria of Graphical Games with Symmetries (Extended Abstract)...	198
<i>Felix Brandt, Felix Fischer, and Markus Holzer</i>	
Equilibrium Points in Fear of Correlated Threats	210
<i>Spyros C. Kontogiannis and Paul G. Spirakis</i>	
Performance Evaluation of a Descent Algorithm for Bi-matrix Games...	222
<i>Haralampos Tsaknakis, Paul G. Spirakis, and Dimitrios Kanoulas</i>	
Worst-Case Nash Equilibria in Restricted Routing	231
<i>Pinyan Lu and Changyuan Yu</i>	

Session B.2: Network Games I

Stackelberg Routing in Arbitrary Networks	239
<i>Vincenzo Bonifaci, Tobias Harks, and Guido Schäfer</i>	
Computational Aspects of a 2-Player Stackelberg Shortest Paths Tree Game	251
<i>Davide Bilò, Luciano Gualà, Guido Proietti, and Peter Widmayer</i>	
Local Two-Stage Myopic Dynamics for Network Formation Games	263
<i>Esteban Arcaute, Ramesh Johari, and Shie Mannor</i>	
Interference Games in Wireless Networks	278
<i>Vincenzo Auletta, Luca Moscardelli, Paolo Penna, and Giuseppe Persiano</i>	
Taxing Subnetworks	286
<i>Martin Hoefer, Lars Olbrich, and Alexander Skopalik</i>	

Session C.2: Solution Concepts

Anonymity-Proof Voting Rules	295
<i>Vincent Conitzer</i>	
Overlapping Coalition Formation	307
<i>Georgios Chalkiadakis, Edith Elkind, Evangelos Markakis, and Nicholas R. Jennings</i>	

A Network-Based Asymmetric Nash Bargaining Solution	322
<i>Edoardo Gallo</i>	
How Public Opinion Forms	334
<i>Fang Wu and Bernardo A. Huberman</i>	
A Game-Theoretic Analysis of Games with a Purpose	342
<i>Shaili Jain and David C. Parkes</i>	

Session A.3: Algorithms and Optimization I

Inapproximability of Combinatorial Public Projects	351
<i>Michael Schapira and Yaron Singer</i>	
Algorithms for Optimal Price Regulations	362
<i>Alexander Grigoriev, Joyce van Loon, and Marc Uetz</i>	
Improving the Efficiency of Load Balancing Games through Taxes	374
<i>Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos</i>	
Network Formation and Routing by Strategic Agents Using Local Contracts	386
<i>Elliot Anshelevich and Gordon Wilfong</i>	
Network Creation Games with Disconnected Equilibria	394
<i>Ulrik Brandes, Martin Hoefer, and Bobo Nick</i>	

Session B.3: Mechanism Design I

Randomized Truthful Mechanisms for Scheduling Unrelated Machines	402
<i>Pinyan Lu and Changyuan Yu</i>	
Optimal Mechanisms for Single Machine Scheduling	414
<i>Birgit Heydenreich, Debasis Mishra, Rudolf Müller, and Marc Uetz</i>	
Welfare Undominated Groves Mechanisms	426
<i>Krzysztof Apt, Vincent Conitzer, Mingyu Guo, and Evangelos Markakis</i>	
Redistribution of VCG Payments in Assignment of Heterogeneous Objects	438
<i>Sujit Gujar and Narahari Yadati</i>	
Bin Packing of Selfish Items	446
<i>Guosong Yu and Guochuan Zhang</i>	

Session C.3: Network Games II

Restricted Core Stability of Flow Games	454
<i>Mao-cheng Cai and Qizhi Fang</i>	
Three Selfish Spanning Tree Games	465
<i>Laurent Gourvès and Jérôme Monnot</i>	
Stochastic Submodular Maximization	477
<i>Arash Asadpour, Hamid Nazerzadeh, and Amin Saberi</i>	
On Pure and (Approximate) Strong Equilibria of Facility Location Games	490
<i>Thomas Dueholm Hansen and Orestis A. Telelis</i>	
Efficiency, Fairness and Competitiveness in Nash Bargaining Games	498
<i>Deeparnab Chakrabarty, Gagan Goel, Vijay V. Vazirani, Lei Wang, and Changyuan Yu</i>	

Session A.4: Equilibrium

Computing an Extensive-Form Correlated Equilibrium in Polynomial Time	506
<i>Wan Huang and Bernhard von Stengel</i>	
Homogeneous Interference Game in Wireless Networks	514
<i>Joseph (Seffi) Naor, Danny Raz, and Gabriel Scalosub</i>	
A Network Coloring Game	522
<i>Kamalika Chaudhuri, Fan Chung Graham, and Mohammad Shoaib Jamall</i>	

Session B.4: Mechanism Design II

Asynchronous Best-Reply Dynamics	531
<i>Noam Nisan, Michael Schapira, and Aviv Zohar</i>	
Fault Tolerance in Distributed Mechanism Design	539
<i>Ronen Gradwohl</i>	
Bargaining Solutions in a Social Network	548
<i>Tanmoy Chakraborty and Michael Kearns</i>	

Session C.4: Online Advertisement

Sharing Online Advertising Revenue with Consumers	556
<i>Yiling Chen, Arpita Ghosh, Randolph Preston McAfee, and David Pennock</i>	

Budget Constrained Bidding in Keyword Auctions and Online Knapsack Problems 566
Yunhong Zhou, Deeparnab Chakrabarty, and Rajan Lukose

Position Auctions with Bidder-Specific Minimum Prices 577
Eyal Even-Dar, Jon Feldman, Yishay Mansour, and S. Muthukrishnan

Session A.5: Sponsored Search Auctions

A Cascade Model for Externalities in Sponsored Search 585
David Kempe and Mohammad Mahdian

Sponsored Search Auctions with Reserve Prices: Going Beyond Separability 597
Rica Gonen and Sergei Vassilvitskii

Auctions for Share-Averse Bidders 609
Mahyar Salek and David Kempe

Sponsored Search Auctions with Markovian Users 621
Gagan Aggarwal, Jon Feldman, S. Muthukrishnan, and Martin Pál

On the Equilibria and Efficiency of the GSP Mechanism in Keyword Auctions with Externalities 629
Ioannis Giotis and Anna R. Karlin

Session B.5: Voting Problem

Biased Voting and the Democratic Primary Problem 639
Michael Kearns and Jinsong Tan

Frequent Manipulability of Elections: The Case of Two Voters 653
Shahar Dobzinski and Ariel D. Procaccia

The Power of Small Coalitions in Cost Sharing 665
Florian Schoppmann

Social Context Games 675
Itai Ashlagi, Piotr Krysta, and Moshe Tennenholtz

Session C.5: Algorithms and Optimization II

Approximability and Parameterized Complexity of Minmax Values 684
Kristoffer Arnsfelt Hansen, Thomas Dueholm Hansen, Peter Bro Miltersen, and Troels Bjerre Sørensen

An “Ethical” Game-Theoretic Solution Concept for Two-Player Perfect-Information Games	696
<i>Joshua Letchford, Vincent Conitzer, and Kamal Jain</i>	
The Secretary Problem with a Hazard Rate Condition	708
<i>Mohammad Mahdian, Randolph Preston McAfee, and David Pennock</i>	
Impact of QoS on Internet User Welfare	716
<i>Galina Schwartz, Nikhil Shetty, and Jean Walrand</i>	
Nonlinear Pricing with Network Externalities	724
<i>Dawen Meng, Guoqiang Tian, and Lei Sun</i>	
Author Index	733

Mechanism Design Theory: How to Implement Social Goals

Eric Maskin

Institute for Advanced Study, Princeton
maskin@ias.edu

Abstract. The theory of mechanism design can be thought of as the engineering side of economic theory. One begins by identifying a social or economic goal. The theory then addresses the question of whether or not an appropriate institution or procedure (that is, a mechanism) could be designed to attain that goal.

Thirty Years of Chinese Economic Reform: Reasons for Its Success and Future Directions

Lawrence J. Lau^{1,2}

¹ President and Ralph and Claire Landau Professor of Economics
The Chinese University of Hong Kong

² Kwoh-Ting Li Professor in Economic Development, Emeritus, Stanford University
lawrence1au@cuhk.edu.hk

Abstract. What are the principal reasons for the highly successful Chinese economic reform that began in 1978? One may say that they are the strong Chinese economic fundamentals-surplus labor, abundant savings, huge domestic market, etc. But the strong fundamentals have always been there, at least since the 1950s. Why did the Chinese economy not take off earlier?

The introduction of the market system, first in the rural area, and then in the urban area, must be regarded as the primary reason for the success of the economic reform. But the former Soviet Union and subsequently Russia also introduced the market system, with disastrous economic results for the entire first decade. Why was China able to do it while others failed?

Three important reasons can be identified: First, Chinese economic reform is characterized by openness-China welcomed international trade with and direct investment from all countries and regions, including Hong Kong, Taiwan, and the United States, and with trade and direct investment came technology, business models, and ideas that were new to China. Second, the Chinese economic reformers are characterized by their pragmatism-they are willing to try almost anything-whatever works-but they will just as readily abandon whatever that proves not to work. Third, Chinese economic reform has been implemented in such a way that it is mostly Pareto-improving, that is, almost everyone is made better off by the economic reform and no one is made worse off, which maximizes support, minimizes opposition and preserves social harmony.

What are some future directions of reform? They should consist of various ways to perfect the market mechanism in China. First, China has reached a stage of development that it needs to make and keep the markets truly competitive, through anti-monopoly laws and other means-and this applies to the both the goods market and the factors (including capital) market. When markets are not competitive, they may result in outcomes that are worse than those under central planning. Second, the markets can also be made more competitive, and hence more efficient, if information asymmetry can be reduced or eliminated. Thus, the Chinese Government can set standards for goods and services and assure quality through government-mandated and operated testing agencies. Third,

markets frequently fail when there is moral hazard. The Chinese Government can reduce the incidence of moral hazard by limiting leverage and requiring bonding. Fourth, the Chinese Government can also make the market system more complete by establishing and maintaining socially desirable markets that do not arise naturally without government intervention, for example, a long-term market for bonds backed by qualified long-term owner-occupied residential mortgages. Finally, the market system is not equipped to redistribute, but redistribution is often necessary on grounds of fairness and social harmony. The Chinese Government should design an equitable tax system as well as undertake public investments in education, health care, environmental protection and mass transportation so that the benefits of the continuing economic reform can be shared by the majority of the people.

Average Distance, Diameter, and Clustering in Social Networks with Homophily*

Matthew O. Jackson

Department of Economics, Stanford University and the Santa Fe Institute

jacksonm@stanford.edu

<http://www.stanford.edu/~jacksonm/>

Abstract. I examine a random network model where nodes are categorized by type and linking probabilities can differ across types. I show that as homophily increases (so that the probability to link to other nodes of the same type increases and the probability of linking to nodes of some other types decreases) the average distance and diameter of the network are unchanged, while the average clustering in the network increases.

Keywords: Networks, Random Graphs, Homophily, Friendships, Social Networks, Diameter, Average Distance, Clustering, Segregation.

1 Introduction

Communication advances and the social networking via the Internet have made it much easier for individuals to locate others with similar backgrounds and tastes. This can affect the formation of social networks. How do such changes in the ability of individuals to locate other similar individuals affect social network structure? Answering this question requires having models of how homophily, the tendency of nodes to be linked to other nodes with similar characteristics, affects social network structure. Homophily is a well-studied and prevalent phenomenon that is observed across all sorts of applications and attributes including ethnicity, age, religion, gender, education level, profession, political affiliation, and other attributes (e.g., see Lazarsfeld and Merton (1954), Blau (1977), Blalock (1982), Marsden (1987, 1988), among others, or the survey by McPherson, Cook and Smith-Lovin (2001)). Despite the extensive empirical research on homophily, there is little that is known about how homophily changes a network's basic characteristics, such as the average distance between nodes, diameter, and clustering.

This paper examines the following questions. Given is a society of nodes that are partitioned into a number of different groups where nodes within a group are of the same “type” and nodes in different groups are of different types. A network formation process is examined that can embody various forms of homophily: the probability of links between pairs of nodes can depend on their respective

* Financial support from the NSF under grant SES-0647867 is gratefully acknowledged. I thank Ben Golub for helpful conversations.

types. Holding the degree distribution constant, how does such a network that is formed with substantial homophily compare to a network formed when types are ignored? One conjecture is that as homophily increases so that the probability of links among nodes of similar types increases and the probability of links across less similar types falls, the average distance and diameter of the network will increase since the density of links across different types of nodes will be falling. This conjecture turns out to be false. Even as the probability of links across types falls, the average distance and diameter are not changed even in some extreme cases where the relative probability a link between nodes of the same type is arbitrarily more likely than a link among nodes of different types, provided some non-vanishing fraction of a node's links are still formed to nodes of other types. In contrast, homophily can have a significant impact on clustering. It is shown that substantial homophily can lead to nontrivial clustering, while a process with the same expected degrees but no homophily exhibits no clustering.

2 A Model of Network Formation with General Forms of Homophily and Degree Sequences

A network $G = (N, g)$ is a graph that consists of a set $N = \{1, \dots, n\}$ of a finite number n of nodes along with a list of edges, g ^[1] which are the pairs of nodes that linked to each other.

Given that the network might not be connected, I follow Chung and Lu (2002) in defining average distance in the network to be the average across pairs of path-connected nodes. In particular, let $\ell_g(i, j)$ be the number of links in the shortest path connecting nodes i and j if there is such a path, and let $\ell_g(i, j)$ be infinity if there is no path between i and j in g ^[2]. Thus, the average distance in the network is defined as^[3]

$$AD(g) = \frac{\sum_{\{i,j\}:\ell_g(i,j)\neq\infty} \ell_g(i,j)}{|\{\{i,j\} : \ell_g(i,j) \neq \infty\}|}.$$

The diameter of the network is $diam(g) = \max_{\{i,j\}:\ell_g(i,j)\neq\infty} \ell_g(i,j)$.

For the network formation processes considered here, the largest component contains all but at most a vanishing fraction of nodes and so these definitions are effectively the same whether we defined them as above, or just work with the largest component of g which is either the whole network or almost all of it.

The clustering of a node i with degree of at least 2 is

$$CL_i(g) = \frac{|\{\{j,j'\} : i \neq j \neq j' \neq i; \{i,j\} \in g, \{i,j'\} \in g, \{j,j'\} \in g\}|}{|\{\{j,j'\} : i \neq j \neq j' \neq i; \{i,j\} \in g, \{i,j'\} \in g\}|}.$$

¹ Formally, $g \subset 2^N$ such that each element in g has cardinality 2.

² Standard definitions, such as path, are omitted. See Jackson (2008) for such definitions.

³ Self-loops are allowed here, and so under these definitions if there is a self-loop then a node is a distance of 1 away from itself. This is irrelevant to the results and simply for convenience. It is easily seen that the results are the same if self-loops are ignored or if self-distance is set to 0. If there are no links in the network, the AD expression is 0/0 which can be set to take any value.

The average clustering is the average of CL_i across nodes i that have degree at least 2⁴

2.1 A General Random Network Model with Homophily

The following model is a generalization of the random network model from Chung and Lu (2002) to allow nodes to be of different types and to allow heterogeneous probabilities of linking across different types.

A set of nodes $N = \{1, \dots, n\}$ is partitioned into K groups or types N_1, \dots, N_K . This partition captures the characteristics of the nodes, so that all nodes with the same characteristics are in the same group N_k . Depending on the application a type might embody ethnicity, gender, age, education, profession, etc. in a social setting, or might involve characteristics of a business in a market network, or might involve some physical characteristics of a node in a physical network.

Also given is a degree sequence $\{d_1, \dots, d_n\}$ which indicates the expected degree or number of connections of each node. Let

$$D = \sum_i d_i$$

and

$$\tilde{d} = \sum_i d_i^2 / D.$$

Note that if $d_i = d$ for all i , then $\tilde{d} = d$.

Let $D_k = \sum_{i \in N_k} d_i$ be the total degree of all nodes of type k .

A random network is formed according to the following process. For each pair of types k and k' there is a parameter $h_{kk'} \geq 0$. This parameter captures the relative proclivity of groups k and k' to link to each other. The parameters satisfy $\sum_{k'} D_{k'} h_{kk'} = D$ for each k . A link between nodes i in group k and j in group k' is formed with probability

$$h_{kk'} d_i d_j / D.$$

Conditions defined below ensure that this expression does not exceed 1.

In the case where $h_{kk} > h_{kk'}$ for all k and $k' \neq k$, then there is homophily, so that nodes are relatively more likely to form their links to their own types than to other types. If $h_{kk'} = 1$ for all k and k' then types are irrelevant and the model reduces to the usual Chung and Lu model. Otherwise, this allows for different patterns of linkings between different types. If $d_i = d$ for all i , then this is a generalization of Erdős-Renyi random graphs where links are type-dependent.⁵ More generally, the degree distribution could vary across nodes, and power-law networks are the special case where the frequency distribution of $\{d_1, \dots, d_n\}$

⁴ Set clustering to 0 if there are no such nodes.

⁵ Note, however, that this process allows for self-loops i may connect to i , although the probability of this for any node i vanishes as n grows provided d_i^2/D vanishes.

has a power distribution where the frequency of degree d is of the form $cd^{-\gamma}$ for some range of d .

An interesting case is where types have some social or spatial geography and type k can be represented as a vector $x_k \in \mathbb{R}^m$ for some m and then $h_{kk'}$ is decreasing in the distance between k and k' ; for example of the form $c - f(|x_k - x_{k'}|)$ where c is a constant and f is an increasing function. One can also consider some hierarchy among the k 's with the relative probabilities depending on the hierarchy (e.g., see Clauset, Moore and Newman (2008)). Another case of interest is where types have a given probability of forming links to their own type and a different probability of forming links all other types (e.g., see Copic, Jackson and Kirman (2005) and Currarini, Jackson and Pin (2007)).

2.2 Admissible Models

The main results consider a growing sequence of network formation models, and so all parameters are indexed by n , the number of nodes. The results use some restrictions on variation in expected degrees across nodes and a minimum bound on the proclivity to link across groups. A sequence of network formation processes is said to be *admissible* if the following conditions are satisfied.

First, there exists $h > 0$ such that $h_{kk'}(n) > h$ for all k and k' for all large enough n . This condition does *not* require that nodes of different types have a probability of linking that is bounded below, as a node's degree could be a fixed number independent of n . This lower bound simply implies that any given node spreads some of its links on types other than its own type. This still allows for extreme homophily, as it can still be that $h_{kk}(n) \rightarrow \infty$ and that the probability of links with own type is becoming infinitely more likely than links with some other types.

Second, the degree sequence satisfies the following:

- $\tilde{d}(n) \geq \log^a(n)$ for some $a > 1$ and $\log(\tilde{d}(n)) / \log(n) \rightarrow 0$
- there exists $c > 0$ such that $hc > 1$, and $M > 0$, such that $d_i \leq M\tilde{d}(n)$ for all i and n , and $d_i \geq c$ for all but $o(n)$ nodes,⁶

The first restriction is that the second-order average degree is growing with n , but more slowly than n . The second requires that no node have an expected degree that explodes relative to the average expected degree and that all but a vanishing fraction of nodes have a lower bound on expected degree that is larger than 1.

3 Diameter and Average Distance in the Model

Let $AD(n, \mathbf{d}(n), \mathbf{h}(n))$ and $diam(n, \mathbf{d}(n), \mathbf{h}(n))$ be the average distance and diameter, respectively, of a graph randomly drawn according to the process above

⁶ Here, h is as defined in the restrictions on proclivity to link across types. These conditions ensure that the degree sequence satisfies (i) and (ii) in Chung and Lu (2000). They also guarantee (iii) setting $U = N$ and noting that $\tilde{d}(n) \leq M^2 D(n)/n$.

with n nodes, degree sequence $\mathbf{d}(\mathbf{n}) = (d_1(n), \dots, d_n(n))$, and homophily parameters $\mathbf{h}(n) = (h_{kk'}(n))_{kk'}$. This average distance and diameter are random variables for each n . Similarly, let $AD(n, \mathbf{d}(n))$ and $diam(n, \mathbf{d}(n))$ be the average distance and diameter, respectively, of a graph randomly drawn according to the process above with n nodes, degree sequence $\mathbf{d}(\mathbf{n}) = (d_1(n), \dots, d_n(n))$, and without any homophily (so that $h_{kk'}(n) = 1$ for all k and k').

Theorem 1. *Consider an admissible sequence of network formation processes $(n, \mathbf{d}(n), \mathbf{h}(n))$. Asymptotically almost surely*

- $AD(n, \mathbf{d}(n), \mathbf{h}(n)) = (1 + o(1)) \log(n) / \log(\tilde{d}(n))$, and so $\frac{AD(n, \mathbf{d}(n), \mathbf{h}(n))}{AD(n, \mathbf{d}(n))} \rightarrow 1$,
- $diam(n, \mathbf{d}(n), \mathbf{h}(n)) = \Theta\left(\log(n) / \log(\tilde{d}(n))\right)$ and so $diam(n, \mathbf{d}(n), \mathbf{h}(n)) = \Theta(diam(n, \mathbf{d}(n)))$.

Thus, the average distance and diameter of the admissible processes are not affected by homophily. Even though there can be an arbitrarily increased density of links within types, and substantial decrease in the density of links across types, this does not impact average distance or the diameter in the network. In order for homophily to affect these aspects of the network, one would have to have the density of links across most types decrease at a level which vanishes relative to overall degree. That is, suppose instead that nodes are grouped into evenly sized groups (up to integer constraints) so that $h_{kk'}(n) \leq f(n)$ for all k and k' with $k' \neq k$ for some $f(n)$ such that $f(n)n\tilde{d}(n)/K(n)$ is bounded above and where $K(n)/n$ is bounded away from 0. Then, it is easy to check that \square almost surely, $\frac{AD(n, \mathbf{d}(n), \mathbf{h}(n))}{AD(n, \mathbf{d}(n))} \rightarrow \infty$ and so $\frac{diam(n, \mathbf{d}(n), \mathbf{h}(n))}{diam(n, \mathbf{d}(n))} \rightarrow \infty$.

Proof of Theorem \square : Consider a network formation process such that each node has expected degree hd_i and $h_{kk'} = 1$ for all kk' . This is the process $(n, h\mathbf{d}(n))$, and the process $(n, \mathbf{h}(n), \mathbf{d}(n))$ is equivalent to a first running the process $(n, h\mathbf{d}(n))$ and then adding some additional links. Under the admissibility requirement here, $(n, h\mathbf{d}(n))$ is admissible and specially admissible under the definitions of Chung and Lu (2002). By Lemma 5 in Chung and Lu (2002), almost surely the largest component of a random graph under the process $(n, h\mathbf{d}(n))$ contains all but at most $o(n)$ of the nodes. By Theorems 1 and 2 in Chung and Lu (2002) the average distance and diameter of this process are almost surely

$$(1 + o(1)) \log(n) / \log(h\tilde{d}(n)) = (1 + o(1)) \log(n) / \log(\tilde{d}(n)),$$

⁷ A lower bound on the average distance is that of a graph where all nodes of a given type are agglomerated to become a single node. There are $K(n)$ nodes in this graph and each of these type-nodes has degree of at most $\tilde{d}Mf(n)n/K(n)$ which is bounded above by some C . The average distance is at least order $\log(K(n))/\log(C)$ which is proportional to $\log(n)$, provided this network has a giant component containing all but at most a vanishing fraction of nodes. The average distance could only be smaller than this if the connectivity across types drops so low so that the network fragments to smaller components.

and

$$\Theta\left(\log(n)/\log(h\tilde{d}(n))\right) = \Theta\left(\log(n)/\log(\tilde{d}(n))\right),$$

respectively. Since the process $(n, \mathbf{h}(n), \mathbf{d}(n))$ is equivalent to a first running the process $(n, h\mathbf{d}(n))$ and then adding some additional links, it then follows directly that a random graph generated in this way contains all but at most $o(n)$ of the nodes and has average distance and diameter of this process are almost surely bounded above by $(1 + o(1))\log(n)/\log(\tilde{d}(n))$, and some factor times $\log(n)/\log(\tilde{d}(n))$, respectively.

Next, let us show that these are also lower bounds. Consider any network where all nodes have degree no more than $M\tilde{d}(n)$. Consider any node i . The T -the neighborhood of i includes fewer than

$$\sum_{t=1}^T \left(M\tilde{d}(n)\right)^t = \frac{\left(M\tilde{d}(n)\right)^{t+1} - M\tilde{d}(n)}{M\tilde{d}(n) - 1}$$

nodes. Thus, in order to reach all nodes in the largest component from some node in the largest component (which as argued above contains at least $(1 - o(n))n$ nodes) it takes at least $T(n) = \log((1 - o(1))n)/\log\left(M\tilde{d}(n)\right)$ steps to reach every other node in the largest component, almost surely. Given that $\tilde{d}(n) \rightarrow \infty$, it follows that $T(n) \geq (1 - o(1))\log(n)/\log\left(\tilde{d}(n)\right)$. The average distance is thus almost surely at least

$$\sum_{t=1}^{T(n)} \left(M\tilde{d}(n)\right)^t t/n.$$

This is at least $(1 - o(1))T(n)$, almost surely. Thus, the lower bound on average distance is $(1 - o(1))\log(n)/\log\left(\tilde{d}(n)\right)$. The diameter is at least the average distance, and so this is also a lower bound on diameter. The result follows by bounding the realized degrees on nodes asymptotically almost surely. ■

4 Clustering

Note that in the model with no homophily if $(\max_i d_i(n))^2/D(n) \rightarrow 0$, then the average clustering almost surely tends to 0 simply because the most probable link has a probability that tends to 0. In contrast, if groups are relatively small (of the order of average degree) and there is substantial homophily, then average clustering does not vanish. Thus, homophilistic networks exhibit the characteristics of the “small worlds” discussed by Watts and Strogatz (1998): nontrivial clustering at the same time as having a diameter on the order of a uniformly random graph.

Theorem 2. *Consider a setting such that (i) there is some $m > 0$ such that for large enough n , $h_{kk}(n)D_k(n)/D(n) > m$ for all k , (ii) $\max_i d_i(n)/\max_k |N_k|$ and $\min_i d_i(n)/\max_i d_i(n)$ are each $\Omega(1)$, and $\max_i d_i(n) > 2$. Almost surely, average clustering is $\Omega(1)$.*

The proof of Theorem 2 is straightforward and so only sketched here. Let $\max_i d_i(n)/\max_k |N_k| > m_1 > 0$ and $\min_i d_i(n)/\max_i d_i(n) > m_2 > 0$ for all large enough n . The probability of a link between any two nodes of the same type is at least

$$\begin{aligned} \frac{(m_2 \max_i d_i(n))^2 \min_k h_{kk}}{D} &> \frac{(m_2 \max_i d_i(n))^2 m}{\max_k D_k(n)} \\ &> \frac{(m_2 \max_i d_i(n))^2 m}{\max_k |N_k(n)| \max_i d_i(n)} \\ &> m_2^2 m_1 m > 0 \end{aligned}$$

for all large enough n . Given that there is a bound $m_3 > 0$ so that each node has an expectation of forming a fraction of at least m_3 of its links within its own group, and the clustering among pairs of nodes that it is linked to of own type is at least $m_2^2 m_1 m > 0$, it follows that the expected clustering of any node is bounded away from 0 (conditional on it having degree at least 2). Given that the expected clustering of all nodes are bounded away from 0 (conditional on having at least degree 2), and all nodes have expected degree bounded away from 0 and so a non-vanishing fraction almost surely end up with degree of at least 2, it can then be shown that the average clustering is almost surely above 0.

5 Discussion

The results here show that substantial homophily and bias in the way that different types of nodes link to each other can be introduced without altering the average distance or diameter of a network. On one level this might not have been expected, and yet the proof of this is very simple and basically relies on the fact that some rescaling of the degree of a node up to a fixed factor does not alter the asymptotic average distance and diameter of the resulting networks. This does not mean that this leaves the properties of the network unchanged, as we have seen with clustering parameters. Also, as shown in Golub and Jackson (2008), networks with substantial homophily can still behave quite differently, so that even though diameter and average distance remain unchanged, the speed of learning can decrease by orders of magnitude and mixing time on such networks can correspondingly increase by orders of magnitude.

References

- [1] Blalock Jr., H.M.: Race and Ethnic Relations. Prentice-Hall, Englewood Cliffs (1982)
- [2] Blau, P.M.: Inequality and Heterogeneity: A Primitive Theory of Social Structure. Free Press, New York (1977)
- [3] Clauset, A., Moore, C., Newman, M.E.J.: Hierarchical Structure and the Prediction of Missing Links in Networks. *Nature* 453(1), 98–101 (2008)

- [4] Chung, F., Lu, L.: The Average Distances in Random Graphs with Given Expected Degrees. *Proceedings of the National Academy of Sciences* 99(25), 15879–15882 (2002)
- [5] Copic, J., Jackson, M.O., Kirman, A.: Identifying Community Structures from Network Data via Maximum Likelihood Methods. Stanford University (preprint, 2005)
- [6] Currarini, S., Jackson, M.O., Pin, P.: An Economic Model of Friendship: Homophily, Minorities and Segregation. *Econometrica* (forthcoming, 2007)
- [7] Golub, B., Jackson, M.O.: How Homophily Affects Communication in Networks. Stanford University (preprint, 2008)
- [8] Jackson, M.O.: *Social and Economic Networks*. Princeton University Press, Princeton (2008)
- [9] Lazarsfeld, P.F., Merton, R.K.: Friendship as a social process: a substantive and methodological analysis. In: Berger, M. (ed.) *Freedom and Control in Modern Society*, pp. 18–66. Van Nostrand, New York (1954)
- [10] Marsden, P.V.: Core Discussion Networks of Americans. *American Sociological Review* 52, 122–131 (1987)
- [11] Marsden, P.V.: Homogeneity in confiding relations. *Social Networks* 10, 57–76 (1988)
- [12] McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of a Feather: Homophily in Social Networks. *Annu. Rev. Sociol.* 27, 415–444 (2001)
- [13] Watts, D.J., Strogatz, S.: Collective dynamics of ‘small-world’ networks. *Nature* 393, 440–442 (1998)

Assignment Exchanges

Paul Milgrom

Stanford University
pmilgrom@gmail.com

Abstract. We analyze “assignment exchanges”- auction and exchange mechanisms which are tight simplifications of direct Walrasian exchange mechanisms. These simplifications are distinguished by their use of assignment messages, which parameterize certain substitutable preferences. The “basic” assignment exchanges respect integer constraints, generalizing the Shapley-Shubik mechanism for indivisible goods. Connections are reported between the assignment exchanges, ascending multi-product clock auctions, uniform price auctions for a single product, and Vickrey auctions. The exchange mechanisms accommodate bids by buyers, sellers and swappers and can support trading for certain kinds of complementary goods.

Search Engine Ad Auctions

Hal R. Varian

Google Inc.
University of California at Berkeley
hal@google.com

Abstract. Auctions for search engine advertising have been one of the most successful examples of economic mechanism design, at least in the private sector. This talk will review some of the history, theory, and practical issues surrounding these auctions.

Computational Economy Equilibrium and Application

Yinyu Ye

Stanford University
yinyu-ye@stanford.edu

Abstract. The rise of the Internet and the emerging E-Commerce applications has created new economic markets of unprecedented scale. They have introduced many cross-disciplinary challenges in mathematics and computer scientists, and engineering, one of which is the algorithmic and complexity issue of economy market equilibrium theory. In this talk, we examine the mathematical connections as well as the computational equivalences between equilibrium and optimization, between game equilibrium and market equilibrium, existence and NP-hardness, and between exact computation and approximation. Being able to compute equilibria numerically also significantly expands the applicability of game/economy equilibrium theory to a wide range of decision problems. We present applications of computational equilibrium from developing communication network protocols in spectrum management and resource allocation to adopting free trade policies in international trade between nations.

Four Graph Partitioning Algorithms

Fan Chung Graham

University of California, San Diego
fan@math.ucsd.edu

Abstract. We will discuss four partitioning algorithms using eigenvectors, random walks, PageRank and their variations. In particular, we will examine local partitioning algorithms, which find a cut near a specified starting vertex, with a running time that depends on the size of the small side of the cut, rather than on the size of the input graph (which can be prohibitively large). Three of the four partitioning algorithms are local algorithms and are particularly appropriate for applications arising in connection with Webgraphs and Internet economics.

Dynamic Spectrum Management: Optimization and Game Theoretic Formulations

Zhi-Quan (Tom) Luo

University of Minnesota
luozq@umn.edu

Abstract. Achieving efficient spectrum usage is a major challenge in the management of a complex communication system. With multiple users having conflicting objectives who share a common spectrum, some of whom may be hostile, careful resource allocation is essential for the effective utilization of the available frequency. Conventionally, spectrum sharing is achieved via orthogonal transmission schemes whereby the available frequency band is divided into multiple tones (or bands) which are pre-assigned to all the users on a non-overlapping basis. However, such “static orthogonal spectrum sharing” approach can lead to low bandwidth utilization. In fact, various recent spectrum occupancy studies have demonstrated that a typical geographical region has wide swathes of frequencies (up to 2/3 of the allocated radio spectrum) that are not used at any given time. While the utilization of spectrum varies with time, a significant amount of spectrum is available for opportunistic wireless applications among secondary users.

Spectrum-sensing cognitive radio technology allows devices to dynamically and automatically seek out and use the optimum frequencies and bandwidth. To take advantage of the unused spectrum capacity, the users dynamically adapt to the spectral environment and change transmission or reception parameters on the fly. This allows for more efficient wireless communication without causing harmful interference with legacy systems or other devices using the same frequency bands. In these systems all users are allowed to use all the tones simultaneously. In comparison with the static spectrum sharing policies, this setup offers significantly greater freedom in utilizing the spectrum. A major challenge in the development of opportunistic spectrum sharing technology is to devise efficient algorithms for the distributed management of frequency slots and transmit power.

This tutorial will describe various optimization and game theoretic formulations of the dynamic spectrum management and present some recent results on its complexity, duality and approximation.

Some Recent Results in Algorithmic Game Theory

Christos Papadimitriou

UC Berkeley

`christos@cs.berkeley.edu`

Abstract. There are three major trends in the field of Algorithmic Game Theory: computational mechanism design, the price of anarchy, and the computation of equilibria; this talk describes one recent result in each. We show computational complexity lower bounds on truthful and approximately efficient mechanisms; we revisit the Roughgarden-Tardos result on selfish routing when routing decisions are made by the nodes, not the flows; and we show that Nash equilibria can be approximated well in several broad, unexpected, and useful classes of games. (Joint work with Costis Daskalakis, Michael Schapira, Yaron Singer, and Greg Valiant).

The Elements of General Equilibrium Theory

Herbert E. Scarf

Yale University
herbert.scarf@yale.edu

Abstract. The lecture will be an introduction to the model of economic equilibrium. The basic concepts: preferences, initial endowments and market clearing prices will be discussed - in general and by means of examples. I will indicate how fixed point theorems are used to demonstrate the existence of equilibrium prices and sketch an algorithm for Brouwer's theorem. If time permits, there will be some remarks on equilibrium models with production.

A Fast and Simple Algorithm for Computing Market Equilibria

Lisa Fleischer¹, Rahul Garg², Sanjiv Kapoor³, Rohit Khandekar²,
and Amin Saberi⁴

¹ Dartmouth College

lkf@cs.dartmouth.edu

² IBM T.J. Watson Research Center

{grahul,rohitk}@us.ibm.com

³ Illinois Institute of Technology

kapoor@iit.edu

⁴ Stanford University

saberi@stanford.edu

Abstract. We give a new mathematical formulation of market equilibria using an *indirect utility function*: the function of prices and income that gives the maximum utility achievable. The formulation is a *convex program* and can be solved when the indirect utility function is convex in prices. We illustrate that many economies including

- Homogeneous utilities of degree $\alpha \in [0, 1]$ in Fisher economies — this includes Linear, Leontief, Cobb-Douglas
- *Resource allocation utilities* like multi-commodity flows

satisfy this condition and can be efficiently solved.

Further, we give a natural and decentralized price-adjusting algorithm in these economies. Our algorithm, mimics the natural tâtonnement dynamics for the markets as suggested by Walras: it iteratively adjusts a good's price upward when the demand for that good under current prices exceeds its supply; and downward when its supply exceeds its demand. The algorithm computes an approximate equilibrium in a number of iterations that is independent of the number of traders and is almost *linear* in the number of goods. Interestingly, our algorithm applies to certain classes of utility functions that are not *weak gross substitutes*.

1 Introduction

The market equilibrium model, common in economics, is that of a market with m traders and n goods, where the traders are endowed with money or/and goods and wish to optimize their utilities. Market equilibrium is defined by a price and an allocation such that no trader has any incentive to trade and there is no excess demand of any good. While the problem was originally formulated by Walras [29] in 1874, the existence of such an equilibrium was established by Arrow and Debreu [1] in 1954 using a fixed-point argument.

The result of Arrow and Debreu does not give much insight into the dynamics of the market. How does market find the equilibrium prices? What is

the complexity of finding these prices? Interested in answering the first question, economists focused on decentralized dynamics that converge to equilibrium. Most notably, Samuelson [26] formalized Walras' idea of tâtonnement as a set of differential equations relating the adjustment of the price with excess demand. Later, Arrow et al. and Nikaido and Uzawa [2,23] showed that in markets with gross substitute property, the process proposed by Samuelson converges to an equilibrium. The number of iterations of such a process depends on the utility functions of the traders.

In computer science literature, the focus has been on designing polynomial-time algorithms for several special cases using techniques such as primal-dual, auctions algorithms and convex programming [9,7,17,13,15,10,30]. The surveys of Vazirani [28] and Codenotti and Varadarajan [4] discuss these results. These algorithms (with the notable exception of [8]¹) are typically centralized.

This paper attempts to combine the advantages of the both approaches for a restricted class of markets. We present a fast and relatively natural algorithm for computing approximate equilibrium prices. The number of iterations required by our algorithm to converge to approximate equilibrium prices is almost linear in the number of goods and is independent of the number of traders. Another desirable feature of our algorithm is its distributed nature: it does not need to gather the information on utility functions and endowments of the traders in a central place to compute the prices. It only offers the sellers a procedure for updating the prices based on the difference of demand and supply of their good that converges to market equilibria. In fact, except a normalization variable, the only information passed between buyer and seller of a good is the current price of the goods and the demand corresponding to the current price.

From an algorithm design perspective, our procedure is different from primal-dual or auction algorithms in the sense that the prices (dual variables) do not approach the equilibrium from below. The process may underestimate or overshoot the equilibrium prices several times before it converges. In that sense, our algorithm is closest to the results of [12,24]. The analysis uses a new convex program for characterizing equilibria. For that reason the class of markets for which we can analyze our procedure is slightly more restricted than the class of markets comprising weakly gross substitute goods. At the same time, it include resource allocation markets, which are in fact not gross-substitute markets.

In particular, our algorithm applies to the market model for network congestion control as a part of a larger class of resource allocation markets [18]. For the case of multiple sources and sinks, the problem of determining, or discovering, equilibrium prices using a tâtonnement or combinatorial process, appears rather challenging, especially since there are no known combinatorial polynomial time algorithms for solving the feasibility of multi-commodity flows in networks. Fortunately, approximate solutions are tractable as we illustrate in this paper.

¹ The result of [8] has a running time that is independent of both number of traders and number of goods, but is dependent on some other market parameters. For example, when all traders share linear utilities, the procedure in [8] may not converge.

1.1 Results

The new convex program. We give a new formulation of the market equilibrium problem using *indirect utility function*. An indirect utility function \tilde{u} of price $\pi \in \mathfrak{R}_+^n$ and budget (or income) $e \in \mathfrak{R}_+$ gives the maximum utility achievable under those prices and budget as follows:

$$\tilde{u}(\pi, e) = \max\{u(x) \mid x \in \mathfrak{R}_+^n, \pi \cdot x \leq e\}$$

where u is the utility function defined on allocation of goods. Although indirect utility functions have been extensively used in Economics to study the behavior of aggregate demand [20,27], here we use them to formulate and solve the market equilibrium problem. Our formulation becomes a convex program if the indirect utility functions are convex on a suitably defined set of prices and income. This enables polynomial-time computation of (approximate) market equilibrium using standard convex programming techniques.

We show that, in the Fisher setting, the indirect utility functions are convex if the utility functions are homogeneous of degree 1. Such utility functions include linear, Leontief, Cobb-Douglas, CES, resource allocation markets. If the utility function u is increasing in all its components, then a necessary and sufficient condition for convexity of the corresponding indirect utility function is (see Proposition 2.4 in [25]): $-\frac{x \cdot \partial^2 u(x) x}{\partial u(x) x} \leq 2$ for all x . Surprisingly, this condition has the same form as those for monotone utilities [7]. They turn out to be a special case of monotone utilities for which market equilibrium can be computed using ellipsoid method [7]. However, note that polynomial time convergent tâtonnement processes are not known for monotone utilities.

The algorithm. A natural approach to computing the equilibrium price (as originally envisaged by Walras) is an iterative algorithm termed as *tâtonnement* process where the prices of goods are updated locally as a function of excess demand. Stability of these processes have been studied extensively in the literature [2,21] (see [22] for a survey). It has been shown that if the utility functions satisfy the *weak gross substitute* (WGS) property then the continuous process is stable and converges to market equilibrium. Polynomial-time convergence of such a process was only recently established in exchange economies with WGS utilities by the works of Codenotti et al. [7].

Our formulation enables us to design efficient processes similar to tâtonnement that converge close to a market equilibrium in polynomial time whenever the indirect utility functions of traders are convex. This partially answers the question raised in [18,19,7] on convergence of tâtonnement processes for a class of utility functions that do not satisfy WGS, for example, Leontief and resource allocation utilities. In order to obtain a $(1 + \epsilon)$ (weak) approximate market equilibrium, our process requires every trader to perform at most $O(\epsilon^{-2} n \log n)$ computations of its demand. For multi-commodity flow resource allocation market, for example, the demand oracle is the shortest-path computation under the given edge-lengths (prices). Thus our algorithm needs $\tilde{O}(kn)$ shortest path computations for a market with k commodities and n edges. This contrasts against the

algorithm of [18] for single-source multi-sink markets that needs $O(k^2)$ max-flow computations. We point out, however, that the algorithm of [18] computes an exact equilibrium while we compute only an approximate equilibrium.

Organization. The rest of the paper is organized as follows. In Section 2, we define the market equilibrium problem and formulate a mathematical program using indirect utility functions. We also outline a convex programming technique for solving this formulation if the indirect utility functions are convex. Section 3 then presents the prominent cases where we consider several utilities in Fisher economy under which the indirect utility functions turn out to be convex. In Section 4, we present our algorithm for computing approximate market equilibria assuming convexity of indirect utility functions. Section 5 concludes with some open directions.

2 An Alternate Formulation Using Indirect Utility Functions

We first describe the exchange market model. Let us consider m economic agents who represent traders of n goods. Let \mathfrak{R}_+^n (resp. \mathfrak{R}_{++}^n) denote the subset of \mathfrak{R}^n where the coordinates are non-negative (resp. strictly positive). The j th coordinate will stand for good j . Each trader i ($i = 1, \dots, m$) is associated with

- a non-empty convex set $\mathcal{K}_i \subseteq \mathfrak{R}^n$ which is the set of all “feasible” allocations that trader i may receive (in many cases, $\mathcal{K}_i = \mathfrak{R}_+^n$),
- a *concave* utility function $u_i : \mathcal{K}_i \rightarrow \mathfrak{R}_+$ which represents her preferences for the different bundles of goods, and
- an initial endowment of goods $w_i = (w_{i1}, \dots, w_{in})^\top \in \mathcal{K}_i$.

At given prices $\pi \in \mathfrak{R}_+^n$, the trader i sells her endowment, and gets the bundle of goods $x_i = (x_{i1}, \dots, x_{in})^\top \in \mathcal{K}_i$ which maximizes $u_i(x)$ subject to budget constraint² $\pi \cdot x \leq \pi \cdot w_i$. A market equilibrium is a price vector $\pi \in \mathfrak{R}_+^n$ and bundles $x_i \in \mathcal{K}_i$ such that: $x_i \in \operatorname{argmax}\{u_i(x) \mid x \in \mathcal{K}_i, \pi \cdot x \leq \pi \cdot w_i\}$ for all i , and $\sum_i x_i \leq \sum_i w_i$. The above described market model is called an *exchange economy*.

We make the following standard assumption on the utility functions:

Assumption 1. For $\pi \in \mathfrak{R}_+^n$, any $x_i \in \operatorname{argmax}\{u_i(x) \mid x \in \mathcal{K}_i, \pi \cdot x \leq \pi \cdot w_i\}$ satisfies $\pi \cdot x_i = \pi \cdot w_i$.

We now define a notion of indirect utility function induced by a utility function.

Definition 2 (Indirect utility function). For trader i , the indirect utility function $\tilde{u}_i : \mathfrak{R}_+^n \times \mathfrak{R}_+ \rightarrow \mathfrak{R}_+$ gives the maximum utility achievable at given price and income:

$$\tilde{u}_i(\pi, e) = \max\{u_i(x) \mid x \in \mathcal{K}_i, \pi \cdot x \leq e\}.$$

The following theorem characterizes the set of all equilibria.

² For two vectors x and y , we use $x \cdot y$ to denote their inner product.

Theorem 3. *The following program gives precisely the set of all market equilibria in the exchange economy.*

$$\begin{aligned} \sum_i x_i &\leq \sum_i w_i \\ \tilde{u}_i(\pi, \pi \cdot w_i) &\leq u(x_i) \quad \text{for all } i \\ \pi &\in \mathfrak{R}_+^n \\ x_i &\in \mathcal{K}_i \quad \text{for all } i. \end{aligned} \tag{1}$$

Proof. From the definition, it follows that a market equilibrium satisfies the above inequalities. Now for converse, consider a solution (π, x_1, \dots, x_m) of the above program. From the second constraint and Assumption [1](#), it follows that $\pi \cdot x_i \geq \pi \cdot w_i$ for all i . Furthermore from the first constraint, it follows that $\sum_i \pi \cdot x_i \leq \sum_i \pi \cdot w_i$. This implies that $\pi \cdot x_i = \pi \cdot w_i$ for all i and hence the solution (π, x_1, \dots, x_m) is a market equilibrium.

Note that the program [\(1\)](#) is convex when, for all i , the function $\tilde{u}_i(\pi, \pi \cdot w_i)$ is a convex function of $\pi \in \mathfrak{R}_+^n$ and the utility function u_i is concave. Unfortunately, for many interesting utility functions u_i , the corresponding indirect utility function \tilde{u}_i is *not* convex. It turns out, however, that in *many* cases (as illustrated later in the paper), if we restrict the prices π to a carefully chosen convex set $\Pi \subseteq \mathfrak{R}_+^n$ that is guaranteed to contain an equilibrium price, the function \tilde{u}_i becomes convex in π . Therefore the program [\(1\)](#) reduces to the following convex program.

$$\begin{aligned} \sum_i x_i &\leq \sum_i w_i \\ \tilde{u}_i(\pi, \pi \cdot w_i) &\leq u(x_i) \quad \text{for all } i \\ \pi &\in \Pi \\ x_i &\in \mathcal{K}_i \quad \text{for all } i. \end{aligned} \tag{2}$$

In order to solve the above convex program using an ellipsoid algorithm, the convex set Π needs to be given in terms of a membership oracle.

Solving Program [\(2\)](#). Assuming that the convex sets Π and \mathcal{K}_i are bounded and full dimensional^{[3](#)} the convex program [\(2\)](#) can be solved to an arbitrary degree of precision by an ellipsoid-like algorithm using the evaluation oracle for the functions u_i and \tilde{u}_i and membership oracles for Π and \mathcal{K}_i . We omit details here and refer the reader to Theorem 4.3.13 in [\[16\]](#).

3 Convexity of the Indirect Utility Functions

In this section, we give a class of Fisher economies in which the indirect utility function \tilde{u}_i is convex in π over a set Π . The Fisher economy is a special case of the exchange economy when $\mathcal{K}_i = \mathfrak{R}_+^n$ and the endowments w_i of the traders are *proportional*, i.e.,

$$w_i = \lambda_i w$$

³ The economies considered in this paper have unbounded Π and \mathcal{K}_i in their description. However one can usually obtain bounds on the largest value that an allocation or a price can take. Moreover the cases that Π is not full dimensional can be handled using standard projection techniques.

where $w \in \mathfrak{R}_{++}^n$ and $\lambda_1, \dots, \lambda_m \in \mathfrak{R}_{++}$. In this case we let $\Pi = \{\pi \in \mathfrak{R}_+^n \mid \pi \cdot w = 1\}$. Thus under any prices $\pi \in \Pi$, the income of trader i is fixed at λ_i .

We now quote a theorem of K.-H. Quah [25] which gives *necessary* and *sufficient* conditions on the utility functions u_i under which the indirect utility functions $\tilde{u}_i(\pi, \lambda_i)$ are convex in $\pi \in \Pi$. We drop the subscript i to simplify the notation.

Proposition 1 (K.-H. Quah [25], Proposition 2.4). *Assume that the utility function $u : \mathfrak{R}_{++}^n \rightarrow \mathfrak{R}$ is continuous, quasi-concave, increasing in all arguments, and has the property that for any $\bar{x} \in \mathfrak{R}_{++}^n$, the set $\{x \in \mathfrak{R}_{++}^n \mid u(x) \geq u(\bar{x})\}$ is closed. Let $\lambda \in \mathfrak{R}_{++}$ be a constant.*

1. *Then, $\tilde{u}(\pi, \lambda)$ is convex in prices π if and only if the functions μ_x are convex for all x , where $\mu_x : \mathfrak{R}_{++} \rightarrow \mathfrak{R}$ is defined by $\mu_x(s) = u(x/s)$.*
2. *Suppose, in addition, that u is \mathcal{C}^2 , a twice differentiable function. Then μ_x is convex if and only if $\psi(x) = -\frac{x \cdot \partial^2 u(x) x}{\partial u(x) x} \leq 2$ for all x .*

Remark 1. Contrast the condition $\psi(x) \leq 2$ above with the condition $\psi(x) < 4$ which is sufficient to guarantee that the induced demand function is *monotone* [7]. Recall that the demand function $x(\pi)$ is monotone if for any π, π' , we have $(\pi - \pi') \cdot (x(\pi) - x(\pi')) \leq 0$. Thus if \tilde{u} is convex, the induced demand function is monotone.

Corollary 4

1. *A concave homogeneous utility function u of degree α where $0 \leq \alpha \leq 1$, i.e., $u(sx) = s^\alpha u(x)$, results in convex indirect utility function \tilde{u} if u satisfies the conditions in Proposition [1].*
2. *If utility functions u_1 and u_2 satisfy the conditions in Proposition [1] and induce convex indirect utility functions, then so does $u_1 + u_2$.*

Proof. For (1), note that $\mu_x(s) = s^{-\alpha} u(x)$ is a convex function of s . For (2), note that if $\mu_{1,x}$ and $\mu_{2,x}$ are convex functions then so is $\mu_{1,x} + \mu_{2,x}$.

Note, however, that some natural homogeneous utility functions of degree one (e.g., Leontief utilities and resource allocation utilities, defined later) do not satisfy the conditions in Proposition [1], in particular, the condition that the utility function is increasing in all arguments. However in the next theorem we show that the homogeneous utilities induce a convex indirect utility function even when they are not increasing in all arguments.

Theorem 5. *If the utility function $u : \mathfrak{R}_+^n \rightarrow \mathfrak{R}$ is homogeneous (of degree one), i.e., $u(\alpha x) = \alpha u(x)$ for all $\alpha \in \mathfrak{R}_+$ and $x \in \mathfrak{R}_+^n$, then the indirect utility function $\tilde{u}(\pi, \lambda)$ is convex in π for all $\lambda \in \mathfrak{R}_{++}$.*

Proof. Let price vectors $\pi, \pi_1, \pi_2 \in \mathfrak{R}_+^n$ satisfy $\pi = \alpha \pi_1 + (1 - \alpha) \pi_2$ for some $0 \leq \alpha \leq 1$. Let $x \in \mathfrak{R}_+^n$ be such that $\pi \cdot x = \lambda$ and $u(x) = \tilde{u}(\pi, \lambda)$. Define $x_1 = \frac{\lambda x}{\pi_1 \cdot x}$ and $x_2 = \frac{\lambda x}{\pi_2 \cdot x}$. Note that $\pi_1 \cdot x_1 = \pi_2 \cdot x_2 = \lambda$ and hence $\tilde{u}(\pi_1, \lambda) \geq u(x_1)$ and $\tilde{u}(\pi_2, \lambda) \geq u(x_2)$. Using the homogeneity of u , we also get that $u(x) = \frac{\pi_1 \cdot x}{\lambda} u(x_1) \leq \frac{\pi_1 \cdot x}{\lambda} \tilde{u}(\pi_1, \lambda)$ and $u(x) = \frac{\pi_2 \cdot x}{\lambda} u(x_2) \leq \frac{\pi_2 \cdot x}{\lambda} \tilde{u}(\pi_2, \lambda)$.

Note that $\alpha(\pi_1 \cdot x) + (1 - \alpha)(\pi_2 \cdot x) = \lambda$. Now

$$\begin{aligned} \left(\frac{\alpha\lambda}{\pi_1 \cdot x} + \frac{\lambda(1 - \alpha)}{\pi_2 \cdot x} \right) &= \left(\frac{\alpha\lambda}{\pi_1 \cdot x} + \frac{\lambda(1 - \alpha)}{\pi_2 \cdot x} \right) \left(\frac{\alpha(\pi_1 \cdot x)}{\lambda} + \frac{(1 - \alpha)(\pi_2 \cdot x)}{\lambda} \right) \\ &= \alpha^2 + \alpha(1 - \alpha) \left(\frac{\pi_1 \cdot x}{\pi_2 \cdot x} + \frac{\pi_2 \cdot x}{\pi_1 \cdot x} \right) + (1 - \alpha)^2 \\ &\geq \alpha^2 + 2\alpha(1 - \alpha) + (1 - \alpha)^2 \\ &= 1. \end{aligned}$$

To complete the proof we now observe

$$\begin{aligned} \tilde{u}(\pi, \lambda) = u(x) &\leq u(x) \left(\frac{\alpha\lambda}{\pi_1 \cdot x} + \frac{\lambda(1 - \alpha)}{\pi_2 \cdot x} \right) \\ &\leq \left(\frac{\pi_1 \cdot x}{\lambda} \tilde{u}(\pi_1, \lambda) \right) \frac{\alpha\lambda}{\pi_1 \cdot x} + \left(\frac{\pi_2 \cdot x}{\lambda} \tilde{u}(\pi_2, \lambda) \right) \frac{\lambda(1 - \alpha)}{\pi_2 \cdot x} \\ &= \alpha \tilde{u}(\pi_1, \lambda) + (1 - \alpha) \tilde{u}(\pi_2, \lambda). \end{aligned}$$

The set of **homogeneous utility functions** of degree one includes the following well-studied utility functions. Here let $a \in \mathfrak{R}_+^n$. Linear utilities $u(x) = a \cdot x$, Leontief utilities $u(x) = \min_{j \in S} a_j x_j$ where $S \subseteq \{1, \dots, n\}$, Cobb-Douglas utilities $u(x) = \prod_j x_j^{a_j}$ assuming $\sum_j a_j = 1$, CES utilities $u(x) = (\sum_j a_j x_j^\rho)^{1/\rho}$ for $-\infty < \rho < 1$ and $\rho \neq 0$, and nested CES utilities [5] [6].

It also includes the **resource allocation utilities** defined as follows. Let k be a positive integer and let $A \in \mathfrak{R}_+^{n \times k}$ be a matrix and $c \in \mathfrak{R}_+^k$ be a vector. The resource allocation utility $u : \mathfrak{R}_+^n \rightarrow \mathfrak{R}$ is defined as

$$u(x) = \max\{c \cdot y \mid y \in \mathfrak{R}_+^k, Ay \leq x\}. \quad (3)$$

The columns of matrix A can be thought of as “objects” that the trader wants to “build”. A unit of an object l needs A_{jl} units of resource (or good) j and accrues c_l units of utility. The trader then builds y_l units of object l such that the total need for resources is at most x and the total utility $c \cdot y$ is maximized. This framework includes interesting markets like

1. Multi-commodity flow markets (in directed or undirected capacitated networks). Here trader i wants to send maximum amount of flow from source s_i to sink t_i such that the total cost of routing the flow under the prices π is at most her budget. The objects here are s_i - t_i paths and the resources are the edges.
2. Steiner-tree markets in undirected (resp. directed) capacitated networks. Here trader i is associated with a subset S_i of nodes and wants to build maximum fractional packing of Steiner trees connecting S_i (resp. fractional arborescences rooted at some $r_i \in S_i$ connecting S_i to r_i) such that the total cost of building under the prices π is at most her budget. The objects here are Steiner trees (resp. arborescences). Note that computing a profit maximizing demand in undirected Steiner-tree market is NP-hard. Therefore the running times of the algorithms are only oracle-polynomial.

From Corollary [4](#) the **additive separable concave** utilities also induce a convex indirect utility functions: (1) $u(x_1, \dots, x_n) = \sum_j a_j x_j^{\rho_j}$ where $a_j \in \mathfrak{R}_{++}$ and $0 \leq \rho_j \leq 1$; (2) $u(x_1, \dots, x_n) = \sum_j \log(1 + a_j x_j)$ where $a_j \in \mathfrak{R}_{++}$ [3](#) — follows from the fact that $\log(1 + \frac{a_j x_j}{s})$ is a convex function of s .

4 The Algorithm

In this section, we present our algorithm to compute a *weak approximate market equilibrium* defined as follows. To simplify the definition, we assume that $\mathcal{K}_i \subseteq \mathfrak{R}_+^n$, i.e., we let x_{ij} take only non-negative values. For some technical reason, we assume that the set Π satisfies the following property: for any vector $p \in \mathfrak{R}_+^n$, there exists $\alpha \in \mathfrak{R}_{++}$ such that $\alpha p \in \Pi$. Note that this requirement is satisfied by the sets Π for the utilities in Fisher markets.

Definition 6 (Weak $(1 + \epsilon)$ -approximate market equilibrium). *A price vector $\pi \in \Pi$ and bundles $x_i \in \mathcal{K}_i$ for each trader i are called a weak $(1 + \epsilon)$ -approximate market equilibrium in the exchange economy if*

1. *The utility of x_i to trader i is at least that of the utility-maximizing bundle under prices π : $u_i(x_i) \geq \tilde{u}_i(\pi, \pi \cdot w_i)$ for each i ,*
2. *The total demand is at most $(1 + \epsilon)$ times the supply: $\sum_i x_i \leq (1 + \epsilon) \sum_i w_i$, and*
3. *The market clears: $\pi \cdot \sum_i w_i \leq \pi \cdot \sum_i x_i$.*

Note that item 3 above follows directly from item 1 and Assumption [1](#). If $\mathcal{K}_i \not\subseteq \mathfrak{R}_+^n$, we use a standard technique of “shifting” the space so that x_{ij} are non-negative. This, however, needs that \mathcal{K}_i is bounded below and we know these bounds. It also weakens the notion of approximate market equilibrium and we omit the details from this extended abstract. Shifting has also been used to address similar problems arising while solving linear programs with negative entries [24](#).

Without loss of generality, we scale the endowments w_i so that $\sum_i w_i = \mathbf{1}$, the vector of all ones. This also implies that we scale the vectors in \mathcal{K}_i . We emphasize that the algorithm also works without scaling; however the scaling simplifies the presentation. The algorithm is given in Figure [1](#). Here $\delta > 0$ is a constant to be fixed later. The algorithm goes in N iterations. In each iteration, we first scale the current price vector p so that it lies in Π . We then “announce” this price vector and receive in response the utility-maximizing bundles $x_i \in \mathcal{K}_i$. We then update the price vector p according to the aggregate demands X_j of goods j as given in Step [2d](#).

Note that this update is essentially same as (within a $(1 + \delta)$ factor) to the following natural update in terms of *excess demand*. Let $Z_j = X_j - \sum_i w_{ij} = X_j - 1$ be the excess demand of good j . We can update p as:

$$p_j \leftarrow p_j(1 + \delta\sigma Z_j).$$

This is so because $(1 + \delta\sigma Z_j) \approx (1 + \delta\sigma X_j)(1 - \delta\sigma)$, which is in turn true since $Z_j = X_j - 1$ and $\delta\sigma$ is small. The extra factor $(1 - \delta\sigma)$ is common to all goods j

1. Initialize $p_j = 1$ for $1 \leq j \leq n$.
2. Repeat for $N = \frac{n}{\delta} \log_{1+\delta} n$ iterations:
 - (a) Find $\alpha > 0$ such that $\alpha p \in \Pi$. Announce prices $\pi = \alpha p$.
 - (b) Each trader i computes $x_i \in \operatorname{argmax}\{u_i(x) \mid x \in \mathcal{K}_i, \pi \cdot x \leq \pi \cdot w_i\}$.
 - (c) Compute the aggregate demand $X = \sum_i x_i$ and let $\sigma = \frac{1}{\max_j X_j}$ where X_j denotes the aggregate demand of good j .
 - (d) Update for each good j : $p_j \leftarrow p_j (1 + \delta \sigma X_j)$.
3. Output for each i : $\bar{x}_i = \frac{\sum_{r=1}^N \sigma(r) x_i(r)}{\sum_{r=1}^N \sigma(r)}$ where $x_i(r)$ and $\sigma(r)$ are the values of x_i and σ in the r th iteration.
4. Output $\bar{\pi} = \frac{\sum_{r=1}^N \sigma(r) \pi(r)}{\sum_{r=1}^N \sigma(r)}$ where $\pi(r)$ and $\sigma(r)$ are the values of $\hat{\pi}$ and σ in the r th iteration.

Fig. 1. Algorithm for the convex program (2)

and is factored away in the price scaling step. The algorithm in the end outputs, $\bar{\pi}$ and \bar{x}_i for all i , the weighted average of the prices and allocations computed in N iterations.

Lemma 1. *The outputs \bar{x}_i and $\bar{\pi}$ satisfy $u_i(\bar{x}_i) \geq \tilde{u}_i(\bar{\pi}, \bar{\pi} \cdot w_i)$ for each i .*

Proof. Since $\tilde{u}_i(\pi, \pi \cdot w_i)$ is convex when $\pi \in \Pi$ and $u_i(x_i)$ is concave when $x_i \in \mathcal{K}_i$, we have $\tilde{u}_i(\bar{\pi}, \bar{\pi} \cdot w_i) \leq \frac{\sum_r \sigma(r) \tilde{u}_i(\pi(r), \pi(r) \cdot w_i)}{\sum_r \sigma(r)} = \frac{\sum_r \sigma(r) u_i(x_i(r))}{\sum_r \sigma(r)} \leq u_i(\bar{x}_i)$.

The following main lemma about the output is proved below. The proof is based on the standard application of “experts theorem” or “multiplicative update” technique used previously in solving packing and covering linear programs [24, 12, 11].

Lemma 2. *The outputs \bar{x}_i satisfy $\sum_i \bar{x}_i \leq \frac{1}{1-2\delta} \sum_i w_i$.*

We set $\delta = \frac{\epsilon}{2(1+\epsilon)}$ so that $\frac{1}{1-2\delta} = 1 + \epsilon$. The proof of Theorem 7 now follows from Lemmas 1, 2, and Assumption 1 on the utility functions.

The main result of this section is summarized in the following theorem.

Theorem 7. *Our algorithm computes a weak $(1 + \epsilon)$ -approximate market equilibrium in an economy for which a set Π containing an equilibrium price is known such that for each i , the indirect utility function $\tilde{u}_i(\pi, \pi \cdot w_i)$ is a convex function of π when restricted to $\pi \in \Pi$.*

In the algorithm, each trader i makes $O(\epsilon^{-2} n \log n)$ calls to her “demand” oracle: given prices $\pi \in \Pi$, compute $x_i \in \operatorname{argmax}\{u_i(x) \mid x \in \mathcal{K}_i, \pi \cdot x \leq \pi \cdot w_i\}$.

Proof of Lemma 2. Let $\bar{x} = \sum_i \bar{x}_i$ and let $(\bar{x})_j$ denote the j th coordinate of \bar{x} . To this end, let us define a potential $\Phi(r) = \sum_j p_j(r)$ where $p_i(r)$ denote the value of p_i in the beginning of r th iteration. From the step 2d) in the algorithm, we have

$$\Phi(r+1) = \Phi(r) + \delta \sigma(r) \sum_j p_j(r) X_j(r)$$

where $X_j(r)$ denotes the value of X_j in the r th iteration. Thus

$$\frac{\Phi(r+1)}{\Phi(r)} = 1 + \delta\sigma(r) \sum_j \frac{p_j(r)}{\Phi(r)} X_j(r) = 1 + \delta\sigma(r) \leq \exp(\delta\sigma(r)).$$

The second equality follows from the fact that $\sum_j p_j(r) X_j(r) = \frac{1}{\alpha(r)} \sum_j \pi_j(r) X_j(r)$ which is, by Assumption [II](#), equal to $\frac{1}{\alpha(r)} \sum_j \pi_j(r) \sum_i w_{ij} = \frac{1}{\alpha(r)} \sum_j \pi_j(r) = \sum_j p_j(r) = \Phi(r)$. Here $\alpha(r)$ is the value of α in r th iteration.

Thus after taking telescoping product, we get

$$\Phi(N+1) \leq \Phi(1) \cdot \exp\left(\delta \sum_r \sigma(r)\right) = n \cdot \exp\left(\delta \sum_r \sigma(r)\right). \quad (4)$$

On the other hand, observe that

$$\begin{aligned} \Phi(N+1) &= \sum_j p_j(N+1) = \sum_j \prod_{r=1}^N (1 + \delta\sigma(r) X_j(r)) \\ &\geq \sum_j \exp\left(\delta(1-\delta) \sum_r \sigma(r) X_j(r)\right) \\ &\geq \max_j \exp\left(\delta(1-\delta) \sum_r \sigma(r) X_j(r)\right) \\ &= \exp\left(\delta(1-\delta) \max_j \sum_r \sigma(r) X_j(r)\right). \end{aligned}$$

The first inequality follows from the elementary fact that $1 + \mu \geq \exp(\mu(1-\delta))$ for all $0 < \mu < \delta < \frac{1}{2}$. Combining the above observation with [\(4\)](#), we get

$$\delta(1-\delta) \max_j \sum_r \sigma(r) X_j(r) \leq \log \Phi(N+1) \leq \log n + \delta \sum_r \sigma(r).$$

Therefore,

$$\max_j(\bar{x})_j = \max_j \frac{\sum_r \sigma(r) X_j(r)}{\sum_r \sigma(r)} \leq \frac{1}{1-\delta} + \left(\frac{\log n}{\delta(1-\delta) \sum_r \sigma(r)}\right). \quad (5)$$

Now we “charge” the second term on the right-hand-side in [\(5\)](#) to $\max_j(\bar{x})_j$ as follows. Note that at least one p_j increases by a factor $(1+\delta)$ in any iteration. Thus after $N = \frac{n}{\delta} \log_{1+\delta} n$ iterations, $\max_j p_j(N+1) \geq n^{1/\delta}$. Also

$$(\bar{x})_j = \frac{\sum_r \sigma(r) X_j(r)}{\sum_r \sigma(r)} = \frac{\log \prod_r \exp(\delta\sigma(r) X_j(r))}{\delta \sum_r \sigma(r)} \geq \frac{\log p_j(N+1)}{\delta \sum_r \sigma(r)}.$$

Thus $\max_j(\bar{x})_j \geq \frac{\log n}{\delta^2 \sum_r \sigma(r)}$. Putting all pieces together, we get

$$\max_j(\bar{x})_j \leq \frac{1}{1-\delta} + \left(\frac{\delta \max_j(\bar{x})_j}{1-\delta}\right).$$

Thus $\max_j(\bar{x})_j \leq \frac{1}{1-2\delta}$.

5 Future Work

Our definitions of approximate market equilibrium is weak because the budget constraints of traders are satisfied only in the aggregate sense. Some of the traders may be spending significantly more than their budget. Moreover, some positively priced items may not be fully allocated. A notion of strongly approximate market equilibrium may be defined on the lines of [13], where budget constraints of no trader may exceed by a factor more than $(1 + \epsilon)$ and no item with positive price is under-demanded. Under this definition it might be possible to prove the “closeness” of the discovered prices to the equilibrium prices (see e.g., [14]). If we set $\delta = O(\frac{\epsilon \min_i \lambda_i}{\sum_i \lambda_i})$, where λ_i is the income of trader i in a Fisher economy, our tâtonnement algorithm obtains a strong $(1 + \epsilon)$ approximate market equilibrium in the above sense in $O((\frac{\epsilon \min_i \lambda_i}{\sum_i \lambda_i})^{-2} n \log n)$ iterations. It will be very interesting to develop a tâtonnement algorithm that converges to a strong approximate market equilibrium in near linear number of iterations. Finally, it is interesting to note that the continuous time version of our process can be described as $\frac{d\pi_j}{dt} = \pi_j Z_j$ where $Z_j = \sum_j x_{ij} - \sum_i w_{ij}$ is the excess demand of good j . Under what conditions is this process or its “time-average” $\hat{\pi}_j = \frac{1}{t} \int_{\tau=0}^t \pi_j d\tau$ stable and does converge to the equilibrium?

References

1. Arrow, K., Debreu, G.: Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22, 265–290 (1954)
2. Arrow, K.J., Block, H.D., Hurwicz, L.: On the stability of the competitive equilibrium, II. *Econometrica* 27(1), 82–109 (1959)
3. Chen, N., Deng, X., Sun, X., Yao, A.: Fisher Equilibrium Price with a class of Concave Utility Functions. In: Albers, S., Radzik, T. (eds.) *ESA 2004*. LNCS, vol. 3221, pp. 169–179. Springer, Heidelberg (2004)
4. Codenotti, B., Varadarajan, K.: Computation of market equilibria by convex programming. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
5. Codenotti, B., McCune, B., Pemmaraju, S.V., Raman, R., Varadarajan, K.: An experimental study of different approaches to solve the market equilibrium problem. In: *ALENEX/ANALCO*, pp. 167–179 (2005)
6. Codenotti, B., McCune, B., Raman, R., Varadarajan, K.: Computing equilibrium prices: Does theory meet practice? In: Brodal, G.S., Leonardi, S. (eds.) *ESA 2005*. LNCS, vol. 3669. Springer, Heidelberg (2005)
7. Codenotti, B., McCune, B., Varadarajan, K.: Market equilibrium via the excess demand function. In: *STOC* (2005)
8. Cole, R., Fleischer, L.: Fast-converging tatonnement algorithms for the market problem. Technical Report TR2007-602, Dept. of Computer Science, Dartmouth College (2007), <http://www.cs.dartmouth.edu/reports/>
9. Deng, X., Papadimitriou, C., Safra, S.: On the Complexity of Equilibria. In: *STOC* (2002)
10. Devanur, N., Papadimitriou, C., Saberi, A., Vazirani, V.: Market Equilibrium via a Primal-Dual-Type Algorithm. In: *FOCS*, pp. 389–395 (2002); Journal version to appear in the *Journal of the ACM*

11. Fleischer, L.: Approximating fractional multicommodity flow independent of the number of commodities. *SIAM J. Discrete Math.* 13, 505–520 (2000)
12. Garg, N., Könemann, J.: Faster and simpler algorithms for multicommodity flow and other fractional packing problems. In: *FOCS*, pp. 300–309 (1998)
13. Garg, R., Kapoor, S.: Auction algorithms for market equilibrium. *Math. Oper. Res.* 31(4), 714–729 (2006)
14. Garg, R., Kapoor, S.: Price roll-backs and path auctions: An approximation scheme for computing the market equilibrium. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) *WINE 2006*. LNCS, vol. 4286, pp. 225–238. Springer, Heidelberg (2006)
15. Garg, R., Kapoor, S., Vazirani, V.: An Auction-Based Market Equilibrium Algorithm for the Separable Gross Substitutability Case. In: Jansen, K., Khanna, S., Rolim, J.D.P., Ron, D. (eds.) *RANDOM 2004 and APPROX 2004*. LNCS, vol. 3122, pp. 128–138. Springer, Heidelberg (2004)
16. Grötschel, M., Lovász, L., Schrijver, A.: *Geometric Algorithms and Combinatorial Optimization*. Springer, Berlin (1988)
17. Jain, K., Mahdian, M., Saberi, A.: Approximating Market Equilibrium. In: Arora, S., Jansen, K., Rolim, J.D.P., Sahai, A. (eds.) *RANDOM 2003 and APPROX 2003*. LNCS, vol. 2764, pp. 98–108. Springer, Heidelberg (2003)
18. Jain, K., Vazirani, V.: Eisenberg-gale markets: Algorithms and structural properties. In: *STOC* (2007)
19. Kelly, F.P., Vazirani, V.: Rate control as a market equilibrium (manuscript, 2002)
20. Mas-Colell, A.: *The Theory of General Economic Equilibrium: A Differential Approach*. Cambridge University Press, Cambridge (1985)
21. Negishi, T.: A note on the stability of an economy where all goods are gross substitutes. *Econometrica* 26(3), 445–447 (1958)
22. Negishi, T.: The stability of a competitive economy: A survey article. *Econometrica* 30(4), 635–669 (1962)
23. Nikaido, H., Uzawa, H.: Stability and non-negativity in a Walrasian process. *International Econ. Review* 1, 50–59 (1960)
24. Plotkin, S., Shmoys, D., Tardos, E.: Fast approximation algorithms for fractional packing and covering problems. *Math. Oper. Res.* 20, 257–301 (1995)
25. Quah, J.K.-H.: The monotonicity of individual and market demand. *Econometrica* 68(4), 911–930 (2000)
26. Samuelson, P.A.: *Foundations of Economic Analysis*. Harvard University Press, Cambridge (1947)
27. Varian, H.: *Microeconomic Analysis*. W. W. Norton, New York (1992)
28. Vazirani, V.: Combinatorial algorithms for market equilibria. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
29. Walras, L.: *Elements of Pure Economics, or the Theory of Social Wealth*. Lausanne, Paris (1874) (in French)
30. Ye, Y.: A path to the Arrow-Debreu competitive market equilibrium. *Mathematical Programming* (2006)

A FPTAS for Computing a Symmetric Leontief Competitive Economy Equilibrium

Zhisu Zhu^{1,*}, Chuangyin Dang², and Yinyu Ye^{3,**}

¹ Institute for Computational and Mathematical Engineering, Stanford University,
Stanford, California 94305-4042

zhuzhisu@stanford.edu

² Department of Manufacturing Engineering & Engineering Management, City
University of Hong Kong, Kowloon, Hong Kong SAR

mecdang@cityu.edu.hk

³ Department of Management Science and Engineering, Stanford University,
Stanford, California 94305-4042

yinyu-ye@stanford.edu

Abstract. We consider a linear complementarity problem (LCP) arisen from the Arrow-Debreu-Leontief competitive economy equilibrium where the LCP coefficient matrix is symmetric. We prove that the decision problem, to decide whether or not there exists a complementary solution, is NP-complete. Under certain conditions, an LCP solution is guaranteed to exist and we present a fully polynomial-time approximation scheme (FPTAS) for computing such a solution, although the LCP solution set can be non-convex or non-connected. Our method is based on solving a quadratic social utility optimization problem (QP) and showing that a certain KKT point of the QP problem is an LCP solution. Then, we further show that such a KKT point can be approximated with running time $\mathcal{O}((\frac{1}{\epsilon}) \log(\frac{1}{\epsilon}) \log(\log(\frac{1}{\epsilon})))$ in accuracy $\epsilon \in (0, 1)$ and a polynomial in problem dimensions. We also report preliminary computational results which show that the method is highly effective.

1 Introduction

Given a real n by n matrix A , consider the linear complementarity problem (LCP) to find u and v such that

$$A^T u + v = e, \quad u^T v = 0, \quad (u \neq 0, v) \geq 0, \quad (1)$$

where e is the vector of all ones. Note that $u^T v = 0$ implies that $u_i v_i = 0$ for all $i = 1, \dots, n$. Also, $u = 0$ and $v = e$ is a trivial complementary solution. But we look for a non-trivial solution where $u \neq 0$ (see Cottle et al. [5] for more literature on linear complementarity problems).

In this note, we focus on the case that A is symmetric. We first prove that the decision problem, to decide whether or not there exists such a complementary

* Zhisu Zhu is supported by a Henry Fan Stanford Graduate Fellowship.

** Research supported in part by NSF grant DMS-0604513.

solution, is NP-complete. Under certain conditions, for example, that all entries of A is non-negative, an LCP solution is guaranteed to exist. Then, we present a fully polynomial-time approximation scheme (FPTAS) for computing a solution, although the LCP solution set can be non-convex or non-connected.

Our method is based on solving a quadratic social utility optimization problem (QP) and showing that a certain KKT point of the QP problem is an LCP solution. Then, we further show that such a KKT point can be approximated with running time $\mathcal{O}((\frac{1}{\epsilon}) \log(\frac{1}{\epsilon}) \log(\log(\frac{1}{\epsilon})))$ in accuracy $\epsilon \in (0, 1)$ and a polynomial in problem dimensions. We also report preliminary computational results which show that the method is highly effective in comparison with other well known LCP solvers.

2 Connection to Competitive Market and Bimatrix Game Equilibria

The LCP (II) rises from the Arrow-Debreu exchange competitive economy equilibrium problem where it was first formulated by Léon Walras in 1874. In this equilibrium problem everyone in a population of m traders has an initial endowment of a divisible goods and a utility function for consuming all goods—their own and others'. Every trader sells the entire initial endowment and then uses the revenue to buy a bundle of goods such that his or her utility function is maximized. Walras asked whether prices could be set for everyone's goods such that this is possible. An answer was given by Arrow and Debreu in 1954 [1] who showed that, under mild conditions, such equilibrium would exist if the utility functions were concave. In general, it is unknown whether or not an equilibrium can be computed efficiently.

Consider a special class of Arrow-Debreu's problems, where each of the n traders has exactly one unit of a divisible and distinctive good for trade, and let trader i , $i = 1, \dots, n$, bring good i , which class of problems is called the *pairing class* [13]. For given prices p_j on good j , consumer i 's maximization problem is

$$\begin{aligned} & \text{maximize } u_i(x_{i1}, \dots, x_{in}) \\ & \text{subject to } \sum_j p_j x_{ij} \leq p_i, \\ & \quad x_{ij} \geq 0, \quad \forall j. \end{aligned} \tag{2}$$

Let x_i^* denote a maximal solution vector of (2). Then, vector p is called the Arrow-Debreu price equilibrium if there exists an x_i^* for consumer i , $i = 1, \dots, n$, such that

$$\sum_i x_i^* = e$$

where e represents available amount of goods on the exchange market.

The Leontief exchange economy problem is the Arrow-Debreu equilibrium when the utility functions are in the Leontief form:

$$u_i(x_i) = \min_{j: a_{ij} > 0} \left\{ \frac{x_{ij}}{a_{ij}} \right\},$$

where the Leontief coefficient matrix is given by

$$A = \begin{pmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \dots & \dots & \dots & \dots \\ a_{n1} & a_{n2} & \dots & a_{nn} \end{pmatrix}.$$

It was proved that

Theorem 1. (Ye [13]) *Let $B \subset \{1, 2, \dots, n\}$, $N = \{1, 2, \dots, n\} \setminus B$, A_{BB} be irreducible, and u_B satisfy the linear system*

$$A_{BB}^T u_B = e, \quad A_{BN}^T u_B \leq e, \quad \text{and} \quad u_B > 0.$$

Then the (right) Perron-Frobenius eigen-vector p_B of $U_B H_{BB}$ together with $p_N = 0$ will be a Leontief economy equilibrium. And the converse is also true.

Theorem 1 has thus established a combinatorial algorithm to compute a Leontief economy equilibrium by finding a right block $B \neq \emptyset$, which is precisely a (non-trivial) complementary solution to the LCP problem (1).

The LCP (1) is also connected to the bimatrix game equilibrium problem specified by a pair of $n \times m$ pay-off matrices C and R , with positive entries, one can construct a Leontief exchange economy with $n + m$ traders and $n + m$ goods as follows.

Theorem 2. (Codenotti et al. [4]) *Let (C, R) denote an arbitrary bimatrix game, where assume, w.l.o.g., that the entries of the matrices C and R are all positive. Let*

$$A^T = \begin{pmatrix} 0 & C \\ R^T & 0 \end{pmatrix}$$

describe the Leontief utility coefficient matrix of the traders in a Leontief economy. There is a one-to-one correspondence between the Nash equilibria of the game (C, R) and the market equilibria A of the Leontief economy.

Therefore, computing a bimatrix game equilibrium is also equivalent to computing a complementary solution of LCP (1). The reader may want to read Brainard and Scarf [2], Gilboa and Zemel [8], Chen, Deng and Teng [3], Daskalakis, Goldberg and Papadimitriou [7], and Tsaknakis and Spirakis [11] on hardness and approximation results of computing a bimatrix game equilibrium.

3 Decision of the Existence of an LCP Solution

In general, it's difficult to decide if LCP (1) has a complementary solution or not, even when A is symmetric.

Theorem 3. *Let A be a real symmetric matrix. Then, it is NP-complete to decide whether or not LCP (1) has a complementary solution such that $u \neq 0$.*

Proof. Given a symmetric matrix A , it's NP-complete (see Murty and Kabadi [10]) to decide if

$$\exists u \geq 0 \quad \text{such that} \quad u^T A u > 0? \quad (3)$$

The complement problem is to decide if or not for all $u \geq 0$ one has $u^T A u \leq 0$, or $-A$ is co-positive plus.

We now prove that the decision problem (3) is equivalent to the problem that if or not LCP (1) has a complementary solution $u \neq 0$.

If (1) has a complementary solution $u \neq 0$, then

$$0 = u^T (e - Au) = e^T u - u^T A u.$$

Since $u \geq 0$ and $u \neq 0$, we have $u^T A u = e^T u > 0$.

On the other hand, if the answer to the decision problem (3) is “yes”, then the maximal value of the following bounded quadratic problem:

$$\begin{aligned} (QP) \quad & \text{maximize} \quad u^T A u \\ & \text{subject to} \quad e^T u = 1, \quad u \geq 0, \end{aligned} \quad (4)$$

is strictly positive. Let u^* be the global maximizer of the problem. Then, u^* must satisfy the Karush-Kuhn-Tucker (KKT) conditions:

$$\begin{aligned} -2Au + \lambda e &= v \\ u^T v &= 0, \\ e^T u &= 1, \\ (u, v) &\geq 0, \lambda \text{ free.} \end{aligned} \quad (5)$$

The first two equations in (5) imply that $\lambda = \frac{2(u^*)^T A u^*}{e^T u^*} = 2(u^*)^T A u^* > 0$. Thus, $\bar{u} = \frac{2u^*}{\lambda} \geq 0$ is complementary solution of LCP (1) and $\bar{u} \neq 0$.

The question remains: given symmetric A , is it easy to compute one if LCP (1) is known to have a complementary solution? Note that, the complementary solution set of (1), even non-empty, is not convex nor even connected. For example, let

$$A^T = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}. \quad (6)$$

Then, there are three isolated non-trivial complementary solutions:

$$u^1 = (1/2; 0), \quad u^2 = (0; 1/2), \quad u^3 = (1/3; 1/3).$$

In the next section, however, we develop a fully polynomial-time approximation scheme (FPTAS) to compute ϵ -approximate complementary solution for LCP (1) when A is symmetric and $\sum_{i,j} a_{ij} > 0$, that is, the sum of all entries of A is positive. Here, an ϵ -approximate complementary solution is a pair $(u \neq 0, v)$ such that

$$A^T u + v = e, \quad (u \neq 0, v) \geq 0, \quad \frac{u^T v}{\bar{a}} \leq \epsilon,$$

where \bar{a} is the largest entry in A :

$$\bar{a} = \max_{i,j} \{a_{ij}\} (> 0). \quad (7)$$

In most applications, one can scale A such that $\bar{a} = 1$.

4 A Social Optimization and FPTAS

We consider a quadratic “social” utility function $u^T A u$, which we like to maximize over the simplex $\{u : e^T u = 1, u \geq 0\}$. This can be written as the quadratic programming problem of QP (4) in the previous section.

Since $e^T A e > 0$ so that LCP (II) has at least one non-trivial complementary solution. Further more, the maximal value of QP (4) is strictly greater than 0 but bounded above by \bar{a} (recall that \bar{a} is the largest entry of A). These facts, together with the proof of Theorem 3, lead to

Lemma 1. *Let A be symmetric. Then, every KKT point u of problem (4), with $u^T A u > 0$, is a (non-trivial) complementary solution for LCP (I).*

In [14], an interior-point potential reduction algorithm was proved to be a FPTAS for computing an ϵ -approximate KKT point of general quadratic programming with bounded feasible region. It can be adapted in solving QP (4) in a running time bounded by $\mathcal{O}((\frac{n^4}{\epsilon} \log \frac{1}{\epsilon} + n^4 \log n)(\log \frac{1}{\epsilon} + \log n))$ arithmetic operations. The algorithm reduces the potential function

$$P(u) = \rho \log(\bar{a} - u^T A u) - \sum_{j=1}^n \log(u_j),$$

where $\rho = (2n + \sqrt{n})/\epsilon$, by a constant each iteration from the initial point $u^0 = \frac{1}{n}e$, till u becomes an ϵ -approximate KKT point.

Note that

$$P(u^0) = \rho \log\left(\bar{a} - \frac{1}{n^2} e^T A e\right) + n \log(n),$$

and for any $u \in \{u : e^T u = 1, u > 0\}$,

$$- \sum_{j=1}^n \log(u_j) \geq n \log(n).$$

Thus, $P(u) < P(u^0)$ implies that

$$\rho \log(\bar{a} - u^T A u) < \rho \log\left(\bar{a} - \frac{1}{n^2} e^T A e\right)$$

or

$$u^T A u > \frac{1}{n^2} e^T A e > 0,$$

that is, any KKT point u generated by the algorithm must have $u^T A u > 0$. To conclude, we have

Theorem 4. *There is a FPTAS to compute an ϵ -approximate non-trivial complementary solution of LCP (1) when A is symmetric and $e^T Ae > 0$. Moreover, such a solution is an ϵ -approximate equilibrium of the symmetric Leontief economy when all entries of A are positive.*

5 Preliminary Computational Results

Here, we computationally compare three type methods to solve the complementarity problem of (1): 1) the QP-based potential reduction algorithm (referred as QP) presented in this paper; 2) a homotopy-based path-following algorithm method (referred as HOMOTOPY) developed in Dang et al. [6]; 3) Mixed Complementarity Problem (MCP) general solvers PATH (Ferris and Munson, <http://www.gams.com/dd/docs/solvers/path.pdf>) and MILES (Rutherford <http://www.gams.com/dd/docs/solvers/miles.pdf>), where both solvers use a Lemke type algorithm that is based on a sequence of pivots similar to those generated by the simplex method for linear programming; see Lemke [9].

If one applies Lemke’s algorithm (PATHS or MILES) directly to solving LCP (1), it will return the trivial solution $u = 0$, $v = e$. To exclude it, we rewrite LCP (1) into an equivalent homogeneous LCP as follows:

$$Mz + q = w, \quad z^T w = 0, \quad (z, w) \geq 0, \quad (8)$$

where $z, w \in R^{n+1}$,

$$M = \begin{pmatrix} -A^T & e \\ e^T & 0 \end{pmatrix} \in M^{n+1}, q = \begin{pmatrix} 0_n \\ -1 \end{pmatrix}.$$

Then, we can obtain a solution for LCP (1) from a complementary solution of LCP (8). However, the standard Lemke algorithm may not be able to solve LCP (8) either, since it may terminate at the second iteration with a non-complementary “secondary-ray” solution. Thus, as shown below, commonly used LCP solver PATH or MILES seems cannot successfully solve LCPs (8) most of times.

Both QP and HOMOTOPY are coded in MATLAB script files, and all solvers are run in the MATLAB environment on a desktop PC (2.8GHz CPU). For the QP-based potential reduction algorithm, we set $\epsilon = 1.e - 8$. After the termination, we use the support of u , $\{i : u_i \geq 1.e - 5\}$, to recalibrate an “exact” solution (to the machine accuracy) for LCP (1).

For different size n ($n = 20 : 20 : 100, 100 : 100 : 1000, 1500 : 500 : 3000$), we randomly generate 15 symmetric and sparse matrices A of two different types (uniform in $[0, 1]$ or binary $\{0, 1\}$) and solve them by the three methods. In the following tables, “mean_sup” the average support size of u and “max_sup” the maximum support size of u in the 15 problems, “mean_iter” the average number of iterations of QP and Homotopy algorithms (each iteration solves a system of linear equations), and “mean_time” the average computing CPU time in seconds.

From our preliminary computational results, we can draw few conclusions. First, LCP (1), although the matrix A is symmetric, seems not an easy problem

Table 1. QP for solving uniform symmetric matrix LCP

n	mean_sup	mean_iter	mean_time	max_sup
20	4.1	39.5	0.1	5
40	4.5	46.0	0.1	5
60	4.5	47.9	0.1	5
80	4.9	47.5	0.2	6
100	5.3	48.2	0.3	7
200	5.5	53.5	1.2	6
300	5.6	59.3	3.4	8
400	5.7	55.1	5.9	7
500	5.9	62.5	11.3	7
600	5.7	58.8	16.0	7
700	5.8	58.8	23.4	7
800	5.8	62.6	33.8	8
900	5.7	65.1	47.3	7
1000	6.3	65.0	60.2	7
1500	6.1	71.5	187.2	8
2000	5.9	73.5	411.9	7
2500	6.4	74.6	774.5	8
3000	6.2	78.7	1404.2	8

Table 2. HOMOTOPY for solving uniform symmetric matrix LCP

n	mean_sup	mean_iter	mean_time	max_sup
20	4.1	37.7	0.2	5
40	4.4	52.7	0.4	5
60	4.4	58.3	0.8	6
80	4.6	68.2	1.4	6
100	5.3	72.6	2.2	7
200	4.9	108.9	14.0	6
300	5.5	127.7	49.3	8
400	5.5	160.5	111.9	7
500	5.7	159.7	181.6	7
600	5.5	182.5	317.0	6
700	5.9	202.9	515.6	7
800	5.5	208.9	706.3	6
900	5.7	231.7	1039.2	7
1000	5.9	267.2	1644.0	7
1500	5.9	305.5	4726.4	7
2000	5.7	307.1	10105.2	6

Table 3. PATH for solving uniform symmetric matrix LCP

n	mean_sup	mean_time	max_sup
20	8.7	0.1004	12
40	13.8	0.3406	23
n \geq 60		fail to solve	

Table 4. QP for solving binary symmetric matrix LCP

n	mean_sup	mean_iter	mean_time	max_sup
20	11.8	35.2	0.1	13
40	16.6	43.3	0.1	20
60	21.1	44.4	0.2	23
80	22.1	46.9	0.3	25
100	23.9	53.3	0.5	27
200	30.0	54.5	1.7	34
300	32.5	66.9	5.2	35
400	34.1	65.1	9.5	38
500	35.4	67.1	16.1	39
600	36.0	82.9	31.4	39
700	37.9	68.0	35.4	42
800	37.8	74.9	55.4	41
900	37.8	78.1	76.5	43
1000	38.7	82.1	106.6	42
1500	40.0	84.9	305.3	43
2000	42.4	91.4	702.2	45
2500	42.9	94.7	1382.8	47
3000	43.9	99.5	1959.4	48

Table 5. HOMOTOPY for solving binary symmetric matrix LCP

n	mean_sup	mean_iter	mean_time	max_sup
20	11.7	48.6	0.2	14
40	16.2	68.3	0.5	21
60	20.6	75.3	0.9	24
80	22.9	84.0	1.7	26
100	24.3	92.9	2.9	27
200	31.3	111.1	14.6	39
300	32.3	130.4	51.1	39
400	32.4	108.2	79.9	34
500	34.8	153.6	263.7	41
600	34.4	144.8	451.3	37
700	35.6	184.0	572.3	38
800	36.5	208.0	1628.1	37
900	37.2	261.2	4733.4	41
1000	37.2	502.8	5370.1	38

Table 6. PATH for solving binary symmetric matrix LCP

n	mean_sup	mean_time	max_sup
20	8.2	0.0445	12
40	10.2	0.3229	17
n _≥ 60		fail to solve	

to solve. Secondly, the QP-based FPITAS algorithm lives up with its theoretical expectation and it is numerically effective. Thirdly, the homotopy-based algorithm seems able to solve sizable problems, although its computational complexity is not proven to be a PITAS. Finally, as mentioned earlier, the general LCP solvers, PATH and MILES, may terminate with a “secondary-ray” solution at the second Lemke pivot, therefore fail to solve LCP (8). As a result, in our numerical experiments MILES can solve none of our test problems, and PATH can only solve a small number of test problems with size no more than 50. (PATH use an alternative default pivoting rule and it switches to original Lemke’s pivot rule only when the default rule fails or the users force to do so.)

In particular, for the simple example (6) with three isolated non-trivial complementary solutions $u^1 = (1/2; 0)$, $u^2 = (0; 1/2)$, $u^3 = (1/3; 1/3)$, all three methods above get the same solution $(1/3; 1/3)$.

6 Further Remarks

We make few final remarks and open questions.

First, is symmetric LCP (II) in the PPAD class described by [3] and [7]?

Secondly, by restricting A being symmetric for bimatrix game setting described in Section 2, we must have $R = C$, that is, the two payoff matrices are identical. But in this case, a trivial, pure-strategic, and Pareto-optimal bimatrix game equilibrium is to simply play the largest entry in C . Thus, it remains to be seen if the QP-based approach offer a PTAS for computing a bimatrix equilibrium with a larger support. Note that the constant-approximation result of Tsaknakis and Spirakis [11] was indeed based on computing a KKT point of a social QP problem.

Thirdly, an important direction is to study the LCP problem (II) where A is not necessarily symmetric. In this case, even all entries of A being non-negative may not guarantee the existence of a (non-trivial) complementary solution; see example:

$$A^T = \begin{pmatrix} 0 & 2 \\ 0 & 1 \end{pmatrix}.$$

Finally, the computational results based on randomly generated data show that the support of u is relative small. Is there a theoretical justification for this fact or observation?

References

1. Arrow, K.J., Debreu, G.: Existence of an equilibrium for competitive economy. *Econometrica* 22, 265–290 (1954)
2. Brainard, W.C., Scarf, H.E.: How to compute equilibrium prices in 1891. Cowles Foundation Discussion Paper 1270 (August 2000)
3. Chen, X., Deng, X., Teng, S.: Computing Nash Equilibria: Approximation and Smoothed Complexity. In: 47th Annual IEEE Symposium on Foundations of Computer Science, pp. 603–612 (2006)

4. Codenotti, B., Saberi, A., Varadarajan, K., Ye, Y.: Leontief Economies Encode Nonzero Sum Two-Player Games. In: SODA 2006 (2006); Theoretical Computer Science (to appear)
5. Cottle, R., Pang, J.S., Stone, R.E.: The Linear Complementarity Problem. Academic Press, Boston (1992)
6. Dang, C., Ye, Y., Zhu, Z.: A path-following algorithm for computing a Leontief economy equilibrium (in preparation, 2008)
7. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. In: 47th Annual IEEE Symposium on Foundations of Computer Science, pp. 71–78 (2006)
8. Gilboa, I., Zemel, E.: Nash and Correlated equilibria: Some Complexity Considerations. *Games and Economic Behavior* 1, 80–93 (1989)
9. Lemke, C.E.: On complementary pivot theory. In: Dantzig, Veinott (eds.) *Mathematics of Decision Sciences, Part 1*, pp. 95–114. American Mathematical Society, Providence (1968)
10. Murty, K.G., Kabadi, S.N.: Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming* 39, 117–129 (1987)
11. Tsaknakis, H., Spirakis, P.: An Optimization Approach for Approximate Nash Equilibria. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 42–56. Springer, Heidelberg (2007)
12. Walras, L.: *Elements of Pure Economics, or the Theory of Social Wealth* (1874) (1899, 4th edn.1926, rev ed., 1954, Engl. Transl.)
13. Ye, Y.: Exchange Market Equilibria with Leontief’s Utility: Freedom of Pricing Leads to Rationality. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 14–23. Springer, Heidelberg (2005); *Theoretical Computer Science* 378(2), 134–142 (2007)
14. Ye, Y.: On The Complexity of Approximating a KKT Point of Quadratic Programming. *Mathematical Programming* 80, 195–212 (1998)

Online and Offline Selling in Limit Order Markets

Kevin L. Chang^{1,*} and Aaron Johnson^{2,**}

¹ Yahoo Inc.

klchang@yahoo-inc.com

² Yale University

ajohnson@cs.yale.edu

Abstract. Completely automated electronic securities exchanges and algorithms for trading in these exchanges have become very important for modern finance. In [4], Kakade *et al.* introduced the limit order market model, which is a prevalent paradigm in electronic markets. In this paper, we consider both online and offline algorithms for maximizing revenue when selling in limit order markets. We first prove that the standard reservation price algorithm has an optimal competitive ratio for this problem. This ratio is not constant, and so we consider computing solutions offline. We show that the offline optimization problem is **NP**-hard, even for very restricted instances. We complement the hardness result by presenting an approximation scheme that runs in polynomial time for a wide class of instances.

1 Introduction

Electronic exchanges are very important venues for trading many different classes of financial securities. With the widespread use of such markets, a technique known as *algorithmic trading* has become popular among financial institutions and institutional investors who seek to buy or sell large amounts of a particular security. Consider, for instance, a pension fund that would like to sell many shares of a particular stock. Typically, such an investor would not sell large amounts of stock himself by just submitting an order to a stock exchange, but rather would seek the expertise of a broker (*e.g.* an investment bank) to perform the transaction on his behalf. A standard strategy for the broker is to break the large order into many smaller orders and to submit the orders gradually over a given time horizon, with the obvious goal of maximizing his total revenue.

Advances in financial technology have automated this process, and the pension fund's entire trade can now be effected by computer with little human intervention: computer programs can choose how to divide the large order into smaller orders, choose at what time and at what price to submit the smaller orders, and

* Work performed while affiliated with the Max Planck Institute for Computer Science.

** Supported by NSF grants 0428422, 0534052, and 0716223, ONR grants N00014-01-1-0795 and N00014-04-1-0725, and ARO grant W911NF-05-1-0417.

then actually execute the transactions on an electronic market. Sophisticated algorithms for such large trades are offered by most large investment banks and by many smaller firms specializing in financial technology.

The manner in which an electronic market itself executes the various orders from clients can also be complex. One of the most important electronic exchange paradigms is the *electronic communication network* (ECN), which implements a *limit order market*. Prominent ECNs for trading equities and equity derivatives are operated by NASDAQ, Instinet, and NYSE-Euronext.

In a limit order market, buyers and sellers submit limit orders to buy or sell a commodity; these orders have semantics such as “I would like to buy v shares, but I am only willing to pay p USD or less for each share.” The market matches buyers and sellers based on a transparent function of the orders submitted. Often, the input order stream is available to market participants, and so algorithms should be conscious of this *market microstructure* in order to effectively trade.

In this paper, we will study natural theoretical computer science questions motivated by algorithmic trading in limit order markets. We consider both on-line and offline algorithms for placing sell orders in limit order markets, with the goal of maximizing the revenue generated by selling volume N of a particular commodity (*e.g.* a stock). This problem domain was introduced into the theoretical computer science literature by Kakade, Kearns, Mansour and Ortiz in [4], although online trading algorithms in much simpler market models have been studied for many years.

1.1 The Trading Model: The Mechanics of Limit Order Markets

In a limit order market, market participants submit limit orders, which consist of three-tuples: $\sigma = \langle \theta, p, v \rangle$. The parameter θ specifies whether the order is to buy or to sell; p denotes the least competitive price the market participant is willing to accept, that is, the lowest price *per share* that is acceptable for a sell order, or the highest price *per share* for a buy order; and v denotes the volume, that is, the number of shares to transact.

A sell order $\sigma_1 = \langle \text{S}, p_1, v_1 \rangle$ can be “matched” to a buy order $\sigma_2 = \langle \text{B}, p_2, v_2 \rangle$ if $p_1 \leq p_2$. If σ_1 and σ_2 transact and $v_1 > v_2$, then σ_2 will be filled, but σ_1 will only be partially filled: its volume will be reduced to $v_1 \leftarrow v_1 - v_2$. If $v_1 < v_2$, then σ_1 will be filled, and for σ_2 , we have $v_2 \leftarrow v_2 - v_1$. Observe that when $p_1 < p_2$, any price $p \in [p_1, p_2]$ would be acceptable to both parties. Limit order markets use the convention that the transaction will occur at the limit price of the order that arrived first. With this convention, it is advantageous to be the second of the two matched orders to arrive.

At any time step, buy and sell orders that have been submitted but that have not yet been paired with suitable counterparties are stored in the *buy order book* and the *sell order book*, respectively. Orders in the book are sorted according to their prices, with the most competitive orders at the “top” of the book (*i.e.* in the buy order book, orders with the highest price are at the top; in the sell order book, orders with the lowest price are at the top). Ties are broken by placing the order that arrives first higher in the book. An important property of the two

books is that the prices of all orders in the buy book are lower than the prices of all orders in the sell book, since the books consist of orders that have not yet been matched.

When a new sell order $\sigma_s = \langle S, p_s, v_s \rangle$ arrives, it is compared with the top order in the buy book, say $\sigma_b = \langle B, p_b, v_b \rangle$. If $p_s > p_b$, then no transaction can occur and σ_s will be placed in the sell book according to the rules specified in the previous paragraph. If $p_b \geq p_s$, then $\min\{v_b, v_s\}$ shares are sold at price p_b per share. If $v_s \geq v_b$, then σ_b has been filled and is removed from the buy book, while the volume of σ_s is adjusted accordingly and a new matching buy order is sought. If $v_s < v_b$, then the order σ_s is filled, and the volume of σ_b is adjusted accordingly. An arriving buy order would be processed in an analogous fashion. The state of the book is public knowledge in many ECNs and can therefore be exploited by sophisticated traders.

Table 1 is an example of buy and sell order books. If a new sell order arrived at price 102.00 and volume 50, then it would transact with the buy order at price 102.20. The volume of this buy order would be reduced by 50.

Table 1.

Buy orders		Sell orders	
Price	Volume	Price	Volume
102.20	100	102.55	200
102.00	500	102.93	300

1.2 The Trading Problem

The trading over the course of the time horizon is represented by time steps $t = 1, \dots, n$. At time step t , order σ_t is placed by some market participant (not the algorithm). As each order arrives, executions occur and the book is updated as described above.

The problem that we consider is to design an algorithm that inserts sell orders into this stream in order to maximize its revenue. At each time step t , the algorithm may place orders before the arrival of σ_t . It cannot sell a total volume of more than N , nor can it submit buy orders.

We will consider both online and offline algorithms for this problem. In the online problem, the algorithm observes the market orders over time and, as they come in, must insert its own orders. The standard measure of quality of an online algorithm is defined by its *competitive ratio*. A randomized profit-maximizing algorithm has a competitive ratio at most $c > 0$ if, for all input sequences Σ , $E[Rev] \geq \frac{1}{c} \cdot \text{OPT}$, where OPT is the revenue that results from an optimum, offline placement of orders on input Σ and $E[Rev]$ is the expected revenue of the algorithm. We refer the reader to the book by Borodin and El-Yaniv [1] for an introduction to online algorithms.

1.3 Our Contributions

We show that the standard online *reservation price* algorithm yields a competitive ratio of $e \log R \leq 2.72 \log R$ in the order book model, where $R = p_{\max}/p_{\min}$ is the *price ratio* between the highest and lowest possible prices in the order stream. This competitive ratio is optimal, in the sense that any *randomized* algorithm must have competitive ratio at least $\frac{1}{2} \log R$. All logarithms are base e .

Our bound on the competitive ratio is an improvement over the $O(\log R \log N)$ bound for the reservation price algorithm by Kakade *et al.* [4], because our bound has no dependence on N .

Since no online selling algorithm can achieve a competitive ratio that is $o(\log R)$, we examine maximizing revenue in the offline case. We prove that this optimization problem is **NP**-hard by reducing from **KNAPSACK**, even for the restricted case where market participants only submit orders with three distinct prices. We also show that a simple dynamic programming algorithm will find the optimum solution. Like **KNAPSACK**, the running time is *pseudopolynomial*: polynomial in N and the largest volume of any order (which in general can be exponential in the size of the input) when the number of distinct prices that market participants may submit is a *fixed constant* k . We then prove that the input volumes can be adjusted so that running the dynamic programming algorithm on the adjusted input will yield a polynomial-time approximation scheme when there is a fixed constant k of distinct prices. The running time with approximation ratio $1 - \epsilon$ is polynomial in n , $1/\epsilon$ and R . We note that all securities indeed only have a constant number of possible prices (*e.g.* US equities are traded in multiples of \$.01), although this constant is admittedly large.

The offline optimization problem is a natural and theoretically interesting question to consider. Although it cannot immediately be applied to the real online problem faced by a trader, nonetheless these results could be of interest to a broader audience. The **NP**-hardness of computing the offline optimum *even with knowledge of the entire order stream* is a very strong statement about the intractability of optimization, for any conceivable context.

Offline algorithms have many potential applications. For instance, an approximation algorithm can be used for studying historical data, and the output of the approximation algorithm can be used to compare the realized performance of a trading algorithm to its theoretical optimum. Also, an offline algorithm could be coupled with an appropriate statistical model for generating sample paths of the future evolution of market microstructure in order to design realizable trading strategies.

Finally, we generalize some of our results to the buying case.

1.4 Related Work in Theoretical Computer Science

Online algorithms for selling in limit order markets were first introduced by Kakade, Kearns, Mansour and Ortiz in [4]. Kakade *et al.* considered selling algorithms that seek to optimize revenue, as well as selling algorithms that seek to sell shares at the *average price* of the market (the *Volume Weighted Average Price*, which is a popular benchmark for commercially available trading algorithms). Even-Dar, Kakade, Kearns, and Mansour also considered the stability of limit order dynamics in [3].

The limit order market is a generalization of simpler online trading models. El-Yaniv, Fiat, Karp, and Turpin in [2] considered the *one-way trading*, *time series search* and *two-way trading* problems in this framework. This work was later extended by Lorenz, Panagiotou, and Steger in [5].

1.5 Preliminaries

The input order stream will be denoted by $\Sigma = \sigma_1, \dots, \sigma_n$, where order $\sigma_t = \langle \theta_t, p_t, v_t \rangle$ arrives at time t . The algorithm's sequence of sell orders is represented by a set $\{\sigma_i^A\}$, where σ_i^A denotes a sell order that is placed right before the arrival of $\sigma_t \in \Sigma$ at time t . If the algorithm places a set of sell orders Σ^A in the order sequence Σ , we will denote the new order stream by $\Sigma \cdot \Sigma^A$.

If the limit prices of all orders in the input stream fall into the interval $[p_{\min}, p_{\max}]$, which is known to the algorithm, then $R = p_{\max}/p_{\min}$ is the *price ratio*. In our paper, we will refer to Lemma 5.3 by Evan-Dar *et al.* [3]. The lemma is most easily stated for the case where all orders have unit volume:

Lemma 1 (Stability Lemma [3]). *Suppose all orders have unit volume, and the order stream Σ' is derived from Σ by inserting a single order σ . (1) If σ is not executed in Σ' , then the sets of executed orders in Σ and in Σ' are identical. (2) If σ is executed in Σ' , then at most 1 order (as specified by id) was executed in Σ but not in Σ' . Conversely, at most 1 order (other than σ) was executed in Σ' but not in Σ .*

From a high-level perspective, the stability lemma states that if we insert an extra order into the order sequence, it will not affect the set of executed orders by very much.

2 Optimal Online Algorithms

We show in this section that the *reservation price* algorithm for selling in the order book model has an optimal competitive ratio. This algorithm was originally considered for the *max-search problem* (also called the *time-series search problem*) by El-Yaniv *et al.* [2] and later by Lorenz, Panagiotou and Steger in [5] and by Kakade *et al.* in [4] for selling in the order book model.

The reservation price algorithm is: pick an integer i uniformly at random between 0 and $\lfloor \log R \rfloor$ and place an order to sell all N shares of stock at price $e^i p_{\min}$ at time $t = 0$, where R is the price ratio.

Theorem 2. *The reservation price algorithm for selling in the order book model has competitive ratio $e \log R$. Furthermore, any randomized algorithm must have competitive ratio at least $\frac{1}{2} \log R$.*

Proof. We first prove that the reservation price algorithm has competitive ratio at most $e \log R$. Let $p_{\text{res}} = e^i p_{\min}$ be the reservation price randomly chosen by the algorithm. Suppose p_1, \dots, p_N are the prices realized for the sale of the N shares by the optimal solution, Σ^A , and let $\text{OPT} = \sum_1^n p_i$ be the optimal revenue. Let $P_i = \{p_j : p_{\text{res}} \leq p_j \leq e \cdot p_{\text{res}}\}$ be the set of prices of executed orders that are within a factor of e of the price p_{res} .

Claim. The reservation price algorithm will sell at least $|P_i|$ shares at price p_{res} .

The main idea of the proof, which is omitted for space, is that if $\sigma_0^R = \langle S, p_{\text{res}}, N \rangle$ is the order placed by the reservation price algorithm, then in the set of executions of $\Sigma \cdot (\{\sigma_0^R\} \cup \Sigma^A)$, σ_0^R will execute $|P_i|$ shares, since it will have a higher priority than the orders in P_i . It follows that if σ_0^R is placed by itself, it will realize at least as much revenue.

With probability $1/\log R$, the reservation price algorithm will choose reservation price $p_{\text{res}} \leftarrow e^i p_{\text{min}}$. Then the expected revenue of the algorithm is

$$\sum_{i=0}^{\lfloor \log R \rfloor} \frac{1}{\log R} |P_i| e^i p_{\text{min}} \geq \sum_{i=1}^{\lfloor \log R \rfloor} \sum_{p_j \in P_i} \frac{p_j}{e \log R} = \frac{\text{OPT}}{e \log R}.$$

Thus, the competitive ratio of the reservation price algorithm is at most $e \log R$.

In order to establish the lower bound, we observe that it is straightforward to reduce the online *max-search* problem to our problem of selling in a limit order market. In this problem, a player observes a sequence of prices and tries to select the highest one. The $\frac{1}{2} \log R$ lower bound on the competitive ratio for any randomized algorithm for max-search proved by Lorenz *et al.* in [5] establishes the lower bound for limit order markets. \square

3 NP-Hardness of the Offline Problem

We prove the **NP**-completeness of optimal selling by reduction from **KNAPSACK**. Our reduction is to the special case of instances in which only three different prices occur, p_h, p_m, p_l , where $p_h \geq p_m \geq p_l$.

We will use several facts about the structure of the optimum solution in the three-price case.

Lemma 3. *There exists an optimal solution that places all its high-price sell orders at time $t = 0$. The number of such orders can be arbitrarily large.*

Lemma 4. *There exists an optimal solution that places all its low-price orders immediately after the last execution it realizes at either price p_h or price p_m .*

Lemma 5. *There exists an optimal solution that only places a medium-price order immediately after an execution occurs at a high price.*

Lemma 6. *There exists an optimal solution, such that if we attempt to insert another medium-price order after an execution at a high price, then the algorithm will achieve one less execution at a high price, except when this occurs after the last execution at a high price.*

The proofs of these lemmas show that any optimal solution can easily be converted to satisfy these properties.

We will reduce the **NP**-complete **KNAPSACK** problem to offline selling. An input to **KNAPSACK** $\mathcal{I}_{\mathcal{K}}$ is a set of n pairs, (w_i, v_i) , a capacity C , and a value V . It is in the language if there exists a subset $S \subseteq [n]$ such that $\sum_{i \in S} v_i \geq V$ and

$\sum_{i \in S} w_i \leq C$. Let $v^* = \max_i(v_i)$ and let $W = \sum_i w_i$. We can assume that all numbers in the instance are positive integers, and that $v^* > 1$.

The three prices in the selling instance we will create are $p_l = 1$, $p_m = (C + 1)v^*$, and $p_h = (C + 1)v^* + 1$. Let $a_i = (C + 1)(w_i v^* - v_i)$ and $b_i = ((C + 1)(w_i v^* - v_i) + w_i)(C + 1)(v^*)$. For $1 \leq i \leq n$, let Σ_i be the following sequence of limit orders:

1. $\langle \mathbf{B}, p_m, a_i + w_i \rangle$
2. $\langle \mathbf{S}, p_l, w_i \rangle$
3. $\langle \mathbf{B}, p_l, w_i \rangle$
4. $\langle \mathbf{S}, p_m, a_i \rangle$
5. $\langle \mathbf{B}, p_h, a_i + b_i \rangle$

Let Ω be the following order sequence:

1. $\langle \mathbf{S}, p_l, W - C \rangle$
2. $\langle \mathbf{B}, p_m, p_m W \rangle$

Let Σ be the concatenated order sequence $(\Sigma_1, \dots, \Sigma_n, \Omega)$. The total volume of buy orders in Σ is $N = \sum_i (2a_i + 2w_i + w_i p_m + b_i)$. N will be the number of shares to sell. The revenue to raise will be $R = p_h \sum_i (a_i + b_i) + p_m^2 W + (C + 1)V$. Then $\mathcal{I}_S = (\Sigma, N, R)$ is an input to the offline selling problem. It is in the language if R revenue can be obtained by selling at most N shares.

Lemma 7. *If there exists a solution $S \subseteq [n]$ to \mathcal{I}_K with total value V , then there exists a solution to \mathcal{I}_S with revenue of at least R .*

Proof Sketch: Given S , first insert the order $\sigma_0^A = \langle \mathbf{S}, p_h, \sum_i (a_i + b_i) \rangle$ at the beginning of Σ . σ_0^A executes with every high-price buy order, which yields revenue $p_h \sum_i (a_i + b_i)$. If $i \in S$, insert the order $\sigma_i^A = \langle \mathbf{S}, p_m, a_i + w_i \rangle$ at the beginning of Σ_i . When σ_i^A is added, then $a_i + w_i$ sales are made by σ_i^A at the medium price, and a_i less sales are made from the high-price buy order at the end of Σ_i . Therefore the change in revenue from σ_i^A is $p_m(a_i + w_i) - p_h(a_i) = (C + 1)v_i$. Finally, insert the order $\sigma_{n+1}^A = \langle \mathbf{S}, p_m, p_m W \rangle$ after subsequence Σ_n , and insert $\sigma_{n+2}^A = \langle \mathbf{S}, p_l, C \rangle$ at the end of Ω . Regardless of the previous insertions, the revenue from σ_{n+1}^A and σ_{n+2}^A is $p_m^2 W + p_l(C - \sum_{i \in S} w_i) \geq p_m^2 W$. Then the total revenue obtained is at least $p_h \sum_i (a_i + b_i) + p_m^2 W + (C + 1)V = R$. \square

Next we prove the converse of Lemma 7. By Lemma 3, we can assume that the optimal solution for \mathcal{I}_S places a large sell order at price p_h at the beginning of Σ . Observe that in the resulting execution, that sell order executes at the high price after every subsequence Σ_i . By Lemma 5, we can assume that all medium-price orders in the optimal solution are each inserted at the beginning either of some Σ_i or of Ω . Let $S \subseteq [n]$ be the set of Σ_i subsequences for which this happens. Finally, by Lemma 4, we can assume that any low-price order is inserted at some point after the last high-price or medium-price order is inserted.

Lemma 8. *For $i \in S$, there is an optimal solution in which the medium-price order inserted at the beginning of Σ_i has volume $a_i + w_i$.*

Proof Sketch: Let $\sigma_i^A = \langle \mathbf{S}, p_m, v \rangle$, $i \in S$, be the medium-price order inserted at the beginning of Σ_i . If $v > a_i + w_i$ or $v \leq a_i$, decreasing v increases high-price sells and decreases medium-price sells, for a net gain in revenue. Otherwise, v can be increased to $a_i + w_i$ without reducing high sells, contradicting Lemma 6. \square

Lemma 9. *There is an optimal solution in which any low-price order is inserted after the Σ_n subsequence.*

Proof Sketch: Suppose instead that the low-price order σ_i^A is inserted before Σ_i . Consider moving σ_i^A to the beginning of Σ_i . If σ_i^A reduces the volume of high-prices transactions before Σ_i , this move would increase net revenue. Otherwise, the move maintains revenue. \square

Lemma 10. *There exists an optimal solution with the order $\langle S, p_m, p_m W \rangle$ inserted at the beginning of Ω and the order $\langle S, p_l, C - \sum_{i \in S} w_i \rangle$ inserted at the end.*

Proof Sketch: By Lemmas 4, 5 and 10, the solution places a low-price order in Ω that may be preceded by a medium-price order. Inspection shows that including the medium-price order is optimal. \square

We can see by inspecting Ω that the payoff of the orders described in Lemma 10, expressed as a function of the volume of the low-price buy book, is:

$$\rho(l) = \begin{cases} p_m^2 W + p_l(C - l) & 0 \leq l \leq C \\ p_m^2 W - p_m(l - C) & C < l \leq W. \end{cases} \quad (1)$$

Lemma 11. *We can assume that $\sum_{i \in S} w_i \leq C$.*

Proof. Suppose instead that $\sum_{i \in S} w_i > C$. Consider removing the medium-price order at the beginning of Σ_i , $i \in S$. The sequence loses the $p_m(a_i + w_i) - p_h(a_i) = (C + 1)v_i$ revenue from that order. The transactions in Σ_j subsequences are unaffected, but at the end of each there is an additional w_i volume in the buy book at the low price. Equation 1 shows that this volume increases the revenue obtained by at least $(C + 1)v^*$. Therefore the total change in revenue is nonnegative, and we can convert this sequence into an optimum such that $\sum_{i \in S} w_i \leq C$.

Finally, we can show that the optimal revenue of \mathcal{I}_S can give us a lower bound on the value of optimal subsets in \mathcal{I}_K .

Lemma 12. *If there exists a solution to \mathcal{I}_S with revenue of at least R , then there exists a solution to \mathcal{I}_K with total value at least V .*

Proof. If $\sum_{i \in S} v_i \geq V$, then, by Lemma 11, S is a solution to \mathcal{I}_K of value at least V . Suppose otherwise, that $\sum_{i \in S} v_i < V$. The solution receives $p_h \sum_i (a_i + b_i)$ revenue from the initial high-price sell order. It receives $(C + 1) \sum_{i \in S} v_i \leq (C + 1)(V - 1)$ revenue from the medium-price orders in S . We can see from Equation 1 that it receives no more than $p_m^2 W + p_l(C)$ revenue from the two orders in Ω . This accounts for all of the revenue R' of the solution. However, $R' - R = (C + 1)(V - 1) + p_l(C) - (C + 1)V < 0$, contradicting the assumption of the lemma. \square

Theorem 13. *The decision version of optimal offline selling is NP-complete.*

Proof. The decision version of the offline selling problem is in **NP**, because we can calculate the revenue of a solution by running the limit order market algorithm on the total sequence. Also, the previous lemmas show that the input to KNAPSACK is in the language if and only if its reduction is a member of the offline selling problem language. \square

In contrast to the three-price case, it can be shown that if there are only two prices in the order sequence, the problem can be solved exactly in $O(n)$ time. The algorithm simply places one high-price sell order at the beginning and then tries all positions for the low-price order.

4 Offline Algorithms and Approximation Schemes

In this section, we present approximation schemes for the offline selling problem. The general approach to our algorithm is similar to the FPTAS for KNAPSACK, but the technical details are more involved. We first give a pseudopolynomial dynamic programming algorithm. Then we show that this algorithm can be used in an approximation scheme by reducing and rounding the order volumes. The approximation scheme will have running time polynomial in $1/\epsilon$, the maximum price ratio R , and n , if the number of prices at which market participants can submit orders is a fixed constant.

4.1 Pseudopolynomial Time Dynamic Programming Algorithm

A simple dynamic programming algorithm can compute the optimal placement of sell orders in polynomial time, under the assumption that the volume of each limit order is 1, and that the number of distinct prices at which the market participants can place orders is at most a fixed constant k .

The input to the dynamic programming subproblem is given by: (1) Times t_1 and t_2 , such that $t_1 \leq t_2$. (2) The initial buy and and sell order books at time t_1 : B_{t_1}, S_{t_1} . (3) The final order books at time t_2 : B_{t_2}, S_{t_2} . (4) m , the number of shares to be sold by the algorithm between times t_1 and t_2 .

Each subproblem is then: Given buy and sell order books B_{t_1} and S_{t_1} at time t_1 (prior to the arrival of order σ_{t_1}), find the optimum placement of orders between times t_1 and t_2 (inclusive), such that the buy and sell order books at the end of time t_2 are B_{t_2} and S_{t_2} and that the number of shares sold by the algorithm between times t_1 and t_2 is at most m .

Theorem 14. *For the case where each order has unit volume, the dynamic programming algorithm will find an optimal solution in time $O(N^{3k+5}n^{3k+5})$, where k is the number of distinct prices.*

The details of the algorithm and its analysis follow standard dynamic programming techniques. This algorithm can be used with orders of arbitrary volume, which adds an additional factor in the runtime.

Corollary 15. *If the volume of each order is unrestricted, the dynamic programming algorithm will run in time $O(n^2 N^{3k+5} (nV)^{3k+3})$, where V is the maximum volume of any order.*

4.2 PTAS for the Arbitrary Volume Case

We now show how the input can be preprocessed in two steps so that the dynamic programming algorithm from the previous section can compute a solution with revenue at least $(1 - \epsilon)\text{OPT}$ in time $O\left(n^{12k+16} (R/\epsilon)^{6k+8}\right)$, for input sequences with arbitrary volumes in each order.

Step 1: Reduce to the Significant Volume Case.

Our first step will be to modify the instance to ensure that N is at least a fraction of the total volume of all orders, which we call the *significant volume* condition. Σ satisfies this condition if $(n + 1) \cdot N \geq V$, where V is the maximum volume of any order in Σ .

Lemma 16. *Given an order stream Σ , we can construct an order stream Π such that*

1. *If V is the maximum volume of any order $\pi_i \in \Pi$, then $(n + 1) \cdot N \geq V$.*
2. *If Σ^A is any set of sell orders with total volume at most N placed by the algorithm, it will realize the same revenue in $\Pi \cdot \Sigma^A$ as in the original input $\Sigma \cdot \Sigma^A$.*

We assume that there is at least one order in the original sequence Σ , $\sigma_i = \langle t_i, p_i, v_i \rangle \in \Sigma$ such that $v_i \geq N \cdot (n + 1)$. Since a selling algorithm will only transact N shares, intuitively its action should have very little effect on order σ_i , which contains many more shares.

Let $\text{trans}_\Sigma(\sigma_i)$ denote the set of orders that are matched with σ_i in the evolution of the order sequence Σ . Let $\text{unex}_\Sigma(\sigma_i)$ denote the volume of σ_i that is unexecuted in the evolution of Σ . Let $\text{match}_\Sigma(\sigma_i, \sigma_j)$ denote the number of shares of σ_i that are matched with σ_j in the order sequence Σ .

Lemma 17. *Let Σ^A be any set of orders placed by the algorithm with an aggregate volume of at most N . Then for any $\sigma_i, \sigma_j \in \Sigma$,*

$$\text{match}_{\Sigma \cdot \Sigma^A}(\sigma_i, \sigma_j) \geq \text{match}_\Sigma(\sigma_i, \sigma_j) - N.$$

The lemma follows from applying the Stability Lemma to the transactions between σ_i and σ_j .

Lemma 17 implies that σ_i and σ_j have excess volume that is, in some sense, superfluous to the problem of selling at most N shares. We eliminate these in a new order sequence Π .

For each order $\sigma_i \in \Sigma$, where $\sigma_i = \langle \theta_i, p_i, v_i \rangle$, we create order π_i such that $\pi_i = \langle \theta_i, p_i, v'_i \rangle$, where $v'_i = v_i - \sum_{\sigma_j \in \text{trans}_\Sigma(\sigma_i)} \max(\text{match}_\Sigma(\sigma_i, \sigma_j) - N, 0) - \max(\text{unex}_\Sigma(\sigma_i) - N, 0)$. With these new volumes,

$$\text{vol}(\pi_i) \leq \text{vol}(\sigma_i) - \sum_j \text{match}_\Sigma(\sigma_i, \sigma_j) - \text{unex}_\Sigma(\sigma_i) + (n + 1)N = (n + 1)N.$$

Thus, the modified input sequence $\Pi = \pi_1, \dots, \pi_n$ satisfies Condition 1 of Lemma 16. Condition 2 follows from Lemma 17.

Step 2: Round Volumes.

Recall that V is the maximum volume of any order. We now assume that our input has been reduced to the significant volume case, where $(n + 1) \cdot N \geq V$. Let $M = \epsilon \cdot N / (nR)$. In the second preprocessing step, we round the volume of every order in Σ to the nearest multiple of M . The volume of each order will be changed by at most $M/2$ and the number of possible values for the volume of an order will be $V/M = nVR / (N\epsilon) \leq n(n + 1)R/\epsilon$. Let Σ' be the input sequence of orders with rounded volumes.

Lemma 18. *Finding the optimum solution to Σ' will induce a solution with revenue at least $(1 - \epsilon)\text{OPT}$ for the original input sequence Σ .*

Proof. We may assume that N is at most the aggregate volume of all buy orders in Σ . Then, $\text{OPT} \geq p_{\min} \cdot N$, because the algorithm could place the order $\langle \mathcal{S}, p_{\min}, N \rangle$ at the beginning and sell to the first N buy orders. We first prove that there exists a solution with revenue at least $(1 - \epsilon/2)\text{OPT}$ for Σ' . Let Σ^A be an optimum set of sell orders for the input sequence Σ . We may assume that the realized price for every order in Σ^A is the same as the price of the order.

Recall that our rounding scheme changed the volume of each order by at most $M/2$. The Stability Lemma therefore implies that the total volume of shares that are executed in $\Sigma \cdot \Sigma^A$ but not in $\Sigma' \cdot \Sigma^A$ is at most $n \cdot M/2$. Therefore, there are at most $n \cdot M/2$ shares that the algorithm sold in $\Sigma \cdot \Sigma^A$ but did not sell in $\Sigma' \cdot \Sigma^A$. The total revenue lost is at most

$$p_{\max} \cdot \frac{nM}{2} = \frac{\epsilon}{2} \cdot p_{\min} \cdot N \leq \frac{\epsilon}{2} \cdot \text{OPT}.$$

An analogous argument will prove that the optimum solution on Σ' with revenue OPT' will induce a solution on Σ with revenue at least $(1 - \epsilon/2)\text{OPT}'$. It can then be inferred that the optimum solution for Σ' will induce a solution on Σ with revenue at least $(1 - \epsilon)\text{OPT}$. \square

We combine the two preprocessing steps with the dynamic programming algorithm to obtain an approximation scheme that runs in polynomial time when the number of price levels k is constant.

Theorem 19. *For any $\epsilon > 0$, dynamic programming with preprocessing will yield an algorithm with approximation ratio at least $1 - \epsilon$ that runs in time $O\left(n^{12k+16} (R/\epsilon)^{6k+8}\right)$.*

5 Extension to Buying

In the buying case, the algorithm's task is to insert buy orders into the order sequence in order to buy at least N shares, with the goal of minimizing the total cost of the trade.

We note that there is an asymmetry between the profit maximization (selling case) and the cost minimization (buying case) online trading problems. The

results of Steger *et al.* [5] for the *min-search* problem imply that no algorithm can achieve a competitive ratio better than $O(\sqrt{R})$. For improved guarantees, we consider offline algorithms. The dynamic programming algorithm can be easily modified to the buying case. During the rounding step, however, we must set $M = \alpha N/n$, for any $\alpha > 0$.

Theorem 20. *Let OPT be the cost of the offline optimum solution that buys exactly N shares. For any $\alpha > 0$, the dynamic programming with preprocessing will yield an algorithm that buys at least $(1 - \alpha)N$ shares with cost at most OPT . The algorithm runs in time $O\left(n^{12k+16} (1/\alpha)^{6k+8}\right)$.*

References

1. Borodin, A., El-Yaniv, R.: Online Computation and Competitive Analysis. Cambridge University Press, Cambridge (1998)
2. El-Yaniv, R., Fiat, A., Karp, R.M., Turpin, G.: Optimal search and one-way trading online algorithms. *Algorithmica* 30(1), 101–139 (2001)
3. Even-Dar, E., Kakade, S.M., Kearns, M.S., Mansour, Y.: (In)Stability properties of limit order dynamics. In: ACM Conference on Electronic Commerce, pp. 120–129 (2006)
4. Kakade, S., Kearns, M.J., Mansour, Y., Ortiz, L.E.: Competitive algorithms for VWAP and limit order trading. In: ACM Conference on Electronic Commerce, pp. 189–198 (2004)
5. Lorenz, J., Panagiotou, K., Steger, A.: Optimal algorithms for k -search with applications in option pricing. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) *ESA 2007*. LNCS, vol. 4698, pp. 275–286. Springer, Heidelberg (2007)

Predictive Pricing and Revenue Sharing

Bobji Mungamuru and Hector Garcia-Molina

Stanford University, Stanford, CA, USA 94305

bobji@i.stanford.edu,

hector@cs.stanford.edu

Abstract. *Predictive pricing* (e.g., Google’s “Smart Pricing” and Yahoo’s “Quality Based Pricing”) and *revenue sharing* are two important tools that online advertising networks can use in order to attract content publishers and advertisers. We develop a simple model of the pay-per-click advertising market to study the market effects of these tools. We then present an algorithm, PRICINGPOLICY, for computing an advertising network’s best response i.e., given the predictive pricing and revenue sharing policies used by its competitors, what policy should an advertising network use in response? Using PRICINGPOLICY, we gain insight into the structure of optimal predictive pricing and revenue sharing policies.

1 Introduction

Google’s “Smart Pricing” [4] and Yahoo’s “Quality-Based Pricing” [9] are examples of a practice we refer to as *predictive pricing*. The idea behind predictive pricing in pay-per-click advertising is to charge the same advertiser different prices for click-throughs, depending on which publisher the click-through originated from. For example, an advertiser who bid on the keyword “camera” might be charged less for a click-through from a travel website than one from a photographer’s blog, since the latter would (ostensibly) be more targeted to potential camera purchasers than the former. Advertising networks use predictive pricing to attract publishers and advertisers to their network.

Revenue sharing, which is the practice of paying out a fraction of earned revenues to the publishers where click-throughs originate, is another tool used by advertising networks to attract traffic. Revenue sharing is the reason publishers display advertisements alongside their content in the first place. In this paper, we study how an online advertising network can apply predictive pricing and revenue sharing “optimally” – that is, in a manner that maximizes the advertising network’s profits.

The sheer size of the online advertising market makes this problem interesting and important. Although predictive pricing and revenue sharing can help advertising networks attract and retain lucrative traffic, applying these tools suboptimally can mean that a network is “leaving money on the table” (either by paying out an unnecessarily large revenue share, or by attracting less- or lower-quality traffic than they could be). And in a market that, by most estimates, is worth several billions of dollars, the losses due to suboptimal pricing policies can be tremendous. Advertising networks that currently do not apply predictive pricing should feel compelled to start – our results suggest that they are yielding a significant advantage to their competitors.

The practice of predictive pricing in the pay-per-click advertising market is relatively new. To the authors’ knowledge, there has been no formal analysis thus far of how to

apply predictive pricing and share revenue properly. Recent research on “click quality” has focused on a related (but orthogonal) problem i.e., *click fraud* [2,3,6,8]. Click fraud relates to whether a given click-through is valid or invalid. Predictive pricing, on the other hand, focuses on the probability that a valid click-through becomes a conversion i.e., the *conversion rate*. Also, techniques for fighting click fraud are typically not applied on a per-publisher basis (apart from simple blacklisting). Predictive pricing, on the other hand, allows for very fine-grained publisher-level control.

1.1 Overview

We begin by constructing a model of the online advertising market as a game between content publishers, advertising networks and advertisers. We then derive an expression for an advertising network’s *best-response function*. That is, if an advertising network knows the predictive pricing and revenue sharing policies of its competitors, what policy should the network choose in response, in order to maximize its profits? The expression we derive for the best-response is implicit – it is the solution to a difficult optimization problem. We then present an algorithm, PRICINGPOLICY, for solving this optimization problem, yielding a best-response predictive pricing and revenue sharing policy.

Finally, we apply PRICINGPOLICY toward answering some qualitative questions about predictive pricing:

- Is it always optimal to charge less for lower-quality traffic? (Yes.)
- Should an advertising network always try to attract as much traffic as it can, regardless of traffic quality? (No.)
- If a network is better at targeting, can it offer a lower revenue share? (Yes.)
- Does predictive pricing harm publishers, as has been conjectured in online forums? (Yes and no – it harms low-quality publishers and helps high-quality publishers.)

In principle, the best-response function can be used as a “subroutine” for computing equilibrium policies for advertising networks (an equilibrium is, by definition, a fixed point of the networks’ best-response functions). However, we believe that the practical value of our algorithm lies in computing best responses, rather than equilibria. It prescribes actions that networks can take “today” in response to their competitors, rather than waiting for equilibria to unfold. Thus, our focus will be on finding best responses.

2 Model

For brevity, we present only a brief overview of our model here. For a complete development, we refer the reader to [7]. Table 1 gives the reader a sense of the various quantities involved in our model.

We model the *pay-per-click* (PPC) advertising market as a one-shot dynamic game between three classes of players: content publishers, advertising networks and advertisers. *Content publishers* (or, *publishers*) publish websites and display advertisements alongside their content. *Advertisers* design advertisements (or, *ads*) and bid on keywords that describe the interests of their target market. *Advertising networks* (or, *networks*) act as intermediaries, auctioning off click-throughs (or, *clicks*) to advertisers and delivering relevant ads to publishers upon request.

Table 1. Summary of notation used in our model

Symbol	Description
I, J, K	Number of publishers, networks and advertisers (respectively)
V_i	Volume of clicks on publisher i 's site
β_i^{ub}	Conversion rate of publisher i 's traffic
c_{ij}	Fraction of publisher i 's clicks sent to ad network j
y_k	Advertiser k 's revenue per conversion
R_k	Advertiser k 's target ROI
β_k^{Adv}	Effectiveness of advertiser k 's ads
v_{kj}	Advertiser k valuation of ad network j 's clicks
β_j^{Net}	Network j 's skill at matching publishers and advertisers
θ_j	Network j 's expected auction revenue per click
κ_j	Network j 's "nominal" auction revenue per click
η_j	Network j 's total profit
η_j^{max}	Network j 's maximum possible profit
g_{ij}	Predictive pricing factor applied to publisher i 's traffic by network j
h_j	Revenue share paid out by network j
\mathbf{g}_1	Length- I vector whose i^{th} element is g_{i1}
\mathbf{C}	I -by- J matrix whose (i, j) -element is c_{ij}
\mathbf{G}_{-1}	Predictive prices chosen by all networks other than network 1
\mathbf{h}_{-1}	Revenue shares chosen by all networks other than network 1

Each time a user visits a publisher's site and clicks on an ad, the advertiser pays the network a small amount. The network then pays out a fraction of this amount to the publisher where the click originated. A small fraction of clicks eventually become *conversions* e.g., a product purchase, or a sign-up to an email list. The advertiser earns some revenue each time a click becomes a conversion.

Predictive pricing affects how much the advertiser is billed by the network, whereas the revenue share determines what fraction of this revenue is paid out to the publisher. For concreteness, suppose each publisher i receives V_i clicks on his website. Suppose $c_{ij} \in [0, 1]$ is the fraction of these clicks that publisher i sends to network j . The total number of clicks that publisher i sends to network j is then:

$$V_i c_{ij} \quad (1)$$

For each click coming from publisher i , network j bills advertisers for only a fraction $g_{ij} \in [0, 1]$ of a click i.e., advertisers receive a $(1 - g_{ij})$ discount. The fraction g_{ij} is the *predictive pricing factor*¹ that network j applies to publisher i 's traffic. The *effective* number of clicks publisher i is paid for by network j is then:

$$V_i c_{ij} g_{ij} \quad (2)$$

Of each dollar of revenue from advertisers, network j pays out a fraction $h_j \in [0, 1]$ to publishers. The fraction h_j is referred to as the *revenue share*. Suppose θ_j is the

¹ The term "predictive pricing" alludes to network j 's prediction about the quality of publisher i 's traffic (i.e., accounting for click-through rates, click fraud and conversion rates).

expected auction revenue per click on network j . Then, the total revenue to publisher i from network j is:

$$V_i c_{ij} g_{ij} h_j \theta_j \quad (3)$$

We refer to $\{g_{ij} \forall i\}$ and h_j together as network j 's *pricing policy*.

Our one-shot dynamic game is, therefore, comprised of two steps:

1. In the *first step*, each network j selects and announces its pricing policy (i.e., its revenue share, h_j , and predictive pricing factors, $\{g_{ij} \forall i\}$).
2. In the *second step*, each publisher i decides which networks to sell its clicks on (i.e., its *allocations*, $\{c_{ij} \forall j\}$). Simultaneously, each advertiser k decides how much it is willing to pay for clicks from each network j (i.e., its *valuations*, $\{v_{kj} \forall j\}$).

After the second step, payoffs are realized: a) publishers sell clicks (i.e., display ads) on their chosen networks, and b) advertisers pay the networks, who then pay the publishers.

3 Optimal Pricing Policies

Network j 's goal is to maximize its own profit, η_j . In [7], it is shown that:

$$\eta_j = \beta_j^{\text{Net}} \left(\sum_i V_i \beta_i^{\text{Pub}} c_{ij} \right) (1 - h_j) \kappa_j \quad (4)$$

Clearly, network j 's profit depends on the decisions made by publishers and advertisers (see Table 1). However, the networks act first in our game: publishers and advertisers observe the networks' decisions in the first step before deciding on their allocations and valuations in the second step. In other words, the outcome in the second step (i.e., publishers' allocations and advertisers' valuations) is the market's reaction to the first-step outcome (i.e., networks' pricing policies). Therefore, to maximize profit, each network j will: a) assume that an equilibrium will be played in the second step, and b) choose a pricing policy that induces the most profitable equilibrium in the second step².

Network j 's profit depends not only on j 's pricing policy, but also on the pricing policies chosen by competing networks in the first step. For example, if the revenue share h_j offered by network j is relatively low, then very few publishers may send traffic to j (i.e., $c_{ij} = 0$ for most i), leading to a low η_j . If h_j were relatively high, more publishers may choose network j , but η_j might be low again since j would be paying out too large a fraction of revenues to publishers. Therefore, network j must account for the actions of all other networks when choosing its own pricing policy.

We will now compute the *best response* for network 1, holding the policies of all other networks fixed, and assuming an equilibrium in the second step³. It can be shown that network 1's best response is a solution to the following optimization problem:

² For a given first-step outcome $\{g_{ij} \forall (i, j)\}$ and $\{h_j \forall j\}$, an equilibrium in the second step is defined as a scenario where every advertiser k chooses its valuations $\{v_{kj} \forall j\}$ optimally and every publisher i chooses its allocations $\{c_{ij} \forall j\}$ optimally.

³ Our choice of network 1 is without loss of generality. Obviously we can compute the best response for any network j in a similar manner.

$$\begin{aligned}
& \text{maximize} && \eta_1 \equiv \beta_1^{\text{Net}} \left(\sum_i V_i \beta_i^{\text{Pub}} c_{i1} \right) (1 - h_1) \kappa_1 \\
& \text{subject to} && X_{ij} = V_i g_{ij} h_j \theta_j \quad \forall (i, j) \\
& && \sum_j c_{ij} X_{ij} = \max_j X_{ij} \quad \forall i \\
& && \sum_j c_{ij} = 1 \quad \forall i \\
& && \theta_j = \kappa_j a_j \quad \forall j \\
& && a_j = \beta_j^{\text{Net}} \frac{\left(\sum_i V_i c_{ij} \beta_i^{\text{Pub}} \right)}{\left(\sum_i V_i c_{ij} g_{ij} \right)} \quad \forall j \\
& && 0 \leq g_{i1}, h_1, c_{ij} \leq 1 \quad \forall (i, j)
\end{aligned} \tag{5}$$

The objective in (5) is an expression for network 1's profit (see (4)). The first three constraints encode the assumption that each publisher chooses allocations optimally in the second step. The fourth and fifth constraints say that advertisers also choose valuations optimally. Thus, the first five constraints together imply that there is an equilibrium in the second step between publishers and advertisers. The final constraint gives ranges for the decision variables we are interested in.

Network 1's optimization problem (5) is highly non-convex, so even feasible points are not easy to find. One of our main contributions is an iterative algorithm, which we call PRICINGPOLICY, for finding approximate solutions to (5).

In [7], we show how to construct a *geometric programming (GP) relaxation* of (5) around a given point $(h_1, \mathbf{g}_1, \mathbf{C})$. GPs are log-convex [11], and therefore can be solved globally and efficiently. PRICINGPOLICY works by solving a sequence of these GPs. It outputs a sequence of feasible (but not-necessarily optimal) points, where each point yields weakly higher profits for network 1 than the previous point. The sequence of solutions to the relaxed problem converge to an approximate solution to (5).

4 Experiments

Using PRICINGPOLICY, we can gain some interesting insights into the structure of best-response pricing policies.

Our first experiment examines whether networks that apply predictive pricing gain a competitive edge, compared to networks that do not. Consider a market with $J = 2$ networks and $I = 20$ publishers. Each publisher i receives 100 clicks (i.e., $V_i = 100 \forall i$), and the quality of i 's traffic, β_i^{Pub} , is linear in i with values ranging from 0.25% to 5%⁴ (i.e., $\beta_i^{\text{Pub}} = 0.0025i$). The networks are equally effective at matching up publishers and advertisers i.e., $\beta_1^{\text{Net}} = \beta_2^{\text{Net}} = 1.0$. We assume $\kappa_1 = \kappa_2 = 10$, which means the auction mechanisms used by each network are also equally efficient.

We used PRICINGPOLICY to compute the best-response pricing policy for network 1, assuming network 2 does not use predictive pricing (i.e., $g_{i2} = 1 \forall i$) and offers

⁴ Such a range is realistic – 5% would be considered a high conversion rate in practice.

Algorithm 1. PRICINGPOLICY**Require:** $\mathbf{G}_{-1}, \mathbf{h}_{-1}, T$

- 1: Select arbitrary initializations $h_1^{(0)}$ and $\mathbf{g}_1^{(0)}$
- 2: Use fixed-point iteration to compute second-step equilibrium, $\mathbf{C}^{(0)}$, assuming other networks play $(\mathbf{h}_{-1}, \mathbf{G}_{-1})$ and network 1 plays $(h_1^{(0)}, \mathbf{g}_1^{(0)})$
- 3: **for** $t \in 1, \dots, T$ **do**
- 4: Solve GP-relaxation of (5) to find an optimal point $(h_1', \mathbf{g}_1', \mathbf{C}')$ that is “close to” $(h_1^{(t-1)}, \mathbf{g}_1^{(t-1)}, \mathbf{C}^{(t-1)})$, assuming other networks play $(\mathbf{h}_{-1}, \mathbf{G}_{-1})$
- 5: $(h_1^{(t)}, \mathbf{g}_1^{(t)}, \mathbf{C}^{(t)}) \leftarrow (h_1', \mathbf{g}_1', \mathbf{C}')$
- 6: **end for**
- 7: Use fixed-point iteration to recompute second-step equilibrium, $\mathbf{C}^{(T)}$, assuming other networks play $(\mathbf{h}_{-1}, \mathbf{G}_{-1})$ and network 1 plays $(h_1^{(T)}, \mathbf{g}_1^{(T)})$
- 8: **return** $(h_1^{(T)}, \mathbf{g}_1^{(T)}, \mathbf{C}^{(T)})$

publishers a revenue share of 50% (i.e., $h_2 = 0.5$). To solve the GP-relaxation of (5) in line 4 of PRICINGPOLICY, we used CVX, a package for specifying and solving convex programs [5]. We initialized the algorithm with random choices of \mathbf{g}_1 and h_1 .

Figure 1(a) shows the revenue share $h_1^{(t)}$ output at each iteration t , as well as the market share $\frac{1}{T} \sum_i c_{i1}^{(t)}$, estimated profit $\hat{\eta}_1^{(t)}$ and actual profit $\eta_1^{(t)}$ at each iteration 5. The “estimated profit” is computed using the estimated allocations $\mathbf{C}^{(t)}$ output by iteration t of PRICINGPOLICY, whereas the “actual profit” is computed using the actual second-stage equilibrium allocations that would result if network 1 played $(h_1^{(t)}, \mathbf{g}_1^{(t)})$.

From Figure 1(a), we see that the algorithm converges after roughly $T = 50$ iterations. The estimated profit tracks the actual profit reasonably well – in this case it is an underestimate of the actual profit, but in other experiments we ran it was an overestimate. As iterations progress, $h_1^{(t)}$ steadily decreases – PRICINGPOLICY recommends progressively better predictive prices $\mathbf{g}_1^{(t)}$, allowing network 1 to offer progressively lower revenue shares. Observe that the algorithm converges to a revenue share of 29%, which is much lower than the 50% being offered by network 2. Despite offering a lower revenue share, network 1 manages to attract 74% market share. Thus, the use of predictive pricing is giving network 1 a significant advantage.

It may seem surprising that the market share in Figure 1(a) is also falling across iterations. The lowest-quality (i.e., lowest β_i^{Pub}) publishers are essentially being driven from network 1 to network 2. Figure 1(b), which shows the final set of predictive prices $\mathbf{g}_1^{(T)}$, suggests why these publishers leave network 1. Advertisers are being charged very low prices (i.e., low g_{i1}) for traffic from low-quality publishers (i.e., low β_i^{Pub}). Consequently, network 1 offers to pay these low-quality publishers very little for their traffic, causing them to choose network 2 instead.

⁵ From (4), note that $\eta_1 \leq (\sum_i V_i \beta_i^{\text{Pub}}) \kappa_1 \beta_1^{\text{Net}} \equiv \eta_1^{\text{max}}$, which is the maximum possible profit network 1 can attain in any outcome. Thus, in Figures 1(a) and 2 we normalize profits by η_1^{max} .

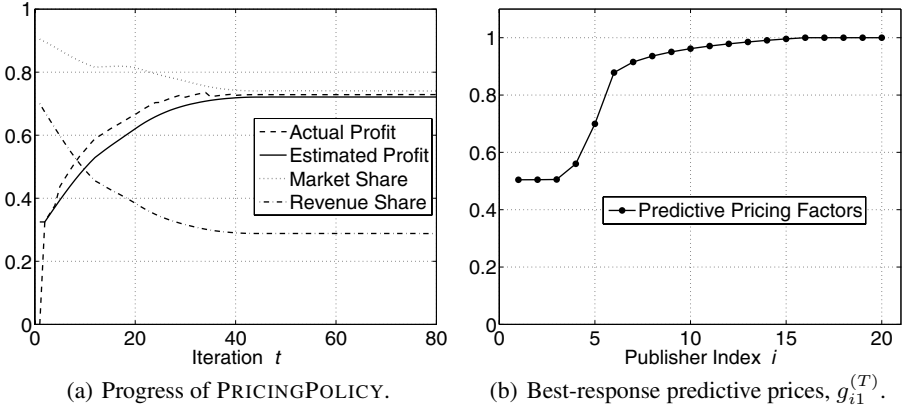


Fig. 1. The effects of predictive pricing




Fig. 2. The effect of network 1’s skill at matching publishers and advertisers (i.e., β_1^{Net})

Observe that the optimal predictive prices in Figure [1\(b\)](#) are increasing in i , and consequently in the conversion rate, β_i^{Pub} . That is, advertisers are being charged less for traffic from publishers whose conversion rate is lower. We ran several other experiments (not discussed here), and found that the optimal g_{i1} was increasing in β_i^{Pub} in every case.

Essentially, a “lemons market” effect is avoided on network 1 as a result of predictive pricing. The lack of low-quality publishers on network 1 raises the average quality of network 1’s traffic, causing advertisers’ bids to increase. The high-quality publishers get paid more per click, and are willing to settle for a lower revenue share as a result.

Our second experiment considers the impact of targeting (i.e., β_j^{Net}) on market outcomes. In particular, if a network is more effective than its competitors at matching publishers with advertisers, does it translate to higher profits for that network? Consider a market with $J = 3$ networks and $I = 20$ publishers. We assume $\beta_i^{\text{Pub}} = 0.000125i^2$ i.e., β_i^{Pub} is quadratic in i , with values ranging from 0.0125% to 5% (there are many low-quality publishers and a few high-quality ones). Networks 2 and 3 are equally skilled at matching i.e., $\beta_2^{\text{Net}} = \beta_3^{\text{Net}} = 1.0$. We assume that $g_{i2} = 20\beta_i^{\text{Pub}}$ (i.e., network 2 uses a

predictive pricing rule that is linear in publisher i 's conversion rate), and network 3 sets $g_{i3} = 1 \forall i$ (i.e., it does not use predictive pricing). Network 2 offers a lower revenue share than network 3, i.e., $h_2 = 0.5$ and $h_3 = 0.6$.

We computed optimal pricing policies for network 1, for various values of β_1^{Net} ranging from 0.7 to 1.3. Recall that β_1^{Net} greater than (less than) 1.0 means that network 1 is better (resp., worse) at targeting than networks 2 and 3. Figure 2 shows network 1's optimal revenue share h_1^* and its resulting profits (normalized by η_1^{max}). As we might expect, network 1 earns higher (lower) profits when β_1^{Net} is higher (resp., lower). From Figure 2 we see that network 1 is able to offer a lower revenue share when β_1^{Net} is higher, since network 1 is generating more conversions for advertisers, causing bids (and consequently publishers' revenues) to increase.

5 Conclusion

Using PRICINGPOLICY, we found that predictive pricing and revenue sharing can be very effective tools for advertising networks to attract publishers and advertisers, especially if their competitors are not using predictive pricing. It is not necessarily optimal to attract as much traffic as possible – quality can be just as important as quantity. Being more effective at matching publishers and advertisers can increase a network's profits, so improving their matching algorithms may be a worthwhile investment for networks.

References

1. Boyd, S., Kim, S.-J., Vandenberghe, L., Hassibi, A.: A tutorial on geometric programming. *Optimization and Engineering* 8(1), 67–127 (2007)
2. Daswani, N., Mysen, C., Rao, V., Weis, S., Gharachorloo, K., Ghosemajumder, S.: Online Advertising Fraud. In: *Crimeware: Understanding New Attacks and Defenses*. Addison-Wesley Professional, Reading (2008)
3. Daswani, N., Stoppelman, M.: The anatomy of clickbot A. In: *Hot Topics in Understanding Botnets (HotBots)*, Usenix (April 2007)
4. Google: The facts about smart pricing. Google Ad Sense Blog (October 2005)
5. Grant, M., Boyd, S.: CVX: Matlab software for disciplined convex programming (web page and software) (June 2008), <http://stanford.edu/~boyd/cvx>
6. Immorlica, N., Jain, K., Mahdian, M., Talwar, K.: Click fraud resistant methods for learning click-through rates. In: Deng, X., Ye, Y. (eds.) *WINE 2005*. LNCS, vol. 3828, pp. 34–45. Springer, Heidelberg (2005)
7. Mungamuru, B., Garcia-Molina, H.: Predictive pricing and revenue sharing. Stanford Infolab Technical Report (July 2008)
8. Mungamuru, B., Weis, S., Garcia-Molina, H.: Should ad networks bother fighting click fraud? (Yes, they should). Stanford Infolab Technical Report (July 2008)
9. Yahoo! A new pricing model rolls out today. Yahoo! Search Marketing Blog (June 2007)

Dual Payoffs, Core and a Collaboration Mechanism Based on Capacity Exchange Prices in Multicommodity Flow Games

Luyi Gui and Özlem Ergun

H. Milton Stewart School of Industrial and System Engineering
Georgia Institute of Technology, Atlanta GA 30332, USA
{lgui3,oergun}@isye.gatech.edu

Abstract. Given a network in which the edge capacities and the commodities are owned by the players, a cooperative multicommodity flow (MCF) game (N, v) can be defined such that $v(S)$, the value of a sub-coalition S , is the maximum profit achievable within S by shipping its commodities through the sub-network owned by its members. In this paper, we study MCF games under a partially decentralized setting where the players make their own routing and resource exchange decisions given a set of capacity prices determined by a central authority.

Keywords: Multicommodity Flow Game, Dual Payoffs, Core, Mechanism Design, Exchange Prices.

1 Introduction

A network is called a collaborative one if its users share with each other the resources on the edges or nodes. Examples of such resource sharing occur in transportation networks where vehicle capacities or in communication networks where bandwidth are shared. Generally speaking, a routing plan which maximizes the social welfare, such as the total network throughput, is desirable. However, the network users aim at maximizing their own revenues. These issues are tackled by cooperative game theory via flow game models. In [4], a flow game is defined on a directed network G in which every edge is owned by a unique player and there exist a unique *source* and *sink*. The value of a sub-coalition is the maximum flow that can be pushed through the network owned by its members. The multicommodity flow (MCF) game is a generalization in that the underlying network has multiple sources and sinks. In this paper, we study a recent model of MCF games with multiple owners of the capacities on a single edge which generalizes the existing models.

A cooperative game theory framework assumes that a centralized overseer manages all the edge capacities and routing decisions, and then distributes the total revenue in a fair manner. However in real life applications which tend to be decentralized to a certain extent, the capacity management and commodity routing decisions are commonly made by individual users. Therefore, it is

important to understand how to design the incentives within the network such that individual players, motivated solely by self-interests, end up following the socially optimal routing plan even under a decentralized setting. In [1], a collaboration mechanism based on capacity exchange prices is proposed to serve this purpose. Players must pay for the edge resources they use at pre-determined unit prices, and at the same time, receive their shares of the money charged on each edge proportional to their ownership levels. These prices alter the final revenue of network users and modify their behaviours. This approach is not only practical, but also has been shown to be effective in [1].

However, cooperation within the grand coalition will be vulnerable if some sub-coalition does not profit enough. In game theory, the notion of the *core* is used to characterize the collection of all fair allocations of the total revenue such that each sub-coalition gets at least as much as it can achieve playing alone. Much effort has been devoted to characterize the core in cooperative games ([2] [14]). One of the well-known results concerns the Linear Production (LP) Game studied by Owen in [7], in which the value of a sub-coalition can be found by solving a linear program. In [7] it is concluded that in a LP-game, an allocation in the core can be obtained by solving a dual linear program, but in general the core is not fully described by such *dual payoffs*.

In this paper, a partially *decentralized* MCF game with multiple owners on a single edge is studied under the collaboration mechanism based on capacity exchange prices introduced in [1]. With a set of *good* exchange prices, each player's net payoff naturally provides an allocation of the value of the grand coalition. This paper focuses on the relationship of such profit allocations arising from a multicommodity network operated under a capacity exchange economy, and the core of the MCF game defined on the network. It shows that the dual payoffs can be achieved by the careful design of the exchange prices but not every fair allocation in the core can be realized in this way. Furthermore, it is proven in [1] that if all edges are uniquely owned, both the social optimum and a fair allocation in the core can be guaranteed under good exchange prices. This paper further generalizes the above result by showing that even with multiple owners of the resources on a single edge, the conclusion still holds under certain conditions. We also give an example that provides insights into how diseconomies arise from multiple ownership of edge capacities and their impact on players' payoffs.

2 Notation and Preliminaries

2.1 A Multicommodity Flow Game

The formulation of the MCF Game used in this paper is mainly from [1]. Formally, a MCF Game is defined on a directed graph $G = (V, E)$. Let N be the set of players. Each commodity type is described as a triplet (o, d, i) such that o and d are the corresponding source and sink, and i is the player who owns demand to be shipped from o to d . Let $d_{(o,d,i)}$ be the amount of the existing demand of commodity (o, d, i) and $r_{(o,d,i)}$ be the unit revenue of it. $D^S = \{(o, d, i) | \forall i \in S\}$ is the demand set of coalition S . Denote D^N as D . Each edge $e \in E$ has a capacity

c_e , which is the maximum amount of flow allowed on e . Player i owns a fraction of γ_e^i of the capacity on edge e . Let $\gamma_e^S = \sum_{i \in S} \gamma_e^i$ be the fraction of capacity on edge e owned by coalition S . Obviously, $\gamma_e^N = \sum_{i \in N} \gamma_e^i = 1 \forall e$.

For modelling convenience, we introduce a fictitious edge (d, o, i) from node d to node o for every commodity type (o, d, i) . Let $f_e^{(o, d, i)}$ be the amount of commodity (o, d, i) shipped through the edge e . Let $IEdges(v) = \{(u, v), \forall u\}$ and $OEdges(v) = \{(v, w), \forall w\}$. The value of the coalition S , $v(S)$, is defined to be the optimal value of the following linear program:

$$P(S) : \quad v(S) = \max \sum_{(o, d, i) \in D^S} f_{(d, o, i)}^{(o, d, i)} r_{(o, d, i)} \quad (1)$$

$$s.t. \quad \sum_{\{e: e \in IEdges(v)\}} f_e^{(o, d, i)} - \sum_{\{e: e \in OEdges(v)\}} f_e^{(o, d, i)} \leq 0 \quad \forall v \in V \quad \forall (o, d, i) \in D^S \quad (2)$$

$$\sum_{(o, d, i) \in D^S} f_e^{(o, d, i)} \leq c_e \gamma_e^S \quad \forall e \in E \quad (3)$$

$$f_{(d, o, i)}^{(o, d, i)} \leq d_{(o, d, i)} \quad \forall (o, d, i) \in D^S \quad (4)$$

$$f \geq 0. \quad (5)$$

After solving the above linear program for every sub-coalition, we define a MCF game (N, \mathbf{v}) , where N is the number of players and \mathbf{v} is the value of coalitions. An allocation in the core of such a game, *i.e.*, the set $C = \{\mathbf{x} \mid \sum_{i \in N} x_i = v(N); \sum_{i \in S} x_i \geq v(S), \forall S \subset N\}$, can be efficiently computed by solving the dual problem of $P(N)$. Mathematically, if we denote the optimal dual solutions associated with constraints (3) and (4) in $P(N)$ as α_e^* , $\forall e \in E$, and $\beta_{(o, d, i)}^*$, $\forall (o, d, i) \in D$, then the dual payoff, as defined in (6), is in the core.

$$\mathbf{x} = \{x_i = \sum_{e \in E} \alpha_e^* c_e \gamma_e^i + \sum_{(o, d, i) \in D^i} \beta_{(o, d, i)}^* d_{(o, d, i)}\}. \quad (6)$$

However, it is a known fact that in a MCF game defined above, the core is not fully characterized by the dual payoffs in general [1]. Moreover, such an allocation scheme requires a central planner to take full control over the operations of the resources, and to allocate the revenue obtained as a result in the grand coalition exactly in the way as described in (6), which is not realistic in most applications.

2.2 A Collaboration Mechanism Based on Exchange Prices

Consider the MCF game defined in 2.1 under a decentralized setting. Players make their own routing decisions and interact with each other via capacity exchanges which are commonly paid in dollars per unit in market applications. In order to achieve the collaborative optimum, a centralized authority intervening minimally may design a price system on the edge capacities to provide enough incentives for the selfish players to choose the social optimal routing for their commodities.

Formally, tag the edge capacities with a set of exchange prices $\bar{cost} = \{cost_e\}$. All players pay for all the edge resources they use, and at the same time get part of the revenue generated by capacity exchanges on every edge proportional to their ownership levels. We model the behaviour of each player by assuming that one makes the routing decision as if he could route all the flow in the network. Such assumption leads to results with desirable properties even in more complex applications. Mathematically, a coalition S solves the following linear program P_{cost}^S to maximize his payoff as in (7) under the same constraints as in $P(N)$.

$$\begin{aligned} \max \quad & \sum_{(o,d,i) \in D^S} f_{(d,o,i)}^{(o,d,i)} r_{(o,d,i)} + \sum_{e \in E} cost_e [\gamma_e^S \sum_{(o,d,i) \notin D^S} f_e^{(o,d,i)} - (1 - \gamma_e^S) \sum_{(o,d,i) \in D^S} f_e^{(o,d,i)}] \quad (7) \\ \text{s.t.} \quad & \text{(2) - (5) in } P(N). \quad (8) \end{aligned}$$

A *good* set of exchange prices makes the social optimum most attractive to every sub-coalition. The following definition captures this idea. Let f^* be the socially optimal routing solution, which is also the optimal solution to $P(N)$.

Definition 1. *A set of exchange prices \bar{cost} is inverse feasible with respect to f^* if f^* is an optimal solution to P_{cost}^S for all sub-coalitions $S \subset N$.*

In [1], it is shown that the set of inverse feasible exchange prices can be identified efficiently by solving a linear program. Specifically, a set of inverse feasible exchange prices must satisfy a set of constraints $\bigcup_{i \in N} I^i$, which contains all the dual constraints and complementary slackness constraints associated with f^* and the problem P_{cost}^i for every player i .¹

Theorem 1. *A set of exchange prices \bar{cost} is inverse feasible if and only if $\bigcup_{i \in N} I^i$ with \bar{cost} as parameter is feasible. Such exchange prices always exist given any MCF game defined in 2.1. (from [1])*

3 Achieving Allocations in the Core by the Mechanism Based on Exchange Prices

With a set of *inverse feasible* exchange prices \bar{cost} , the payoff of player i is

$$x^i = \sum_{(o,d,i) \in D^i} f_{(d,o,i)}^{*(o,d,i)} r_{(o,d,i)} + \sum_{e \in E} cost_e [\gamma_e^i \sum_{(o,d,i) \notin D^i} f_e^{*(o,d,i)} - (1 - \gamma_e^i) \sum_{(o,d,i) \in D^i} f_e^{*(o,d,i)}]. \quad (9)$$

Because the total profits earned from edge capacity exchanges summed over all the players is 0, $\{x^i\}$ as in (9) naturally provides an allocation of the amount $\sum_{(o,d,i) \in D} f_{(d,o,i)}^{*(o,d,i)} r_{(o,d,i)} = v(N)$ as defined in (1)-(5). In other words, inverse feasible exchange prices can serve as a practical tool to realize profit allocations in a MCF game under decentralized settings. This section deals with the problem of how to design the exchange prices such that the resulting profits of players happen to be identical to some allocation in the core.

¹ See the full paper on the author’s webpage for the complete description of $\bigcup_{i \in N} I^i$.

3.1 Dual Payoffs

The dual optimal solutions associated with $P(N)$ are important in the sense that every element of it defines a fair payoff allocation in the core [7]. They are also termed as *market prices* in [3] to emphasize on their economic interpretation. In this section, they are further shown to be inverse feasible in the mechanism based on exchange prices. The full result is introduced in Theorem 2²

Theorem 2. *In a MCF game defined in 2.1, the payoff vector induced under the mechanism based on exchange prices is identical to the dual payoff defined in (6) if the exchange prices are set to be the market prices of edge capacities.*

Theorem 2 is closely related to the economic notion of *competitive equilibrium* studied in the Edgeworth model of exchange economy such as in [2], [12], [13]. Please refer to a recent working paper by the authors for details.

3.2 An Arbitrary Allocation in the Core

Although all the dual payoffs can be achieved under the mechanism based on exchange prices as explained in the previous section, this cannot be slated for an arbitrary allocation in the core. Consider the following simple MCF game.

Example 1. A network has two nodes o and d and an edge e from o to d . There are 2 players. Each player owns 0.5 units of capacity on e . Player I has 0.4 units and Player II has 0.7 units of demands to be shipped from o to d . The unit revenue of either commodity is 1.

By simple calculations we conclude that the exchange price on e is inverse feasible if and only if it is 1. So only the dual payoff $[0.5, 0.5]$ can be achieved under the mechanism based on exchange prices. However, since $v(I) = 0.4$, $v(II) = 0.5$ and $v(I \cup II) = 1$, there are infinitely many solutions in the core that cannot be realized in this way.

To conclude, given a MCF game, let D be the set of dual payoffs as in (6), C be the core, and I be the set of allocations induced by inverse feasible exchange prices. It is always true that $D \subset C$ and $D \subset I$. However, generally speaking, $C \not\subset I$, and it is also easy to find a counter example to $I \subset C$ (see Examples 2 in section 4.2).

4 The Inverse Feasible Exchange Prices and the Core

While inverse feasibility ensures that every selfish player chooses to follow the social optimum when he plays within the grand coalition, the resulting allocation must be in the core to sustain the stability of the grand coalition. In general, it is hard to tell whether the allocation generated by an arbitrary set of inverse feasible exchange prices is in the core. In [1], it is proven that under the assumption of unique ownership of the edge capacities, inverse feasibility automatically guarantees the long-term cooperation of all players, *i.e.*, $I \subset C$.

² See the full paper on the author's webpage for the proof of Theorem 2.

Theorem 3. *If each edge of the network has an unique owner, the payoff vector $x = \{x_i\}$ given by (9) lies in the core of the multicommodity flow game as long as the set of exchange prices in use is inverse feasible. (from [1])*

However, as it is mentioned in [15], diseconomies can arise when an arc is not owned by a unique player, and when diseconomies exist desirable properties of a flow game may be violated. In order to further understand the relationship between the core and the allocations induced by inverse feasible exchange prices under a more general setting, multiple ownership of the capacities on a single edge is allowed here. It is concluded that only a certain type of edges is needed to be uniquely owned in order to reach the conclusion in Theorem 3.

4.1 A Single Sub-coalition Problem

Our analysis begins by studying the payoff to a particular sub-coalition S . First we introduce some notation. let $M_S = \{e | 0 < \gamma_e^S < 1\}$. Let $f(-S)$ denote one feasible solution to $P(N \setminus S)$ (1)-(5) and define the utilized capacity on edge e within the coalition $N \setminus S$ to be $u_e^{-S} = \sum_{(o,d,i) \notin D^S} f(-S)_e^{(o,d,i)}$. If $u_e^{-S} = c_e(1 - \gamma_e^S)$, we say that there are *no excess resources within $N \setminus S$ under $f(-S)$* .

Theorem 4. *If there exists a solution $f(-S)$ under which there are no excess resources within $N \setminus S$ on every edge in M_S , then under every set of inverse feasible exchange prices, $x^S \geq v(S)$.*

Proof. The theorem is proven by showing that, under our assumption, there is a feasible routing for S such that the payoff induced is guaranteed to be no less than $v(S)$. See the full paper on the author’s webpage for the complete proof. \square




Fig. 1. An illustration of the proof of Theorem 4

Fig. 1 illustrates the main idea of the proof to Theorem 4. It also explains the relationship between x^S and $v(S)$, which largely depends on the payments made and gains received from capacity exchanges. Specifically, if a single edge e is owned by multiple players and is also used by at least one of its owners who, after joining the grand coalition, will be paying for his own resources on e which should be free for him, *i.e.*, $ab > 0$ in Fig. 1. Hence diseconomies arise from multiple ownership of edge capacities. The power of the unique ownership

condition lies in the fact that it guarantees any sub-coalition to not pay for edge capacities owned by its members, *i.e.*, $ab = 0$, hence always leads to $I \subset C$.

Meanwhile, it should be noticed that the capacity exchange mechanism enables players to earn extra profits from other parties. The key effect of the no-excess-resources condition in Theorem 4 is to guarantee that there exists a feasible way for the coalition S to route the flows such that the money paid by $N \setminus S$ covers the amount S pays on its own resources, *i.e.*, $ac \geq ab$ in Fig. 1. In this way the diseconomies resulting from multiple ownership can be remedied and $x^S \geq v(S)$ whenever a set of inverse feasible exchange prices is used.

However, no such conclusions can be made in general. In fact, the condition in Theorem 4 is sometimes *necessary* to obtain the conclusion.

4.2 On Excess Resources

The following example illustrates how the payoff to coalition S induced by inverse feasible exchange prices is affected by the excess resources within $N \setminus S$, so that $x^S \geq v(S)$ is violated.

Example 2. A network has only two nodes o and d and an edge e from o to d . Player I owns $c_e \gamma_e^1$ units of the capacity, but his demand exceed the total capacity, *i.e.*, $d_{(o,d,1)} > c_e$. On the other side, Player II owns no shipping demand but only $c_e(1 - \gamma_e^1)$ units of capacity. The unit revenue is $r_{(o,d,1)} = 1$.

Obviously the collaborative optimal solution is to ship c_e units of Player I's demand. Denote the exchange price as $\bar{c}ost$. Consider the objective functions in $P_{\bar{c}ost}^1$ and $P_{\bar{c}ost}^2$. We conclude that $\bar{c}ost$ is inverse feasible iff $0 \leq \bar{c}ost \leq \frac{1}{1-\gamma_e^1}$.

Calculate Player I's payoff under an inverse feasible exchange price $\bar{c}ost$.

$$x^1 = c_e[1 - \bar{c}ost(1 - \gamma_e^1)] \quad (10)$$

$$= v(1) + c_e(1 - \gamma_e^1)(1 - \bar{c}ost). \quad (11)$$

By (10), $x^1 < v(1)$ if $\bar{c}ost > 1$. Since Player I owns a positive fraction of capacity on e , $1 < \frac{1}{1-\gamma_e^1}$. Hence, any inverse feasible exchange price in the nonempty interval $(1, \frac{1}{1-\gamma_e^1}]$ leads to $x^1 < v(1)$.

Fig. 2 illustrates how Player I accumulates his profit as the shipping amount increases if $\bar{c}ost > 1$. The line segment ob demonstrates the situation in which all resources in use are exchanged at a price of $\bar{c}ost$. By (10), Player I earns a unit profit of $1 - \bar{c}ost(1 - \gamma_e^1)$ from his shipping business, hence the slope of ob is $1 - \bar{c}ost(1 - \gamma_e^1)$ which is nonnegative by inverse feasibility. From this point of view, Player I benefits from the excess resources owned by Player II. However, (11) implies that the profit accumulation process can also be understood in another way. At first, Player I plays alone to use up his own resources with a unit profit 1, which is the slope of the line segment oa . Then he joins the coalition and ships more using the capacities owned by Player II. By (11), the unit profit he actually earns after joining the grand coalition, which is the slope of ab , is $1 - \bar{c}ost$. Hence, the excess resources owned by Player II in fact undermines the total profit of Player I since $1 - \bar{c}ost$ is negative when $\bar{c}ost > 1$.




Fig. 2. Payoff accumulation of player I in Example 2, where x denotes the shipping amount as well as the resources used on e , and y denotes the profit earned

Notice the condition in Theorem 4 is *necessary* in Example 2 to guarantee $x^{\{I\}} \geq v(\{I\})$ when an arbitrary inverse feasible exchange price is used.

4.3 A Sufficient Condition for $I \subset C$

By considering every sub-coalition using Theorem 4, we can easily derive the following sufficient condition for every set of inverse feasible exchange prices to induce a payoff vector in the core.

Theorem 5. *If $\forall i \in N, \forall e \in \bigcup_{i \in N} M_i$, there exists a feasible solution $f(i)$ to $P(i)$, under which there are no excess resources within $\{i\}$ on e , then every set of inverse feasible exchange prices leads to a payoff allocation vector in the core.*

Theorem 5 requires that the edges are uniquely owned only if there exists some player who cannot use up his own resources on them by his shipping demands. As the assumption of unique ownership of edge capacities implies $\bigcup_{i \in N} M_i = \emptyset$, Theorem 3 is indeed an extreme case of Theorem 5.

5 Open Problems

In section 3.1, it has been shown that the implementation of the collaborative mechanism with market prices as exchange prices promotes cooperation among selfish players. However, this approach sometimes results in allocations with undesirable properties. Because the market price of every partially used edge should be zero, players who own resources on those edges are very likely to provide *free* service to edge users if the exchange prices are indeed set to be the market price. This phenomenon might deviate players' behaviours from the social optimum, thus is a serious drawback. Our present study finds a cutting plane algorithm to correct the exchange prices so that the problem is tackled. We also consider the resulting allocations and study their properties.

Another open problem concerns the fact that there are many examples in which $C \subset I$ and the core can be fully described by inverse feasibility. The open problem is to characterize the situation under which this desirable result is true.

References

1. Agarwal, R., Ergun, Ö.: Mechanism Design for a Multicommodity Flow Game in Service Network Alliances. *Operations Research Letters* (to appear)
2. Debreu, G., Scarf, H.: A Limit Theorem on the Core of an Economy. *International Economic Review* 4(3), 235–246 (1963)
3. Engelbrecht-Wiggans, R., Granot, D.: On Market Prices in Linear Production Games. *Mathematical Programming* 32, 366–370 (1985)
4. Kalai, E., Zemel, E.: Totally Balanced Games and Games of Flow. *Mathematics of Operations Research* 7(3), 476–478 (1982)
5. Kalai, E., Zemel, E.: Generalized Network Problems Yielding Totally Balanced Games. *Operations Research* 30(5), 998–1008 (1982)
6. Markakis, E., Saberi, A.: On the Core of the Multicommodity Flow Game. *Decision Support Systems* 39, 3–10 (2005)
7. Owen, G.: On the Core of Linear Production Games. *Mathematical Programming* 9, 358–370 (1975)
8. Özener, Ö., Ergun, Ö.: Allocating Costs in a Collaborative Transportation Procurement Network. *Transportation Science* 42(2), 146–165 (2008)
9. Papadimitriou, C.: Algorithms, Games, and the Internet. *Annual ACM Symposium on the Theory of Computing*, 749–753 (2001)
10. Samet, D., Zemel, E.: On the Core and Dual Set of Linear Programming Games. *Mathematics of Operations Research* 9(2), 309–316 (1984)
11. Sánchez-Soriano, J., López, M., García-Jurado, I.: On the Core of Transportation Games. *Mathematical Social Sciences* 41, 215–225 (2001)
12. Shapley, L.S., Shubik, M.: Pure Competition, Coalitional Power and Fair Division. *International Economic Review* 10(3), 337–361 (1969)
13. Shapley, L.S., Shubik, M.: Concepts and Theories of Pure Competition. In: Shubik, M. (ed.) *Honor of Oskar Morgenstern. Essays in Mathematical Economics*, pp. 63–79. Princeton University Press, Princeton (1967)
14. Scarf, H.: The Core of an N Person Game. *Econometrica* 35(1), 50–69 (1967)
15. Sounderpandian, J.: Totally Monotonic Games and Flow Games. *Operations Research Letters* 36, 165–167 (2008)

Graphical Congestion Games^{*}

Vittorio Bilò¹, Angelo Fanelli², Michele Flammini²,
and Luca Moscardelli^{2,3}

¹ Dipartimento di Matematica, University of Salento
Provinciale Lecce-Arnesano, P.O. Box 193, 73100 Lecce, Italy
vittorio.bilo@unile.it

² Dipartimento di Informatica, University of L'Aquila
Loc. Vetoio, Coppito, 67100 L'Aquila, Italy
{angelo.fanelli,flammini,moscardelli}@di.univaq.it

³ Dipartimento di Informatica ed Applicazioni “R. M. Capocelli”,
University of Salerno
Via Ponte don Melillo, 84084 Fisciano (SA), Italy
moscardelli@dia.unisa.it

Abstract. We consider congestion games with linear latency functions in which each player is aware only of a subset of all the other players. This is modeled by means of a social knowledge graph G in which nodes represent players and there is an edge from i to j if i knows j . Under the assumption that the payoff of each player is affected only by the strategies of the adjacent ones, we first give a complete characterization of the games possessing pure Nash equilibria. We then investigate the impact of the limited knowledge of the players on the performance of the game. More precisely, given a bound on the maximum degree of G , for the convergent cases we provide tight lower and upper bounds on the price of stability and asymptotically tight bounds on the price of anarchy. All the results are then extended to load balancing games.

Keywords: Algorithmic Game Theory, Nash Equilibrium, Price of Anarchy, Price of Stability, Congestion Games, Social Knowledge.

1 Introduction

Congestion games constitute a well-known class of non-cooperative games in which a set of facilities E is available to the players and the strategy set of each player i can be any $S_i \subseteq 2^E$. The cost of each facility $e \in E$ (usually called the latency of e) is a function of the number of players using e and the latency experienced by each player i is the sum of the latencies of all the facilities used by i .

Congestion games have been introduced by Rosenthal [17] in 1973. By defining an elegant potential function he showed that they always possess (and always

^{*} Work supported by the European Union under IST FET Integrated Project AEO-LUS (IST-015964).

converge to) pure Nash equilibria [16]. In the last decade, they have come across the analysis of the Computer Science community with the purpose of characterizing the complexity of computing their pure Nash equilibria and evaluating their suboptimality in terms of price of stability and anarchy. One of the most interesting and studied special cases is the class of the linear congestion games, in which the latency of each resource e is defined as a linear function of the number of players using e .

A major concern related to the model of multiplayer games is given by the fact that it is always assumed that each player knows all the context parameters, is aware of the existence of all the other ones and, more important, of the consequences of their choices in the definition of her payoff. Such interdependent effects clearly represent the core of each non-cooperative game, but with a huge number of players in highly dynamic and distributed environments such a global knowledge might be unfeasible. Therefore, Harsanyi in his two pioneering works on games with incomplete information [11][12] introduced Bayesian games, where players may have different types and are uncertain about each others types according to a probability distribution over all possible type profiles. Bayesian congestion games have then been studied in [2][7][8][9][10]. Recently, Koutsoupias et al. applied a similar model to load balancing games [13].

Along the line of incomplete information, a further realistic step is to assume that each player is aware only of the strategies played by a subset of players representing somehow her neighborhood. The idea of exploiting and modeling the locality of mutual influences constitutes the basis of another famous class of games, called graphical games. These games have been introduced in [14] with the main purpose of providing a succinct form for representing non-cooperative games in the cases in which the payoff of a player is influenced only by the choices performed by a relative small subset of the players in the game. General non-cooperative games, in fact, are usually represented in normal form by using n matrices each of size m^n , where n is the number of players and m is the number of strategies available for each player (assuming for simplicity that all the strategy sets have the same cardinality). In graphical games, a game is represented by a *social knowledge graph* G with n nodes and a set of n matrices. Each player corresponds to a node in the graph and the set of her neighbors in G to all players directly influencing with their choices her payoff. When the maximum degree in G is small, say a constant Δ , each of the payoff matrices will have size equal to m^Δ .

In this paper we push this idea a little further by associating a given game a social knowledge graph in such a way that the payoff matrices of the resulting game will be determined on the basis of those of the initial game and of the neighborhoods yielded by the topology of the social knowledge graph.

1.1 Related Work

After Rosenthal's seminal paper [17], in [15] it was shown that congestion games are isomorphic to exact potential games. The price of anarchy of general congestion games is known to be arbitrarily high and cannot be better than n (the

number of players in the game). Thus, a lot of research has been devoted in the study of linear congestion games with respect to two different social functions: the total latency, that is the sum of all the latencies, and the maximum experienced latency. In [14] it is shown, among the various results, that in case of total latency social cost the price of anarchy of pure Nash equilibria is $5/2$, while for mixed Nash equilibria or pure Nash equilibria of weighted players it is 2.618. Moreover, in [4] it is also shown that the price of anarchy of the maximum latency social cost is $\Theta(\sqrt{n})$. The price of stability of linear congestion games has been studied in [5] where it was shown that for the total latency social cost it is between $1 + \frac{1}{\sqrt{3}} \approx 1.577$ and 1.6. Such a value has been fixed to $1 + \frac{1}{\sqrt{3}}$ in [3].

1.2 Our Contribution

We analyze the impact of the social knowledge among the players on congestion games with linear latency functions. Under the assumption that the payoff of each player is affected only by the strategies of the adjacent ones in the social knowledge graph, we first give a complete characterization of the games possessing pure Nash equilibria. Namely, if the social graph G is undirected the game is an exact potential game and thus isomorphic to a classical congestion game. As a consequence, it always converges and possesses Nash equilibria. We then show that if G is directed an equilibrium is not guaranteed to exist, but the game is always convergent and an equilibrium can be found in polynomial time if G is acyclic, even if finding the best equilibrium remains an intractable problem.

We then investigate the impact of the limited knowledge of the players on the performance of the game. More precisely, given a bound Δ on the maximum degree of G , for all the convergent cases we bound the respective prices of stability and anarchy.

Such results are determined for the social cost functions (i.) total presumed latency, that is the one the players believe to pay due to the fact that they are only aware of the existence of their neighbors, (ii.) maximum presumed latency, (iii.) total perceived latency, i.e. actually experience due to all and not only the known players using the same facilities and (iv.) maximum perceived latency.

All the results are then extended to load balancing games, that is congestion games in which every pure strategy consists of a single facility.

We provide tight and asymptotically tight bounds for 31 of the 32 arising cases. Such results are summarized in Tables 1 and 2, where $\mathcal{G}(\Delta)$ and $\vec{\mathcal{A}}\mathcal{G}(\Delta)$

Table 1. Presumed latencies: bounds for congestion and load balancing games

	PoS^{sum}, PoS^{max}	PoA^{sum}, PoA^{max}
$\mathcal{G}(\Delta)$	2, $\Theta(\Delta + 1)$	$\Theta(\Delta + 1), \Delta + 1$
$\vec{\mathcal{A}}\mathcal{G}(\Delta)$	$\Theta(\Delta + 1), \Delta + 1$	$\Theta(\Delta + 1), \Delta + 1$

Table 2. Perceived latencies: bounds for congestion and load balancing games

	Congestion games		Load balancing games	
	PoS^{sum}, PoS^{max}	PoA^{sum}, PoA^{max}	PoS^{sum}, PoS^{max}	PoA^{sum}, PoA^{max}
$\mathcal{G}(\Delta)$	$n, n \div n\sqrt{\Delta + 1}$	$\Theta(n(\Delta + 1))$	$n, \Theta(n)$	$\Theta(n)$
$\vec{\mathcal{A}}\mathcal{G}(\Delta)$	$\Theta(n(\Delta + 1))$	$\Theta(n(\Delta + 1))$	$\Theta(n)$	$\Theta(n)$

are the classes of all the symmetric (or undirected) and directed acyclic social graphs with maximum node degree bounded by Δ , respectively.

In some sense our result seems contradictory: the more players know, the worse the prices of stability and anarchy are. This is actually true for Δ approaching to n , with the worst case being $\Delta = n/2$. Note also that, in all our lower bound constructions, Δ is upper bounded by a constant fraction of the number of players. However, in the case of perceived latencies, it is possible to see that if every player knows at least a certain number of other players Δ , that is if Δ is a lower bound on the minimum degree, then as $\Delta \geq n/2$ approaches n all the prices gradually tend to $O(n)$. For $\Delta < n/2$ all our results coincide with the ones where Δ is the maximum degree, as dummy players can be added using only dummy facilities of null cost so as to induce social graphs of minimum degree at least Δ . Another crucial observation is that better bounds can be obtained for specific social graphs. In fact, for the undirected complete graph constant bounds derive directly from the classical congestion game.

Besides the particular results, our framework is particularly effective in modelling situations in which users choices are done a priori or modified during a preprocessing phase under partial knowledge of the arising system performance in the following operating phase, during which preemption or alternative strategy selections are not allowed or yield excessive costs. As an example, we have particular routing protocols, real traffic networks and in general contexts in which users must subscribe conflicting services in advance. More in general, the framework can be applied to all non-cooperative games in which a complete knowledge among the players cannot be achieved or can be guaranteed up to a limited extent. Thus, we believe that it will possibly capture future research attention.

2 Model

A *graphical congestion game* is defined by a tuple $\mathcal{H} = \langle G = (N, M), E, (S_i)_{i \in N}, (f_e)_{e \in E} \rangle$ where $G = (N, M)$ is a directed graph, called *social knowledge graph*, $N = \{1, \dots, n\}$ is a set of n players, E is a set of facilities, $S_i \subseteq 2^E$ is a set of pure strategies for player i , each consisting of a set of facilities, and f_e is the latency function for the facility e depending on the number of players using e .

Roughly speaking, the graph G defines the social knowledge among the players. In fact, the players j adjacent to a player i in G , that is such that $(i, j) \in M$, are all and only the ones whose strategy choices are known by i and can potentially influence her strategy choice. This models the fact that a player may be not aware of the choices of all the other players, and thus her choices are affected only by the strategies played by a restricted neighborhood. Clearly, any complete symmetric social graph induces a classical congestion game.

Let $\delta_i(G)$ be the (out-)degree of player i in G , and $\vec{\mathcal{G}}(\Delta)$ be that class of all the social graphs with maximum node degree bounded by Δ . Moreover, let $\mathcal{G}(\Delta)$ and $\mathcal{AG}(\Delta)$ be the subclasses of the symmetric (or undirected) graphs and of the directed acyclic graphs in $\vec{\mathcal{G}}(\Delta)$, respectively.

The pure strategy profile (state) set of the game is $S = \prod_{i \in N} S_i$. Given a pure strategy profile $s = (s_1, s_2, \dots, s_n) \in S$, we denote with $G_e(s)$ the subgraph of G induced by the set of players using facility e , i.e. $G_e(s) = (N_e(s), M_e(s))$ where $N_e(s) = \{i \in N : e \in s_i\}$ and $M_e(s) = \{(i, j) \in M : i, j \in N_e(s)\}$. Let $n_e(s) = |N_e(s)|$ and $m_e(s) = |M_e(s)|$ be the number of nodes and arcs in $G_e(s)$ respectively, and $\delta_e^i(s)$ be the degree of node i in $G_e(s)$. The cost of player i in the strategy profile s is $c_i(s) = \sum_{e \in s_i} f_e(n_e^i(s))$, where $n_e^i(s)$ is the number of nodes adjacent to i in $G_e(s)$, i included, that is $n_e^i(s) = \delta_e^i(s) + 1 = |\{j \in N_e(s) : j = i \vee (i, j) \in M_e(s)\}|$.

We focus on the case in which the latencies of the facilities are linear functions with nonnegative coefficients, i.e. $f_e(x) = \alpha_e x + \beta_e$, with $\alpha_e, \beta_e \geq 0$, for any $e \in E$. Moreover, we consider four different social cost functions of a strategy profile s : the *total presumed social cost* $C_{PR}^{sum}(s) = \sum_{i \in N} c_i(s) = \sum_{e \in E} \sum_{i: e \in s_i} f_e(n_e^i(s))$ given by the sum of all the players' costs, the *maximum presumed social cost* $C_{PR}^{max}(s) = \max_{i \in N} c_i(s) = \max_{i \in N} \sum_{e \in s_i} f_e(n_e^i(s))$, that is the maximum players' cost, the *total perceived social cost* $C_{PE}^{sum}(s) = \sum_{e \in E} \sum_{i: e \in s_i} f_e(n_e(s)) = \sum_{e \in E} n_e(s) f_e(n_e(s))$, i.e., the total cost due to the actual congestion of the facilities, and finally the *maximum perceived social cost* $C_{PE}^{max}(s) = \max_{i \in N} \sum_{e \in s_i} f_e(n_e(s))$.

The objective of a player i is to choose the pure strategy minimizing her own cost, given the strategy of the players adjacent to i in G . Given a strategy profile $s = (s_1, s_2, \dots, s_i, \dots, s_n)$, we denote as $s \oplus s'_i = (s_1, s_2, \dots, s'_i, \dots, s_n)$ the strategy profile obtained from s if player i changes her strategy from s_i to s'_i . A (pure) *Nash equilibrium* is a pure strategy profile such that no player can reduce her cost by seceding in favor of a better strategy, given the strategies of the other players. More formally, a Nash equilibrium is a pure strategy profile $s = (s_1, s_2, \dots, s_n)$ such that $\forall i \in N$ and strategy $s'_i \in S_i$, it holds $c_i(s) \leq c_i(s \oplus s'_i)$. Denoting with \mathcal{N} the set of all the possible Nash equilibria, the *price of anarchy (PoA)* of a game \mathcal{H} for the total presumed latency social cost is defined as the worst case ratio among the Nash versus optimal performance, i.e., $PoA_{PR}^{sum}(\mathcal{H}) = \frac{\max_{s \in \mathcal{N}} C_{PR}^{sum}(s)}{OPT_{PR}^{sum}(\mathcal{H})}$, where $OPT_{PR}^{sum}(\mathcal{H}) = \min_{s \in S} C_{PR}^{sum}(s)$. On the other hand the *price of stability (PoS)* of \mathcal{H} is defined as the best case ratio among the Nash versus optimal performance, i.e., $PoS_{PR}^{sum}(\mathcal{H}) = \frac{\min_{s \in \mathcal{N}} C_{PR}^{sum}(s)}{OPT_{PR}^{sum}(\mathcal{H})}$. $PoA_{PR}^{max}(\mathcal{H})$, $PoS_{PE}^{sum}(\mathcal{H})$,

$PoA_{PE}^{max}(\mathcal{H})$, $PoS_{PR}^{max}(\mathcal{H})$, $PoS_{PE}^{sum}(\mathcal{H})$ and $PoS_{PE}^{max}(\mathcal{H})$ can be defined accordingly. In the following, when clear from the context, we will drop the indices *sum* and *max* and the argument \mathcal{H} from the notation.

3 Existence, Convergence and Complexity

In this section, we focus on the existence and convergence to Nash equilibria and completely characterize the complexity of finding a generic Nash equilibrium and an optimal one. We first consider undirected social knowledge graphs.

Theorem 1. *Every graphical linear congestion game defined over an undirected social graph is an exact potential game, and thus always converges to a Nash equilibrium.*

Proof. Given the strategy profile $s = (s_1, s_2, \dots, s_n)$, the potential function establishing the result is $\Phi(s) = \sum_{e \in E} F_e(s)$, where $F_e(s) = \alpha_e (m_e(s) + n_e(s)) + \beta_e n_e(s)$.

Let i be a player reducing her cost by changing her strategy from s_i in s to s'_i , thus yielding a new strategy profile $s' = s \oplus s'_i$. The change of the potential function, $\Phi(s') - \Phi(s)$, is then equal to $\sum_{e \in s'_i \setminus s_i} (F_e(s') - F_e(s)) - \sum_{e \in s_i \setminus s'_i} (F_e(s) - F_e(s')) = \sum_{e \in s'_i \setminus s_i} (\alpha_e (\delta_e^i(s') + 1) + \beta_e) - \sum_{e \in s_i \setminus s'_i} (\alpha_e (\delta_e^i(s) + 1) + \beta_e) = \sum_{e \in s'_i \setminus s_i} (\alpha_e n_e^i(s') + \beta_e) - \sum_{e \in s_i \setminus s'_i} (\alpha_e n_e^i(s) + \beta_e) = c_i(s') - c_i(s)$. \square

We now turn our attention to directed social knowledge graphs, by first showing that each game converges to a Nash equilibrium and an equilibrium can be efficiently determined.

Theorem 2. *Each graphical congestion game defined over a directed acyclic social graph converges to a Nash equilibrium. Moreover, there always exists a sequence of at most n best replies which can be computed in polynomial time ending to a Nash equilibrium.*

Notice that the above theorem holds for any latency function, and this is in contrast with the hardness for the undirected case, where the PLS completeness follows from [6] by restricting to complete social graphs. However, the following theorem shows that for directed social graphs determining an equilibrium with minimum social cost remains an intractable problem. We simply refer to social cost, since slight modifications of the same reduction apply to all the four social functions.

Theorem 3. *Given a graphical linear congestion game with directed acyclic social graph and an integer $k > 0$, determining whether there exists a Nash equilibrium with social cost at most k is an NP-complete problem.*

On the other hand, if the social graph contains cycles, Nash equilibria might not exist.

Theorem 4. *There exists a graphical linear congestion game defined over a directed social graph not admitting any Nash equilibrium.*

4 Presumed Social Cost

In this section we first analyze the prices of stability and anarchy with respect to the social cost $C_{PR}^{sum}(s) = \sum_{i \in N} c_i(s)$, and then sketch how to extend our results to the max presumed latency social cost.

We provide matching and asymptotically matching upper and lower bounds for the games defined over the social graphs always guaranteeing the convergence to Nash equilibria, i.e. undirected and directed acyclic graphs.

We first focus on graphical games defined over undirected social graphs, and, by exploiting the potential function defined in Theorem 4, we prove that the price of stability is equal to 2, regardless of the maximum degree of the social graph.

Theorem 5. *For any graphical linear congestion game \mathcal{H} defined over an undirected social graph, $PoS_{PR}^{sum}(\mathcal{H}) \leq 2$.*

Theorem 6. *For any $\epsilon > 0$ there exists a graphical linear congestion game \mathcal{H} with an arbitrarily large number of players defined over an undirected social graph such that $PoS_{PR}^{sum}(\mathcal{H}) \geq 2 - \epsilon$.*

As the following theorems state, both the price of anarchy of graphical linear congestion games defined over undirected social graphs and the price of stability of graphical linear congestion games defined over directed acyclic social graphs cannot be upper bounded by a constant, but are linear in Δ . Notice that if the social graph has no edges, the price of anarchy is trivially 1, since in any Nash equilibrium all the players experience the lowest possible cost. Thus, for clarity of presentation, in the following we focus on graphical games defined on graphs belonging to $\mathcal{G}(\Delta) \cup \overrightarrow{\mathcal{AG}}(\Delta)$ with $\Delta > 0$.

Theorem 7. *For any $\Delta > 0$ there exists a graphical linear congestion game \mathcal{H} with an arbitrarily large number of players defined over a social graph $G \in \mathcal{G}(\Delta)$ such that $PoA_{PR}^{sum}(\mathcal{H}) \geq \frac{2\Delta+1}{3}$.*

Theorem 8. *For any $\Delta > 0$ there exists a graphical linear congestion game \mathcal{H} with an arbitrarily large number of players defined over a social graph $G \in \overrightarrow{\mathcal{AG}}(\Delta)$ such that $PoS_{PR}^{sum}(\mathcal{H}) \geq \frac{\Delta+1}{2}$.*

The following theorem provides asymptotically matching upper bounds on the price of anarchy for undirected social graphs, and on the prices of stability and anarchy for directed acyclic social graphs.

Theorem 9. *Given any graphical linear congestion game \mathcal{H} defined over a social graph $G \in \mathcal{G}(\Delta) \cup \overrightarrow{\mathcal{AG}}(\Delta)$, $PoA_{PR}^{sum}(\mathcal{H}) \leq 1 + \Delta$.*

By combining Theorem 9 with Theorem 7 and Theorem 8, respectively, we have that the price of anarchy for undirected and directed acyclic social graphs is $\Theta(\Delta + 1)$, as well as the price of stability for directed acyclic social graphs.

Let us finally sketch how to extend our results to the max presumed latency social function.

Exactly matching bounds on the price anarchy can be shown by very similar arguments and constructions; the same holds for the price of stability of games defined over directed acyclic social graphs (see Table I). Different results hold for the price of stability of games defined over undirected social graphs. In fact, the potential function argument of Theorem 5 cannot be applied to establish the same upper bound. However, while a trivial $\Delta + 1$ upper bound is given by the price of anarchy, an asymptotically matching lower bound is established in the following theorem.

Theorem 10. *For any $\Delta > 0$ there exists a graphical linear congestion game \mathcal{H} with an arbitrarily large number of players defined over a social graph $G \in \mathcal{G}(\Delta)$ such that $PoS_{PR}^{max}(\mathcal{H}) \geq \frac{\Delta+1}{2}$.*

5 Perceived Social Cost

In this section we first analyze the prices of stability and anarchy with respect to the social cost $C_{PE}^{sum}(s) = \sum_{e \in E} n_e(s) f_e(n_e(s))$, in which we are interested in minimizing the sum of the latencies actually perceived by the players. We then sketch how to extend our results to the maximum perceived latency social function.

Again we provide matching and asymptotically matching upper and lower bounds for the games defined over the social graphs always guaranteeing the convergence to Nash equilibria, i.e. undirected and directed acyclic.

We first focus on the price of stability of graphical games defined over undirected social graphs. By the same potential function technique of Theorem 5, we prove that it is equal to n regardless of the maximum degree of the social graph.

Theorem 11. *For any graphical linear congestion game \mathcal{H} defined over an undirected social graph, $PoS_{PE}^{sum}(\mathcal{H}) \leq n$.*

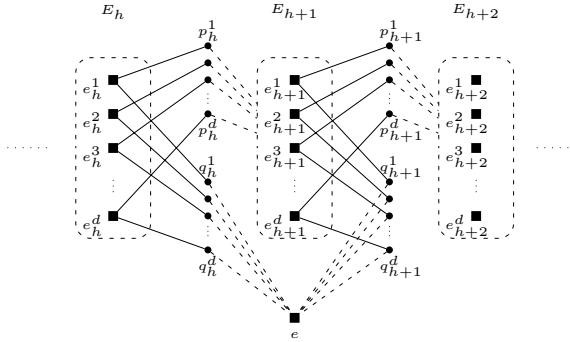
Theorem 12. *For any $\epsilon > 0$ there exists a graphical linear congestion game \mathcal{H} with an arbitrarily large number of players defined over an undirected social graph such that $PoS_{PE}^{sum}(\mathcal{H}) \geq n - \epsilon$.*

We now show that there exist graphical congestion games defined on undirected social graphs G for which the price of anarchy is $\Omega(n(\Delta+1))$, and that there exist graphical congestion games defined on directed acyclic social graphs for which this bound holds even for the price of stability. Notice that if the social graph has no edges, by the same arguments in the proof of Theorem 12, the prices of stability and anarchy are lower bounded by n . Thus, for clarity of presentation, in the following again we focus on graphical games defined on graphs belonging to $\mathcal{G}(\Delta) \cup \overrightarrow{\mathcal{AG}}(\Delta)$ with $\Delta > 0$.


Theorem 13. *For any $\Delta > 0$ there exists a graphical linear congestion game \mathcal{H} with an arbitrarily large number of players defined over a social graph $G \in \mathcal{G}(\Delta)$ such that $PoA_{PE}^{sum}(\mathcal{H}) = \Omega(n(\Delta + 1))$.*

Proof. Let us consider a graphical linear congestion game \mathcal{H} with $n = 2kd$ players, where $d = \lfloor \Delta/3 \rfloor$ and k is an arbitrarily large integer. The social graph of the game is $G = (N_0 \cup N'_0 \cdots \cup N_{k-1} \cup N'_{k-1}, M_0 \cup M'_0 \cup \cdots \cup M_{k-1} \cup M'_{k-1})$ (see Figure 1(b)), where for every $h = 0, \dots, k-1$, $N_h = \{p_h^1, \dots, p_h^d\}$, $N'_h = \{q_h^1, \dots, q_h^d\}$, $M_h = \{\{p_h^i, p_{(h+1)}^j \}_{\text{mod } k} \mid i, j \in \{1, \dots, d\}\}$ and $M'_h = \{\{p_h^i, q_{(h+1)}^j \}_{\text{mod } k} \mid i, j \in \{1, \dots, d\}\}$. The set of facilities is $E = \{e\} \cup E_0 \cup \dots \cup E_{k-1}$, where for every $h = 0, \dots, k-1$ $E_h = \{e_h^1, \dots, e_h^d\}$, and for every $l = 1, \dots, d$ the latency functions are $f_{e_h^l}(x) = x$; moreover, $f_e(x) = dx$. Each player p_h^l has the strategy set $\{\{e_h^l\}, E_{(h+1)} \}_{\text{mod } k}$ and each q_h^l has the strategy set $\{\{e_h^l\}, \{e\}\}$ (see Figure 1(a)). For the sake of clearness, we refer to the first strategy of each strategy set as the *small* strategy, and to the second one as the *big* strategy.

Clearly, if each player chooses her small strategy, the achieved social cost is equal to $4kd$; thus $OPT_{PE} \leq 4kd$.



(a) The optimal assignment and a Nash equilibrium (in dashed lines)



(b) The social graph of \mathcal{H}

Fig. 1. Lower bound of Theorem 13

Moreover, the strategy profile \bar{s} in which each player chooses her big strategy is a Nash equilibrium; in fact, each player p_h^l has a cost equal to d in \bar{s} , and she would experience a presumed cost equal to $d + 1$ by changing her strategy. Moreover, each player q_h^l has a cost equal to d in \bar{s} , and she would experience a presumed cost equal to $d + 1$ by changing her strategy. The perceived social cost of \bar{s} is lower bounded by the sum of the latencies on facility e , that is equal to $d(kd)^2 = k^2d^3$. Therefore, the price of anarchy of \mathcal{H} is at least $\Omega(kd^2) = \Omega(n(\Delta + 1))$. \square

Theorem 14. *For any $\Delta > 0$ there exists a graphical linear congestion game \mathcal{H} with an arbitrarily large number of players defined over a social graph $G \in \overrightarrow{\mathcal{AG}}(\Delta)$ such that $PoS_{PE}^{sum}(\mathcal{H}) = \Omega(n(\Delta + 1))$.*

The following theorem provides asymptotically matching upper bounds on the price of anarchy for undirected social graphs, and on the prices of stability and anarchy for directed acyclic social graphs.

Theorem 15. *Given any graphical linear congestion game \mathcal{H} defined over a social graph $G \in \mathcal{G}(\Delta) \cup \overrightarrow{\mathcal{AG}}(\Delta)$, $PoA_{PE}^{sum}(\mathcal{H}) = O(n(\Delta + 1))$.*

Notice that by combining Theorem 13 and Theorem 14 with Theorem 15 we can derive a price of anarchy $\Theta(n(\Delta + 1))$ for undirected social graphs and prices of stability and anarchy $\Theta(n(\Delta + 1))$ for directed acyclic social graphs. Moreover, for $\Delta = \frac{n}{2}$, Theorem 15 combined with the following theorem provides an exactly matching bound to the price of anarchy for undirected social graphs, expressed in the number of players.

Theorem 16. *There exists a graphical linear congestion game \mathcal{H} with an arbitrarily large number of players defined over an undirected social graph such that $PoA_{PE}^{sum}(\mathcal{H}) = \frac{n^2}{4}$.*

Again, for the maximum perceived latency social function the same bounds on the price anarchy can be shown by very similar arguments and constructions; the same holds for the price of stability of games defined over directed acyclic social graphs (see Table 2). The potential function argument of Theorem 11 cannot be applied to establish the same upper bounds on the the price of stability. However, while an $\Omega(n)$ lower bound still holds by the same construction of Theorem 12, the following theorem improves upon the trivial $O(n(\Delta + 1))$ upper bound given by the price of anarchy. It exploits a novel technique showing that a high price of stability for the maximum perceived social function would imply a high price of stability for the total perceived social function, thus contradicting Theorem 11.

Theorem 17. *Given any graphical linear congestion game \mathcal{H} defined over a social graph $G \in \mathcal{G}(\Delta)$, $PoS_{PE}^{max}(\mathcal{H}) \leq n\sqrt{\Delta + 1}$.*

6 Load Balancing Games

In this section we sketch how to extend our results to load balancing games, that is congestion games in which every pure strategy consists of a single facility.

Almost all the previous results implicitly consider load balancing instances, with the exception of the lower bounds established in Theorem 13 and 14 for the perceived latencies. Both for the total and the maximum perceived latency, lower bounds equal to n come directly from the respective price of stability constructions, that in fact are load balancing ones and hold for games defined over both undirected and directed acyclic social graphs. The following two theorems show corresponding asymptotically matching upper bounds.

Theorem 18. *Given any graphical linear load balancing game \mathcal{H} defined over a social graph $G \in \mathcal{G}(\Delta) \cup \overrightarrow{\mathcal{AG}}(\Delta)$, $PoA_{PE}^{sum}(\mathcal{H}) \leq 8n$.*

Proof. Given any Nash Equilibrium s , for every $i = 1, \dots, n$ let $c_i(s)$, or simply c_i , be the latency of player i in s , and c_i^* be the latency of i in a fixed optimal strategy profile s^* . Without loss of generality, let us assume that players are non-increasingly ordered with respect to the ratio between the latency at equilibrium and the one at optimum, i.e. $\frac{c_1}{c_1^*} \geq \dots \geq \frac{c_n}{c_n^*}$. Consider the largest index r such that $\frac{c_r}{c_r^*} \geq \frac{PoAPE}{2}$. Then $\sum_{j=1}^r c_j \geq \frac{\sum_{j=1}^n c_j}{2} = \frac{C_{PE}(s)}{2}$, as otherwise $\sum_{j=r+1}^n c_j^* \geq \sum_{j=r+1}^n \frac{2c_j}{PoAPE} > \frac{\sum_{j=1}^n c_j}{PoAPE} = \sum_{j=1}^n c_j^*$, getting a contradiction.

Moreover, since $\sum_{j=1}^r c_j \geq \frac{\sum_{j=1}^n c_j}{2}$, there must exist a facility e having latency function $f_e(x) = \alpha_e x + \beta_e$ with at least half of the clients using it at Nash equilibrium belonging to the first r players. Let $J = \{j_1, \dots, j_{n'_e}\}$ be the set of such n'_e players, and n_e be the overall number of players using facility e at Nash equilibrium, with $n'_e \geq \lceil \frac{n_e}{2} \rceil$. Consider all the facilities e_1, \dots, e_k used by at least one player of J in s^* . For every $h = 1, \dots, k$, let $f_{e_h}(x) = \alpha_{e_h} x + \beta_{e_h}$ be the latency function of e_h and $o_h > 0$ be the number of players of J using e_h in s^* .

We consider two distinct cases.

If $(\alpha_e + \beta_e)PoAPE \leq 2(\alpha_e n_e + \beta_e)$, it clearly follows that $PoAPE \leq 2n_e \leq 2n$.

It remains to analyze the case in which $(\alpha_e + \beta_e)PoAPE > 2(\alpha_e n_e + \beta_e)$. Since players in J cannot unilaterally decrease their latencies, $\alpha_e + \beta_e \leq \alpha_{e_h}(n_{e_h} + 1) + \beta_{e_h}$ for every $h = 1, \dots, k$, where n_{e_h} is the number of players using facility e_h in s . Moreover, $\alpha_{e_h} o_h + \beta_{e_h} = c_{j_h}^* \leq \frac{2c_{j_h}}{PoAPE} = \frac{2(\alpha_e n_e + \beta_e)}{PoAPE}$, because $\bar{j} < r$ is a player belonging to J . By combining the last two inequalities, it follows that $n_{e_h} + 1 \geq \frac{PoAPE o_h (\alpha_e + \beta_e - \beta_{e_h})}{2(\alpha_e n_e + \beta_e) - \beta_{e_h} PoAPE}$. Summing up over all $h = 1, \dots, k$, recalling that $(\alpha_e + \beta_e)PoAPE > 2(\alpha_e n_e + \beta_e)$, $2n \geq \sum_{h=1}^k (n_{e_h} + 1) \geq \sum_{h=1}^k \frac{o_h ((\alpha_e + \beta_e)PoAPE - \beta_{e_h} PoAPE)}{2(\alpha_e n_e + \beta_e) - \beta_{e_h} PoAPE} \geq \sum_{h=1}^k \frac{o_h (\alpha_e + \beta_e) PoAPE}{2(\alpha_e n_e + \beta_e)} \geq \sum_{h=1}^k \frac{o_h \alpha_e PoAPE}{2\alpha_e n_e} = \frac{PoAPE}{2n_e} \sum_{h=1}^k o_h \geq \frac{n_e PoAPE}{2n_e} \geq \frac{PoAPE}{4}$, as $n'_e \geq \frac{n_e}{2}$.

Therefore, $PA_{PE} \leq 8n$. \square

By exploiting a similar technique it is possible to prove the following theorem.

Theorem 19. *Given any graphical linear load balancing game \mathcal{H} defined over a social graph $G \in \mathcal{G}(\Delta) \cup \overrightarrow{\mathcal{AG}}(\Delta)$, $PoA_{PE}^{max}(\mathcal{H}) \leq 2n$.*

As a consequence of the above theorems and discussion, while for the presumed latencies all the bounds coincide with the congestion game ones (Table [1](#)), the tight results shown in Table [2](#) hold.

Acknowledgements. We thank Elias Koutsoupias and Tim Roughgarden for helpful comments.

References

1. Awerbuch, B., Azar, Y., Epstein, A.: The Price of Routing Unsplittable Flow. In: Proc. of STOC, pp. 57–66. ACM Press, New York (2005)
2. Beier, R., Czumaj, A., Krysta, P., Vocking, B.: Computing Equilibria for Congestion Games with (Im)perfect Information. In: Proc. of SODA, pp. 746–755. ACM Press, New York (2004)
3. Caragiannis, I., Flammini, M., Kaklamani, C., Kanellopoulos, P., Moscardelli, L.: Tight Bounds for Selfish and Greedy Load Balancing. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 311–322. Springer, Heidelberg (2006)
4. Christodoulou, G., Koutsoupias, E.: The Price of Anarchy of Finite Congestion Games. In: Proc. of STOC, pp. 67–73. ACM Press, New York (2005)
5. Christodoulou, G., Koutsoupias, E.: On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, pp. 59–70. Springer, Heidelberg (2005)
6. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The Complexity of Pure Nash Equilibria. In: Proc. of STOC, pp. 604–612. ACM Press, New York (2004)
7. Facchini, G., van Megan, F., Borm, P., Tijs, S.: Congestion Models and Weighted Bayesian Potential Games. *Theory and Decision* 42, 193–206 (1997)
8. Gairing, M., Monien, B., Tiemann, K.: Selfish Routing with Incomplete Information. In: Proc. of SPAA, pp. 203–212. ACM Press, New York (2005)
9. Garg, D., Narahari, Y.: Price of Anarchy of Network Routing Games with Incomplete Information. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 1066–1075. Springer, Heidelberg (2005)
10. Georgiou, C., Pavlides, T., Philippou, A.: Network Uncertainty in Selfish Routing. In: Proc. of IPDPS. Computer Society (2006)
11. Harsanyi, J.C.: Games with Incomplete Information Played by Bayesian Players, I, II, III. *Management Science* 14, 159–182, 320–332, 468–502 (1967)
12. Harsanyi, J.C.: Games with Randomly Disturbed Payoffs. *International Journal on Game Theory* 21, 1–23 (1973)
13. Koutsoupias, E., Panagopoulou, P.N., Spirakis, P.G.: Selfish Load Balancing Under Partial Knowledge. In: Kučera, L., Kučera, A. (eds.) MFCS 2007. LNCS, vol. 4708, pp. 609–620. Springer, Heidelberg (2007)
14. Kearns, M.J., Littman, M.L., Singh, S.P.: Graphical Models for Game Theory. In: Proc. of UAI, pp. 253–260. Morgan Kaufmann, San Francisco (2001)
15. Monderer, D., Shapley, L.S.: Potential Games. *Games and Economic Behavior* 14, 124–143 (1996)
16. Nash, J.: Equilibrium Points in n -person Games. *Proceedings of the National Academy of Sciences* 36, 48–49 (1950)
17. Rosenthal, R.W.: A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2, 65–67 (1973)

How Hard Is It to Find Extreme Nash Equilibria in Network Congestion Games?

Elisabeth Gassner^{1,*}, Johannes Hatzl¹, Sven O. Krumke², Heike Sperber^{2,**},
and Gerhard J. Woeginger^{3,***}

¹ Graz University of Technology, Institute of Optimization and Discrete
Mathematics, Steyrergasse 30, Graz, Austria
{gassner,hatzl}@opt.math.tu-graz.ac.at

² University of Kaiserslautern, Department of Mathematics, P.O.Box 3049,
67653 Kaiserslautern, Germany
{krumke,sperber}@mathematik.uni-kl.de

³ Eindhoven University of Technology, Department of Mathematics and Computer
Science, P.O.Box 513, 5600 MB Eindhoven, The Netherlands
gwoegi@win.tue.nl

Abstract. We study the complexity of finding extreme pure Nash equilibria in symmetric (unweighted) network congestion games. In our context best and worst equilibria are those with minimum respectively maximum makespan. On series-parallel graphs a worst Nash equilibrium can be found by a Greedy approach while finding a best equilibrium is NP-hard. For a fixed number of users we give a pseudo-polynomial algorithm to find the best equilibrium in series-parallel networks. For general network topologies also finding a worst equilibrium is NP-hard.

Keywords: Network congestion game, unsplittable flow, makespan objective, extreme equilibria, complexity.

1 Introduction

In the last years there has been a lot of interest in algorithmic game theory combining aspects of game theory and computer science. Driven by growing demand for faster and larger communication networks more and more questions were asked: How do non-cooperative users interact in such networks where increasing load on individual parts of the network causes a degradation in service, often in the form of reduced transfer speed? How does this congestion effect influence the whole network? Is there some kind of self-regulation among the

* This research is partially supported by the Austrian Science Fund Project P18918-N18 *Efficiently solvable variants of location problems*.

** Supported by the Rhineland-Palatinate Cluster of Excellence Dependable Adaptive Systems and Mathematical Modeling.

*** Supported by the Netherlands Organisation for Scientific Research (NWO), grant 639.033.403, and by BSIK grant 03018 (BRICKS: Basic Research in Informatics for Creating the Knowledge Society).

users? Classical game theory provides qualitative answers such as existence of equilibria, states of the network in which all users are satisfied, and computer scientists added more quantitative question and concepts. It is a well known fact (cf. Pigou [1]), that in general selfish non-cooperative behaviour does not lead to social optimal outcome. Papadimitriou [2] coined the term price of anarchy for the ratio of the social cost of a worst Nash equilibrium and the minimal social cost. The KP-Model named after Koutsoupias and Papadimitriou [3] describes the situation in which users of possibly different size assign their traffic to parallel links with linear latency functions. For pure assignments this corresponds to uniform/related machines in scheduling. Fotakis et al. [4] came up with the question whether a best or worst pure equilibrium w.r.t. to makespan can be computed efficiently and established that in the KP-Model both problems are strongly NP-hard. Gairing et al. [5] added that it is even hard to approximate the worst equilibrium social cost on identical links while there is a PTAS for the best equilibrium social cost. Fischer and Vöcking [6] considered the worst mixed equilibrium.

The hardness proofs for extreme equilibria stated above are based on the users' different sizes, i.e., the amounts of unsplittable traffic they send through the network and the close relationship to scheduling and bin-packing problems. The question arises whether finding extreme Nash equilibria for unit-size users is substantially easier as for the unit-size case the corresponding scheduling and bin-packing instances become polynomially solvable. We will show that most versions of finding extreme equilibria are still NP-hard even for unit-size users. Up to now the complexity status of finding extreme equilibria with respect to the makespan was only considered for the KP-Model. However, in this case finding extreme equilibria for unit-size users is trivial because even for arbitrary non-decreasing latency functions on parallel links all Nash Equilibria have equal and minimal makespan as shown by Epstein et al. [7].

The game describing unit-size users sending their unsplittable traffic through arbitrary directed networks with latency functions on edges is called network congestion game and was already studied in the 1970's by Rosenthal [8]. He established that the more general congestion games possess pure strategy Nash equilibria. Fabrikant et al. [9] established that for symmetric (single-commodity) network congestion games an arbitrary equilibrium can be computed in polynomial time, but for asymmetric network congestion games or general symmetric congestion games it is PLS-complete to find an equilibrium. Fotakis et al. [10] introduced that the greedy approach yields a pure Nash equilibrium not only on parallel links but also on series-parallel graphs.

Contribution. We consider (unweighted) network congestion games with arbitrary non-decreasing latency functions on edges. Our negative results need only linear latencies $\ell_e(x) = a_e x$.

We establish that finding a best or a worst Nash equilibrium concerning makespan social cost is not equally hard in the following meaning: We prove that on series-parallel graphs finding a best equilibrium is NP-hard. It is strongly NP-hard if the number of users is part of the input and weakly NP-hard

otherwise. Moreover, we suggest a pseudo-polynomial time algorithm that determines a best Nash equilibrium on series-parallel graphs if the number of users is fixed. This indicates that this problem is not strongly NP-hard. In contrast to this we show that a worst pure equilibrium is found by the Greedy approach of Fotakis et al. [10] on these graphs.

In general networks also finding a worst equilibrium is NP-hard. In fact, we prove it to be NP-hard in the strong sense already for two users on an acyclic network with linear latencies.

Road Map. The paper is organized as follows: Section 2 introduces our notation and preliminary results such as existence of pure equilibria and computation of an arbitrary equilibrium. In Section 3 we discuss our results on finding a worst Nash equilibrium and in Section 4 for a best Nash equilibrium, respectively.

2 Preliminaries

We consider N users of the same size, i.e., each routing the same amount of unsplittable flow from a single source s to a single sink t through a directed graph $G = (V, E)$. The edges of G are equipped with *non-decreasing latency functions* $\ell_e : \mathbb{N}_0 \rightarrow \mathbb{R}_0^+$ for all $e \in E$ modelling the congestion effects. An instance of the game is thus given by $(G = (V, E), (\ell_e)_{e \in E}, s \in V, t \in V, N)$. By scaling the latency functions appropriately we assume without loss of generality all users to have *unit size*.

Let \mathcal{P} denote the set of all simple s - t -paths in G and thus the strategy set of all users. In our context a *flow* is a function $f : \mathcal{P} \rightarrow \mathbb{N}_0$ that assigns integer values to paths in the network. The latency on a path is the sum of the latencies on its edges that depends on the total flow on the edge:

$$\ell_P(f) := \sum_{e \in P} \ell_e \left(\sum_{P' \in \mathcal{P}: e \in P'} f_{P'} \right) \quad (1)$$

We denote by $f_e := \sum_{P \in \mathcal{P}: e \in P} f_P$ the flow on edge e uniquely induced by the flow f defined on paths. Note that there may be different so-called flow-decompositions or flows on paths that correspond to the same flow on edges. Example 1 shows that we need the information about paths for modelling the users' behaviour in our game.

A Nash equilibrium is a stable situation in which no user wants to deviate from her chosen path because she cannot decrease her experienced latency this way:

Definition 1 (Nash Equilibrium, Nash Flow). *A flow on paths $f = (f_P)_{P \in \mathcal{P}}$ is at Nash equilibrium, if and only if for all paths P_1, P_2 with $f_{P_1} > 0$ we have*

$$\ell_{P_1}(f) \leq \ell_{P_2}(\tilde{f}) \text{ with } \tilde{f}_P = \begin{cases} f_P - 1 & \text{if } P = P_1 \\ f_P + 1 & \text{if } P = P_2 \\ f_P & \text{otherwise} \end{cases} \quad (2)$$

Existence of Nash Equilibria. Rosenthal [8] used the following potential function $\Pi : \mathcal{F} \rightarrow \mathbb{R}$ defined on the set of feasible flows \mathcal{F} to prove the existence of pure Nash equilibria in network congestion games:

$$\Pi(f) = \sum_{e \in E} \sum_{i=1}^{f_e} \ell_e(i) \quad (3)$$

Flows corresponding to local optima of this potential function constitute Nash equilibria. Fabrikant et al. [9] establish that one equilibrium can be computed in polynomial time because a min-cost flow in the following instance $\text{MCF}(G)$ minimizes Rosenthal's potential function and is thus a Nash flow.

Definition 2 (Min-cost Flow Instance, $\text{MCF}(G)$). *Given a network congestion game $(G = (V, E), (\ell_e)_{e \in E}, s \in V, t \in V, N)$ construct the corresponding min-cost flow instance as follows:*

For every edge $e \in E$ we need N copies with costs $c_{e_i} = \ell_e(i)$, $i = 1, \dots, N$. The capacities of all edges are 1 and we send N units of flow from s to t .

Observe that every path decomposition of every optimal solution of the min-cost flow instance $\text{MCF}(G)$ yields a Nash equilibrium as the negative cycle optimality condition for optimal min-cost flows directly implies that no user wants to deviate from her chosen strategy. However, not every Nash equilibrium is also an optimal solution of the min-cost flow instance (cf. Examples [3] and [4]).

Note that there are instances and Nash flows (not global but local optima of Rosenthal's potential) such that a different path decomposition of the flow on edges induced by a Nash flow is not again Nash (cf. Example [1]). Thus, it is necessary to have the information about the flow on paths as the output of the game.

Example 1 (Nash equilibria and flow decompositions). Consider the instance given in Figure [1] in which two users travel from s to t . The latency functions are given as edge labels. In order to distinguish parallel edges (s, u) (or (u, t)), we call them upper and lower edge between s and u (u and t).

Observe that the flow sending the first user on edge (s, t) , the second user on the path consisting of the upper edge from s to u and the lower edge from u to t and the third user on the path containing the so far unused edges is a Nash equilibrium. The flow on every edge is equal to 1.

But if we change the flow decomposition and send the second user on both upper and the third on the lower edges this last user becomes unsatisfied because she would be better off changing to edge (s, t) . Hence, not every path decomposition of a flow on edges yields a Nash equilibrium.




Fig. 1. Nash equilibrium property might depend on flow decomposition (Example [1])

Social Cost. In this paper, we consider minimizing the makespan as the social objective function. This notion comes from scheduling and is a priori only applicable to parallel link networks. Flows minimizing the following more general makespan definition are sometimes also called min-max flows.

Definition 3 (Makespan, Social Cost). *Given a flow on paths $f = (f_P)_{P \in \mathcal{P}}$ the makespan is given by*

$$C_{\max}(f) := \max_{P \in \mathcal{P}: f_P > 0} \ell_P(f). \tag{4}$$

Epstein et al. [7] showed that on parallel links all Nash equilibria have equal makespan but this does not hold in general:

Example 2 (Nash equilibria with different non-optimal makespans). Consider the instance given in Figure 2 for two users. If every edge is used by exactly one user and the paths are alternating between upper and lower edges then an optimal solution with makespan 12 is achieved. Observe that in any Nash equilibrium there is exactly one user on every edge between s and u_1 and between u_1 and u_2 and there are two users on the upper edges between u_2 and u_3 and between u_3 and t . A best Nash equilibrium with makespan 13 can be obtained if both users alternate between upper and lower connection on the first two edges. However, one user may also choose the lower connections on both first edges. This yields again a Nash equilibrium, which is worst and has makespan 14.




Fig. 2. Instance with several Nash equilibria

As in general the makespan of different Nash equilibria as well as an optimum makespan are not equal, we are now interested in computing two extreme Nash equilibria.

Extreme Nash Equilibria. We introduce the following two problems of finding a best or worst pure equilibrium, respectively.

WORST NASH EQUILIBRIUM (W-NE for short):

Given: Network congestion game $(G = (V, E), (\ell_e)_{e \in E}, s \in V, t \in V, N)$
 Output: Nash equilibrium f with maximal makespan among all Nash equilibria.

BEST NASH EQUILIBRIUM (B-NE for short):

Given: Network congestion game $(G = (V, E), (\ell_e)_{e \in E}, s \in V, t \in V, N)$
 Output: Nash equilibrium f with minimal makespan among all Nash equilibria.

Note that the decision versions of these two problems are in NP for acyclic networks G as in those networks for a given flow f a longest path w.r.t. to the fixed edge lengths $\ell_e(f_e)$ can be computed in polynomial time [11].

Unfortunately, it can be shown that in general neither a best nor a worst Nash equilibrium is an optimal solution of $\text{MCF}(G)$:

Example 3 (Best Nash flow not optimal in $\text{MCF}(G)$). Reconsider the instance of Example 2 and observe that the unique solution of $\text{MCF}(G)$ is the Nash equilibrium with makespan 14 and thus not the best one.

Example 4 (Worst Nash flow not optimal in $\text{MCF}(G)$). In case of the worst Nash equilibrium consider the instance given in Figure 3 for two users:




Fig. 3. The unique worst Nash equilibrium does not imply an optimal min-cost flow (Example 4)

The optimal solution f^* of $\text{MCF}(G)$ for the graph given in Figure 3 is unique and has a unique path decomposition sending one user on $Q_1 = (s, u_1, u_4, t)$ and $Q_2 = (s, u_3, u_2, t)$ each with makespan $C_{\max}(f^*) = 2$. However, f with $f_{P_1} = f_{P_2} = 1$ where $P_1 = (s, u_1, u_2, t)$ and $P_2 = (s, u_3, u_4, t)$ is a Nash equilibrium with $C_{\max}(f) = 3$.

The fact that in general no worst Nash equilibrium is an optimal min-cost flow in $\text{MCF}(G)$ is quite interesting because in the special case of series-parallel graphs there always exists a worst Nash equilibrium that is an optimal solution of the min-cost flow problem $\text{MCF}(G)$. This follows from the result that the Greedy approach determines a worst Nash equilibrium in series-parallel graphs (cf. Section 3).

Series-Parallel Graphs. As already mentioned we consider not only arbitrary network topologies but also series-parallel networks. Series-parallel graphs can be defined inductively. A single edge $e = (s, t)$ is series-parallel with start terminal s and end-terminal t by definition. Let G_i be series-parallel with start-terminal s_i and end-terminal t_i ($i = 1, 2$). Then the graph $S(G_1, G_2)$ obtained by identifying t_1 as s_2 is a series-parallel graph, with s_1 and t_2 as its terminals (series composition). And the graph $G = P(G_1, G_2)$ obtained by identifying s_1 as s_2 and also t_1 as t_2 is a series-parallel graph (parallel composition). This graph has $s_1(=s_2)$ and $t_1(=t_2)$ as its terminals (cf. [10]).

This class of graphs has some very nice properties: Bein et al. [12] established that the Greedy approach solves the min-cost flow problem in series-parallel graphs. Combined with the min-cost flow instance introduced by Fabrikant et al. [9] this yields that the greedy approach of iteratively assigning the users to a shortest path with respect to the latency induced by the current flow plus an

additional user on every edge yields a Nash equilibrium on series-parallel graphs. This result was also obtained by Fotakis et al. [10] who call this algorithm GBR (greedy best response) and we keep this notation.

3 Worst Pure Nash Equilibrium

In this section the complexity status of determining a worst Nash equilibrium is investigated. We prove that a Greedy strategy solves the problem on series-parallel graphs and show strong NP-hardness for the problem on general graphs.

Special Case of Series-Parallel Graphs. In the following we show that the Greedy Best Response (GBR) algorithm introduced by Fotakis et al. [10] always leads to a worst Nash equilibrium in series-parallel graphs. The idea of this algorithm is as follows: If one considers a setting where the users arrive consecutively, a new user routes her path such that her personal latency is minimized given the flow induced by the users currently in the network. This choice is irrevocable, i.e., no user can change the strategy in the future. More formally, let us denote by

$$L^+(f) := \min_{P \in \mathcal{P}} \sum_{e \in P} \ell_e(f_e + 1) \quad (5)$$

the minimum latency for a new $(N+1)^{\text{st}}$ user given a flow f sending N users from s to t . According to GBR the new user chooses her path P_{N+1} such that the latency of P_{N+1} is $L^+(f)$. If a flow f' is obtained by a given flow f where a single user is added according to GBR we use $f' = f \oplus P_{N+1}$. For series-parallel graphs it has been shown in [10] that if f is an arbitrary Nash equilibrium then $f' = f \oplus P_{N+1}$ is again a Nash equilibrium. Note that this property does not hold in general graphs. As a consequence GBR always leads to a Nash equilibrium if all users have the same size and the underlying network is series-parallel. In this paper, we strengthen this result and show that the obtained Nash equilibrium is always a worst Nash equilibrium. This holds for all latency functions that are non-decreasing. The next lemma, which is a key point in order to prove our result, has already been used implicitly in [10]. It states that if we start with a Nash equilibrium and add one more user according to GBR then the latency of the new user is not less than the latency of all the previous users in the new flow.

Lemma 1. *Let $G = (V, E)$ be a series-parallel graph and f a Nash equilibrium for N users. If we choose $P_{N+1} \in \mathcal{P}$ according to GBR we obtain a new Nash equilibrium $f' = f \oplus P_{N+1}$ such that*

$$\ell_{P_{N+1}}(f') = C_{\max}(f').$$

The next two lemmata are dealing with the two compositions in the definition of series-parallel graphs. In fact, we give a characterization of a Nash equilibrium in $S(G_1, G_2)$ and $P(G_1, G_2)$. Before the results are stated the following notation is introduced. Let G_i be a series-parallel graph and $f_i : \mathcal{P}_i \rightarrow \mathbb{N}_0$ a flow in G_i for

$i = 1, 2$. Then the set of all simple s - t -paths in $P(G_1, G_2)$ is given by $\mathcal{P}_1 \cup \mathcal{P}_2$. We define a new flow f in $P(G_1, G_2)$ by $f := f_1 \cup f_2$, where $f : \mathcal{P}_1 \cup \mathcal{P}_2 \rightarrow \mathbb{N}_0$ and $f|_{\mathcal{P}_i} = f_i$ for $i = 1, 2$.

Lemma 2. *Let f_i be a flow in a series-parallel graph G_i for $i = 1, 2$. Then $f = f_1 \cup f_2$ is a Nash equilibrium in $P(G_1, G_2)$ if and only if the following conditions are satisfied:*

1. f_i is a Nash equilibrium in G_i for $i = 1, 2$,
2. $L_{G_1}^+(f_1) \geq C_{\max}(f_2)$ and $L_{G_2}^+(f_2) \geq C_{\max}(f_1)$.

We want to establish a similar result for the series composition. Therefore let G_i be series-parallel and $f_i : \mathcal{P}_i \rightarrow \mathbb{N}_0$ a flow in G_i for $i = 1, 2$ for N users. Let us assume without loss of generality that the users choose the paths P_1, \dots, P_N (Q_1, \dots, Q_N) in G_1 (G_2). For each permutation ϕ of $\{Q_1, \dots, Q_N\}$ we can obtain a new flow f in $S(G_1, G_2)$ if we define a new path for user i by $\bar{P}_i = P_i \cup Q_{\phi(i)}$. The set of all flows that can be obtained this way will be denoted by $f_1 \otimes f_2$.

Lemma 3. *Let f_i be a flow in G_i for $i = 1, 2$. Let $f \in f_1 \otimes f_2$ then f is a Nash flow in $S(G_1, G_2)$ if and only if f_i is a Nash equilibrium in G_i for $i = 1, 2$.*

Using these lemmata we are able to prove the following theorem by induction on the composition steps. The detailed proof is omitted due to space restrictions.

Theorem 1. *If G is a series-parallel graph then the Nash equilibrium obtained by GBR is a worst Nash equilibrium.*

Complexity Status on General Graphs. Before proving NP-hardness of the problem of finding a worst Nash equilibrium, we consider a related problem that is called Blocking Path problem:

BLOCKING PATH PROBLEM (BlockP for short):

Given: Digraph $G = (V, E)$ with source $s \in V$ and sink $t \in V$.

Question: Does there exist an s - t -path $P \in \mathcal{P}$ such that after deleting the edges of P there is no path from s to t ?

Theorem 2. *The Blocking Path Problem is strongly NP-complete even on acyclic networks.*

The proof is a reduction from 3SAT and due to lack of space postponed to the full version of this paper. The Blocking Path Problem is used to show that determining a worst Nash equilibrium in general networks is NP-hard even for two users.

Theorem 3. *Determining a Worst Pure Nash equilibrium is strongly NP-hard even for two users on acyclic networks and with linear latency functions.*

Proof. Consider an instance $I(\text{BlockP})$ of the strongly NP-complete Blocking Path problem. Let $G = (V, E)$ be the acyclic network of instance $I(\text{BlockP})$ with

$s, t \in V$. An instance of $I(\text{W-NE})$ of determining a worst pure Nash equilibrium is defined as follows: $I(\text{W-NE})$ is defined on a graph $G' = (V, E')$ which contains the same vertex set as G and $E' = E \cup \{(s, t)\}$. Since G' is acyclic it is possible to define a bijective function $\pi : V \rightarrow \{1, \dots, n\}$ such that $\pi(i) < \pi(j)$ if $(i, j) \in E$. Given any such bijection π the latency functions are given by

$$\ell_e(x) = (\pi(j) - \pi(i))x, \quad e = (i, j) \in E.$$

Observe that due to this definition of the latency functions of edges in G every path from s to t is a shortest path with respect to the edge lengths $\ell_e(1)$. Let L^* be the length of a shortest path from s to t in G with respect to edge lengths $\ell_e(1)$ for $e \in E$. Then the latency of (s, t) is defined by $\ell_{(s,t)}(x) = (L^* + \frac{1}{2})x$.

We show that there exists a blocking path P^* for $I(\text{BlockP})$ if and only if the answer to the decision problem corresponding to $I(\text{W-NE})$ is “yes” for $K = L^* + \frac{1}{2}$, i.e. there exists a Nash equilibrium f in G' with cost $C_{\max}(f) \geq L^* + \frac{1}{2}$.

Given a blocking path P^* in $I(\text{BlockP})$ we construct a feasible flow f in G' by sending one user on P^* and the other on edge (s, t) inducing $C_{\max}(f) = L^* + \frac{1}{2}$. Observe that indeed both users are satisfied and this flow constitutes a Nash equilibrium.

On the other hand, assume that there exists a Nash equilibrium f with makespan $C_{\max}(f) \geq L^* + \frac{1}{2}$. Analysing the different cases of flow values on (s, t) , the Nash property of f together with the lower bound on $C_{\max}(f)$ tell us that in this setting one user is sent over (s, t) and one on a path P^* in G' . The fact, that the user on (s, t) does not want to change to G' implies that P^* is in fact a blocking path. □

4 Best Pure Nash Equilibrium

In this section, we show several complexity results concerning the problem of determining a best Nash equilibrium. All results given in this section hold even for series-parallel graphs. We show that computing a best Nash equilibrium for N users is strongly NP-hard if N is part of the input. If the number of users is fixed then the problem remains weakly NP-hard. At least for series-parallel graphs this result is best possible because there exists a dynamic programming algorithm with pseudo-polynomial running time.

Strong NP-Hardness Result. In this subsection, we prove that finding a best Nash equilibrium on series-parallel graphs is strongly NP-hard if the number of users is part of the input. We show this by a reduction of the corresponding decision problem to the numerical 3-dimensional matching problem, which is known to be strongly NP-complete (see [11]).

NUMERICAL 3-DIMENSIONAL MATCHING (N3M for short):

Given: Disjoint sets X, Y, Z , each containing m elements, a weight $w(a)$ for all elements $a \in X \cup Y \cup Z$ and a bound $B \in \mathbb{Z}^+$.

Question: Does there exist a partition of $X \cup Y \cup Z$ into m disjoint sets A_1, \dots, A_m such that each A_j contains exactly one element from each of X, Y and Z and $\sum_{a \in A_i} w(a) = B$ for all $1 \leq i \leq m$.

Theorem 4. *Determining a best Nash equilibrium is strongly NP-hard on series-parallel graphs if the number of users is part of the input.*

Proof. Consider an instance $I(\text{N3M})$ of N3M. Observe that we may assume without loss of generality that $w(a) < 2w(b)$ and $w(b) < 2w(a)$ for each pair $a, b \in X$. Otherwise a large number M can be added to all elements in the set X and to B until the desired condition is satisfied. An analogue property holds for Y and Z .

Based on this instance we construct the following series-parallel graph $G = (V, E)$: Let $V = (s, u, v, t)$ and for each element in the set X (Y, Z) we introduce a directed edge from s to u (u to v , v to t). The latency function of an edge e is given by $\ell_e(x) = w(a)x$ where $w(a)$ is the weight of the corresponding element in the instance $I(\text{N3M})$.

Observe that in a best Nash equilibrium every edge is used by exactly one user. Hence, there is a one-to-one correspondence between the subsets A_i $i = 1, \dots, m$ and the paths of the users and therefore there exists a Nash equilibrium with m users in G with social cost at most B if and only if $I(\text{N3M})$ is a YES-instance. \square

Weak NP-Hardness for Fixed Number of Users. This subsection deals with the problem of determining a best Nash equilibrium if the number of users N is fixed. The proof is a reduction from Even-Odd Partition. As it works similar to that of the previous section it is omitted here.

Theorem 5. *Determining a best Nash equilibrium is weakly NP-hard even for two users and on series-parallel graphs.*

A Pseudo-Polynomial Time Algorithm for Series-Parallel Graphs. In this subsection, we discuss a dynamic programming approach to find a best Nash equilibrium in series-parallel graphs if the number N of users is not part of the input. Let f be a Nash equilibrium in a graph G for k users which choose the paths P_1, \dots, P_k . Then we define a multiset

$$C(f) := \{\ell_{P_1}(f), \dots, \ell_{P_k}(f)\}$$

which will be called cost profile of f . Note that several Nash equilibria can have the same cost profile. The idea of the algorithm is to decide if for a given multiset $C = \{c_1, \dots, c_k\}$ with $0 \leq k \leq N$ there exists a corresponding Nash flow f with k users. This is done using the inductive definition of series-parallel graphs. In order to decide if a cost profile can be realized by a Nash flow $f = f_1 \cup f_2$ in $P(G_1, G_2)$ we need to know $L^+(f_i)$. More formally, for a given multiset $C = \{c_1, \dots, c_k\}$ and a graph G we define

$$S_G(C) := \max\{L^+(f) \mid C(f) = \{c_1, \dots, c_k\}, f \text{ is a Nash flow}\}.$$

If such a Nash equilibrium does not exist we set $S_G(C) := -\infty$. Hence, all cost profiles with $S_G(C) \geq 0$ do have a corresponding Nash flow f . Let us discuss the algorithm in more detail.

1. A single edge (s, t)

For the simplest series-parallel graph there is a unique flow for all $0 \leq k \leq N$ and all users have latency $\ell_{(s,t)}(k)$. Thus, we obtain immediately

$$S_G(C) = \begin{cases} \ell_{(s,t)}(k+1) & \text{if } C = \{\ell_{(s,t)}(k), \dots, \ell_{(s,t)}(k)\}, |C| = k \leq N, \\ -\infty & \text{otherwise.} \end{cases}$$

2. The series composition

Let $C = \{c_1, \dots, c_k\}$ be given. Note that this cost profile can only be obtained by a Nash flow $f \in f_1 \otimes f_2$ with $C_i := C(f_i) = \{c_1^i, \dots, c_k^i\}$ for $i = 1, 2$ and $C = \{c_1^1 + c_{\phi(1)}^2, \dots, c_k^1 + c_{\phi(k)}^2\}$ for some permutation ϕ . If such a permutation exists we write $C_1 \otimes C_2 = C$. Moreover, $L^+(f) = L^+(f_1) + L^+(f_2)$ because every $s - t$ path in G has to pass the vertex $t_1 = s_2$. Thus, we obtain

$$S_G(C) = \max_{C_1 \otimes C_2 = C} \{S_{G_1}(C_1) + S_{G_2}(C_2)\}. \tag{6}$$

3. The parallel composition

Let $C = \{c_1, \dots, c_k\}$ be given. A corresponding Nash flow f is of the form $f_1 \cup f_2$ with $C_1 := C(f_1) = \{c_1^1, \dots, c_{k_1}^1\}$, $C_2 := C(f_2) = \{c_1^2, \dots, c_{k_2}^2\}$, $k_1 + k_2 = k$ and $C = C_1 \cup C_2$. Moreover the conditions from Lemma 3 have to be satisfied, i.e., $\max\{c_1^1, \dots, c_{k_1}^1\} \leq S_{G_2}(C_2)$ and $\max\{c_1^2, \dots, c_{k_2}^2\} \leq S_{G_1}(C_1)$. The shortest path in G with respect to the flow f is given by $\min\{L^+(f_1), L^+(f_2)\}$, because the shortest path in $P(G_1, G_2)$ chooses either a path with edges in G_1 or in G_2 . Thus,

$$S_G(C) = \max_{\substack{C=C_1 \cup C_2 \\ |C_1|+|C_2|=k \\ \max\{c|c \in C_1\} \leq S_{G_2}(C_2) \\ \max\{c|c \in C_2\} \leq S_{G_1}(C_1)}} \min\{S_{G_1}(C_1), S_{G_2}(C_2)\} \tag{7}$$

is satisfied.

Note that it is straightforward to get the best Nash flow at the end if the corresponding flows which determine $S_G(C)$ during the algorithm are stored as well. In order to analyze the running time of this algorithm note that for a graph G and a fixed number k of users there are at most $\frac{(|V|L)^k}{k!} = \mathcal{O}((|V|L)^k)$ different multisets, where $L := \max_{e \in E} l_e(N)$ is the maximum latency on an edge and a simple path can have at most $|V|$ edges. Due to the fact that this is needed for all $0 \leq k \leq N$ the number of multisets that have to be stored is at most $N(|V|L)^N = \mathcal{O}((|V|L)^N)$. It is easy to see that for the series and parallel composition (6) and (7) can be done in polynomial time with respect to the number of multisets. Thus, the proposed dynamic programming approach is pseudo-polynomial which implies that B-NE is indeed not NP-hard in the strong sense for series-parallel graphs.

References

1. Pigou, A.C.: The economics of welfare. Macmillan, Basingstoke (1920)
2. Papadimitriou, C.: Algorithms, games, and the internet. In: Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing, pp. 749–753 (2001)

3. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
4. Fotakis, D., Kontogiannis, S., Koutsoupias, E., Mavronicolas, M., Spirakis, P.: The structure and complexity of nash equilibria for a selfish routing game. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) ICALP 2002. LNCS, vol. 2380, pp. 123–134. Springer, Heidelberg (2002)
5. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B., Spirakis, P.: The structure and complexity of extreme nash equilibria. *Theoretical Computer Science* 343(1-2), 133–157 (2005)
6. Fischer, S., Vöcking, B.: On the structure and complexity of worst-case equilibria. *Theoretical Computer Science* 378(2), 165–174 (2007)
7. Epstein, A., Feldman, M., Mansour, Y.: Efficient graph topologies in network routing games. In: Joint Workshop on Economics of Networked Systems and Incentive-Based Computing (2007)
8. Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2(1), 65–67 (1973)
9. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure nash equilibria. In: Proceedings of the 36th Annual ACM Symposium on the Theory of Computing, pp. 604–612 (2004)
10. Fotakis, D., Kontogiannis, S., Spirakis, P.: Symmetry in network congestion games: Pure equilibria and anarchy cost. In: Erlebach, T., Persinao, G. (eds.) WAOA 2005. LNCS, vol. 3879, pp. 161–175. Springer, Heidelberg (2005)
11. Garey, M.R., Johnson, D.S.: Computers and intractability. A guide to the theory of NP-completeness. A Series of Books in the Mathematical Sciences. W. H. Freeman & Co., New York (1979)
12. Bein, W.W., Brucker, P., Tamir, A.: Minimum cost flow algorithm for series-parallel networks. *Discrete Applied Mathematics* 10, 117–124 (1985)

On the Road to \mathcal{PLS} -Completeness: 8 Agents in a Singleton Congestion Game*

Dominic Dumrauf^{1,2} and Burkhard Monien²

¹ Paderborn Institute for Scientific Computation
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
dumrauf@uni-paderborn.de

² Faculty of Computer Science, Electrical Engineering and Mathematics,
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
bm@uni-paderborn.de

Abstract. In this paper, we investigate the complexity of computing *locally optimal* solutions for *Singleton Congestion Games* (SCG) in the framework of \mathcal{PLS} , as defined in Johnson et al. [25]. Here, in an instance *weighted agents* choose links from a set of *identical links*. The *cost of an agent* is the load (the sum of the weights of the agents) on the link it chooses. The agents are selfish and try to minimize their individual cost. Agents may form arbitrary, *non-fixed coalitions*. The cost of a coalition is defined to be the maximum cost of its members. The potential function is defined as the *lexicographical order of the agents' cost*. In each selfish step of a coalition, the potential function decreases. Thus, a local minimum is a *Nash Equilibrium* among coalitions of size at most k —an assignment where no coalition of size at most k has an incentive to unilaterally decrease its cost by switching to different links. The neighborhood of a feasible assignment (every agent chooses a link) are all assignments, where the cost of some arbitrary non-fixed coalition of at most k reallocating agents decreases. We call this problem $\text{SCG-}(k)$ and show that $\text{SCG-}(k)$ is \mathcal{PLS} -complete for $k \geq 8$. On the other hand, for $k = 1$, it is well known that the solution computed by Graham's LPT-algorithm [14,16,22] is locally optimal for $\text{SCG-}(k)$.

We show our result by tight reduction from the $\text{MAXCONSTRAINTASSIGNMENT}$ -problem (p, q, r) -MCA, which is an extension of $\text{GENERALIZED SATISFIABILITY}$ to higher valued variables. Here, p is the maximum number of variables occurring in a constraint, q is the maximum number of appearances of a variable, and r is the valuedness of the variables.

To the best of our knowledge, $\text{SCG-}(k)$ is the first problem, which is known to be solvable in polynomial time for a small neighborhood and \mathcal{PLS} -complete for a larger, but still constant neighborhood.

1 Introduction

Routing on Parallel Links. *Routing games* model large scale networks, like e.g. traffic networks. These networks often lack a central regulation due to their size or the fact

* This work has been partially supported by the European Union within the Integrated Project IST-15964 "AEOLUS" and the German Science Foundation (DFG) Research Training Group GK-693 of the Paderborn Institute for Scientific Computation (PaSCo).

that users may be free to act according to their private interest. Such an environment can be modeled as a *non-cooperative game* [28]. A famous solution concept for non-cooperative games is the concept of *Nash equilibrium*. A Nash equilibrium is a state in which no player can improve its objective by unilaterally changing its strategy. In a *pure* Nash equilibrium, all players choose a pure strategy and in a *mixed* Nash equilibrium, all players choose probability distributions over strategies. Routing games belong to the class of congestion games, introduced by Rosenthal [31]. In a *congestion game*, the strategy set of each player is a subset of the power set of given resources. The cost of a player for some choice of strategy is defined as the sum (over the chosen resources) of functions in the number of players sharing this resource. Routing games are defined on general graphs, but special attention is given to the model of *routing on parallel links*. In terms of a congestion game, strategies are then single resources and the class of games is therefore often referred to as *singleton congestion games*. This model has been intensively studied [7][15][19], starting with the seminal paper of Koutsoupias and Papadimitriou [26]. On parallel links, the degradation of social welfare due to the selfish behavior of the players – usually coined as *price of anarchy* or *coordination ratio* – has been thoroughly investigated, [9][19][26]. In contrast to that, little progress has been made in the investigation of the complexity of computing (pure) Nash Equilibria. For general congestion games, Fabrikant et al. [13] and Ackermann et al. [3] show that the problem of computing a Nash Equilibrium is \mathcal{PLS} -complete. Skopalik and Vöcking [34] prove that even the approximation is \mathcal{PLS} -complete. For parallel links, it is well known that Graham’s LPT-algorithm [14][16][22] computes a pure Nash Equilibrium.

Coalitions and Local Search. A natural and convincing concept to model limited collusion in games is to allow agents to form *coalitions*. This concept has been investigated in several areas of computer science. In mechanism design, agents form a coalition, such that no player decreases its utility and at least one player strictly increases its utility. Differing from that, we allow players in *non-fixed coalitions* to sacrifice some of their own utility for the welfare of the group. Similar to Hayrapetyan et al. [23] and Fotakis et al. [17], we assume full cooperation among the members of a coalition, who aim to minimize their collective cost. Thus, sets of agents of constant size can collaborate and collectively improve. Here, individual deficits may be compensated by e.g. monetary transfers between the members of a coalition. Similar to Fotakis et al. [17], we define the *cost of a coalition* to be the *maximum* cost of its members and in any improving step, the maximum cost of a coalition has to decrease. In contrast to this, Hayrapetyan et al. [23] define the cost of a coalition to be the *sum* of the costs of its members and in any improving step, the sum of the costs of the players has to decrease. To investigate the complexity of computing a pure Nash Equilibrium on parallel links, we formulate the problem as a local search problem, since *pure Nash equilibria* are the local optima for the heuristic of *selfish steps*.

Local search is a natural approach to approximate solutions of hard combinatorial optimization problems. Local search algorithms are well-known to lead to very powerful heuristics for many hard problems [4][21]. Starting from an arbitrary (feasible) solution, a sequence of (feasible) solutions is iteratively generated, such that each solution

is contained in the neighborhood of its predecessor solution and strictly improves the objective function. If no improvement within the neighborhood of a solution is possible, a *local optimum* (or *locally optimal solution*)—in our case a Nash Equilibrium—is found. In practice, local search algorithms often require only a few iterations to compute a local optimum. However, their running time depends on the objective function, and thus is pseudo-polynomial in general and exponential in the worst case. In many papers, local algorithms have been investigated for the MULTIPROCESSOR SCHEDULING-problem (MPS), [5,6]. In an improving step, either the makespan decreases or the number of makespan-machines decreases. Lately, the MPS-problem was shown to be \mathcal{PLS} -complete for the k -move neighborhood for a sufficiently large k , [12], where in the k -move neighborhood up to k jobs may be relocated in an improving step. This local version of the MULTIPROCESSOR SCHEDULING-problem can be viewed also as a Singleton Congestion Game, but has a different cost function than the one we study in this paper. The concept of local search has also been successfully applied to other areas of computer science. For an overview of the application of local search, confer Aarts et al. [2].

Polynomial Time Local Search. Johnson, Papadimitriou, and Yannakakis, [25], introduced the class \mathcal{PLS} (polynomial-time local search) in 1988 to investigate the complexity of local search. Essentially, a problem in \mathcal{PLS} is given by some *minimization* or *maximization problem* over instances with finite sets of feasible solutions together with a non-negative *cost function*. A *neighborhood structure* is superimposed over the set of feasible solutions, with the property that a local improvement in the neighborhood can be found in polynomial time. The objective is to find a *locally optimal* solution. The notion of a \mathcal{PLS} -reduction was also defined in Johnson et al. [25] to establish relationships between \mathcal{PLS} -problems and to further classify them. Similar to reductions from problem A to problem B in \mathcal{NP} , one asks for a mapping from instances of A to instances of B . While in \mathcal{NP} , the question is about the *existence* of a solution with the desired properties, in \mathcal{PLS} the challenge is to actually *compute* locally optimal solutions. By definition of \mathcal{PLS} -reductions, local optima carry over from B to A . Not many problems are known to be \mathcal{PLS} -complete, since reductions are mostly technically involved. Also, since the goal should be to show that \mathcal{PLS} -problems with a *small neighborhood* are \mathcal{PLS} -complete, \mathcal{NP} -reductions are of little help, since the neighborhood is usually unbounded in these reductions. The first problem, which was shown to be \mathcal{PLS} -complete is CIRCUIT/FLIP [25]. In the meantime, only a handful of problems were shown to be \mathcal{PLS} -complete. Our knowledge about \mathcal{PLS} is still very limited and not at all comparable with the rich knowledge which we have about the class \mathcal{NP} .

Our Contribution. In this paper, we investigate the complexity of computing *locally optimal* solutions for *Singleton Congestion Games* (SCG). Here, in an instance *weighted agents* choose links from a set of *identical links*. The *cost of an agent* is the load (the sum of the weights of the agents) on the link it chooses. The agents are selfish and try to minimize their individual cost. Agents may form arbitrary, *non-fixed coalitions*. We consider coalitions up to size k , for some $k \in \mathbb{N}$. The cost of a coalition is defined

to be the maximum cost of its members. In a *selfish step of a coalition*, a coalition of size at most k improves its cost by unilateral deviation. The *potential function* is defined as the *lexicographical order of the agents' cost*. In each selfish step, the potential function decreases and a local minimum is a *Nash Equilibrium* among coalitions of size at most k —an assignment where no coalition of size at most k has an incentive to unilaterally decrease its cost by switching to different links. This problem is contained in \mathcal{PLS} , the objective is to find a local minimum with respect to the lexicographic order of the agents' costs. The neighborhood of a feasible assignment (every agent chooses a link) are all assignments where the cost of some arbitrary non-fixed coalition of at most k reallocating agents decreases. We call this problem $\text{SCG-}(k)$ and show that $\text{SCG-}(k)$ is \mathcal{PLS} -complete for $k \geq 8$. On the other hand, for $k = 1$, it is well known that the solution computed by Graham's LPT-algorithm [16,22] is locally optimal for both models and also the standard algorithm problem can be solved for both models in polynomial time [14,24]. We show our result by reduction from the $\text{MAXCONSTRAINT-ASSIGNMENT}$ problem (p, q, r) -MCA, which is an extension of weighted, $\text{GENERALIZED SATISFIABILITY}$ (confer problem [L06] in [20] for a formal description) to higher valued variables. Here, p is the maximum number of variables occurring in a predicate, q is the maximum number of appearances of a variable, and r is the valuedness of the variables. The problem (p, q, r) -MCA is known to be \mathcal{PLS} -complete for triples $(3,2,16)$, $(2,3,18)$, $(3,3,3)$, and $(6,3,2)$, [10,27]. In detail, we use a *tight* reduction from $(3, 2, r)$ -MCA and we want to stress that the parameter r does not have a negative influence on the size of the neighborhood in the proof of \mathcal{PLS} -completeness of $\text{SCG-}(k)$. The tightness of our reduction implies that there exist instances of $\text{SCG-}(8)$ with assignments such that every sequence of selfish steps of coalitions starting in such an assignment has exponential length. Furthermore, this implies that it is \mathcal{PSPACE} -complete for $\text{SCG-}(8)$ to compute a Nash equilibrium among coalitions of size at most k reachable by successive selfish steps of coalitions from a given initial assignment. To the best of our knowledge, $\text{SCG-}(k)$ is the first problem that establishes the \mathcal{PLS} -completeness of computing a Nash Equilibrium for the class of routing games on parallel links. Structurally, it is the first type of a problem, which is known to be solvable in polynomial time for a small neighborhood and \mathcal{PLS} -complete for a larger, but still constant neighborhood. Also, it is one of the first purely numerical problems shown to be \mathcal{PLS} -complete and it contributes to the narrow class of known \mathcal{PLS} -complete problems.

Further Related Work. Survey articles about local search algorithms can be found in several books [12,21]. \mathcal{PLS} was defined in [25] and the fundamental definitions and results are presented in [25,32]. Further findings on the complexity of computing a locally optimal solution are presented in [3,8,13,27]. Results considering the approximation of \mathcal{PLS} -problems can be found in [8,29,34]. The book of Aarts et al. [1] contains a list of \mathcal{PLS} -complete problems known so far. Computing Nash equilibria for coalitions of agents has been considered in [17]. Local search has been applied to a large number of scheduling problems (see chapters in [12,21]). Besides the complexity of computing a locally optimal solution, the quality of the obtained solution has also been investigated [5,6,18,33].

2 Notation and Contribution

In this section, we describe the notation, classes and problems used in the paper. For all $j, k \in \mathbb{N}$ with $j \geq k$, denote $[j : k] = \{j, \dots, k\}$, $[k] = \{1, \dots, k\}$, and $[k]_0 = [k] \cup \{0\}$.

$\mathcal{P}\mathcal{L}\mathcal{S}$, Reductions and Completeness, [25]. A $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem $L = (D_L, F_L, c_L, N_L, \text{INIT}_L, \text{COST}_L, \text{IMPROVE}_L)$ is defined as follows: The set of instances is given by $D_L \subseteq \{0, 1\}^*$, membership in D_L can be decided in polynomial time. Every instance $I \in D_L$ has a finite set of feasible solutions $F_L(I)$, where feasible solutions $\mathbf{s} \in F_L(I)$ have length bounded by a polynomial in the length of I . Every feasible solution $\mathbf{s} \in F_L(I)$ has a non-negative real cost $c_L(\mathbf{s}, I)$ and a neighborhood $N_L(\mathbf{s}, I) \subseteq F_L(I)$. The three polynomial-time algorithms are as follows: Algorithm $\text{INIT}_L(I)$ computes an initial feasible solution $\mathbf{s} \in F_L$. Algorithm $\text{COST}_L(\mathbf{s}, I)$ computes the cost of a solution $\mathbf{s} \in F_L(I)$. Algorithm $\text{IMPROVE}_L(\mathbf{s}, I)$, given an instance $I \in D_L$ and a feasible solution $\mathbf{s} \in F_L(I)$, finds a better solution in $N_L(\mathbf{s}, I)$ or returns that there is no better one.

We consider *maximization* and *minimization problems*. A solution $\mathbf{s} \in F_L(I)$ is *locally optimal*, if it holds for every neighboring solution $\mathbf{s}' \in N_L(\mathbf{s}, I)$ that $c_L(\mathbf{s}', I) \leq c_L(\mathbf{s}, I)$ in case L is a maximization $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem and $c_L(\mathbf{s}', I) \geq c_L(\mathbf{s}, I)$ in case L is a minimization $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem. A *search problem* R is given by a relation over $\{0, 1\}^* \times \{0, 1\}^*$. An algorithm “solves” R , when given $I \in \{0, 1\}^*$ it computes an $\mathbf{s} \in \{0, 1\}^*$, such that $(I, \mathbf{s}) \in R$ or it correctly outputs that such an \mathbf{s} does not exist. Given a $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem L , let the according search problem be $R_L := \{(I, \mathbf{s}) \mid I \in D_L, \mathbf{s} \in F_L(I) \text{ is a local optimum}\}$. Then, the *class* $\mathcal{P}\mathcal{L}\mathcal{S}$ is defined as $\mathcal{P}\mathcal{L}\mathcal{S} := \{R_L \mid L \text{ is a } \mathcal{P}\mathcal{L}\mathcal{S}\text{-problem}\}$. A $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem L_1 is $\mathcal{P}\mathcal{L}\mathcal{S}$ -reducible to a $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem L_2 (written $L_1 \leq_{\text{pls}} L_2$), if there exist two polynomial-time computable functions $\Phi : D_{L_1} \mapsto D_{L_2}$ and Ψ defined for $\{(I, \mathbf{s}) \mid I \in D_{L_1}, \mathbf{s} \in F_{L_2}(\Phi(I))\}$ with $\Psi(I, \mathbf{s}) \in F_{L_1}(I)$, such that for all $I \in D_{L_1}$ and for all $\mathbf{s} \in F_{L_2}(\Phi(I))$ it holds that, if $(\Phi(I), \mathbf{s}) \in R_{L_2}$, then $(I, \Psi(I, \mathbf{s})) \in R_{L_1}$. A $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem L is $\mathcal{P}\mathcal{L}\mathcal{S}$ -complete if every $\mathcal{P}\mathcal{L}\mathcal{S}$ problem is $\mathcal{P}\mathcal{L}\mathcal{S}$ -reducible to L .

Let L be a $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem and $I \in D_L$ be an instance of L . The *transition graph* $\text{TG}(I)$ of the instance I is a directed graph with one vertex for each feasible solution to I and with an arc $s \rightarrow t$, whenever $t \in N_L(\mathbf{s}, I)$ and $c_L(\mathbf{t}, I)$ is strictly better than $c_L(\mathbf{s}, I)$ (i.e., greater if L is a maximization problem, and smaller if L is a minimization problem). Schäffer and Yannakakis [32] define a $\mathcal{P}\mathcal{L}\mathcal{S}$ -reduction (Φ, Ψ) from $\mathcal{P}\mathcal{L}\mathcal{S}$ -problem L_1 to L_2 to be *tight* if for any instance $I \in D_{L_1}$, there exists a subset \mathcal{R} of feasible solutions for the image instance $J = \Phi(I)$ of L_2 , so that the following properties are satisfied: (1) \mathcal{R} contains all local optima of J . (2) For every feasible solution $\mathbf{s} \in F_{L_1}(I)$, we can construct in polynomial time a solution $\mathbf{t} \in \mathcal{R}$ of J such that $\Psi(\mathbf{t}, I) = \mathbf{s}$. (3) Suppose that the transition graph of J , $\text{TG}(J)$, contains a directed path $q \rightarrow \dots \rightarrow q'$, such that $q, q' \in \mathcal{R}$, but all internal path vertices are outside of \mathcal{R} , and let $p = \Psi(q, I)$ and $p' = \Psi(q', I)$ be the corresponding feasible solution of I . Then, either $p = p'$ or $\text{TG}(I)$ contains an arc from p to p' . The *standard algorithm problem* is to compute a locally optimal solution reachable by a sequence of local improvement steps from a given initial solution.

Problems (p, q, r) -MCA and SCG- (k) . We next describe the \mathcal{PLS} -problems that we consider in this paper. For sake of readability, we write limitations to a problem as a prefix and the size of the neighborhood as a suffix. For the \mathcal{PLS} -problems L studied in this paper, the algorithms INIT_L , COST_L , and IMPROVE_L are straightforward and polynomial-time computable. The size of the neighborhood is limited by a constant that is independent of the size of the input and therefore algorithm $\text{IMPROVE}_L(\mathbf{s}, I)$ can search the neighborhood of $N_L(\mathbf{s}, I)$ in polynomial time.

Problem 1. We first present the base of our reduction – $\text{MAXCONSTRAINTASSIGNMENT}$ – which is an extension of $\text{GENERALIZED SATISFIABILITY}$ (confer problem [L06] in [20] for a formal description) to higher valued variables. We allow variables to take values from a set $[r]$ with $r \in \mathbb{N}$ and we replace the sequence of weighted clauses by a sequence of functions (*constraints*), where each function returns the weight of the constraint for the given assignment. We consider the subclass of instances by limiting the maximum *length* of any constraint (where the length of a constraint is the number of variables it depends on), the *number of appearances* of any variable in all constraints and the *values* any variable can take. The problem is then to compute an assignment maximizing the sum of the weights. A feasible solution \mathbf{a} is locally optimal for an instance I in the change neighborhood (one variable changes its assignment), if $(I, \mathbf{a}) \in R_{(p,q,r)\text{-MCA}}$.

Definition 1 ((p, q, r)-MCA). An instance $I \in D_{(p,q,r)\text{-MCA}}$ of problem (p, q, r) - $\text{MAXCONSTRAINTASSIGNMENT}$ is a set $\mathbf{C} = \{C_1, \dots, C_m\}$ of constraints, where each constraint has length at most p , over a set of variables $\mathbf{X} := \{x_1, \dots, x_n\}$, where each variable appears in at most q constraints, and variables can take values from $[r]$. For every constraint $C_i(x_{i_1}, \dots, x_{i_{p_i}}) \in \mathbf{C}$, there is a function $w_{C_i} : [r]^{p_i} \mapsto \mathbb{R}_{\geq 0}$. The set of feasible solutions $F_{(p,q,r)\text{-MCA}}(I)$ to instance I consists of all assignments $a : \mathbf{X} \mapsto [r]$ of values to variables. Given an assignment $\mathbf{a} \in F_{(p,q,r)\text{-MCA}}(I)$, the cost is $c_{(p,q,r)\text{-MCA}}(\mathbf{a}, I) := \sum_{C_i(x_{i_1}, \dots, x_{i_{p_i}}) \in \mathbf{C}} w_{C_i}(a(x_{i_1}), \dots, a(x_{i_{p_i}}))$. The neighborhood of assignment \mathbf{a} consists of all assignments, where the value of one variable is changed.

Problem 2. The main problem we study in this paper is the $\text{SINGLETONCONGESTIONGAME}$ -problem (SCG-problem). Here, in an instance *weighted agents* choose links from a set of *identical links*. The *cost of an agent* is the load (the sum of the weights of the agents) on the link it chooses. The agents are selfish and try to minimize their individual cost. Agents may form arbitrary, *non-fixed coalitions*. We consider coalitions up to size k , for some $k \in \mathbb{N}$. The cost of a coalition is defined to be the maximum cost of its members. In a *selfish step of a coalition*, a coalition of size at most k improves its cost by unilateral deviation. The *potential function* is defined as the *lexicographical order of the agents' cost*. In each selfish step, the potential function decreases and a local minimum is a *Nash Equilibrium* among coalitions of size at most k . The neighborhood of a feasible assignment (every agent chooses a link) are all assignments, where the cost of some arbitrary non-fixed coalition of at most k reallocating agents decreases. Note that any solution in the neighborhood of a solution \mathbf{a} has a better cost than \mathbf{a} . Denote by *swap* the special neighborhood operation, where coalitions of two or more agents mutual exchange their choice of links.

Definition 2 (SCG- (k)). An instance $I \in D_{\text{SCG-}(k)}$ to problem SINGLETONCONGESTIONGAME is a set of selfish agents $\mathbf{A} = \{a_1, \dots, a_n\}$, a function $w : \mathbf{A} \mapsto \mathbb{N}$ and a number $m \in \mathbb{N}$ of identical links. The set of feasible solutions $F_{\text{SCG-}(k)}(I)$ consists of all assignments of all agents to links. Let $\mathbf{a} : \mathbf{A} \mapsto [m]$ be a function that assigns agents to links. The cost of an agent is the load on the link to which it is assigned. Here, the load on some link $j \in [m]$ is the sum of the weights of the agents assigned to j in \mathbf{a} . The cost of a coalition is the maximum cost of one of its members. The cost $c_{\text{SCG-}(k)}(\mathbf{a}, I)$ of a solution $\mathbf{a} \in F_{\text{SCG-}(k)}(I)$ is the n -vector of the costs of all agents. The order is given by the lexicographical order on vectors. The neighborhood $N_{\text{SCG-}(k)}(\mathbf{a}, I)$ of solution $\mathbf{a} \in F_{\text{SCG-}(k)}(I)$ consists of all assignments, where coalitions of up to k agents can relocate, such that the cost of the coalition decreases.

2.1 Our Contribution

We prove the following theorem where we omitt some parts of the proof due to lack of space. They can be found in the full version, [11].

Theorem 1. $(3, 2, r)$ -MCA \leq_{pls} SCG- (8) via a tight reduction for all $r \in \mathbb{N}$.

3 The General Method

In this section, we present the general method that our reduction relies on. We build on the \mathcal{PLS} -completeness of $(3, 2, r)$ -MCA, [10], and we may assume that every constraint has length three and every r -valued variable appears in two constraints. Our reduction is constructed such that the parameter r does not have a negative influence on the size of the neighborhood. We model the set of variables \mathbf{X} and the set of constraints \mathbf{C} with agents and links.

The Links. We introduce $r - 1$ links for every variable, one link for every constraint and one repository-link. All links have identical speed.

The Agents and Their Weights. Using the weights of the agents, we create a framework where the given MCA-instance is simulated in. Here, the actual weights of the constraints from the MCA-instance play a secondary role compared to the weights of the agents in the SCG-instance that ensure the framework. Of course, there is some similarity between our construction and known \mathcal{NP} -reductions; confer the proof of strong \mathcal{NP} -completeness of BINPACKING (see pp.204 in [30]). We introduce *three types of agents*: The variable-agents simulate the double appearance of every variable and its r values in the set of constraints \mathbf{C} . The frame-agents simulate the assignment of values to variables in every constraint $C_i \in \mathbf{C}$. The base-agents, which are sub-divided into variable-base-agents, constraint-base-agents and one repository-base-agents, create some large base load. The weight of a variable-base-agent for some variable $x \in \mathbf{X}$ is tailored such that the two variable-agents for variable $x \in \mathbf{X}$ and some *arbitrary* valuedness fit on a link and create a load, which is equal to the average load on every link. The weight of the constraint-base-agent is tailored to fit one frame-agent and the three variable-agents for the respective assignment and their load is the average load on every

link *plus* the cost of the assignment for the constraint. Furthermore, the weight of the repository-base-agent is tailored to fit on a link with all remaining frame-agents, which are not on a link with one constraint-base-agent or a variable-base-agent.

The Core Ideas. In the design of our reduction, we bear the size of the neighborhood in mind at all times. A key is to use a *direct reduction* from the MCA-problem, where *every variable appears exactly twice*. On the upside, this double appearance of variables significantly lowers the size of the neighborhood required to simulate the MCA-problem and is conceptionally crucial in the design of the variable-agents. We want to stress that for triple appearances, our construction would not work. On the downside, it leaves us with a higher valuedness and here, a core idea is to *mask the actual valuedness* of variable-agents with matching numbers. Furthermore, we only require *three types of agents*. The variable-agents and the frame-agents naturally arise from the MCA-problem and the only newly introduced base-agents serve as a *classification* of the identical links. With these three types of agents, we are able to ensure the intended *framework* in every locally optimal solution with the exchange of at most 6 agents. We prove this technical main result in Lemma 5. Furthermore, the simulation of the local search in the MCA-instance is possible by exchanging at most 8 agents.

4 SCG-(8) Is \mathcal{PLS} -Complete

In this section, we prove Theorem 1. We present the reduction function Φ and the solution mapping Ψ . Given an instance $I \in D_{(3,2,r)\text{-MCA}}$, we construct an instance $\Phi(I) = (\mathbf{A}, w, m) \in D_{\text{SCG-}(k)}$ consisting of a set of agents \mathbf{A} , a weight function $w : \mathbf{A} \mapsto \mathbb{N}$ that maps weights to agents and a number $m \in \mathbb{N}$ of identical links. We assume that in instance $I \in D_{(3,2,r)\text{-MCA}}$, every constraint $C_i \in \mathbf{C}$ has length 3, where clauses are given in natural order, every variable $x \in \mathbf{X}$ appears in 2 constraints and takes values from $[r]$. Furthermore, we may assume that the sum of the weights of two constraints is larger than the maximum weight of a single constraint (otherwise, we can add the weight of the largest constraint to all other constraints and this does not modify the set of local optima). Let $|\mathbf{C}| = m$ and $|\mathbf{X}| = n$. We create an instance of SCG-(8) with $N = (r - 1) \cdot n + m + 1$ identical links, $r - 1$ links for every variable, one link for every constraint, and one repository-link, and the following agents:

- For every $1 \leq i \leq n, 1 \leq k \leq r$, we introduce two *variable-agents* v_{ik1} and v_{ik2} with $w(v_{ik1}) = M^{2i} + k \cdot M^{2i-1}$ and $w(v_{ik2}) = M^{2i} + (2r + 1 - k) \cdot M^{2i-1}$.
- For every constraint $C_j(x_{i1}, x_{i2}, x_{i3})$, with $C_j \in \mathbf{C}, 1 \leq j \leq m; k_1, k_2, k_3 \in [r], i1 < i2 < i3$, we introduce *frame-agents* $\alpha(j, k_1, k_2, k_3)$ with

$$w(\alpha(j, k_1, k_2, k_3)) = M^{2n+2} - j \cdot M^{2n+1} - M^{2i_1} - \tilde{k}_1 \cdot M^{2i_1-1} - M^{2i_2} \\ - \tilde{k}_2 \cdot M^{2i_2-1} - M^{2i_3} - \tilde{k}_3 \cdot M^{2i_3-1} + w_{C_j}(k_1, k_2, k_3),$$

where

$$\tilde{k}_t = \begin{cases} k_t, & \text{if variable } x_{i_t} \text{ occurs in constraint } C_j \text{ for the first time} \\ (2r + 1 - k_t), & \text{otherwise.} \end{cases}$$

- We also introduce the following *base-agents*: For all $i \in [n], k \in [r - 1]$, we introduce *variable-base-agents* b_{ik} with $w(b_{ik}) = B + A - 2 \cdot M^{2i} - (2r + 1)M^{2i-1}$. For all $j \in [m]$, we introduce *constraint-base-agents* \hat{b}_j with $w(\hat{b}_j) = B + A - M^{2n+2} + j \cdot M^{2n+1}$. We also introduce a *repository-base-agent* b_0 , with $w(b_0) = B$.

Here, the constants A, B, M, W, \hat{M} are defined as follows:

$$\begin{aligned}
 W &= \sum_{j=1}^m \sum_{k_1=1}^r \sum_{k_2=1}^r \sum_{k_3=1}^r w_j(k_1, k_2, k_3) \\
 \hat{M} &= m \cdot r^3 \cdot M^{2n+2} + W \\
 B &= 2 \cdot \hat{M} \\
 A &= A_1 - A_2 \\
 A_2 &= m \cdot M^{2n+2} - \frac{m \cdot (m+1)}{2} M^{2n+1} - 2 \sum_{i=1}^n M^{2i} - (2r+1) \sum_{i=1}^n M^{2i-1} \\
 A_1 &= \sum_{j=1}^m \sum_{k_1=1}^r \sum_{k_2=1}^r \sum_{k_3=1}^r w(\alpha(j, k_1, k_2, k_3)) = W + r^3 \cdot A_2
 \end{aligned}$$

M is chosen large enough to ensure correctness of Lemmas [4.5](#). We first need to make the following definitions, before we are able to present the solution mapping:

Definition 3. We call an assignment admissible, if the base-agents are assigned to different links. For an admissible assignment, we use the following notation: A link to which b_{ik} is assigned for $i \in [n], k \in [r - 1]$, is called an i -variable-link. The link to which \hat{b}_j is assigned for $j \in [m]$ is called j -constraint-link. The link to which b_0 is assigned is called repository-link.

Definition 4. Let $i \in [n]$. An admissible assignment \mathbf{a} of agents to links is i -regular iff there exists a mapping $\pi : [i : n] \mapsto [r]$ such that the following conditions are fulfilled:

1. For each $t \in [i : n]$ and for each $k \in [r] \setminus \{\pi(t)\}$ the two agents v_{tk_1}, v_{tk_2} have chosen the same t -variable-link.
2. For $t \in [m]$ let $C_{j_1}, C_{j_2} \in \mathcal{C}$ with $j_1 < j_2$, be the two constraints containing x_t in their variable list. Then, for each $t \in [i : n]$, agent $v_{t\pi(t)1}$ is assigned to the j_1 -constraint-link and $v_{t\pi(t)2}$ is assigned to the j_2 -constraint-link.
3. For all $j \in [m]$ one frame-agent $\alpha(j, k_1, k_2, k_3)$ is assigned to the j -constraint-link. Furthermore, for each $j \in [m]$ and for each $t \in [i : n]$ if x_t is contained in the variable list of $C_j \in \mathcal{C}$ in position μ with $\mu \in [3]$, then $k_\mu = \pi(t)$.

Note that the number A_2 defined above is the sum of the weights of all constraint-agents assigned to the constraint-links in a 1-regular assignment.

Solution mapping. For a feasible and 1-regular solution $\mathbf{a} \in F_{\Phi(I)}$, function $\Psi(I, \mathbf{a})$ returns the mapping π . Otherwise, function $\Psi(I, \mathbf{a})$ returns some assignment of values to variables, which is not locally optimal for $I \in D_{(3,2,r)\text{-MCA}}$.

Correctness of the Construction. In the following, we prove properties of locally optimal solutions \mathbf{a} for $\Phi(I)$ in different move-neighborhoods up to size 6. We prove that in every locally optimal solution the intended framework is obeyed. In detail, we show that every locally optimal assignment \mathbf{a} is admissible (Lemma 1), exactly one frame-agent is assigned to every constraint-link (Lemma 2) and eventually that \mathbf{a} is 1-regular (Lemmas 4 and 5).

Lemma 1. *In every locally-optimal assignment \mathbf{a} for $\Phi(I)$ in the 1-move-neighborhood, the base-agents are assigned to different links.*

Lemma 2. *In every locally-optimal assignment \mathbf{a} for $\Phi(I)$ in the 2-move-neighborhood, exactly one frame-agent is assigned to each j -constraint-link for every $j \in [m]$. All other frame-agents are assigned to the repository-link.*

Lemma 3. *For every locally-optimal assignment \mathbf{a} for $\Phi(I)$ in the 2-move-neighborhood, the load on every link $\lambda \in [N]$ can be written as*

$$A + B + \sum_{t=1}^{2n+1} \gamma_t(\lambda) M^t + \gamma_0(\lambda),$$

with $\gamma_0(\lambda) \leq W$ and $\gamma_t(\lambda) \leq \sqrt{M}$ for all $t \in [2n + 1]$.

We next prove a property of i -regular assignments, which we need in the proof of Lemma 5.

Lemma 4. *For all $i \in [n + 1]$, every locally-optimal assignment \mathbf{a} for $\Phi(I)$ in the 2-move-neighborhood is an i -regular assignment if $\gamma_t(\lambda) = 0$ for all $t \in [2i - 1 : 2n + 1]$, $\lambda \in [N]$.*

Lemma 5. *Every locally optimal solution \mathbf{a} for $\Phi(I)$ in the 6-move-neighborhood is a 1-regular assignment.*

Proof. We will show for all $t \in [2n + 2]$ by downward induction on t that $\gamma_t(\lambda) = 0$ holds for all $\lambda \in [N]$. From this, the claim follows because of Lemma 4. We introduce a new notation: For some agent a , denote by $w_l(a)$ the factor k that M^l is multiplied with in the weight $w(a)$ of agent a .

Basis for $t = 2n + 2$. Follows from Lemma 1 and Lemma 2.

Induction Step. Next consider $t = 2n + 1$. Because of Lemma 2, no frame-agent is assigned to a variable-link and therefore $\gamma_{2n+1}(\lambda) = 0$ holds for all variable-links λ . We assume now that $\gamma_{2n+1}(\lambda) \neq 0$ holds for some link λ . We distinguish three cases defined by the load on the repository-link λ_0 and show that in each case the assignment is not locally optimal in the 6-move-neighborhood.

- (a) If $\gamma_{2n+1}(\lambda_0) > 0$ holds, then there exists some constraint-link λ_1 with $\gamma_{2n+1}(\lambda_1) < 0$. Let λ_1 be a j -constraint-link and let j be the largest number with this property. Consider some frame-agent $\alpha = \alpha(j, \cdot, \cdot, \cdot)$. α has chosen some link λ which is either a constraint-link or the repository-link λ_0 . If λ is a constraint-link then $\gamma_{2n+1}(\lambda) > 0$ due to the choice of j . In both cases, the coalition formed by α and the frame-agent that has chosen link λ_1 improves by a swap.

- (b) For the case $\gamma_{2n+1}(\lambda_0) < 0$ a contradiction is shown in the same way. In this case, some j -constraint-link λ_1 exists with $\gamma_{2n+1}(\lambda_1) > 0$. j is chosen to be the smallest number with this property. Let α be some frame-agent $\alpha = \alpha(j, \cdot, \cdot, \cdot)$. Then, the coalition formed by α and the frame-agent that has chosen link λ_1 improves by a swap.
- (c) In the third case $\gamma_{2n+1}(\lambda_0) = 0$ holds and there exist some j -constraint-links λ_μ , where $\mu \in \{1, 2\}$, with $\gamma_{2n+1}(\lambda_1) > 0$ and $\gamma_{2n+1}(\lambda_2) < 0$. Let j_1 be the smallest number with $\gamma_{2n+1}(\lambda_1) > 0$ and let j_2 be the largest number with $\gamma_{2n+1}(\lambda_2) < 0$. Consider some arbitrary frame-agents $\alpha_1 = \alpha(j_1, \cdot, \cdot, \cdot)$ and $\alpha_2 = \alpha(j_2, \cdot, \cdot, \cdot)$. If α_1 has chosen some constraint-link λ , then $\gamma_{2n+1}(\lambda) < 0$ holds due to the definition of j_1 . The coalition formed by α_1 and the frame-agent that has chosen λ_1 can improve by a swap. If α_2 has chosen some constraint-link λ , then $\gamma_{2n+1}(\lambda) > 0$ holds due to the definition of j_2 . The coalition formed by α_2 and the frame-agent that has chosen λ_2 can improve by a swap. So, there remains to consider the case that α_1 and α_2 both have chosen the repository-link λ_0 . In this case, the coalition formed by α_1, α_2 , frame-agent $\hat{\alpha}_1$ that has chosen λ_1 , and frame-agent $\hat{\alpha}_2$ that has chosen λ_2 can improve by a simultaneous swap between α_1 and $\hat{\alpha}_1$ and between α_2 and $\hat{\alpha}_2$.

Now, assume that the claim holds for $t + 1 \leq 2n + 1$ and we have to show that it holds also for t . By induction hypothesis $\gamma_r(\lambda) = 0$ holds for all $r \in [t + 1, 2n + 1]$ and for all links λ . We distinguish two cases:

1. t is an *even* number, i.e. $t = 2i$ for some i . Because of Lemma [4](#), \mathbf{a} is an $(i + 1)$ -regular assignment. In order to prove the claim, we have to look mainly at the placements of the agents $v_{ik1}, v_{ik2}, k \in [r]$. Let I_i be the set of these agents. Furthermore, let Ω_i be the set containing all i -variable-links, $k \in [r - 1]$, and those j -constraint-links such that x_i occurs in $C_j \in \mathbf{C}$. Then, in an i -regular assignment, all agents from I_i are assigned to links from Ω_i . Up to now, we only know that \mathbf{a} is an $(i + 1)$ -regular assignment. This implies that $\gamma_{2i}(\lambda) \geq 0$ holds for all links $\lambda \notin \Omega_i$ and if $\gamma_{2i}(\lambda) > 0$ holds for some link $\lambda \notin \Omega_i$ then the following two properties hold:

- There exists a link $\lambda' \in \Omega_i$ with $\gamma_{2i}(\lambda') < 0$.
- An agent $\beta \in I_i$ is assigned to λ and even after removing β from λ , still $\gamma_{2i}(\lambda) \geq 0$ holds.

Therefore, β can improve by moving from λ to λ' . Thus, we have shown that $\gamma_{2i}(\lambda) = 0$ holds for all links $\lambda \notin \Omega_i$. Now, consider $\lambda \in \Omega_i$. If $\gamma_{2i}(\lambda) > 0$ holds for some $\lambda \in \Omega_i$ then there exists also $\lambda' \in \Omega_i$ with $\gamma_{2i}(\lambda') < 0$ and an agent $\beta \in I_i$ is assigned to λ . The agent β can improve by moving from λ to λ' .

2. t is an *odd* number, i.e. $t = 2i - 1$ for some i . Let I_i, Ω_i be defined as in [\(1\)](#). Let $C_{j_1}, C_{j_2} \in \mathbf{C}$, with $j_1 < j_2$, be the two constraints containing variable i . By induction hypothesis \mathbf{a} is an $(i + 1)$ -regular assignment. To each variable-link in Ω_i there are assigned two variable-agents from I_i and the remaining two variable-agents from I_i are assigned to the two constraint-links (one to each) in Ω_i . This implies that $\gamma_{2i-1}(\lambda) = 0$ holds for $\lambda \notin \Omega_i \cup \{\lambda_0\}$ and that no variable agent from I_i is assigned to the repository-link λ_0 . Thus, we know that $\sum_{\lambda \in \Omega_i \cup \{\lambda_0\}} \gamma_{2i-1}(\lambda) = 0$ holds and we have to show that $\gamma_{2i-1}(\lambda) = 0$ holds for all $\lambda \in \Omega_i \cup \{\lambda_0\}$. We

introduce a new notion. Let $p \in [r]_0$. We call an assignment (i, p) -regular iff the following conditions are fulfilled:

- (i) $\gamma_t(\lambda) = 0$ for all $t \in [2i : 2n + 1]$ and all $\lambda \in N$
- (ii) $\gamma_{2i-1}(\lambda) = 0$ for $\lambda \in \Omega_i \cup \{\lambda_0\}$
- (iii) For every $l \in [p]$ the two agents v_{il1}, v_{il2} are assigned to some subset $\Omega_i(p) \subset \Omega_i$ of links with $\gamma_{2i-1}(\lambda) = 0$ for all $\lambda \in \Omega_i(p)$.

Note that property (i, p) -regular implies that for all $1 \leq l \leq p$ the two agents v_{il1}, v_{il2} are assigned to the same variable-link from Ω_i (and no other agent from I_i is assigned to this link) or they are assigned to the j_1 -constraint-link and the j_2 -constraint-link (and no other agent from I_i is assigned to these constraint-links) and some frame-agents $\alpha_1 = \alpha(j_1, \cdot, \cdot, \cdot)$, $\alpha_2 = \alpha(j_2, \cdot, \cdot, \cdot)$ with $w_{2i-1}(\alpha_1) = -l$, $w_{2i-1}(\alpha_2) = -(2r + 1 - l)$ have chosen these constraint-links. We have to show that \mathbf{a} is (i, p) -regular for $p = r$ and we will do it by induction on p . For $p = 0$, the claim follows from the induction hypothesis as seen above. Now, assume that \mathbf{a} is $(p - 1)$ -regular. We will show that \mathbf{a} is also p -regular. Let v_{ipq} be assigned to link $\lambda_q \in \Omega_i$ for $q \in [2]$. Then $\lambda_1, \lambda_2 \notin \Omega_i(p - 1)$. If $\lambda_1 = \lambda_2$, then the claim is proved also for $l = p$. Now, let $\lambda_1 \neq \lambda_2$. We distinguish 4 cases:

- (a) λ_1 and λ_2 are both variable-links. Let $\hat{v}_q, q \in [2]$ be the other i -variable-agent assigned to λ_q . Then, $w(v_{ip1}) < w(\hat{v}_q) < w(v_{ip2})$ for $q \in [2]$. So, $\gamma_{2i-1}(\lambda_2) > 0$ holds and the coalition formed by v_{ip1} and \hat{v}_2 would improve by a swap, a contradiction.
- (b) λ_1 and λ_2 are both constraint-links. Then, $\gamma_{2i-1}(\lambda_1) + \gamma_{2i-1}(\lambda_2) + \gamma_{2i-1}(\lambda_0) = 0$ holds, since exactly all constraint-agents for constraints C_{j_1}, C_{j_2} , and the two variable-agents for variable i with value p have chosen the links $\lambda_0, \lambda_1, \lambda_2$. If $\gamma_{2i-1}(\lambda_1) = \gamma_{2i-1}(\lambda_2) = \gamma_{2i-1}(\lambda_0) = 0$ holds, then the claim is proved also for $l = p$. Otherwise, two cases have to be considered. If λ_1 is a j_2 -constraint-link and λ_2 is a j_1 -constraint-link, then $\gamma_{2i-1}(\lambda_2) > 0$, $\gamma_{2i-1}(\lambda_1) < 0$ and the coalition formed by v_{ip1} and v_{ip2} can improve by a swap. If, on the other hand, λ_μ is j_μ -constraint-link for $\mu \in \{1, 2\}$, then the coalition formed by the two frame-agents assigned to λ_1 and λ_2 and two suitably chosen frame-agents assigned to λ_0 could improve by a simultaneous swap.
- (c) λ_1 is a constraint-link and λ_2 is a variable-link. Let \hat{v}_2 be the other i -variable-agent assigned to λ_2 and let \hat{v}_1 be the i -variable-agent assigned to the other constraint-link λ_3 from Ω_i . Then $w(v_{ip1}) < w(\hat{v}_q) < w(v_{ip2})$ holds for $q \in [2]$ and $\gamma_{2i-1}(\lambda_2) > 0$. Let $\alpha_1 = \alpha(\hat{j}_1, \cdot, \cdot, \cdot)$ be the frame-agent that has chosen λ_1 and let $\alpha_2 = \alpha(\hat{j}_2, \cdot, \cdot, \cdot)$ be the frame-agent that has chosen λ_3 . Note that $\{\hat{j}_1, \hat{j}_2\} = \{j_1, j_2\}$. If $\gamma_{2i-1}(\lambda_1) < 0$, then the coalition formed by v_{ip1} and \hat{v}_2 can improve by a swap. It is $w_{2i-1}(\alpha(j_2, \cdot, \cdot, \cdot)) \leq -(r + 1)$ and therefore, $\hat{j}_1 = j_2$ implies $\gamma_{2i-1}(\lambda_1) < 0$. So, we assume now $\gamma_{2i-1}(\lambda_1) \geq 0$, $\hat{j}_1 = j_1$, and $\hat{j}_2 = j_2$. If $\gamma_{2i-1}(\lambda_0) \geq 0$, then $w_{2i-1}(\alpha_1) + w_{2i-1}(\alpha_2) \leq -(2r + 1)$ and together with $w(v_{ip1}) + w(\hat{v}_1) < 2r + 1$, this implies $\gamma_{2i-1}(\lambda_1) + \gamma_{2i-1}(\lambda_3) < 0$ and therefore $\gamma_{2i-1}(\lambda_3) < 0$. If $\gamma_{2i-1}(\lambda_0) > 0$, then the coalition formed by α_2 and some frame-agent $\beta = \alpha(j_2, \cdot, \cdot, \cdot)$, which has chosen λ_0 , with $w_{2i-1}(\beta) = w_{2i-1}(\alpha_2) + 1$ improves by a swap. If $\gamma_{2i-1}(\lambda_0) = 0$, then the coalition formed by $\alpha_1, \alpha_2, \hat{v}_1, v_{ip2}$, and the two frame-agents $\hat{\alpha}_1 = \alpha(\hat{j}_1, \cdot, \cdot, \cdot)$ and $\hat{\alpha}_2 = \alpha(j_2, \cdot, \cdot, \cdot)$ with $w_{2i-1}(\hat{\alpha}_1) = -p$ and $w_{2i-1}(\hat{\alpha}_2) =$

$-(2r + 1 - p)$ improves by a simultaneous swap between agents α_1 and $\hat{\alpha}_1$, between agents α_2 and $\hat{\alpha}_2$, and between agents v_{ip2} and \hat{v}_i . The last case to consider is that $\gamma_{2i-1}(\lambda_0) < 0$. If $\gamma_{2i-1}(\lambda_1) > 0$, then the coalition formed by α_1 and some frame-agent $\alpha = \alpha(j_1, \cdot, \cdot, \cdot)$ with $w_{2i-1}(\alpha) = w_{2i-1}(\alpha_1) - 1$ improves by a swap. If $\gamma_{2i-1}(\lambda_1) = 0$, then the coalition formed by v_{ip2} , \hat{v}_1 , α_2 , and some frame-agent $\alpha = \alpha(j_2, \cdot, \cdot, \cdot)$ with $w_{2i-1}(\alpha) = -(2r + 1 - p)$ improves by a simultaneous swap between agents α_2 and α , and between agents v_{ip2} and \hat{v}_1 .

- (d) λ_1 is a variable-link and λ_2 is a constraint-link. Let \hat{v}_1 be the other i -variable-agent assigned to λ_1 and let \hat{v}_3 be the i -variable-agent assigned to the other constraint-link λ_3 from Ω_i . Then $\gamma_{2i-1}(v_{ip1}) < \gamma_{2i-1}(w(\hat{v}_q)) < \gamma_{2i-1}(w(v_{ip2}))$ holds for $q \in \{1, 3\}$ and $\gamma_{2i-1}(\lambda_1) < 0$ holds. Furthermore, let α_q , $q \in \{2, 3\}$, be the frame-agent assigned to λ_q . If $\gamma_{2i-1}(\lambda_2) > 0$, then the coalition formed by \hat{v}_1 and v_{ip2} could improve by a swap. Especially, it is $\gamma_{2i-1}(\lambda_2) > 0$, if λ_2 is the j_1 -constraint-link. So, we can assume now that $\gamma_{2i-1}(\lambda_2) \leq 0$, where λ_2 is the j_1 -constraint-link and λ_1 is the j_2 -constraint-link. This implies $\gamma_{2i-1}(\lambda_2) + 2r + 1 - p \leq 0$. We distinguish now two cases:

If $w_{2i-1}(\alpha_2) + w_{2i-1}(\alpha_3) < -(2r + 1)$, then $\gamma_{2i-1}(\lambda_0) > 0$ and the coalition formed by \hat{v}_1 , $v_1(i, p, 2)$, α_2 and some frame-agent $\hat{\alpha}_2$ assigned to λ_0 with $w_{2i-1}(\hat{\alpha}_2) = w_{2i-1}(\alpha_2) + 1$ improves by a simultaneous swap between \hat{v}_1 and v_{ip2} and between α_2 and $\hat{\alpha}_2$.

If, on the other hand, $w_{2i-1}(\alpha_2) + w_{2i-1}(\alpha_3) \geq -(2r + 1)$ then $\gamma_{2i-1}(\lambda_0) \leq 0$. $\gamma_{2i-1}(\lambda_2) \leq 0$ implies $w_{2i-1}(\alpha_3) \geq -(2r + 1) - w_{2i-1}(\alpha_2) \geq -p$. This implies that $\gamma_{2i-1}(\lambda_3) = w_{2i-1}(\alpha_3) + w_{2i-1}(\hat{v}_3) > w_{2i-1}(\alpha_3) + w_{2i-1}(v_{ip1}) \geq 0$. It is $\gamma_{2i-1}(\lambda_3) = w_{2i-1}(\alpha_3) + w_{2i-1}(\hat{v}) > 0$ and therefore there exists a frame-agent $\hat{\alpha}_3$ assigned to λ_0 with $w_{2i-1}(\hat{\alpha}_3) = w_{2i-1}(\alpha_3) - 1$. Furthermore, $\gamma_{2i-1}(\lambda_2) = w_{2i-1}(\alpha_2) + w_{2i-1}(v_{ip2}) = w_{2i-1}(\alpha_2) + 2r + 1 - p \leq 0$ and therefore, there exists a frame-agent $\hat{\alpha}_2$ assigned to λ_0 with $w_{2i-1}(\hat{\alpha}_2) = w_{2i-1}(\alpha_2) + 1$. Then, the coalition formed by \hat{v}_1 , v_{ip2} , α_2 , α_3 and $\hat{\alpha}_2$, $\hat{\alpha}_3$ improves by a simultaneous swap between \hat{v}_1 and v_{ip2} , between α_2 and $\hat{\alpha}_2$, and between α_3 and $\hat{\alpha}_3$. \square

Theorem 11. $(3, 2, r)$ -MCA \leq_{pls} SCG-(8) via a tight reduction for all $r \in \mathbb{N}$.

Proof. Assume that a feasible solution $\mathbf{a} \in F_{\text{SCG-(8)}}$ is locally optimal for $\Phi(I)$, but $\Psi(I, \mathbf{a})$ is not locally optimal for $I \in D_{(3,2,r)\text{-MCA}}$. This implies that there exists a variable $x \in \mathbf{X}$, which can be set from value $k \in [r]$ to a value $l \in [r]$ in some constraints $C_{j_1}, C_{j_2} \in \mathbf{C}$, with $j_1 < j_2$, such that $c_{(3,2,r)\text{-MCA}}$ strictly *increases* by some $\Delta > 0$. By Lemma 5 \mathbf{a} is a 1-regular assignment and all frame-agents that do not chose a constraint-link are on the repository-link, which also has the largest load. Then, there exists a coalition of 8 agents that can improve their coalitional cost. Frame-agents $\alpha(j_1, x_{ik1}, \cdot, \cdot)$ and $\alpha(j_1, x_{il1}, \cdot, \cdot)$ and frame-agents $\alpha(j_2, x_{ik2}, \cdot, \cdot)$ and $\alpha(j_2, x_{il2}, \cdot, \cdot)$ swap links and variable-agents v_{ik1} and v_{il2} and variable-agents v_{ik2} and v_{il2} swaps links. Thereby, the makespan *decreases* by Δ on the repository-link. Thus, \mathbf{a} is not locally optimal, a contradiction.

We define \mathcal{R} to be the set $F_{\text{SCG-}(8)}$. It is obvious to see that our reduction is tight, since the assignment of new values to two variables would require at least 12 agents to swap. \square

References

1. Aarts, E., Korst, J., Michiels, W.: Theoretical Aspects of Local Search. Monographs in Theoretical Computer Science. An EATCS Series. Springer, New York (2007)
2. Aarts, E., Lenstra, J. (eds.): Local Search in Combinatorial Optimization. John Wiley & Sons, Inc., New York (1997)
3. Ackermann, H., Röglin, H., Vöcking, B.: On the Impact of Combinatorial Structure on Congestion Games. In: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), pp. 613–622. IEEE Computer Society, Los Alamitos (2006)
4. Ahuja, R., Ergun, Ö., Orlin, J., Punnen, A.: A Survey of Very Large-Scale Neighborhood Search Techniques. *Discrete Appl. Math.* 123(1-3), 75–102 (2002)
5. Brucker, P., Hurink, J., Werner, F.: Improving Local Search Heuristics for Some Scheduling Problems. Part II. *Discrete Appl. Math.* 72(1-2), 47–69 (1997)
6. Brüggemann, T., Hurink, J., Vredeveld, T., Woeginger, G.: Very Large-Scale Neighborhoods With Performance Guarantees for Minimizing Makespan on Parallel Machines. In: Kaklamani, C., Skutella, M. (eds.) WAOA 2007. LNCS, vol. 4927, pp. 41–54. Springer, Heidelberg (2007)
7. Caragiannis, I., Flammini, M., Kaklamani, C., Kanellopoulos, P., Moscardelli, L.: Tight Bounds for Selfish and Greedy Load Balancing. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 311–322. Springer, Heidelberg (2006)
8. Chien, S., Sinclair, A.: Convergence to Approximate Nash Equilibria in Congestion Games. In: Proceedings of the 18th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA 2007) (2007)
9. Czumaj, A., Vöcking, B.: Tight Bounds for Worst-Case Equilibria. In: Proceedings of the 13th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA 2002), pp. 413–420 (2002); Also accepted *Journal of Algorithms* as Special Issue of SODA 2002
10. Dumrauf, D., Monien, B.: On the \mathcal{PLS} -Complexity of Maximum Constraint Assignment (submitted, 2008)
11. Dumrauf, D., Monien, B.: On the Road to \mathcal{PLS} -Completeness: 8 Agents in a Singleton Congestion Game. Technical report, University of Paderborn (2008)
12. Dumrauf, D., Monien, B., Tiemann, K.: Multiprocessor Scheduling is \mathcal{PLS} -Complete. In: Proceedings of the 42nd Annual Hawaii International Conference on System Sciences (HICSS 42) (2008)
13. Fabrikant, A., Papadimitriou, C., Talwar, K.: The Complexity of Pure Nash Equilibria. In: Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of computing (STOC 2004), pp. 604–612. ACM Press, New York (2004)
14. Feldmann, R., Gairing, M., Lücking, T., Monien, B., Rode, M.: Nashification and the Coordination Ratio for a Selfish Routing Game. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 514–526. Springer, Heidelberg (2003)
15. Fiat, A., Kaplan, H., Levy, M., Olonetsky, S.: Strong Price of Anarchy for Machine Load Balancing. In: Arge, L., Cachin, C., Jurdziński, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 583–594. Springer, Heidelberg (2007)

16. Fotakis, D., Kontogiannis, S., Koutsoupias, E., Mavronicolas, M., Spirakis, P.: The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) ICALP 2002. LNCS, vol. 2380, pp. 123–134. Springer, Heidelberg (2002)
17. Fotakis, D., Kontogiannis, S., Spirakis, P.: Atomic Congestion Games Among Coalitions. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 572–583. Springer, Heidelberg (2006)
18. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B.: Computing Nash Equilibria for Scheduling on Restricted Parallel Links. In: Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC 2004), pp. 613–622 (2004)
19. Gairing, M., Monien, B., Tiemann, K.: Selfish Routing With Incomplete Information. *Theory of Computing Systems* 42(1), 91–130 (2007)
20. Garey, M., Johnson, D.: *Computers and Intractability; A Guide to the Theory of NP-Completeness*. Mathematical Sciences Series. W. H. Freeman & Co., New York (1990)
21. Gonzalez, T.: *Handbook of Approximation Algorithms and Metaheuristics*. Chapman & Hall/CRC Computer & Information Science Series. Chapman & Hall/CRC, Boca Raton (2007)
22. Graham, R.L.: Bounds on Multiprocessing Timing Anomalies. *SIAM Journal on Applied Mathematics* 17(2), 416–429 (1969)
23. Hayrapetyan, A., Tardos, É., Wexler, T.: The Effect of Collusion in Congestion Games. In: Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing (STOC 2006), pp. 89–98. ACM, New York (2006)
24. Hurkens, C., Vredeveld, T.: Local Search for Multiprocessor Scheduling: How Many Moves Does It Take to a Local Optimum? *Oper. Res. Lett.* 31(2), 137–141 (2003)
25. Johnson, D., Papadimitriou, C., Yannakakis, M.: How Easy is Local Search? *Journal of Computer and System Science* 37(1), 79–100 (1988)
26. Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
27. Krentel, M.: Structure in Locally Optimal Solutions (Extended Abstract). In: 30th Annual Symposium on Foundations of Computer Science (FOCS 1989), pp. 216–221 (1989)
28. Mas-Colell, A., Whinston, M., Green, J.: *Microeconomic Theory*. Oxford University Press, Inc., Oxford (1995)
29. Orlin, J., Punnen, A., Schulz, A.: Approximate Local Search in Combinatorial Optimization. In: Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2004), pp. 587–596. Society for Industrial and Applied Mathematics, Philadelphia (2004)
30. Papadimitriou, C.: *Computational Complexity*. Addison Wesley, Reading (1993)
31. Rosenthal, R.W.: A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
32. Schäffer, A., Yannakakis, M.: Simple Local Search Problems That Are Hard to Solve. *SIAM J. Comput.* 20(1), 56–87 (1991)
33. Schuurman, P., Vredeveld, T.: Performance Guarantees of Local Search for Multiprocessor Scheduling. In: Proceedings of the 8th International IPCO Conference on Integer Programming and Combinatorial Optimization, London, UK, pp. 370–382. Springer, Heidelberg (2001)
34. Skopalik, A., Vöcking, B.: Inapproximability of Pure Nash Equilibria. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC 2008), pp. 355–364. ACM, New York (2008)

Conflicting Congestion Effects in Resource Allocation Games

Michal Feldman¹ and Tami Tamir²

¹ School of Business Administration and Center for the Study of Rationality,
Hebrew University of Jerusalem

`mfeldman@cs.huji.ac.il`

² School of Computer Science, The Interdisciplinary Center, Herzliya, Israel
`tami@idc.ac.il`

Abstract. We consider resource allocation games with heterogeneous users and identical resources. Most of the previous work considered cost structures with either negative or positive congestion effects. We study a cost structure that encompasses both the resource's load and the job's share in the resource's activation cost.

We consider the *proportional* sharing rule, where the resource's activation cost is shared among its users proportionally to their lengths. We also challenge the assumption regarding the existence of a fixed set of resources, and consider settings with an unlimited supply of resources.

We provide results with respect to equilibrium existence, computation, convergence and quality. We show that if the resource's activation cost is shared equally among its users, a pure Nash equilibrium (NE) might not exist. In contrast, under the proportional sharing rule, a pure NE always exists, and can be computed in polynomial time. Yet, starting at an arbitrary profile of actions, best-response dynamics might not converge to a NE. Finally, we prove that the price of anarchy is unbounded and the price of stability is between 18/17 and 5/4.

1 Introduction

In resource allocation applications, tasks are assigned to resources to be performed. For example, in job scheduling models, jobs are assigned to servers to be processed, and in network routing models, traffic is assigned to network links to be routed. In the last decade, algorithmic game theory has introduced game theoretic considerations to many of these problems [17,13,21,3,2]. At the heart of the game theoretic view is the assumption that the players have *strategic* considerations and act to minimize their own cost, rather than optimizing the global objective. In resource allocation settings, this would mean that the jobs *choose* a resource instead of being assigned to one by a central designer.

The literature is divided into two main approaches with respect to the cost function. The first class of models emphasizes the negative congestion effect, and assumes that the cost of a resource is some non-decreasing function of its load. Job scheduling [11,23] and selfish routing [10,21] belong to this class of models.

The second class assumes that each resource has some *activation cost*, which should be covered by its users, thus a user wishes to share its resource with additional users in attempt to decrease its share in the activation cost. Roughly speaking, the cost of using a resource in this class is some decreasing function of its load. Positive congestion effects have been considered in network design games [8,5,2].

We claim that in practice both the positive and the negative congestion effect take place. On the one hand, a heavy-loaded resource might be less preferred due to negative congestion effects; on the other hand, resources do have some activation cost, and sharing this cost with other users releases the burden on a single user. Our goal is to combine these two components into a unified cost function. Consequently, the cost function in our model is composed of (i) the load on its resource, and (ii) its share in the activation cost of its chosen resource.

An additional assumption we wish to challenge is the existence of an *a priori* given set of resources. In many practical settings a set of users controlling some jobs have the opportunity to utilize a new resource at their own cost. For example, a user might be able to purchase a dedicated server for his job if he is willing to cover its cost. Consequently, we consider settings in which the number of resources is unlimited *a priori*. (Obviously, the number of resources will never exceed the number of users.)

In our model, each resource is associated with some fixed activation cost, which should be jointly incurred by the set of jobs using it. A crucial question in this setting is how to divide the resource cost among its users. Sharing of joint costs among heterogeneous players is a common problem, and a large number of sharing rules have been proposed for this problem, each associated with different efficiency and fairness properties [15,16,12]. Here, our focus is not on the mechanism design point of view. Rather, we analyze two specific sharing rules with respect to equilibrium existence, computation, convergence and quality. The first rule is the *uniform* sharing rule, under which the resource's cost is shared evenly among its users. The second rule is the *proportional* sharing rule, under which the resource's cost is shared among its users in proportion to their sizes. Note that under both sharing rules, for a sufficiently small activation cost, the unique NE will be one in which each job is processed by a different resource. In the other extreme, for a sufficiently large activation cost (in a sense that will be formalized below), the unique NE will be one in which all the jobs will be assigned to a single resource.

1.1 Our Results

Equilibrium existence: Our game in its general form does not comply with the family of *potential games* (or congestion games), which always admit a NE in pure strategies [19,14]. Thus we need to pursue new techniques for proving equilibrium existence. In particular, as we show, the cost sharing method strongly affects the equilibrium existence. Specifically, in the uniform sharing model a pure NE might not exist, while in the proportional sharing model a pure NE always exists. This motivates the use of this sharing model in our study of the remaining aspects.

Computational complexity: Under a job scheduling model with a fixed number of machines and where a user's cost is the load of its chosen machine, the longest processing time (LPT) algorithm always results in a NE [10]. Here, we devise an algorithm that computes a NE for our setting in polynomial time. The main challenge of the algorithm is to determine the number of active machines.

Convergence to equilibrium: Even if a NE exists, it is not necessarily the case that natural dynamics (like best-response dynamics (BRD), where each job, in turn, performs a best-response to the current profile) always lead to a NE. Yet, in potential games [14], BRD is guaranteed to converge to a NE. BRD is known to converge to a NE both in resource allocation games that ignore the negative congestion effects and in those ignoring the activation costs [6]. However, as we show, this is not the case in our unified model, that is, BRD might not converge to a NE. Yet, if all the jobs are of equal size, the game is a congestion game (as in [2]), and convergence of BRD is guaranteed.

Equilibrium quality: A NE may not be socially optimal. In order to quantify the inefficiency we define an objective function, and compare its value under the optimal solution and its value under some NE.

We quantify the inefficiency according to well-established measurements, namely the *price of anarchy* (PoA) [13,18] and the *price of stability* (PoS) [2]. The PoA is defined as the ratio between the cost of the worst NE and the cost of the optimal solution, while the PoS is defined as the ratio between the cost of the best NE and the cost of the optimal solution. These metrics have been studied in a variety of applications, such as selfish routing [20], job scheduling [13,4], network formation [7,12], facility location [22] and more. The objective function we consider is the egalitarian one, i.e., we wish to minimize the cost of the job that incurs the highest cost. We show that the PoA is not bounded. For the PoS we give an upper bound of 5/4 and a lower bound of 18/17.

All missing proofs are given in the full version of this paper [9].

2 Model and Preliminaries

An instance of our game, $G = \langle I, B \rangle$, consists of a set of n jobs, each associated with length p_j (processing time, bandwidth requirement, etc.). Let $I = \{p_1, \dots, p_n\}$ denote the job lengths. Also given is a set of identical resources $M = \{M_1, M_2, \dots\}$ (machines, links, etc.), each associated with an activation cost B . If the set of machines is limited, we denote $m = |M|$. While our model is general, we use terminology of job scheduling for simplicity of presentation.

The action space S_j of player j is defined as all the individual resources, i.e., $S_j = M$. The joint action space is $S = \times_{j=1}^n S_j$. In a joint action $s \in S$, player j selects machine s_j as its action. We denote by R_i^s the set of players on machine M_i in the joint action $s \in S$, i.e., $R_i^s = \{j : s_j = M_i\}$. The load of M_i in s , denoted by $L_i(s)$, is the sum of the weights of the players that chose machine M_i . In particular, a player can chose to be on a *dedicated* machine (i.e., assigned to a machine with no additional jobs). In this case, $L_i(s) = p_j$.

The cost function of player j , denoted by c_j , maps a joint action $s \in S$ to a real number, and is composed of two components; one depends on the total load on the chosen resource, and the other is its share in the resource’s activation cost. Formally, the cost of player j under a joint action s in which $s_j = M_i$ is $c_j(s) = f(L_i(s), b_j(s))$, where $L_i(s) = \sum_{j \in R_i^s} p_j$ is the total load of players served by M_i , and $b_j(s)$ is j ’s share in the cost B . The function f is increasing in both $L_i(s)$ and $b_j(s)$. In this paper, we assume that $c_j(s) = L_i(s) + b_j(s)$.

The resource’s activation cost may be shared among its users according to different sharing rules, two of which we consider in this paper. Under the *uniform sharing* rule, all the jobs assigned to a particular resource share its cost equally. Formally, a job assigned to M_i under joint action s pays $b_j(s) = B/|R_i^s|$. Under the *proportional sharing* rule, the jobs assigned to a particular resource share its cost proportionally to their sizes. Formally, a job assigned to M_i under joint action s pays $b_j(s) = \frac{p_j B}{L_i(s)}$. For example, let $G = \langle I = \{1, 2\}, B = 12 \rangle$, and let s be the schedule in which both jobs are assigned to the same machine. Then, under uniform sharing $c_1(s) = c_2(s) = 3 + 12/2 = 9$, while under proportional sharing, $c_1(s) = 3 + 12/3 = 7$, $c_2(s) = 3 + 2 \cdot 12/3 = 11$.

Nash Equilibrium (NE): A joint action $s \in S$ is a *pure* Nash Equilibrium if no player $j \in N$ can benefit from unilaterally switching his action.

Let $g(s)$ denote the social cost function under the joint action s . The optimal social cost is $OPT = \min_{s \in S} g(s)$. We consider the egalitarian objective function, in which the goal is to minimize the highest cost some player incurs. Formally, $g(s) = \max_j c_j(s)$. Let $\Phi(G)$ be the set of Nash equilibria of the game G . If $\Phi(G) \neq \emptyset$ then the *PoA (PoS)* is the ratio between the maximal (minimal) cost of a Nash equilibrium and the social optimum, i.e., $\max_{s \in \Phi(G)} g(s)/OPT$ ($\min_{s \in \Phi(G)} g(s)/OPT$).

2.1 Proportional Sharing Rule – Useful Observations

In this section we present several observations that provide some intuition regarding proportional sharing. These observations will be used repeatedly in the sequel. The first observation specifies the conditions under which a job prefers to migrate from one machine to another. Note that in the standard model (where a job’s cost depends only on the load on its chosen machines), the equivalent condition is simply $L_{i'}(s) + p_j > L_i(s)$. In our model, however, a migration might be beneficial even if it involves an increase of load.

Lemma 1. *Consider a schedule s . Suppose $j \in R_i^s$, and let $\rho = \frac{L_i(s)(L_{i'}(s) + p_j)}{p_j}$. Job j reduces its cost by a migration to machine i' if and only if $L_{i'}(s) + p_j > L_i(s)$ and $B > \rho$ or $L_{i'}(s) + p_j < L_i(s)$ and $B < \rho$.*

Proof. The cost of job j under schedule s is $c_j(s) = L_i(s) + p_j B / L_i(s)$. Let s' be the obtained schedule after j ’s migration to machine $M_{i'}$. It holds that $c_j(s') = L_{i'}(s) + p_j + p_j B / (L_{i'}(s) + p_j)$. The assertion follows immediately from comparing $c_j(s)$ and $c_j(s')$.

The following observations provide lower and upper bounds for an agent’s individual cost.

Observation 1. *In any joint action s , for every job j , $c_j(s) \geq 2\sqrt{p_j B}$. Additionally, for every j s.t. $p_j \geq B$, $c_j(s) \geq p_j + B$.*

Observation 2. *In any NE, s , for every job j , $c_j(s) \leq p_j + B$.*

The following observation, whose proof can be easily derived by Lemma 1, provides some insight into beneficial and non-beneficial migrations of jobs.

Observation 3. *(i) A job j of length $p_j < B$ which is assigned to a machine with load smaller than B cannot reduce its cost by migrating to a machine with load greater than B or to a dedicated machine. (ii) Given an assignment s of jobs of lengths smaller than B s.t. $L_{i'}(s) + p_j \geq L_i(s)$ for every i, i' and j assigned to machine M_i , if $L_i(s) + L_{i'}(s) > B$, then no migration is beneficial.*

2.2 Longest Processing Time (LPT) Rule

LPT is a well-known scheduling heuristic [11]. The LPT rule sorts the jobs in a non-increasing order of their lengths and greedily assigns each job to the least loaded machine. In the traditional load-balancing problem, the LPT rule is known to produce a NE [10]. However, the stability of an LPT assignment in our setting is not clear since LPT cares about the machines’ loads solely and does not consider the activation costs. Obviously, under an unlimited supply of resources, LPT will simply assign each job to a new machine, and the resulting schedule is not necessarily a NE. A natural generalization of LPT, in which each job is assigned to a machine minimizing its cost, does not necessarily lead to a NE either, even with unit-size jobs (consider for example $G = \langle I = \{1, 1, 1, 1\}, B = 4 - \epsilon \rangle$). In this paper we use a variant of LPT (see Sections 3 and 4). The next lemma provides an important non-trivial property of the LPT algorithm, to be used in the sequel.

Lemma 2. *Let I be a set of jobs s.t. $p_j < C$ for every j . Let m be the minimal number of machines s.t. an LPT-schedule of I on m machines has makespan at most C . The total load on any two machines in the LPT-schedule on m machines is greater than C .*

3 Equilibrium Existence and Computation

3.1 No Equilibrium under the Uniform Sharing Rule

Under the uniform sharing rule a pure NE might not exist. Consider for example the instance $G = \langle I = \{1, 10\}, B = 4 \rangle$. On dedicated machines, the jobs’ costs are 5 and 14 respectively. If they are assigned together, each job pays 13. Thus, no schedule is stable: the short job will escape to a dedicated machine, while the long job will join it. This example motivates the use of the proportional sharing rule.

3.2 Equilibrium under the Proportional Sharing Rule

In this section we prove that under the proportional sharing rule and unlimited supply of resources a pure NE always exists. Moreover, a NE can be found in time $O(n \log^2 n)$. Our algorithm, denoted LPT*, uses as a subroutine the assignment rule *Longest Processing Time* (LPT) [11]. Given an instance I , let $I_{short} \subseteq I$ be the subset of jobs having length less than B , and let $I_{long} = I \setminus I_{short}$.

Algorithm LPT*:

1. Schedule each of the jobs in I_{long} on a dedicated machine.
2. The jobs of I_{short} are scheduled by algorithm LPT. The number of machines, m , is the minimal number of machines such that LPT produces a schedule having makespan at most B (i.e., LPT on $m-1$ machines produces a schedule having makespan more than B).

Note that the number of machines used in the second step is well defined, since all the participating jobs are shorter than B , therefore, a schedule having makespan less than B exists. The running time of LPT* is $O(n \log^2 n)$. Long jobs are identified and scheduled in time $O(n)$, the short jobs are sorted in time $O(n \log n)$ and then LPT is executed at most $\log n$ times (binary search for the right value of m - which is an integer in the range $[1, n]$).

Theorem 4. *The profile \bar{s} obtained by LPT* is a NE.*

Minimal Lexicographic Assignment: In the traditional load balancing game with a fixed number of machines, the minimal lexicographic profile is known to be a NE [10]. In our model, this profile is not well-defined as the number of machines is not fixed. Let \hat{s}_k^* be the lexicographically minimal assignment of I_{short} on k machines. Let m be such that the makespan under \hat{s}_m^* is smaller than B whereas the makespan under \hat{s}_{m-1}^* is at least B . Let \hat{s}^* be the profile in which: (i) every long job is assigned to a dedicated machine, and (ii) the jobs of I_{short} are assigned according to \hat{s}_m^* . The proof of Theorem 4 can be easily tuned to show that \hat{s}^* is a NE. However, this profile cannot be found efficiently. Moreover, as shown in Theorem 9, both \bar{s} and \hat{s}^* might incur arbitrarily large cost compared to the social optimum.

Identical Jobs: A simpler case is when all the jobs have the same length. Note that for this case the uniform and the proportional sharing rule coincide.

Theorem 5. *If all jobs have the same length, a NE can be computed in linear time.*

Limited Supply of Resources: Assume that the number of machines that can be used is limited. Let $m = |M|$ be the given number of machines, and let m^* be the number of machines required by algorithm LPT*. If $m^* \leq m$ then clearly LPT* produces a NE. Otherwise, it can be seen that the assignment according to LPT rule on m machines results produces a NE. Thus,

Theorem 6. *Every resource allocation game under the proportional sharing rule and a limited supply of resources admits a Nash equilibrium in pure strategies. The NE can be computed efficiently*

3.3 Convergence of Best-Response Dynamics

In this section we show that unlike other job scheduling games, in our model best-response-dynamics (BRD) do not necessarily converge to a Nash equilibrium. BRD is a local-search method in which, starting from an arbitrary joint action, in each step, some player is chosen and plays its best-response strategy (i.e., the strategy that minimizes its cost, given the strategies of the other players). By considering the instance $\langle I = \{10, 10, 10, 20\}, B = 72 \rangle$, and the initial joint action $\{(10, 10); (10, 20)\}$, we get:

Theorem 7. *Under proportional sharing, BRD might not converge to a NE.*

Yet, with unit-size jobs the resulting game is a congestion game [19], thus BRD is guaranteed to converge to a NE (note that while the set of resources is not given, a game with a fixed set of n resources is equivalent to our game, thus it is a congestion game). Moreover, one can easily verify that the function $P(s) = \sum_i B \cdot H_{x_i} + \frac{1}{2}x_i^2$ where x_i denotes the number of jobs on machine i , $H_0 = 0$, and $H_k = 1 + 1/2.. + 1/k$, is a potential function for the game.

4 Equilibrium Quality

In this section we provide bounds for the price of anarchy (PoA) and the price of stability (PoS). In particular, we present sufficient condition for having $PoA = PoS = 1$, we show that the PoA is unbounded, and finally, we prove that the PoS is less than $5/4$ and provide an example in which the PoS is $18/17$.

Theorem 8. *If there exists a job j s.t. $p_j \geq B$, then $PoA = PoS = 1$.*

Therefore, we would like to analyze the PoA and PoS for instances in which all the jobs have load less than B . We first present an upper bound for the PoA which depends on the length of the longest job. Let $p = \alpha B$ be the length of the longest job in the instance, for some $\alpha < 1$,

Lemma 3. $PoA \leq \frac{1+\alpha}{2\sqrt{\alpha}}$.

However, α can be arbitrarily small, therefore, the PoA is not bounded, as we show below.

Theorem 9. *For any given r , there exist instances for which $PoA > r$, even with unit-size jobs.*

Proof. Given r , let $B = 4 \lceil r \rceil^2$ and consider an instance with B unit-length jobs. An optimal schedule groups the jobs in sets of $\sqrt{B} = 2 \lceil r \rceil$, each paying $2\sqrt{B}$. A possible NE is to schedule all the jobs on a single machine. This is a NE because each job incurs a cost of $B + 1$ which cannot be reduced by migrating to a new machine. In particular, this is the NE produced by LPT*, and by finding the minimal lexicographic assignment. For this instance, $\alpha = 1/B$, and the analysis in the proof of Theorem 3 is tight. Moreover, the above construction can be repeated with B^{z+1} jobs, each of length $1/B^z$ to get $PoA = \Omega(B^{O(z/2)})$.

For standard load balancing games, it is well-known that the price of stability is 1, even for the model of unrelated machines [23]. We show that this is not the case in our model. By analyzing the instance $G = \langle I = \{2, 1, 1\}, B = 4 \rangle$, we get

Theorem 10. *In the resource allocation game under the proportional sharing rule, $PoS \geq \frac{18}{17}$.*

On the other hand, the price of stability is bounded by a small constant:

Theorem 11. *In any resource allocation game under the proportional sharing rule, $PoS \leq \frac{5}{4}$.*

Proof. Let αB be the length of the longest job in the instance, for $\alpha < 1$. If $\alpha > 0.25$ then by Theorem 3, $PoS < \frac{5}{4}$, and the assertion follows since $PoS \leq PoA$.

Thus, assume that $\alpha \leq 0.25$, and let $c = \sqrt{\alpha}$. Let m be the minimal number of machines such that algorithm LPT on m machines produces a schedule whose makespan is at most $2cB$. Let s be the profile obtained by LPT on m machines. We show that s is a NE: Note that for any $\alpha \leq 0.25$, $c \leq 0.5$ and thus the makespan is at most B . Therefore, by Observation 3(i), no job will migrate to a dedicated machine. Also, by Lemma 2 (applied with $C = 2cB$), the total load on any two machines is at least $2cB$, and since the maximal gap in the load between any two machines is at most αB , we have that for any two machines having loads $L_i, L_{i'}$, it holds that $L_i L_{i'} \geq (c - \frac{\alpha}{2})B(c + \frac{\alpha}{2})B = (\alpha - \alpha^2/4)B^2$. Finally, the load on any machine is at least $(c - \alpha)B$. A known property of schedules produced by LPT is that any migration involves increase in the load. By Lemma 1 such a migration is profitable for a job of length p migrating from load L_i into load $L_{i'}$ only if $B > L_i(L_{i'} + p)/p$. However $L_i(L_{i'} + p)/p = (L_i L_{i'}/p) + L_i \geq ((\alpha - \alpha^2/4)B^2/\alpha B) + (c - \alpha)B = (1 - \alpha/4 + c - \alpha)B > B$ for any $\alpha \leq 0.25$ (since $\sqrt{\alpha} > \frac{5}{4}\alpha$).

The maximal cost of a job in s is at most $2cB + \alpha BB/2cB = \frac{5}{2}\sqrt{\alpha}B$. By Observation 1 the cost of the longest job is at least $2\sqrt{\alpha}B$, thus $PoS \leq \frac{5}{4}$.

References

1. Albers, S., Elits, S., Even-Dar, E., Mansour, Y., Roditty, L.: On Nash Equilibria for a Network Creation Game. In: SODA (2006)
2. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The PoS for Network Design with Fair Cost Allocation. In: FOCS (2004)
3. Anshelevich, E., Dasgupta, A., Tardos, E., Wexler, T.: Near-Optimal Network Design with Selfish Agents. In: STOC (2003)
4. Czumaj, A., Vöcking, B.: Tight Bounds for Worst-case Equilibria. In: SODA (2002)
5. Epstein, A., Feldman, M., Mansour, Y.: Strong Equilibrium in Cost Sharing Connection Games. In: ACMEC (2007)
6. Even-Dar, E., Mansour, Y.: Fast Convergence of Selfish Rerouting. In: SODA (2005)
7. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C., Shenker, S.: On a Network Creation Game. In: PODC (2003)

8. Feigenbaum, J., Papadimitriou, C., Shenker, S.: Sharing the Cost of Multicast Transmissions. *J. of Computer and System Sciences* 63, 21–41 (2001)
9. Feldman, M., Tamir, T.: Conflicting Congestion Effects in Resource Allocation Games, <http://www.faculty.idc.ac.il/tami/Papers/coco.pdf>
10. Fotakis, D., Kontogiannis, S., Mavronicolas, M., Spiraklis, P.: The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) *ICALP 2002*. LNCS, vol. 2380. Springer, Heidelberg (2002)
11. Graham, R.: Bounds on Multiprocessing Timing Anomalies. *SIAM J. Appl. Math.* 17, 263–269 (1969)
12. Herzog, S., Shenker, S., Estrin, D.: Sharing the "Cost" of Multicast Trees: An Axiomatic Analysis. *IEEE/ACM Transactions on Networking* (1997)
13. Koutsoupias, E., Papadimitriou, C.H.: Worst-case Equilibria. In: Meinel, C., Tison, S. (eds.) *STACS 1999*. LNCS, vol. 1563. Springer, Heidelberg (1999)
14. Monderer, D., Shapley, L.S.: Potential Games. *Games and Economic Behavior* 14, 124–143 (1996)
15. Moulin, H., Shenker, S.: Serial Cost Sharing. *Econometrica* 60, 1009–1037 (1992)
16. Moulin, H., Shenker, S.: Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Journal of Economic Theory* 18, 511–533 (2001)
17. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
18. Papadimitriou, C.: Algorithms, Games, and the Internet. In: *STOC* (2001)
19. Rosenthal, R.W.: A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
20. Roughgarden, T.: The Price of Anarchy is Independent of the Network Topology. In: *STOC* (2002)
21. Roughgarden, T., Tardos, E.: How Bad is Selfish Routing? *Journal of the ACM* 49(2), 236–259 (2002)
22. Vetta, A.R.: Nash Equilibria in Competitive Societies with Applications to Facility Location, Traffic Routing and Auctions. In: *FOCS* (2002)
23. Vöcking, B.: Selfish Load Balancing. In: *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)

The Price of Malice in Linear Congestion Games

Aaron Roth*

Department of Computer Science
Carnegie Mellon University
alroth@cs.cmu.edu

Abstract. We study the *price of malice* in linear congestion games using the technique of no-regret analysis in the presence of Byzantine players. Our assumptions about the behavior both of rational players, and of malicious players are strictly weaker than have been previously used to study the price of malice. Rather than assuming that rational players route their flow according to a Nash equilibrium, we assume only that they play so as to have no *regret*. Rather than assuming that malicious players myopically seek to maximize the social cost of the game, we study Byzantine players about whom we make no assumptions, who may be seeking to optimize any utility function, and who may engage in an arbitrary degree of counter-speculation. Because our assumptions are strictly weaker than in previous work, the bounds we prove on two measures of the price of malice hold also for the quantities studied by Babaioff et al. [2] and Moscibroda et al. [15]. We prove tight bounds both for the special case of parallel link routing games, and for general congestion games.

1 Introduction

The price of anarchy measures the deterioration of performance in a system due to selfishness and lack of coordination. It is a brittle measure however, since it assumes that all agents in the system are perfectly rational and adeptly seek to minimize their own cost. In real systems, agents vary in their rationality, computational power, access to information, and objectives. In the case of malicious users, they may seek to harm particular individuals or general social welfare, and may be myopic or able to engage in a high degree of counter-speculation. We would therefore like to be able to characterize the deterioration of performance in a system containing both selfish but rational agents, as well as Byzantine agents. We have a choice as to how to model both the rational agents and the Byzantine agents, and in both cases, we make very weak assumptions: we assume that the rational agents play so as to experience no *regret*, and we make no assumptions at all about the behavior of the Byzantine agents.

We bound the degradation in social welfare due to Byzantine players for the class of non-atomic congestion games with linear edge costs. In non-atomic congestion games, there are a set of source-sink pairs, and for each source-sink pair (s_i, t_i) there exists a continuum of players who each choose among $s_i \rightarrow t_i$

* Partially Supported by an NSF Graduate Research Fellowship.

paths, which induces a flow along the edges of the paths. Each edge has a load-dependent latency function, which in this paper takes the form $\ell_e(x) = a_e x + b_e$ for $a_e, b_e \geq 0$. In a game with a set of agents of measure 1, we model a set of measure $(1 - v)$ rational agents who wish to minimize their own latency, and a set of measure v Byzantine agents about whom we make no assumptions.

We define social cost to be the average latency experienced by the rational players, and we consider two measures of the degradation of social welfare due to the presence of the Byzantine players. The *price of malice* measures the ratio of the social cost in the presence of v Byzantine flow to the optimal social cost without Byzantine flow, and is the analogue of the quantity studied by Moscibroda, Schmid, and Wattenhofer [15] (also termed “price of malice”). The *differential price of malice* measures the marginal cost to the rational players incurred by introducing ϵ Byzantine flow – in effect the brittleness of the Nash flow to Byzantine players – and is the analogue of the quantity studied by Babaioff, Kleinberg, and Papadimitriou [2] (also termed “price of malice”). Upper bounding this quantity was posed in [2] as an important open problem. Our definitions of the price of malice and the differential price of malice allow for a far wider range of adversarial behavior than those defined by Moscibroda et al. [15] and Babaioff et al. [2], and the upper bounds we prove hold also for the quantities studied in the more restricted settings of [15] and [2].

We model Byzantine players who may behave arbitrarily by using the no-regret framework recently introduced by Blum et al. [5] to bound the *price of total anarchy*. The price of total anarchy compares the average social cost over T rounds of repeated play to the cost of the optimal flow, when the rational players have no *regret*. This is a strictly more general assumption than that rational players play according to a Nash equilibrium, since players in a Nash equilibrium all experience no regret. Studying the price of total anarchy instead of the price of anarchy has the advantage that it allows one naturally to model a game in which only a fraction of the players are rational, allowing the others to behave arbitrarily. Moreover, it is known that in both nonatomic and atomic congestion games, the price of total anarchy exactly matches the price of anarchy [4, 5]. Finally, bounding the price of malice in terms of the price of total anarchy has the attraction that there exist simple and efficient algorithms that guarantee regret quickly approaching 0, even in the case that the number of paths is exponential in the description length of the game, and even in the case when players receive information only about their own costs, and not the costs of other paths [12, 9, 11, 13, 10]. Therefore, bounds on the price of malice proven in terms of the price of total anarchy can plausibly be achieved by rational agents with limited computational power and informational awareness.

We consider both the special case in which the congestion game is defined over a graph consisting of m parallel links, and also the general case of congestion games in which the path set of the game need not correspond to any graph. In the case of parallel links, we prove tight bounds on both the price of malice and the differential price of malice, and show that Byzantine flow cannot hurt social welfare at all. In the general case, we prove a tight bound on the price of

malice and a tight bound the differential price of malice for congestion games with scalar latency functions of the form $\ell_e(x) = a_e x$. We omit all proofs for space; they can be found in the full version of this paper.

1.1 Related Work

Moscibroda et al. study a virus inoculation game in which a certain fraction of players are malicious and seek to maximize the sum costs of the rational players. [15]. They define an equilibrium concept in which rational players are extremely risk-averse, and assume that all malicious players are playing a worst-case strategy profile with respect to their own utility. They then define the price of malice with k malicious players to be the ratio of the social cost in equilibria with k malicious players to the social cost in Nash equilibria without any malicious players, which is akin to our definition of the price of malice. Moscibroda et al. also observe that malicious play can improve social welfare, by causing rational players to cooperate [15].

Two papers by Karakostas and Viglas [8] and Babaioff, Kleinberg, and Papadimitriou [2] initiate the study of malicious users in non-atomic congestion games. Both papers consider congestion games in which a fraction of players are rational and wish to minimize their own costs, and a fraction are malicious, and wish to maximize the sum costs of the rational players. They then study (slightly different) notions of equilibria among these rational and malicious players. Babaioff et al. [2] show lower bounds for an alternative definition for price of malice [2]. They also observe that malicious players can improve social welfare (even in the case of linear edge costs), and term this phenomenon the ‘windfall of malice’.

Blum et al. [5] define the price of total anarchy as an alternative to the price of anarchy in quantifying the degradation of social welfare in the presence of selfish players. They show that in many classes of games, the price of total anarchy exactly matches the price of anarchy, and they analyze the price of total anarchy in the presence of Byzantine players in several games. [5].

Chung et al. [7] study the *price of stochastic anarchy* in which players are imperfect and play random actions rather than best responses with some probability. They show that imperfect play can actually improve social welfare, by showing that the price of stochastic anarchy in the load balancing game on unrelated machines is a bounded function of the number of players and machines, whereas the price of anarchy can be unboundedly large.

Our results are most similar to those from Blum et al. [5] and differ from other previous work [6,14,7,15,8,2] in that we make no assumptions about how irrational or malicious agents should behave. As a result, in our model there cannot exist a windfall of malice as there does in the models of malicious but myopic adversaries from [2,15,8], since if nothing else, an adversary can behave like a selfish, rational player. However, since we are modeling more general adversaries, the bounds we prove on the price of malice and the differential price of malice also hold for equilibrium models of adversarial behavior.

2 Preliminaries

2.1 Nonatomic Congestion Games

A nonatomic congestion game is defined by a four-tuple $\mathcal{G} = (E, \{\ell_e\}, \{\mathcal{P}_i\}, \{R_i\})$. E is a finite set of elements which we will refer to as *edges*. There are k player types, and for each player type i there is a set of feasible paths \mathcal{P}_i where for each $P_j \in \mathcal{P}_i$, P_j is a subset of E . R_i is a Lebesgue measurable continuum of agents of type i represented by the interval $[0, \rho_i]$. In total, we say that a congestion game has $s = \sum_{i=1}^k \rho_i$ units of flow. In this paper we will generally assume without loss of generality that $s = 1$. Finally, associated with each edge is a traffic-dependent latency function $\ell_e(x)$, which in this paper will take the form $\ell_e(x) = a_e x + b_e$ for $a_e, b_e \geq 0$. The names ‘edge’ and ‘path’ suggest a graph, and indeed, we often think of congestion games as *traffic routing games*, in which there is an underlying graph G for which E is the edge set, each player type i corresponds to a source sink pair (s_i, t_i) , and \mathcal{P}_i corresponds to the set of simple $s_i \rightarrow t_i$ paths. However, our results hold for general congestion games which need not correspond to any underlying graph.

A *flow* f partitions the set of players according to the set of paths (we say that players in the partition corresponding to path P_i play on path P_i). We denote by A_i^f the set of players who play on path P_i in flow f , and write $f_{P_i} = \int_{A_i^f} 1$. Note that $\sum_{i=1}^k \sum_{P_i \in \mathcal{P}_i} f_{P_i} = 1$. A flow f induces a unique flow on edges: we write that the flow on edge e is $f(e) = \sum_{P_i: e \in P_i} f_{P_i}$. Given a flow f , the latency of each edge e is $\ell_e(f(e))$, and the latency of each path P_i is $\ell_{P_i}(f) = \sum_{e \in P_i} \ell_e(f(e))$. We say that a player who plays on a path P_i experiences cost $\ell_{P_i}(f)$. We will let $\mathcal{F}(\mathcal{G})$ denote the set of all possible flows in a game \mathcal{G} .

The social cost of a flow is the aggregate of player costs. We define a social cost function γ , and say that the cost of a flow f is: $\gamma(f) = \left(\sum_{i=1}^k \sum_{P_j \in \mathcal{P}_i} \int_{A_j^f} \ell_{P_j}(f) \right) = \frac{1}{s} \left(\sum_{e \in E} f(e) \ell_e(f(e)) \right)$. We write $f^* \in \operatorname{argmin}_{f \in \mathcal{F}(\mathcal{G})} \gamma(f)$ to denote an optimal flow, and write $\mathbf{OPT} = \gamma(f^*)$ to denote the cost of the optimal flow. When the game instance is not clear from context, we will write $f_{\mathcal{G}}^*$ and $\mathbf{OPT}_{\mathcal{G}}$.

We will often speak of flows in which a portion of flow of measure v is controlled by (possibly adversarial) Byzantine players, and the remaining $1 - v$ flow is controlled by rational players. In this case, we write $f(e) = f^r(e) + f^b(e)$ where $f^r(e)$ represents the portion of flow on edge e due to rational players, and $f^b(e)$ represents the portion of flow on edge e due to Byzantine players. The Byzantine players can be of any player type. In the presence of Byzantine players, the social cost that we are concerned with is simply the aggregate of rational player costs: $\gamma(f) = \frac{1}{1-v} \left(\sum_{e \in E} f^r(e) \ell_e(f(e)) \right)$.

Definition 2.1. *A flow f in a congestion game \mathcal{G} is a Nash equilibrium if for each player type i and for all $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$, $\ell_{P_1}(f) \leq \ell_{P_2}(f)$.*

Intuitively, a flow f is a Nash equilibrium if no player would like to change his path. In an equilibrium flow, all paths of each type have the same latency.

Proposition 2.2 (Beckmann et al. [3]). For f, \hat{f} two Nash equilibrium flows of \mathcal{G} , $\gamma(f) = \gamma(\hat{f})$.

Therefore we may refer to the cost of a Nash flow of \mathcal{G} which we will write as $\gamma(\mathcal{G})$.

2.2 Anarchy, Regret, and Malice

In this section we define quantities that we will use to characterize the loss of efficiency due to selfishness and “malice.”

Definition 2.3 ([11]). The price of anarchy of an instance of a congestion game \mathcal{G} is defined to be: $\text{PoA}(\mathcal{G}) = \frac{\gamma(\mathcal{G})}{\text{OPT}_{\mathcal{G}}}$. The price of anarchy of the class of congestion games is: $\text{PoA} = \max_{\mathcal{G}} \text{PoA}(\mathcal{G})$.

In this paper, we will assume that rational players play so as to have no regret. Play proceeds in a series of T timesteps, and at time t each player chooses a path, which results in a flow f^t .

Definition 2.4. A player who has played on paths P_{n_1}, \dots, P_{n_T} after T timesteps experiences ϵ -regret if his average cost is no more than that of his best fixed path in hindsight plus an additive ϵ . That is, for a player of type i : $\frac{1}{T} \sum_{t=1}^T \ell_{P_{n_t}}(f^t) \leq \frac{1}{T} \min_{P_i^* \in \mathcal{P}_i} \sum_{t=1}^T \ell_{P_i^*}(f^t) + \epsilon$. If $\epsilon = 0$, we say that the player satisfies the no regret property.

Assuming that rational players play so as to have no regret is a strictly weaker assumption than that they play according to a Nash equilibrium, since in a Nash equilibrium, players experience no regret. A number of efficient algorithms can guarantee players ϵ regret with ϵ quickly approaching 0 with T , even in the case when the number of paths is exponential in the description length of the game, and even when players receive information only about their own costs [12,9,11,13,10]. For simplicity in our paper, we will assume that rational players actually satisfy the no regret property, but all of our results can be carried through with players who experience $\epsilon(T)$ regret with $\epsilon(T) = o(1)$.

Throughout this paper, we study the time averaged cost of the rational players in the presence of Byzantine players. We write $\text{COST}(v) = \frac{1}{T} \sum_{t=1}^T \gamma(f^t)$.

Definition 2.5 (Blum et al. [5]). The price of total anarchy in a game instance \mathcal{G} with v Byzantine flow is the ratio of the worst case average social cost (among the rational players) over T rounds of repeated play to OPT , when $1-v$ flow corresponds to players with the no-regret property, and the remaining v flow behaves arbitrarily. $\text{PoTA}(\mathcal{G}, v) = \max_{f^1, \dots, f^T} \frac{\text{COST}(v)}{\text{OPT}_{\mathcal{G}}}$. where the max is taken over flows $(f^1, \dots, f^T) \in \mathcal{F}(\mathcal{G})^T$ such that a set of players of measure $1-v$ satisfy the no-regret property and the remaining players behave arbitrarily. The price of total anarchy with v Byzantine flow of the class of congestion games is $\text{PoTA}(v) = \max_{\mathcal{G}} \text{PoTA}(\mathcal{G}, v)$.

Observation 2.6 (Blum et al. [5]). Since when playing a Nash equilibrium all players satisfy the no regret property, for any class of games, $\text{PoTA}(0) \geq \text{PoA}$.

In many classes of games, the price of total anarchy matches the price of anarchy exactly, including in congestion games [54].

Proposition 2.7 (Blum et al. [4]). *For the class of non-atomic congestion games, $PoTA(0) = PoA$.*

We now define the price of malice. Our definition is parallel to the quantity studied by Moscibroda et al. [15] (also termed price of malice). In particular, any upper bound that applies to our definition of price of malice also applies to the price of malice in [15].

Definition 2.8. *The price of malice in an instance of a congestion game \mathcal{G} with v Byzantine flow is the ratio of the price of total anarchy with v Byzantine flow and the price of anarchy. $PoM(\mathcal{G}, v) = \frac{PoTA(\mathcal{G}, v)}{PoA(\mathcal{G})} = \frac{PoTA(\mathcal{G}, v)}{PoTA(\mathcal{G}, 0)}$. The price of malice of the class of congestion games is $PoM(v) = \max_{\mathcal{G}} PoM(\mathcal{G})$.*

Finally, we define the differential price of malice, which parallels the quantity studied by Babaioff et al. [2] (also called price of malice). Any upper bound that applies to the differential price of malice also applies to the price of malice as defined in [2].

Definition 2.9. *The differential price of malice is the maximum marginal cost incurred in any game instance when an ϵ fraction of flow is converted from rational to Byzantine: $DPoM = \max_{\mathcal{G}} \frac{d}{d\epsilon}(PoM(\mathcal{G}, \epsilon))|_{\epsilon=0}$.*

In principle, a game may have a large price of total anarchy and a small price of malice or vice versa, although in linear congestion games the two quantities differ only by a factor of 4/3 [16].

It is *not* sufficient to upper bound $PoTA(v)$ to find an upper bound to $DPoM$, since the slope of the price of total anarchy is measured on an instance by instance basis for $DPoM$. We require further conditions:

Observation 2.10. *If the following conditions are met: 1. $g(v) \geq PoTA(v)$ for all non-negative v 2. $g(0) = PoA(\mathcal{G}, 0)$ for all game instances \mathcal{G} then: $DPoM \leq \frac{d}{d\epsilon}(g(\epsilon)/PoA)|_{\epsilon=0}$*

3 Parallel Links

We first consider the case in which the underlying graph G consists of two vertices s and t (the source and sink for all players), and m $s \rightarrow t$ edges with linear latency functions of the form $\ell_e(x) = a_e x + b_e$. This is an interesting special case because instances of parallel link congestion games can have a price of anarchy as high as in the general case [16], and it also serves as a model of the load balancing game on related machines. We bound the price of total anarchy in terms of $\gamma(\mathcal{G})$, the social cost at Nash equilibrium of the instance in question.

Theorem 3.1. *In the parallel links congestion game with linear edge costs, $PoM(v) = 1$ and $DPoM = 0$.*

Since in the Byzantine adversary model, $PoM(v) \geq 1$ and $DPoM \geq 0$, Theorem 3.1 is tight.

4 General Congestion Games

In this section, we consider the general case of linear congestion games. Instances of these congestion games may or may not be defined over an underlying (arbitrary) graph, although we will continue using the language of paths and edges. The game is played over T timesteps, where at time t , the flow on edge e is $f^t(e) = (f^{rt}(e) + f^{bt}(e))$ where $f^{rt}(e)$ is the flow on edge e due to the rational players and $f^{bt}(e)$ is the flow on edge e due to the Byzantine players. For simplicity of presentation, in this section, we consider *adding* v units of Byzantine flow, rather than converting rational flow to Byzantine flow (and so we always have one unit of rational flow). The case in which Byzantine flow replaces rational flow is similar (but leads to more unwieldy equations). We first prove a tight bound on the price of malice for congestion games with linear edge costs of the form $\ell_e(x) = a_e x + b_e$ for $a_e, b_e \geq 0$. We then consider congestion games with scalar edge costs of the form $\ell_e(x) = a_e x$ for $a_e \geq 0$, and bound both the price of malice and the differential price of malice in such games.

The bounds given here are asymptotically tight; Proofs appear in the full version.

Theorem 4.1. *In non-atomic congestion games with linear edge costs: $PoM(v) \leq PoTA(v) \leq \frac{4}{3} + \sqrt{\frac{a \cdot r(v^2 + v)}{OPT}}$ where $a = \max_{e \in E} a_e$ and $r = \max_{P_i} |\{e \in P_i : \ell_e(x) \neq 0\}|$ is the length of the longest path (not including edges with no latency cost).*

We now consider congestion games with scalar edge costs of the form $\ell_e(x) = a_e x$ for some $a_e \geq 0$.

Theorem 4.2. *In non-atomic congestion games with scalar edge costs: $PoM(v) = PoTA(v) \leq 1 + \sqrt{\frac{a \cdot r(v^2 + v)}{OPT}}$ where $a = \max_{e \in E} a_e$ and $r = \max_{P_i} |\{e \in P_i : \ell_e(x) \neq 0\}|$.*

Theorem 4.3. *In non-atomic single-source single-sink congestion games with scalar edge costs, the differential price of malice is at most $DPoM \leq r = \max_{P_i} |\{e \in P_i : \ell_e(x) \neq 0\}|$.*

Acknowledgements. The author would like to thank Mike Dinitz for useful discussions, and Katrina Ligett for her time and patience in reading an early draft of this paper, and her invaluable comments and suggestions.

References

1. Awerbuch, B., Kleinberg, R.: Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches. In: Proceedings of the 36th ACM Symposium on Theory of Computing (STOC) (2004)
2. Babaioff, M., Kleinberg, R., Papadimitriou, C.H.: Congestion games with malicious players. In: EC 2007: Proceedings of the 8th ACM conference on Electronic commerce, pp. 103–112. ACM, New York (2007)
3. Beckmann, M., McGuire, C.B., Winsten, C.B.: Studies in the Economics of Transportation. Yale University Press, New Haven (1956)
4. Blum, A., Even-Dar, E., Ligett, K.: Routing without regret: on convergence to nash equilibria of regret-minimizing algorithms in routing games. In: PODC 2006: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing, pp. 45–52. ACM, New York (2006)
5. Blum, A., Hajiaghayi, M., Ligett, K., Roth, A.: Regret minimization and the price of total anarchy. In: STOC 2008: Proceedings of the fortieth annual ACM symposium on Theory of computing (2008)
6. Brandt, F., Sandholm, T., Shoham, Y.: Spiteful bidding in sealed-bid auctions. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) (2007)
7. Chung, C., Ligett, K., Pruhs, K., Roth, A.: The price of stochastic anarchy. In: Monien, B., Schroeder, U.-P. (eds.) SAGT 2008. LNCS, vol. 4997, pp. 303–314. Springer, Heidelberg (2008)
8. Kalai, A., Vempala, S.: Efficient algorithms for on-line optimization. In: Proceedings of the 16th Annual Conference on Learning Theory, pp. 26–40 (2003)
9. Karakostas, G., Viglas, A.: Equilibria for networks with malicious users. Math. Program. 110(3), 591–613 (2007)
10. Kleinberg, R.: Anytime algorithms for multi-armed bandit problems. In: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm, pp. 928–936. ACM Press, New York (2006)
11. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: Proceedings of 16th STACS, pp. 404–413 (1999)
12. Littlestone, N., Warmuth, M.K.: The weighted majority algorithm. Inf. Comput. 108(2), 212–261 (1994)
13. McMahan, B., Blum, A.: Online geometric optimization in the bandit setting against an adaptive adversary. In: Shawe-Taylor, J., Singer, Y. (eds.) COLT 2004. LNCS, vol. 3120, pp. 109–123. Springer, Heidelberg (2004)
14. Morgan, J., Steiglitz, K., Reis, G.: The spite motive and equilibrium behavior in auctions. Contributions to Economic Analysis & Policy 2(1), 1102–1127 (2003)
15. Moscibroda, T., Schmid, S., Wattenhofer, R.: When selfish meets evil: byzantine players in a virus inoculation game. In: PODC 2006: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing, pp. 35–44. ACM, New York (2006)
16. Roughgarden, T., Tardos, É.: How bad is selfish routing? J. ACM 49(2), 236–259 (2002)

Parimutuel Betting on Permutations^{*}

Shipra Agrawal¹, Zizhuo Wang², and Yinyu Ye³

¹ Department of Computer Science, Stanford University
shipra@stanford.edu

² Department of Management Science and Engineering, Stanford University
zzwang@stanford.edu

³ Department of Management Science and Engineering, Stanford University
yinyu-ye@stanford.edu

Abstract. We focus on a permutation betting market under parimutuel call auction model where traders bet on final rankings of n candidates. We present a *Proportional Betting* mechanism for this market. Our mechanism allows traders to bet on any subset of the n^2 ‘candidate-rank’ pairs, and rewards them proportionally to the number of pairs that appear in the final outcome. We show that market organizer’s decision problem for this mechanism can be formulated as a convex program of polynomial size. Further, the formulation yields a set of n^2 *unique* marginal prices that are sufficient to price the bets in this mechanism, and are computable in polynomial-time. These marginal prices reflect the traders’ beliefs about the marginal distributions over outcomes. More importantly, we propose techniques to compute the joint distribution over $n!$ permutations from these marginal distributions. We show that using a maximum entropy criterion, we can obtain a concise parametric form (with only n^2 parameters) for the joint distribution which is defined over an exponentially large state space. We then present an approximation algorithm for computing the parameters of this distribution. In fact, our algorithm addresses a generic problem of finding the maximum entropy distribution over permutations that has a given mean, and is of independent interest.

1 Introduction

Prediction markets are increasingly used as an information aggregation device in academic research and public policy discussions. The fact that traders must “put their money where their mouth is” when they say things via markets helps to collect information. To take full advantage of this feature, however, we should ask markets the questions that would most inform our decisions, and encourage traders to say as many kinds of things as possible, so that a big picture can emerge from many pieces. Combinatorial betting markets hold great promise on this front. Here, the prices of contracts tied to the events have been shown to reflect the traders’ belief about the probability of events. Thus, the pricing or ranking of possible outcomes in a combinatorial market is an important research topic.

^{*} Research supported in part by NSF DMS-0604513 and Boeing.

We consider a permutation betting scenario where traders submit bids on final rankings of n candidates, for example, an election or a horse race. The possible outcomes are the $n!$ possible orderings among the candidates, and hence there are $2^{n!}$ subset of events to bid on. In order to aggregate information about the probability distribution over the entire outcome space, one would like to allow bets on all these event combinations. However, such betting mechanisms are not only intractable, but also exacerbate the thin market problems by dividing participants attention among an exponential number of outcomes [1] [2]. Thus, there is a need for betting languages or mechanisms that could restrict the possible bid types to a tractable subset and at the same time provide substantial information about the traders' beliefs.

1.1 Previous Work

Previous work on parimutuel combinatorial markets can be categorized under two types of mechanisms: a) *posted price mechanisms* including the Logarithmic Market Scoring Rule (LMSR) of Hanson [2] [3] and the Dynamic Pari-mutuel Market-Maker (DPM) of Pennock [4] b) *call auction models* developed by Lange and Economides [5], Peters et al. [6], in which all the orders are collected and processed together at once. An extension of the call auction mechanism to a dynamic setting similar to the posted price mechanisms, and a comparison between these models can be found in Peters et al. [7].

Chen et al. (2008) [8] analyze the computational complexity of market maker pricing algorithms for combinatorial prediction markets under LMSR model. They examine both permutation combinatorics, where outcomes are permutations of objects, and Boolean combinatorics, where outcomes are combinations of binary events. Even with severely limited languages, they find that LMSR pricing is #P-hard, even when the same language admits polynomial-time matching without the market maker. Chen, Goel, and Pennock [9] study a special case of Boolean combinatorics and provide a polynomial-time algorithm for LMSR pricing in this setting based on a Bayesian network representation of prices. They also show that LMSR pricing is NP-hard for a more general bidding language.

More closely related to our work are the studies by Fortnow et al. [10] and Chen et al. (2006) [11] on *call auction* combinatorial betting markets. Fortnow et al. [10] study the computational complexity of finding acceptable trades among a set of bids in a Boolean combinatorial market. Chen et al. (2006) [11] analyze the auctioneer's matching problem for betting on permutations, examining two bidding languages: *subset bets*, which are bets of the form candidate i finishes in positions x , y , or z or candidate i , j , or k finishes in position x , and *pair bets*, which take the form candidate i beats candidate j . They give a polynomial-time algorithm for matching divisible subset bets, but show that matching pair bets is NP-hard.

1.2 Our Contribution

In this paper, we focus on the problem of *pricing* a call auction under permutation betting scenario. We consider a new mechanism called *Proportional Betting*

for betting on permutations, which is a slightly more generalized form of *Subset Betting* [11], and will be shown to include it as a special case (details in Section 3.2). In proportional betting mechanism, the traders bet on one or more of the n^2 ‘candidate-position’ pairs, and receive rewards proportional to the number of pairs that appear in the final outcome. For example, a trader may place an order of the form “Horse A will finish in position 2 OR Horse B will finish in position 4”. He [9] will receive a reward of \$2 if both Horse A & Horse B finish at the specified positions 2 & 4 respectively; and a reward of \$1 if only one horse finishes at the position specified. The market organizer collects all the orders and then decides which orders to accept in order to maximize his worst case profit.

We propose this proportional betting mechanism as a relaxation of *Fixed reward Betting* where a trader receives a fixed reward (say \$1) if *any* of his horse-position pairs appear in the outcome permutation. We show that the market organizer’s problem is NP-hard for fixed reward betting. Note that a further relaxation of proportional betting would be to allow traders to bet only on individual candidate position pairs (or individual columns or rows like in subset betting [11]), and allow each trader to submit multiple bets. Here, a difference from our model is that in the relaxed model, a trader may place different bids for different bets and an arbitrary subset of his bets could be accepted, rather than all or nothing.

Our results for proportional betting model are described as follows:

- We show that the market organizer’s decision problem for this mechanism can be formulated as a convex program with only $O(n^2 + m)$ variables and constraints, where m is the number of bidders. Further we show that we can obtain, in *polynomial-time*, a small set (n^2) of dual ‘marginal prices’ that satisfy the desired price consistency constraints, and are sufficient to price the bets in this mechanism. The polynomial-time computability of marginal prices in our call auction setting seems particularly interesting considering that computing the n^2 marginal prices that correspond to Hanson’s logarithmic market scoring rule is #P-hard, even under a restricted form of “proportional betting” where traders are allowed to bet only on *individual* candidate-position pairs [8].
- In the second, and perhaps more interesting part of our work, we suggest a maximum entropy criteria to obtain a joint distribution over $n!$ outcomes from the n^2 marginal prices. Although defined over an exponential space, this distribution is shown to have a concise parametric form involving only n^2 parameters. Moreover, it is shown to agree with the maximum-likelihood distribution when prices are interpreted as observed statistics from the traders’ beliefs.

We present an approximation algorithm to compute the parameters of the maximum entropy joint distribution to any given accuracy in (pseudo)-polynomial time [2]. In fact, this algorithm can be directly applied to a generic problem of finding the maximum entropy distribution over permutations that has a given expected value, and is of independent interest.

¹ ‘he’ shall stand for ‘he or she’.

² The approximation factors and running time will be established precisely in the text.

To the best of our knowledge, this is the first result on *pricing* a parimutuel call auction under permutation betting scenario.

2 Parimutuel Call Auction Model

In this section, we briefly describe the Convex Parimutuel Call Auction Model (CPCAM) developed by Peters et al. [6] that will form the basis of our betting mechanism. Consider a market with one organizer and m traders or bidders. There are S states of the world in the future on which the traders are submitting bids. For each bid that is accepted by the organizer and contains the realized future state, the organizer will pay the bidder some fixed amount of money, which is assumed to be \$1 without loss of generality. The organizer collects all the bids and decides which bids to accept in order to maximize his worst case profit.

Let $a_{ik} \in \{0, 1\}$ denote the trader k 's bid for state i . Let q_k and π_k denote the limit quantity and limit price for trader k , i.e., the maximum number of orders requested by trader k , and the maximum price he is willing to pay for the contract, respectively. The number of contracts accepted for trader k is denoted by x_k . x_k is allowed to take fractional values, that is, the orders are ‘divisible’ in the terminology of [11]. Also, let p_i denote the price computed for outcome state i . Below is the convex formulation of the market organizer’s problem given by [6]:

$$\begin{aligned}
 & \max_{x,s,r} \pi^T x - r + \sum_{i=1}^S \theta_i \log(s_i) \\
 & \text{s. t. } \sum_k a_{ik} x_k + s_i = r \quad 1 \leq i \leq S \\
 & \quad 0 \leq x \leq q \\
 & \quad s \geq 0
 \end{aligned} \tag{1}$$

The above convex program maximizes the worst case profit of the organizer which is given by the difference between the total amount of money collected ($\pi^T x$) and the worst case payment made (r). A “parimutuel” state price vector $\{p_i\}_{i=1}^S$ is given by the dual variables associated with the first set of constraints. The parimutuel property implies that if the bidders are charged a price of $\{\sum_i a_{ik} p_i\}$, instead of their limit price, the payouts made to the bidders are exactly funded by the money collected from the accepted orders in the worst-case outcome. $\theta > 0$ represents starting orders needed to guarantee uniqueness of the state price vector. They capture the prior belief of the organizer. The market organizer could actually lose this seed money in some outcomes. However, as shown in [6], infinitesimal quantity of starting orders are sufficient. That is, if we reduce θ uniformly to 0, the price vector converges to a unique limit.

3 Permutation Betting Mechanisms

In this section, we propose new mechanisms for betting on permutations under the parimutuel call auction model described above. Consider a permutation betting scenario with n candidates. Traders bet on rankings of the candidates in

the final outcome. The final outcome is represented by an $n \times n$ permutation matrix, where ij^{th} entry of the matrix is 1 if the candidate i takes position j in the final outcome and 0 otherwise. We propose betting mechanisms that restrict the admissible bet types to ‘set of candidate-position pairs’. Thus, trader k ’s bet will be specified by an $n \times n$ $(0, 1)$ matrix A_k , with 1 in the entries corresponding to the candidate-position pairs he is bidding on. We will refer to this matrix as the ‘bidding matrix’ of the trader. If a trader’s bid is accepted, he will receive some payout in the event that his bid is a “winning bid”.

Depending on how this payout is determined, two variations of this mechanism are examined: a) Fixed Reward Betting and b) Proportional Betting. The intractability of fixed reward betting will provide motivation to examine proportional betting more closely, which is the focus of this paper.

3.1 Fixed Reward Betting

In this mechanism, a trader receives a fixed payout (assume \$1 w.l.o.g.) if *any* entry in his bidding matrix matches with the corresponding entry in the outcome permutation matrix. That is, if M is the outcome permutation matrix, then the payout made to trader k is given by $I(A_k \bullet M > 0)$. Here, the operator ‘ \bullet ’ denotes the Frobenius inner product³, and $I(\cdot)$ denotes an indicator function. The market organizer must decide which bids to accept in order to maximize the worst case profit. Using the same notations as in the CPCAM model described in Section 2 for limit price, limit quantities, and accepted orders, the problem for the market organizer in this mechanism can be formulated as follows:

$$\begin{aligned} & \max \pi^T x - r \\ \text{s. t. } & r \geq \sum_{k=1}^m I(A_k \bullet M_\sigma > 0) x_k \quad \forall \sigma \in \mathcal{S}_n \\ & 0 \leq x \leq q \end{aligned} \quad (2)$$

Here, \mathcal{S}_n represents the set of n dimensional permutations, M_σ represents the permutation matrix corresponding to permutation σ . Note that this formulation encodes the problem of maximizing the worst-case profit of the organizer with no starting orders.

Above is a linear program with exponential number of constraints. We prove the following theorem regarding the complexity of solving this linear program.

Theorem 1. *The optimization problem in (2) is NP-hard even for the case when there are only two non-zero entries in each bidding matrix.*

Proof. The separation problem for the linear program in (2) corresponds to finding the permutation that “satisfies” maximum number of bidders. Here, an

³ The Frobenius inner product, denoted as $A \bullet B$ in this paper, is the component-wise inner product of two matrices as though they are vectors. That is,

$$A \bullet B = \sum_{i,j} A_{ij} B_{ij}.$$

outcome permutation is said to “satisfy” a bidder, if his bidding matrix has at least one coincident entry with the permutation matrix. We show that the separation problem is NP-hard using a reduction from maximum satisfiability (MAX-2-SAT) problem. Then, using the result on equivalence of separation and optimization problem from [12], the theorem follows. A detailed proof can be found in our technical report [13].

This result motivates us to examine the following variation of this mechanism which makes payouts proportional to the number of winning entries in the bidding matrix.

3.2 Proportional Betting

In this mechanism, the trader receives a fixed payout (assume \$1 w.l.o.g.) for each coincident entry between the bidding matrix A_k and the outcome permutation matrix. Thus, the payoff of a trader is given by the Frobenius inner product of his bidding matrix and the outcome permutation matrix. The problem for the market organizer in this mechanism can be formulated as follows:

$$\begin{aligned} & \max \pi^T x - r \\ \text{s. t. } & r \geq \sum_{k=1}^m (A_k \bullet M_\sigma) x_k \quad \forall \sigma \in \mathcal{S}_n \\ & 0 \leq x \leq q \end{aligned} \quad (3)$$

The above linear program involves exponential number of constraints. However, the separation problem for this program is polynomial-time solvable, since it corresponds to finding the maximum weight matching in a complete bipartite graph, where weights of the edges are given by elements of the matrix $(\sum_k A_k x_k)$. Thus, the ellipsoid method with this separating oracle would give a polynomial-time algorithm for solving this problem. This approach is similar to the algorithm proposed in [11] for *Subset Betting*. Indeed, for the case of subset betting [11], the two mechanisms proposed here (fixed and proportional) are equivalent. This is because subset betting can be equivalently formulated under our framework, as a mechanism that allows non-zero entries only on a single row or column of the bidding matrix A_k . Hence, the number of entries that are coincident with the outcome permutation matrix can be either 0 or 1, resulting in $I(A_k \bullet M_\sigma > 0) = A_k \bullet M_\sigma$, for all permutations σ . Thus, subset betting forms a special case of the proportional betting mechanism proposed here, and all the results derived in the sequel for proportional betting will directly apply to it.

4 Pricing in Proportional Betting

In this section, we reformulate the market organizer’s problem for Proportional Betting into a compact linear program involving only $O(n^2 + m)$ constraints. The new formulation is not only faster to solve in practice (using interior point methods) but also generates a compact dual price vector of size n^2 . These ‘marginal prices’ will be sufficient to price the bets in Proportional Betting, and are shown

to satisfy some useful properties. The reformulation will also allow introducing n^2 starting orders in order to obtain unique prices.

Observe that the first constraint in (3) implicitly sets r as the worst case payoff over all possible permutations (or matchings). Since the matching polytope is integral [12], r can be equivalently set as the result of following linear program that computes maximum weight matching:

$$\begin{aligned}
 r &= \max_M (\sum_{k=1}^m x_k A_k) \bullet M \\
 \text{s.t. } & M^T e = e \\
 & M e = e \\
 & M_{ij} \geq 0 \quad i, j \in \{1, \dots, n\}
 \end{aligned} \tag{4}$$

Here e denotes the vector of all 1s (column vector). Taking dual, equivalently,

$$\begin{aligned}
 r &= \min_{v,w} e^T v + e^T w \\
 \text{s.t. } & v_i + w_j \geq \sum_{k=1}^m (x_k A_k)_{ij} \quad i, j \in \{1, \dots, n\}
 \end{aligned} \tag{5}$$

Here, $(x_k A_k)_{ij}$ denotes the ij^{th} element of the matrix $(x_k A_k)$. The market organizer’s problem in (3) can now be formulated as:

$$\begin{aligned}
 \max_{x,v,w} & \pi^T x - e^T v - e^T w \\
 \text{s.t. } & v_i + w_j \geq \sum_{k=1}^m (x_k A_k)_{ij} \quad i, j \in \{1, \dots, n\} \\
 & 0 \leq x \leq q
 \end{aligned} \tag{6}$$

Observe that this problem involves only $n^2 + 2m$ constraints.

Let $Q \in \mathbb{R}^{n \times n}$ represent the dual variables corresponding to the first n^2 constraints in the above problem. It is easy to show that the dual matrix Q is well interpreted as a “parimutuel price”. That is, $Q \geq 0$; and, if we charge each trader k a price of $A_k \bullet Q$ instead of their limit price (π_k), then the optimal decision remains unchanged and the total premium paid by the accepted orders will be equal to the total payout made in the worst case. Further, Q satisfies the following extended definition of “price consistency condition” introduced in [5].

Definition 1. *The price matrix Q satisfies price consistency constraints if and only if for all k :*

$$\begin{aligned}
 x_k = 0 & \quad \Rightarrow Q \bullet A_k = c_k \geq \pi_k \\
 0 < x_k < q_k & \Rightarrow Q \bullet A_k = c_k = \pi_k \\
 x_k = q_k & \quad \Rightarrow Q \bullet A_k = c_k \leq \pi_k
 \end{aligned} \tag{7}$$

That is, a trader’s bid is accepted only if his limit price is greater than the calculated price for the order.

These properties can be shown using the KKT conditions for (6), in a manner similar to [6] where a non-combinatorial setting is considered. However, the dual price Q thus computed is not guaranteed to be unique. To ensure uniqueness, we can use starting orders as discussed for the CPCAM model in Section 2. We introduce one starting order $\theta_{i,j} > 0$ for each candidate-position pair (i,j) . These starting orders can be of possibly infinitesimal quantity and represent the prior

belief of organizer. Refer [13] for detailed proofs of properties of price matrix Q and the implications of introducing starting orders.

To summarize, we have shown that:

Theorem 2. *One can compute in polynomial-time, an $n \times n$ marginal price matrix Q which is sufficient to price the bets in the Proportional Betting mechanism. Further, the price matrix is unique, parimutuel, and satisfies the desired price-consistency constraints.*

5 Pricing the Outcome Permutations

There is analytical as well as empirical evidence that prediction market prices provide useful estimates of average beliefs about the probability that an event occurs [14] [15] [16]. Therefore, prices associated with contracts are typically treated as predictions of the probability of future events. The marginal price matrix Q derived in the previous section associates a price to each candidate-position pair. Also, it is easy to observe that Q is a doubly stochastic matrix (use KKT conditions of problem in (6)). Thus, the distributions given by a row (column) of Q could be interpreted as marginal distribution over positions for a given candidate (candidates for a given position). One would like to compute the complete price vector that assigns a price to each of the $n!$ outcome permutations. This price vector would provide information regarding the joint probability distribution over the entire outcome space. In this section, we discuss methods for computing this complete price vector from the marginal prices given by Q .

Let p_σ denote the price for permutation σ . Then, the marginal constraints on the price vector p are represented as:

$$\begin{aligned} \sum_{\sigma \in \mathcal{S}_n} p_\sigma M_\sigma &= Q \\ p_\sigma &\geq 0 \quad \forall \sigma \in \mathcal{S}_n \end{aligned} \quad (8)$$

Finding a feasible solution under these constraints is equivalent to finding a decomposition of doubly-stochastic matrix Q into a convex combination of $n \times n$ permutation matrices. There are multiple such decompositions possible. For example, one such solution can be obtained using Birkhoff-von Neumann decomposition [17] [18]. Next, we propose a criterion to choose a meaningful distribution p from the set of distributions satisfying constraints in (8).

5.1 Maximum Entropy Criterion

Intuitively, we would like to use all the information about the marginal distributions that we have, but avoid including any information that we do not have. This intuition is captured by the ‘Principle of Maximum Entropy’. It states that the least biased distribution that encodes certain given information is that which maximizes the information entropy. Therefore, we consider the problem of finding the maximum entropy distribution over the space of n dimensional permutations, satisfying the above constraints on the marginal distributions. The problem can be represented as follows:

$$\begin{aligned} \min \quad & \sum_{\sigma \in \mathcal{S}_n} p_\sigma \log p_\sigma \\ \text{s.t.} \quad & \sum_{\sigma \in \mathcal{S}_n} p_\sigma M_\sigma = Q \\ & p_\sigma \geq 0 \end{aligned} \tag{9}$$

The maximum entropy distribution obtained from above has many nice properties. Firstly, as we show next, the distribution has a concise representation in terms of only n^2 parameters. This property is crucial for combinatorial betting due to the exponential state space over which the distribution is defined. Let $Y \in R^{n \times n}$ be the Lagrangian dual variable corresponding to the marginal distribution constraints in (9), and s_σ be the dual variables corresponding to non-negativity constraints on p_σ . Then, the KKT conditions for (9) are given by:

$$\begin{aligned} \log(p_\sigma) + 1 - s_\sigma &= Y \bullet M_\sigma \\ \sum_{\sigma} p_\sigma M_\sigma &= Q \\ s_\sigma, p_\sigma &\geq 0 \quad \forall \sigma \\ p_\sigma s_\sigma &= 0 \quad \forall \sigma \end{aligned} \tag{10}$$

Assuming $p_\sigma > 0$ for all σ , this gives $p_\sigma = e^{Y \bullet M_\sigma - 1}$. Thus, the distribution is completely specified by the n^2 parameters given by Y . Once Y is known, it is possible to perform operations like computing the probability for a given set of outcome permutations, or finding the most probable outcomes.

Further, we show that the dual solution Y is a maximum likelihood estimator of distribution parameters under suitable interpretation of Q .

Maximum likelihood interpretation. For a fixed set of data and an *assumed* underlying probability model, maximum likelihood estimation method picks the values of the model parameters that make the data “more likely” than any other values of the parameters would make them. Let us assume in our model that the traders’ beliefs about the outcome come from an exponential family of distributions D_η , with probability density function of the form $f_\eta \propto e^{\eta \bullet M_\sigma}$ for some parameter $\eta \in R^{n \times n}$. Suppose Q gives a summary statistics of s sample observations $\{M^1, M^2, \dots, M^s\}$ from the traders’ beliefs, i.e., $Q = \frac{1}{s} \sum_k M^k$. This assumption is inline with the interpretation of the prices in prediction markets as mean belief of the traders. Then, the maximum likelihood estimator is given by

$$\begin{aligned} \hat{\eta} &= \arg \max_{\eta} \log f_\eta(M^1, M^2, \dots, M^s) \\ &= \arg \max_{\eta} \log \left(\prod_k \frac{e^{\eta \bullet M^k}}{\sum_{\sigma} e^{\eta \bullet M_\sigma}} \right) \end{aligned} \tag{11}$$

The optimality conditions for the above unconstrained convex program are:

$$\frac{1}{Z} \sum_{\sigma} e^{\eta \bullet M_\sigma} M_\sigma = \frac{1}{s} \sum_k M^k \tag{12}$$

where Z is the normalizing constant, $Z = \sum_{\sigma} e^{\eta \bullet M_\sigma}$. Since $\frac{1}{s} \sum_k M^k = Q$, observe from the KKT conditions for the maximum entropy model given in (10) that $\eta = Y$ satisfies the above optimality conditions. Hence, the parameter Y computed from the maximum entropy model is also the maximum likelihood estimator for the model parameters η .

5.2 Complexity of the Maximum Entropy Model

In this section, we analyze the complexity of solving the maximum entropy model in (9). As shown in the previous section, the solution to this model is given by the parametric distribution $p_\sigma = e^{Y \bullet M_\sigma - 1}$. The parameters Y are the dual variables given by the optimal solution to the following dual problem of (9)

$$\max_Y Q \bullet Y - \sum_\sigma e^{Y \bullet M_\sigma - 1} \tag{13}$$

We prove the following result regarding the complexity of computing the parameters Y :

Theorem 3. *It is #P-hard to compute the parameters of the maximum entropy distribution $\{p_\sigma\}$ over n dimensional permutations $\sigma \in \mathcal{S}_n$, that has a given marginal distribution.*

Proof. We make a reduction from the following problem:

Permanent of a (0, 1) matrix. The permanent of an $n \times n$ matrix B is defined as $\text{perm}(B) = \sum_{\sigma \in \mathcal{S}_n} \prod_{i=1}^n B_{i, \sigma(i)}$. Computing permanent of a (0, 1) matrix is #P-hard [19].

We use the observation that $\sum_\sigma e^{Y \bullet M_\sigma} = \text{perm}(e^Y)$, where the notation e^Y is used to mean component-wise exponentiation: $(e^Y)_{ij} = e^{Y_{ij}}$. For complete proof, see [13].

Interestingly, there exists an FPTAS based on MCMC methods for computing the permanent of any non-negative matrix [20]. Next, we derive a polynomial-time algorithm for approximately computing the parameter Y that uses this FPTAS along with the ellipsoid method for optimization.

5.3 An Approximation Algorithm

Here, we give an outline of the algorithm and present main ideas involved in the analysis. The details along with a complete technical proof can be found in [13].

Using the KKT conditions for the problem, we show that computing optimal Y is equivalent to finding a feasible point in the following bounded convex set:

$$\begin{aligned} \mathbf{K}: \quad & Q \bullet Y - 1 \geq t \\ & \sum e^{Y \bullet M_\sigma} M_\sigma \leq Q \\ & 0 \geq Y_{ij} \geq -\gamma \quad \forall i, j \end{aligned} \tag{14}$$

where $\gamma = \frac{n \log n}{q_{min}}$, $q_{min} = \min\{Q_{ij}\}$, and $t \in [-n \log n - 1, 0]$ is a fixed parameter. Showing this equivalence involves proving upper and lower bounds on optimal Y . Next, we use ellipsoid method to solve this feasibility problem. In each iteration, the ellipsoid method requires to determine if the given iterate Y is feasible, or compute a separating hyperplane, if infeasible. The gradient of a violated constraint forms a natural candidate for separating hyperplane. In the above

problem, both these tasks pose a problem due to the intractability of second set of constraints. Checking feasibility requires computing the quantity:

$$f(Y) := \sum_{\sigma} e^{Y \bullet M_{\sigma}} M_{\sigma}$$

And, the gradient takes the form⁴:

$$\nabla f(Y) = \sum_{\sigma} e^{Y \bullet M_{\sigma}} (M_{\sigma} \otimes M_{\sigma})$$

Both these quantities are #P-hard to compute. We use MCMC method for computing permanent [20] to compute an $(1 + \epsilon)$ -approximation of these quantities. For a fixed $\epsilon > 0$, each iteration of the resulting ellipsoid algorithm looks like this:

Algorithm

1. If Y violates any constraints other than the constraint on $f(Y)$, report $Y \notin \mathbf{K}$. The violated inequality gives the separating hyperplane.
2. Otherwise, compute a $(1 \pm \delta)$ -approximation $\tilde{f}(Y)$ of $f(Y)$, where $\delta = \min\{\frac{\epsilon}{12}, 1\}$.
 - (a) If $\tilde{f}(Y) \leq (1 + 3\delta)Q$, then report Y is feasible.
 - (b) Otherwise, say i_j th constraint is violated. Compute a $(1 \pm \gamma)$ -approximation $\hat{\nabla} f_{ij}(Y)$ of the gradient $\nabla f_{ij}(Y)$, where $\gamma = \delta q_{min}/2n^4$. The approximate gradient $C = \hat{\nabla} f_{ij}(Y)$ gives the desired separating hyperplane.

We show that the above algorithm gives an approximate (pseudo-)polynomial time separating oracle for our problem, in the following sense:

Lemma 1. *Given any $Y \in R^{n \times n}$, and any parameter $\epsilon > 0$, the algorithm with runs time polynomial in n , $1/\epsilon$ and $1/\{\min Q_{ij}\}$ and does one of the following:*

- asserts that $Y \in \mathbf{K}_{\epsilon}$, where \mathbf{K}_{ϵ} represents the set \mathbf{K} with relaxed constraints $f(Y) \leq (1 + \epsilon)Q$.
- or, finds $C \in R^{n \times n}$ such that $C \bullet X \leq C \bullet Y$ for every $X \in \mathbf{K}$.

Thus, the ellipsoid algorithm using this oracle will terminate with either $Y \in \mathbf{K}_{\epsilon}$, or declares that there exists no Y in \mathbf{K} . The proof of the lemma involves proving bounds on the diameter of set \mathbf{K} , and gradient $\nabla f_{ij}(Y)$. The details are available in [13]. Overall, we prove the following theorem (refer [13] for proof):

Theorem 4. *Using the proposed approximate ellipsoid method, a distribution $\{p_{\sigma} \sim e^{Y \bullet M_{\sigma}}\}$ over permutations can be constructed in time $\text{poly}(n, \frac{1}{\epsilon}, \frac{1}{q_{min}})$, such that*

- $(1 - \epsilon)Q \leq \sum_{\sigma} p_{\sigma} M_{\sigma} \leq Q$
- p has close to maximum entropy, i.e., $\sum_{\sigma} p_{\sigma} \log p_{\sigma} \leq (1 - \epsilon)OPT_E$, where $OPT_E(\leq 0)$ is the optimal value of (2).

⁴ $A \otimes B$ denotes ‘Kronecker product’ of matrix A and B .

Acknowledgements. We thank Arash Asadpour and Erick Delage for valuable insights and discussions.

References

1. Chen, Y., Fortnow, L., Nikolova, E., Pennock, D.M.: Combinatorial betting. *SIGecom Exch.* 7(1), 61–64 (2007)
2. Hanson, R.: Combinatorial information market design. *Information Systems Frontiers* 5(1), 107–119 (2003)
3. Hanson, R.D.: Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* (2007)
4. Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: *ACM conference on Electronic commerce* (2004)
5. Lange, J., Economides, N.: A parimutuel market microstructure for contingent claims. *European Financial Management* 11(1) (2005)
6. Peters, M., So, A.M.C., Ye, Y.: A convex parimutuel formulation for contingent claim markets. Working Paper (2005), <http://www.stanford.edu/~yyye/cpcam-ec.pdf>
7. Peters, M., So, A.M.C., Ye, Y.: Pari-mutuel markets: Mechanisms and performance. In: *Workshop on Internet and Network Economics* (2007)
8. Chen, Y., Fortnow, L., Lambert, N., Pennock, D.M., Wortman, J.: Complexity of combinatorial market makers. *CoRR abs/0802.1362* (2008)
9. Chen, Y., Goel, S., Pennock, D.M.: Pricing combinatorial markets for tournaments. In: *ACM Symposium on Theory of Computing* (2008)
10. Fortnow, L., Kilian, J., Pennock, D.M., Wellman, M.P.: Betting boolean-style: a framework for trading in securities based on logical formulas. In: *ACM conference on Electronic commerce* (2003)
11. Chen, Y., Fortnow, L., Nikolova, E., Pennock, D.M.: Betting on permutations. *ACM conference on Electronic commerce* (2007)
12. Grötschel, M., Lovász, L., Schrijver, A.: *Geometric Algorithms and Combinatorial Optimization*. Springer, Heidelberg (1988)
13. Agrawal, S., Wang, Z., Ye, Y.: Parimutuel betting on permutations. Technical report (2008), <http://arxiv.org/abs/0804.2288>
14. Manski, C.: Interpreting the predictions of prediction markets. *Economic Letters* 91(3) (2006)
15. Ottaviani, M., Sørensen, P.N.: Aggregation of information and beliefs in prediction markets. Mimeo, London Business School (2006)
16. Wolfers, J., Zitzewitz, E.: Interpreting prediction market prices as probabilities. Working Paper 12200, National Bureau of Economic Research (May 2006), <http://www.nber.org/papers/w12200>
17. Birkhoff, G.: Three observations on linear algebra. *Univ. Nac. Tucuman Rev. A* 5 (1946)
18. Dulmage, L., Halperin, I.: On a theorem of Frobenius-König and J. von Neumann's game of hide and seek. *Trans. Roy. Soc. Canada Sect. III* 49 (1955)
19. Valiant, L.: The complexity of computing the permanent. *Theoretical Computer Science* (1979)
20. Jerrum, M., Sinclair, A., Vigoda, E.: A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries. In: *ACM symposium on Theory of computing* (2001)

Strategies in Dynamic Pari-Mutual Markets

Tian-Ming Bu^{1,*}, Xiaotie Deng^{2,**}, Qianya Lin², and Qi Qi³

¹ Shanghai Key Laboratory of Trustworthy Computing
East China Normal University
Shanghai, P.R. China
tmbu@sei.ecnu.edu.cn

² Department of Computer Science
City University of Hong Kong
Hong Kong SAR

csdeng@cityu.edu.hk, lqianya2@student.cityu.edu.hk

³ Department of Management Science and Engineering
Stanford University
kaylaqi@stanford.edu

Abstract. We present a strategic model for pari-mutual markets by traders using a cumulative utility function. Under this model, we derive guidelines for the traders on how much to buy or sell. Those guidelines can be implemented with three action combinations, called strategies. We prove that those strategies are payoff equivalent for both the involved trader and the others in the current transaction. However, in the long run, their payoffs can be quite different.

We show that the *buy-only strategy* (*BOS*) achieves the highest market capitalization for the current transaction. In addition, simulation results also prove that BOS always yields the fastest growth of market capitalization even when multiple stages are taken into consideration. Simulation results also show that BOS is a better revelation of the traders' personal beliefs, though it exhibits a higher risk in traders' payoffs.

1 Introduction

The Internet has not only made it possible to create a global electronic market but also allowed for creations of new market models by providing a boundary-less testing base through its powerful communication infrastructures. Prediction markets have been among those that have benefited from such global medium. It builds on an idea that combines the characteristics of investing and betting to create a type of financial markets for wagers on different types of activities such as political events, horse racing, sports, entertainments, or other uncertain events. Despite of a short history of its introduction into the Internet life,

* Research is supported by grants of NSF of China (No. 60496321 and No. 90718013) and a grant of 863 program of China (No. 2007AA01Z189).

** This work described in this paper was supported by a grant from a donation project to CityU (9220046) with a matching grant from UGC of Hong Kong SAR, China.

prediction markets have a remarkable list of successful records in probability forecasting of uncertain events.

Created as a type of financial markets to predict future [11,16,15], prediction market models are based on the *efficient markets hypothesis* [6] that states that traders' information about the outcomes is aggregated into prices. In other words, market prices on the predicted outcomes of future events reflect the collective estimation. It has been known that when a *rational expectation equilibrium* is reached [11,7,5,2], the information distributed among traders will indeed be aggregated on the market in the form of the market-clearing price.

Double auctions have been the most used mechanisms in prediction markets where sales take place only when both sides of trades accept the same trading prices and quantities. However, it may suffer from the *thin market* problem [8] when there is a large gap between bid and ask prices because of a low level of participation.

Hanson's market scoring rule [8] (MSR) offered a new approach to solve the above *thin market* problem with an automatic market maker who accepts orders from traders sequentially and determine the prices by a proper scoring rule. Subsequently, Hanson proposed a logarithmic version [9] (LMSR) which he advocated for its advantages in terms of both cost and modularity.

Pennock [13] invented a new mechanism combining the advantages of traditional pari-mutuel markets and continuous double auctions. In this pricing mechanism, prices change dynamically according to a price function. It allows traders to buy or sell securities at any moment from the system according to the price function. Payoff per share is calculated according to the quantity of the winning security and the amount of losing money or of total money. Pennock proposed and studied several types of price functions. A share-ratio version [3] is now the most commonly used pricing model.

1.1 Related Work

Nikolova and Sami [12] introduced the method of projection games for the design and analysis of prediction markets. The projection game was shown to serve as a strategic model of DPM to capture the essence of strategics in MSR. Their studies concluded that DPM and MSR are deeply connected to each other such that they may be regarded as two different interfaces to the same underlying game.

Chen et al. [3], Dimitrov and Sami [4] independently studied traders' untruthful betting behaviors to mislead the next trader in LMSR markets. [3] found out that, in LMSR, traders with joint probability distributions on signals have the incentive to bet against their own information. [4] used a projection game to study non-myopic strategies in LMSR in an infinite number of periods of plays. [3] also showed that when there're two players in DPM market, the penultimate trader will withhold information.

Peters et al. [14] gave a performance comparison among MSR, DPM and sequential convex pari-mutuel mechanism (SCPM) under the purely pari-mutuel, full charge and tax penalty situations. They established that LMSR has a less stable pricing function and outperforms DPM in the pure pari-mutuel setting in which the total money pool is redistributed to the traders.

1.2 Our Contributions

The utility functions in the previous works [13] and [12] consider a trader's utility for one transaction. However, a trader may trade in the market for several times and the previous lost and gain should be also taken into account. We introduce a cumulative utility function in which the profit not only concerns the shares bought in the current transaction, but also the previous one. So a trader can maximize his utility in total by this utility function. Furthermore, our model is based on the independent belief distribution while previous ones are dependent.

Strategic analysis of DPM market [12] requires the knowledge of the true probability of the event which is hard to obtain in reality. Our analysis is based on the traders' personal beliefs. Moreover, in [12], the utility function is also used to maximize a trader's payoff in the current transaction, which is different from our work.

We provide actions for traders in general cases giving concrete guidelines about how much to trade and what type of actions should be taken. From these actions, we study three action combinations, called strategies, including *first-prior strategy*, *second-prior strategy* and *buy-only strategy*. We prove that these three strategies yield the same expected payoffs for all traders in the current transaction. We should call this property the payoff equivalence property.

We also prove that the *buy-only strategy* achieves the highest market capitalization for the current trader. In addition, our simulation results show that it also yields the fastest growth of market capitalization in the long run. By simulations, we find that the market capitalization has an impact on traders' payoffs. The higher market capitalization is, the higher risk traders may suffer. The *buy-only strategy*, which is most commonly used in pari-mutual markets, exhibits a riskier performance than the other two implying traders using this strategy tend to win more or lose more.

On the other hand, from simulations, higher market capitalization leads to a better fitting of market probability into traders' beliefs, resulting in a better revelation of traders' private information. Hence, the market capitalization is a double-edged sword for market designers.

2 Dynamic Parimutuel Markets

The dynamic parimutuel market is first proposed by Pennock [13] and implemented in an on-line prediction market named Yahoo! Buzz market [10].

Suppose there're n securities in the market. Each security i represents a mutual exclusive outcome i and \mathbf{I} is the collection of all outcomes. So $i \in \mathbf{I}$ and $|\mathbf{I}| = n$. The market is initialized with a number of outstanding shares on all securities which in fact is a subsidy from the market maker. Traders trade with the market maker by choosing appropriate securities to buy or sell according to their personal beliefs. Prices vary dynamically all the time as the total money pool changes. After the true outcome is revealed, the market is liquidated and the winning security is cashed by re-distributing the money pool.

Let $\mathbf{q} = \{q_1, q_2, \dots, q_n\}$ be the quantities of outstanding shares in the market. In a share-ratio version of DPM market [13,3], the price function is related to the ratio of quantities of securities. So the spot price for shares on outcome i is :

$$\hat{p}_i = \frac{q_i}{\sqrt{\sum_{j=1}^n q_j^2}}.$$

The aggregating market estimation (i.e. market probability) on outcome i is:

$$Pr_i = \frac{q_i^2}{\sum_{j=1}^n q_j^2}.$$

The market capitalization (i.e. cost function) is:

$$C(\mathbf{q}) = \sqrt{\sum_{j=1}^n q_j^2}.$$

The transition cost (or the trader’s payment) is:

$$C = C(\mathbf{q}^{after}) - C(\mathbf{q}^{bef}).$$

Here \mathbf{q}^{bef} . and \mathbf{q}^{after} denote the quantity of securities before and after transaction respectively.

As the word “parimutuel” implies, traders who wager on the true outcome win the money re-distributed by the total pool. Due to the dynamic changes of shares on different events, the winning security’s return money, which we called redemption price in such a market is not fixed. This adds some difficulties for traders to report their beliefs since traders can not simply buy or sell securities until prices reach their personal estimations such as in MSR, but they should take the redemption price into account also. So the traders’ actions of maximizing their expected payoffs a bit more complicated.

The redemption price of outcome i in DPM is the total market capitalization divided by the quantity of outstanding shares of i if it happens,

$$\tilde{p}_i = \frac{\sqrt{\sum_{j=1}^n q_j^2}}{q_i}.$$

$\forall i$, the spot price for shares on outcome i , always varies between $(0, 1)$, while the redemption price for shares on outcome i always greater than 1. So traders will be guaranteed positive utilities if they holds the securities of true outcome.

In order to simplify the model, in the rest of the paper we follow the same assumption as in Pennock’s paper [13].

Assumption 2.1. [13] The current value for the payoff per share of security i is the same as the expected final value of the payoff per share of i given that i occurs. That is,

$$E[\rho_i|i] = \rho_i \tag{2.1}$$

3 Strategic Model of DPM

We consider two mutual exclusive outcomes A and B with x and y outstanding shares respectively, i.e., $I = \{A, B\}$, $|I| = 2$ and $\mathbf{q} = \{x, y\}$. All the results in this paper can be generalized to multi-event cases easily.

3.1 Symmetry Property

Nikolova and Sami [12] give some insight about actions when traders have no possession on their hands. But after several rounds of transactions, traders who have securities will keep an closer eye on the market, waiting for the decisions to sell or to buy more. So the strategies for traders who have involved in the market seem to play more important roles for the reason that those traders are more active and incentive.

We seek to propose a general action model for traders in all situations. Given a trader having π_A shares of security A and π_B shares of security B, and the total payment of these shares is C . There are x outstanding shares of A and y of B in the market currently. After purchasing Δx and Δy extra shares on outcome A and B, the trader's total utility will be:

$$L = p(\pi_A + \Delta x) \frac{\sqrt{(x + \Delta x)^2 + (y + \Delta y)^2}}{x + \Delta x} + (1 - p)(\pi_B + \Delta y) \frac{\sqrt{(x + \Delta x)^2 + (y + \Delta y)^2}}{y + \Delta y} - (\sqrt{(x + \Delta x)^2 + (y + \Delta y)^2} - \sqrt{x^2 + y^2}) - C \quad (3.1)$$

Note that C is the total payment in the previous transactions, which is independent of Δx and Δy .

In order to maximize his utility, we take the partial derivative by Δx ,

$$\frac{\partial L}{\partial \Delta x} \Rightarrow \Delta x = (y + \Delta y) \sqrt[3]{\frac{p}{1 - p} \frac{x - \pi_A}{y - \pi_B}} - x \quad (3.2)$$

Similarly,

$$\frac{\partial L}{\partial \Delta y} = 0 \Rightarrow \Delta y = (x + \Delta x) \sqrt[3]{\frac{1 - p}{p} \frac{y - \pi_B}{x - \pi_A}} - y \quad (3.3)$$

Theorem 3.1. *In DPM, the way to purchase shares on outcome A and outcome B to maximize trader's utility is not unique.*

Corollary 3.2. *In DPM, in order to maximize the payoff, a trader can always purchase securities only on one side.*

Similarly, we can get the following corollary.

Corollary 3.3. *In DPM, in order to maximize the payoff, a trader can always sell securities only on one side.*

From Corollary 3.2 and Corollary 3.3 we assume that a trader never buys or sells shares simultaneously on both sides.

3.2 Strategic Actions

In DPM, the redemption price depends on the quantity of outstanding shares and changes dynamically, a trader couldn't simply compare his personal belief to the market probability when deciding which action to take. The number of shares he possesses currently also has an impact on the decision. Hence, we introduce a new concept called *virtual market probability* which depends on the number of shares the current holding.

Definition 3.4 (Virtual Market Probability \tilde{Pr}). *If a trader has π_A shares of security A and π_B of B on hand, his virtual market probability on A is $\tilde{Pr} = \frac{x^3(y-\pi_B)}{x^3(y-\pi_B)+y^3(x-\pi_A)}$ and on B is $1 - \tilde{Pr}$.*

Likewise, in the rest of the paper, we use Pr to denote the market probability on outcome A and $1 - Pr$ on outcome B. Now we can summarize the action model.

Theorem 3.5. *Given a trader with probability estimation p on A and $1 - p$ on B, suppose he has π_A shares of security A and π_B of B on hand. The current market outstanding shares on A are x , and on B are y . He will compare his personal belief to the virtual market probability to maximize his expected payoff.*

1. If $p > \tilde{Pr}$:

he should purchase $\Delta x = y \sqrt[3]{\frac{p}{1-p} \frac{x-\pi_A}{y-\pi_B}} - x$ on outcome A, or,

sell $\Delta y = y - x \sqrt[3]{\frac{1-p}{p} \frac{y-\pi_B}{x-\pi_A}}$ on outcome B.

2. If $p < \tilde{Pr}$:

he should sell $\Delta x = x - y \sqrt[3]{\frac{p}{1-p} \frac{x-\pi_A}{y-\pi_B}}$ on outcome A, or,

purchase $\Delta y = x \sqrt[3]{\frac{1-p}{p} \frac{y-\pi_B}{x-\pi_A}} - y$ on outcome B.

Theorem 3.5 provides a set of actions for general context. When a trader has no shares on hand as follows, we can simplify the formula by setting $\pi_A = 0$ and $\pi_B = 0$. Furthermore, short sell is forbidden. In this case, things become so straightforward that a trader just needs to compare his personal estimation to the market probability and choose one type of securities to buy.

Remark 3.6. If a trader has no shares on hand, his virtual market probability equals the market probability, i.e., $\tilde{Pr} = Pr$. Thus,

1. If $p > Pr$:

he should purchase $\Delta x = \sqrt[3]{\frac{p}{1-p} y^2 x} - x$ on outcome A.

2. If $p < Pr$:

he should purchase $\Delta y = \sqrt[3]{\frac{1-p}{p} x^2 y} - y$ on outcome B.

3.3 Definition of Strategies

According to Theorem 3.5, there're two alternative actions under each case. We define three strategies which are combinations of these actions as Figure 4 (See Appendix B).

Definition 3.7

1. *First-prior Strategy (FPS):* If a trader’s personal belief about A is higher than his virtual market probability on A , he will buy $\Delta x = y \sqrt[3]{\frac{p}{1-p} \frac{x-\pi_A}{y-\pi_B}} - x$ on A . Otherwise, he will sell $\Delta x = x - y \sqrt[3]{\frac{p}{1-p} \frac{x-\pi_A}{y-\pi_B}}$ on A when $\Delta x \leq \pi_A$. And he will sell π_A on A , then buy $\Delta y' = y - (x - \pi_A) \sqrt[3]{\frac{1-p}{p} \frac{y-\pi_B}{x-\pi_A}}$ on B when $\Delta x > \pi_A$.
2. *Second-prior Strategy (SPS):* If a trader’s personal belief about A is higher than his virtual market probability on A , he will sell $\Delta y = y - x \sqrt[3]{\frac{1-p}{p} \frac{y-\pi_B}{x-\pi_A}}$ on B when $\Delta y \leq \pi_B$. And he will sell π_B on B , then buy $\Delta x' = (y - \pi_B) \sqrt[3]{\frac{p}{1-p} \frac{x-\pi_A}{y-\pi_B}} - x$ on A when $\Delta y > \pi_B$. Otherwise, he will choose buying B .
3. *Buy-only Strategy (BOS):* A trader will choose buying $\Delta x = y \sqrt[3]{\frac{p}{1-p} \frac{x-\pi_A}{y-\pi_B}} - x$ on A when his personal estimation about A is higher than his virtual market probability on A and buying $\Delta y = x \sqrt[3]{\frac{1-p}{p} \frac{y-\pi_B}{x-\pi_A}} - y$ on B otherwise. He will never sell his shares in this case.

So far we’ve generalized the strategic model of DPM. In the following sections, we will focus on the properties of these strategies.

4 Strategies Comparison

4.1 Payoff Equivalence

Lemma 4.1. *[Payoff Equivalence For Others] FPS, SPS and BOS yield the same expected payoff for the other traders in the market.*

Intuitively, payoff of these three strategies for the involved trader himself should be equal too. Now we proof this conjecture.

Lemma 4.2. *[Payoff Equivalence of the Involved Trader] FPS, SPS and BOS yield the same expected payoff for the involved trader himself.*

As Lemma 4.1 and 4.2 conclude, from the myopic point of view, three strategies are equivalent to the involved trader himself and to the others which is summarized by Theorem 4.3.

Theorem 4.3. *FPS, BOS and SPS yield the same expected payoff for all traders in the current transaction.*

The expected utilities under these three strategies currently is the same too. Even so, these three strategies have different impact on the total capitalization of the market and the outstanding shares on events. Analysis below gives a closer insight to this influence.

Proposition 4.4. *If a trader chooses BOS rather than FPS and SPS, market capitalization will raise the highest among the three strategies.*

Proposition 4.4 proves that BOS achieves the highest market capitalization for the current transaction. Now we turn to see what happens if all traders choose BOS. Intuitively, in BOS markets, which use BOS as the dominant strategies, capitalization should also exceed the one that in FPS and SPS markets in the long run. We design simulations to validate our conjecture.

5 Simulations and Observations

The market is open with two mutual exclusive events A and B, and initialized with 100 shares on each event. A number of traders with a normal distribution of beliefs are prepared. At each round, a trader is chosen randomly and enters the market with the goal of achieving a maximal payoff by using his private estimation. The choice of the trader draws from a uniform distribution so that traders are selected with equal chances. After a number of trading rounds the market closes and each trader attains a profit (may be negative). We always assume event A comes out to be true in the end.

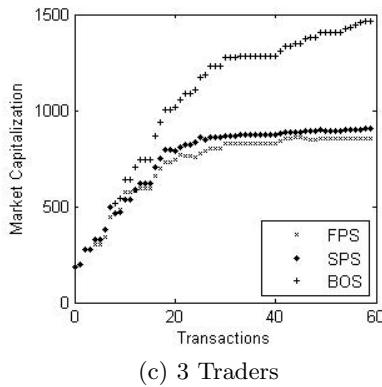



Fig. 1. Market Capitalization of FPS, SPS and BOS

5.1 Observation: Market Capitalization

In our experiments, we found that no matter what belief distribution is, BOS always yields the highest market capitalization among the three strategies. Three concrete examples are shown in Figure 11. Each curve represents one case that all traders choose one type of the strategies as their dominant strategy.

5.2 Observation: Payoff

We put 10 traders into the market, with $\mu = 0.5$ and $\sigma = 0.2$. In our all experiments, the higher market capitalization suffers higher risk of payoff at all time. Take Figure 2 for example. In (a) it shows the growth of market capitalization. BOS, which always achieves the highest capitalization among the three tends to win more but lose more in Figure 2 (b). The reason may be that traders have to give higher investment to report their estimations when market capitalization is higher. Hence, one may win more if his report is close to the true probability of the outcome but lose more on the other hand.

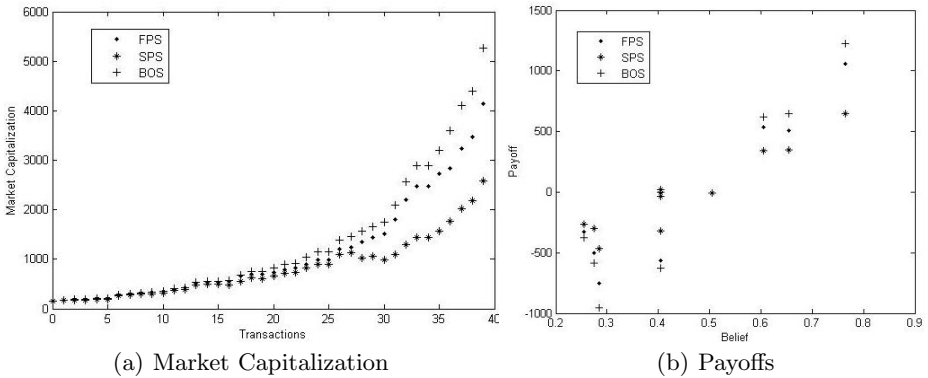


Fig. 2. Capitalization and Payoffs of FPS, SPS and BOS with 10 Traders and 40 Rounds

5.3 Observation: Share Ratio and Market Probability

In fact, the market probability can be looked on as a function of share ratio(x/y), because $Pr = \frac{x^2}{x^2+y^2} = \frac{(x/y)^2}{(x/y)^2+1}$. However, in DPM, a trader compares his personal estimation to his virtual market probability, not market probability, thus the share ratio doesn't simply change along with the beliefs of the traders. We try to keep track of the share ratios in these three strategies.

¹ (a) 10 traders whose beliefs draw from a normal distribution between (0, 1) with mathematical expectation $\mu = 0.5$, standard deviation $\sigma = 0.2$. (b) 20 Traders whose beliefs draw from a normal distribution between (0.4, 1) with mathematical expectation $\mu = 0.8$, standard deviation $\sigma = 0.1$. (c) 3 Traders with $p_1 = 0.2$, $p_2 = 0.4$ and $p_3 = 0.8$.

For simplicity, we only make comparisons between FPS and BOS. The analysis for SPS is quite similar. Suppose the initial state of the market is x_0 shares on outcome A and y_0 shares on outcome B with $x_0 = y_0$. At the first k rounds, assume no trader has the chance to take his second transaction in the market. So traders all take buying actions and the share ratio after round k equals to $\frac{x_k}{y_k} = \prod_{j=1}^k (\frac{p^{(j)}}{1-p^{(j)}})^{\frac{1}{3j}} (\frac{x_0}{y_0})^{\frac{1}{3k}} = \prod_{j=1}^k (\frac{p^{(j)}}{1-p^{(j)}})^{\frac{1}{3j}}$. At $k + 1st$ round, assume trader m ($m < k$) taking his second transaction. If his action is buying A at his 1st transaction, then $\pi_A^{(m)} > 0$ and $\pi_B^{(m)} = 0$ currently. For FPS and BOS, he will buy Δx when his personal estimation is higher than his virtual market probability. So in this case FPS and BOS yield the same share ratio, even the same number of shares on both sides. For the case when his personal estimation is lower than his virtual market belief, FPS will sell A and BOS will buy B.

$$FPS, \text{ selling } \Delta x_{k+1} : \frac{x'_{k+1}}{y'_{k+1}} = \frac{x_k - \Delta x_{k+1}}{y_k} = (\frac{p^{(m)}}{1-p^{(m)}})^{1/3} (\frac{x_k - \pi_A^{(m)}}{y_k})^{1/3}$$

$$BOS, \text{ buying } \Delta y_{k+1} : \frac{x''_{k+1}}{y''_{k+1}} = \frac{x_k}{y_k + \Delta y_{k+1}} = (\frac{p^{(m)}}{1-p^{(m)}})^{1/3} (\frac{x_k - \pi_A^{(m)}}{y_k})^{1/3}$$

So, $\frac{x'_{k+1}}{y'_{k+1}} = \frac{x''_{k+1}}{y''_{k+1}}$ but $x'_{k+1} < x''_{k+1}$, $y'_{k+1} < y''_{k+1}$.

At $k + 2nd$ round, assume trader n taking his second transaction. Suppose $\pi_A^{(n)} > 0$ and $\pi_B^{(n)} = 0$, $\frac{p^{(n)}}{1-p^{(n)}} = \epsilon^{(n)}$. In the case $p^{(n)} > \tilde{P}r$, FPS and BOS will both buy A. If $p^{(n)} < \tilde{P}r$, FPS will sell A and BOS will buy B. As a result, $(\frac{x'_{k+2}}{y'_{k+2}})^3 = \epsilon^{(n)} (\frac{x'_{k+1}}{y'_{k+1}} - \frac{\pi_A^{(n)}}{y'_{k+1}})$, $(\frac{x''_{k+2}}{y''_{k+2}})^3 = \epsilon^{(n)} (\frac{x''_{k+1}}{y''_{k+1}} - \frac{\pi_A^{(n)}}{y''_{k+1}})$. So $\frac{x'_{k+2}}{y'_{k+2}} < \frac{x''_{k+2}}{y''_{k+2}}$. When $\pi_A^{(n)} = 0$ and $\pi_B^{(n)} > 0$, FPS will be $(\frac{x'_{k+2}}{y'_{k+2}})^3 = \epsilon^{(n)} (\frac{x'_{k+1}}{y'_{k+1} - \pi_B^{(n)}})$, and BOS is $(\frac{x''_{k+2}}{y''_{k+2}})^3 = \epsilon^{(n)} (\frac{x''_{k+1}}{y''_{k+1} - \pi_B^{(n)}})$, so $\frac{x'_{k+2}}{y'_{k+2}} > \frac{x''_{k+2}}{y''_{k+2}}$. From the above deduction we can see share ratio depends on the securities one already possessed and may be different for FPS and BOS after a few rounds of transactions.

Our simulations conform to our analysis. We put three traders into the market, with beliefs $p^{(1)} = 0.1$, $p^{(2)} = 0.3$, and $p^{(3)} = 0.8$. Figure 3 is the evolution of share ratio from 6th to 18th round. It shows the irregular change. When other variables are fixed, share ratio still depends on the type and amount of shares the current trader holding.

Next we analyze the strategies' performances from another aspect. We record the market probabilities of the three strategies in each transaction, then compare them with the involved trader's personal belief. As we mentioned above, during each transaction, the trader will try to maximize his expected payoff by changing virtual market probability to his personal belief. After each transaction, virtual market belief equals to the involved trader's belief regardless of which strategy he chose. Table 1 computes the mean value μ and the

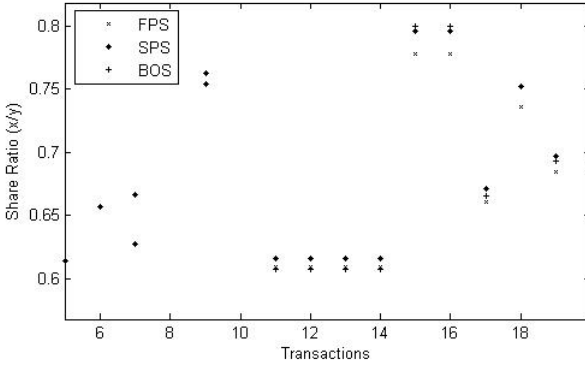


Fig. 3. Share Ratio Evolution of 3 Traders: A Closer Look

Table 1. Market Probabilities VS Personal Beliefs

	FPS	SPS	BOS
$\mu =$	0.0700	0.0703	0.0680
$\sigma =$	0.0897	0.0902	0.0866

standard deviation σ by taking the personal beliefs as the benchmark². The market probability of BOS always varies most closely to personal belief in our experiments, suggesting Pr in the BOS market is closest to \tilde{Pr} . This may be due to the huge amount of market capitalization of BOS market, since $\lim_{\frac{\pi_A}{x} \rightarrow 0, \frac{\pi_B}{y} \rightarrow 0} \tilde{Pr} = \lim_{\frac{\pi_A}{x} \rightarrow 0, \frac{\pi_B}{y} \rightarrow 0} \frac{x^3(y - \pi_B)}{x^3(y - \pi_B) + y^3(x - \pi_A)} = \frac{x^2}{x^2 + y^2} = Pr$. While the higher of the market capitalization, the weaker role of a trader’s personal possession plays, so \tilde{Pr} tends to converge to Pr .

In DPM market, traders can’t obtain the current trader’s belief directly because the virtual market probability depends on the involved trader’s possession which is private information to the others. For this reason, BOS, which is better fitting in with traders’ personal beliefs, has the advantage of revealing aggregating market belief. However, markets adopting BOS result in a huge amount of capitalization so that traders in such markets suffer higher risk than the others. Market designer should make the risk-fitting trade-off when designing a DPM market.

6 Conclusions and Discussions

The previous work usually focuses on maximizing traders’ expected payoffs at the current step. In our paper, we take traders previous possessions into consideration, seeking to get a maximal payoff in its entirety.

² $\mu = \frac{\sum_{i=1}^n (Pr_i - p_i)}{n}$, where Pr_i is the market probability on A in transaction i and p_i is the involved trader’s personal belief. $\sigma = \sqrt{\frac{\sum_{i=1}^n (Pr_i - p_i)^2}{n}}$.

By using this cumulative utility function, we summarize actions for traders in general cases. There're three strategies covering all actions a trader may take in different situations. They yield the same expected payoffs for all traders in myopic, called payoff equivalence.

We observe in experiments that BOS, most commonly used in pari-mutual markets, achieves the highest market capitalization than the other two strategies. Traders in such a rapid growing capitalization market tend to win more and lose more, exhibiting a riskier performance. But higher capitalization drives the market probability to be a better indication of traders' beliefs. Market designers have to take this double-sided effect into account when open a DPM market.

In reality, traders' behaviors may be some variations since people adjust their estimations as they observe others' behaviors and as more and more information is revealed to the public. Moreover, people may have budget constraint, which limit their buying power to report their personal beliefs. In our experiments, traders' beliefs remain unchanged during the transactions. Our future work will introduce the dynamic changes of beliefs and budget control to the analysis of the strategies.

References

1. Berg, J., Forsythe, R., Nelson, F., Rietz, T.: Results from a dozen years of election futures markets research (2001)
2. Chen, Y., Mullen, T., Chu, C.-H.: An in-depth analysis of information markets with aggregate uncertainty, vol. 38, pp. 94–97. IEEE Computer Society, Los Alamitos (2005)
3. Chen, Y., Reeves, D.M., Pennock, D.M., Hanson, R.D., Fortnow, L., Gonen, R.: Bluffing and strategic reticence in prediction markets. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 70–81. Springer, Heidelberg (2007)
4. Dimitrov, S., Sami, R.: Non-myopic strategies in prediction markets. In: The 9th ACM conference on Electronic commerce, Chicago, USA, pp. 200–209 (2008)
5. Feigenbaum, J., Fortnow, L., Pennock, D.M., Sami, R.: Computation in a distributed information market. *Theoretical Computer Science* 343, 114–132 (2005)
6. Fama, E.F.: Efficient capital markets: A review of theory and empirical work. *Journal of Finance* 25, 383–417 (1970)
7. Grossman, S.J.: An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies* 48, 541–559 (1983)
8. Hanson, R.: Combinatorial information market design. *Information Systems Frontiers* 5, 107–119 (2003)
9. Hanson, R.: Logarithmic market scoring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets* 1, 3–15 (2007)
10. Mangold, B., Dooley, M., Flake, G.W., Hoffman, H., Kasturi, T., Pennock, D.M., Dornfest, R.: The tech. buzz game. *Electronic Commerce Research* 6, 201–221 (2006)
11. Muth, J.F.: Rational expectations and the theory of price movements. *Econometrica* 29, 315–335 (1961)
12. Nikolova, E., Sami, R.: A strategic model for information markets. In: The 8th ACM conference on Electronic commerce, pp. 316–325 (2007)

13. Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. *Electronic Commerce*, 170–179 (May 2004)
14. Peters, M., So, A.M.-C., Ye, Y.: Pari-mutuel markets: Mechanisms and performance. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 82–95. Springer, Heidelberg (2007)
15. Schmidt, C., Werwatz, A.: How accurate do markets predict the outcome of an event? the euro, soccer championships experiment. Technical Report, number 09-2002 (2000)
16. Wolfers, J., Zitzewitz, E.: Prediction markets. *Journal of Economic Perspectives* 18, 107–126 (Spring, 2004)

A Appendix: Proofs

A.1 Proof of Theorem 3.1

Proof. In fact, Equation (3.2) and (3.3) are the same equations, we re-arrange them and get

$$\frac{x + \Delta x}{y + \Delta y} = \sqrt[3]{\frac{p}{1-p} \frac{x - \pi_A}{y - \pi_B}} \quad (\text{A.1})$$

Equation (A.1) implies that for an arbitrary Δx , we can find one corresponding Δy which maximizes the utility. \square

A.2 Proof of Corollary 3.2

Proof. We can prove it directly from Equation (A.1). By either setting $\Delta x = 0$ or $\Delta y = 0$, maximal utility by purchasing shares on one side only could be obtained. \square

A.3 Proof of Theorem 3.5

Proof. Proof of Case 1:

In Equation (3.1), let $\Delta y = 0$, then,

$$\begin{aligned} L &= p(\Delta x + \pi_A) \frac{\sqrt{(x + \Delta x)^2 + y^2}}{x + \Delta x} + (1-p)\pi_B \frac{\sqrt{(x + \Delta x)^2 + y^2}}{y} \\ &\quad - (\sqrt{(x + \Delta x)^2 + y^2} - \sqrt{x^2 + y^2}) - C \\ \frac{\partial L}{\partial \Delta x} &= 0 \implies \Delta x = y \sqrt[3]{\frac{p}{1-p} \frac{x - \pi_A}{y - \pi_B}} - x \end{aligned}$$

When $\Delta x > 0$, we have:

$$\begin{aligned} y \sqrt[3]{\frac{p}{1-p} \frac{x - \pi_A}{y - \pi_B}} - x > 0 &\implies \frac{x}{y} \sqrt[3]{\frac{y - \pi_B}{x - \pi_A}} < \sqrt[3]{\frac{p}{1-p}} \\ &\implies p > \frac{x^3(y - \pi_B)}{x^3(y - \pi_B) + y^3(x - \pi_A)} = \tilde{P}r \end{aligned}$$

Similarly, if the trader sells Δy shares of event B instead of buying A,

$$L = p\pi_A \frac{\sqrt{x^2 + (y - \Delta y)^2}}{x} + (1 - p)(\pi_B - \Delta y) \frac{\sqrt{x^2 + (y - \Delta y)^2}}{y} - (\sqrt{x^2 + (y - \Delta y)^2} - \sqrt{x^2 + y^2}) - C$$

Let $\frac{\partial L}{\partial \Delta y} = 0$, then, $\Delta y = y - x^3 \sqrt{\frac{1-p}{p} \frac{y-\pi_B}{x-\pi_A}}$.

When $\Delta y > 0$, we also have $p > \tilde{P}r$.

Proof of Case 2:

The trader should sell security A or buy security B when his estimation is lower than the virtual market probability. We first consider selling A.

The expected utility of selling Δx shares of A is:

$$L = p(\pi_A - \Delta x) \frac{\sqrt{(x - \Delta x)^2 + y^2}}{x - \Delta x} + (1 - p)\pi_B \frac{\sqrt{(x - \Delta x)^2 + y^2}}{y} - (\sqrt{(x - \Delta x)^2 + y^2} - \sqrt{x^2 + y^2}) - C$$

Take the first order derivative, we have, $\Delta x = x - y^3 \sqrt{\frac{p}{1-p} \frac{x-\pi_A}{y-\pi_B}}$.

When $\Delta x > 0$, we have $p < \tilde{P}r$.

Similarly, if the trader buys Δy shares of event B instead of selling A,

$$L = p\pi_A \frac{\sqrt{x^2 + (y + \Delta y)^2}}{x} + (1 - p)(\pi_B + \Delta y) \frac{\sqrt{x^2 + (y + \Delta y)^2}}{y + \Delta y} - (\sqrt{x^2 + (y + \Delta y)^2} - \sqrt{x^2 + y^2}) - C$$

So we get $\Delta y = x^3 \sqrt{\frac{1-p}{p} \frac{y-\pi_B}{x-\pi_A}} - y$ and $p < \tilde{P}r$. □

A.4 Proof of Lemma 4.1

Proof. We first consider the case when the involved trader's personal belief is higher than his virtual market probability on event A. If he chooses FPS or BOS, he will buy Δx on A. Redemption prices on security A and B will become:

$$\tilde{p}_A = \frac{\sqrt{(x + \Delta x)^2 + y^2}}{x + \Delta x} \quad \tilde{p}_B = \frac{\sqrt{(x + \Delta x)^2 + y^2}}{y}$$

If he chooses SPS, his strategy is selling Δy on B, in which there're two cases.

Case I: If $\Delta y \leq \pi_B$, trader will sell Δy shares on B. Redemption prices become:

$$\tilde{p}'_A = \frac{\sqrt{x^2 + (y - \Delta y)^2}}{x} \quad \tilde{p}'_B = \frac{\sqrt{x^2 + (y - \Delta y)^2}}{y - \Delta y}$$

Since $\frac{\tilde{p}_A^2}{\tilde{p}'_A^2} = \frac{1 + \frac{y^2}{(x + \Delta x)^2}}{1 + \frac{(y - \Delta y)^2}{x^2}}$, replacing Δx and Δy with $\Delta x = \sqrt[3]{\frac{p}{1-p} \frac{y^3}{y - \pi_B} (x - \pi_A)} - x$ and $\Delta y = y - \sqrt[3]{\frac{1-p}{p} \frac{x^3}{x - \pi_A} (y - \pi_B)}$ in the above equation, it's easy to verify that $\tilde{p}_A = \tilde{p}'_A$ and $\tilde{p}_B = \tilde{p}'_B$.

Case II: If $\Delta y > \pi_B$, trader will sell π_B shares on B, then buy $\Delta x'$ shares on A. Redemption prices become:

$$\tilde{p}'_A = \frac{\sqrt{(x + \Delta x')^2 + (y - \pi_B)^2}}{x + \Delta x'} \quad \tilde{p}'_B = \frac{\sqrt{(x + \Delta x')^2 + (y - \pi_B)^2}}{y - \pi_B}$$

Replacing $\Delta x'$ with $\Delta x' = (y - \pi_B) \sqrt[3]{\frac{p}{1-p} \frac{x - \pi_A}{y - \pi_B}} - x$ we can also obtain $\tilde{p}_A = \tilde{p}'_A$ and $\tilde{p}_B = \tilde{p}'_B$.

That means, no matter which strategy the involved trader chooses, there's just buying or selling for option. And no matter which action he takes, the expected redemption prices of the events are the same. Since the expected payoff of another trader, for example j , is $p^{(j)} \pi_A^{(j)} \tilde{p}_A + (1 - p^{(j)}) \pi_B^{(j)} \tilde{p}_B - C^{(j)}$, the equivalence of redemption prices implies the expected payoffs of all other traders in the market currently are also the same regardless of which strategy the involved trader chooses.

The proof for the case that the involved trader's personal belief is lower than the virtual market probability on event A is quite similar, so we omit here. \square

A.5 Proof of Lemma 4.2

Proof. First consider the case the involved trader's personal estimation is higher than his virtual market probability on event A. He may buy A or sell B for maximizing the profit in this case. According to Lemma 4.1, redemption prices of two actions are even.

$$\begin{aligned} \tilde{p}_A = \tilde{p}'_A &= \frac{\sqrt{(x + \Delta x)^2 + y^2}}{x + \Delta x} = \frac{\sqrt{x^2 + (y - \Delta y)^2}}{x} = \frac{\sqrt{(x + \Delta x')^2 + (y - \pi_B)^2}}{x + \Delta x'} \\ \tilde{p}_B = \tilde{p}'_B &= \frac{\sqrt{(x + \Delta x)^2 + y^2}}{y} = \frac{\sqrt{x^2 + (y - \Delta y)^2}}{y - \Delta y} = \frac{\sqrt{(x + \Delta x')^2 + (y - \pi_B)^2}}{y - \pi_B} \end{aligned} \tag{A.2}$$

If he chooses BOS or FPS, he will buy shares on A, his expected payoff is $\tilde{L}_1 = p(\pi_A + \Delta x)\tilde{p}_A + (1 - p)\pi_B\tilde{p}_B - (\sqrt{(x + \Delta x)^2 + y^2} - \sqrt{x^2 + y^2}) - C$. If he chooses SPS, he will sell shares on B, in which there're two cases.

Case I: If $\Delta y \leq \pi_B$, he will sell Δy shares on B. And his expected payoff is $\tilde{L}_2 = p\pi_A\tilde{p}'_A + (1 - p)(\pi_B - \Delta y)\tilde{p}'_B - (\sqrt{x^2 + (y - \Delta y)^2} - \sqrt{x^2 + y^2}) - C$, where C is the total payment of previous transactions and C in the above two equations are obviously the same. $\tilde{L}_1 - \tilde{L}_2 = p\Delta x\tilde{p}_A - \sqrt{(x + \Delta x)^2 + y^2} - (-(1 - p)\Delta y\tilde{p}_B - \sqrt{x^2 + (y - \Delta y)^2})$. We combine it with Equation (A.2) and get $\tilde{L}_1 = \tilde{L}_2$.

Case II: If $\Delta y > \pi_B$, trader will sell π_B shares on B, then buy $\Delta x'$ shares on A. His expected payoff is: $\tilde{L}_2 = p(\pi_A + \Delta x')\tilde{p}'_A - (\sqrt{x^2 + (y - \Delta y)^2} - \sqrt{x^2 + y^2}) - C$. $\tilde{L}_1 - \tilde{L}_2 = p\Delta x\tilde{p}'_A - \sqrt{(x + \Delta x)^2 + y^2} - (-\sqrt{(x + \Delta x')^2 + (y - \pi_B)^2})$. Combining with Equation (A.2), we can also get $\tilde{L}_1 = \tilde{L}_2$.

Similarly, we can prove the expected payoff equivalence of three strategies in the case $p < \tilde{P}r$. □

A.6 Proof of Proposition 4.4

Proof. Given an arbitrary trader entering the market, there're x and y outstanding shares on outcome A and B respectively. The trader has π_A and π_B shares on hand. Firstly we consider the case $p > \tilde{P}r$, if he chooses FPS or BOS, he should purchase Δx shares on outcome A, so the market capitalization is: $M_{BOS} = M_{FPS} = \sqrt{(x + \Delta x)^2 + y^2}$. If he chooses SPS, he should sell Δy shares on outcome B.

Case I: The trader has enough shares on B to sell ($\pi_B \geq \Delta y$). $M_{SPS} = \sqrt{x^2 + (y - \Delta y)^2}$. Since $y \geq y - \Delta y$ and $x + \Delta x \geq x$, we have $M_{BOS} = M_{FPS} \geq M_{SPS}$

Case II: The trader has not enough shares on B to sell ($\pi_B < \Delta y$), so he has to sell all shares on B first ($\Delta y' = \pi_B$), then buy $\Delta x'$ on A in order to get a maximal profit. The market capitalization of SPS becomes $M_{SPS} = \sqrt{(x + \Delta x')^2 + (y - \pi_B)^2}$. From $x + \Delta x' = (y - \pi_B)(\frac{p}{1-p})^{1/3}(\frac{x - \pi_A}{y - \pi_B})^{1/3}$, $x + \Delta x = y(\frac{p}{1-p})^{1/3}(\frac{x - \pi_A}{y - \pi_B})^{1/3}$ and $y \geq y - \pi_B$, we can conclude $\sqrt{(x + \Delta x)^2 + y^2} \geq \sqrt{(x + \Delta x')^2 + (y - \pi_B)^2}$, implying $M_{BOS} = M_{FPS} \geq M_{SPS}$.

The proof of $M_{BOS} = M_{SPS} \geq M_{FPS}$ in the case $p < \tilde{P}r$ is similar. □

B Appendix: Figure of Subsection 3.3

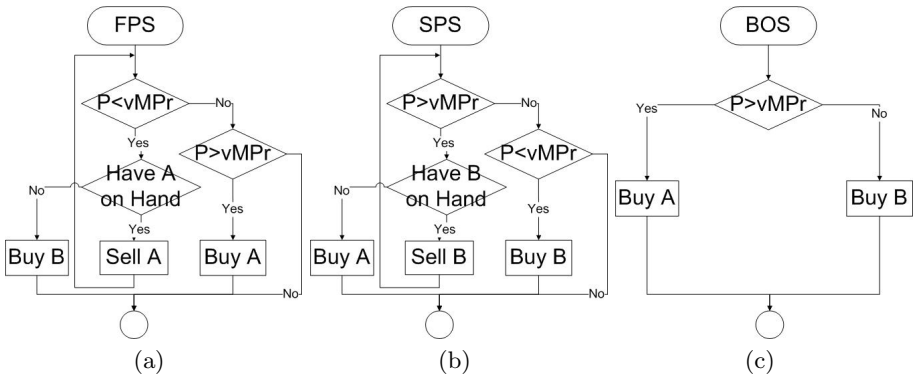


Fig. 4. Work Flows of FPS, SPS and BOS

Truthful Surveys

Nicolas Lambert* and Yoav Shoham

Department of Computer Science
Stanford University
Stanford, CA 94305

Abstract. We consider the problem of truthfully sampling opinions of a population for statistical analysis purposes, such as estimating the population distribution of opinions. To obtain accurate results, the surveyor must incentivize individuals to report unbiased opinions. We present a rewarding scheme to elicit opinions that are representative of the population. In contrast with the related literature, we do not assume a specific information structure. In particular, our method does not rely on a common prior assumption.

1 Introduction

Online surveys, opinion polls and questionnaires are primary tools to gather information on a population and have been growing at a fast pace over the past few years. There already exists an extensive literature on the construction of questionnaires and their statistical processing and analysis (see, for example, Montgomery [1] or Kish [2]). However, to derive meaningful results, it is also imperative to get accurate samples. To induce honest behavior, the surveyor should reward participants appropriately. This paper focuses on the design of survey mechanisms that incentivize participants to provide true samples of opinions.

This problem falls under the broader umbrella of information elicitation, for which there exists several available solutions in various settings. For example, when one is interested in gathering information regarding the uncertainty of an upcoming event, such as the probability of a political candidate winning an election, one may use scoring rules and score functions [3,4]. These induce honest participation by setting rewards as a function of the outcome of the event. However, these methods rely on the verifiability of the outcomes, and do not apply to more general information, such as subjective opinions.

Miller et al. [5] show that eliciting non-verifiable information is nonetheless possible if we make some assumptions on the information being retrieved, and

* Part of this work was done while the author was visiting Yahoo! Research, New York.

the knowledge available to the individuals. They consider the problem of gathering ratings of the quality of a product or service. They assume that products have a true quality, distributed according to a prior $P(\omega)$, and that each individual experiencing a particular product of quality ω^* gets a noisy signal t of the quality, distributed according to $P(t|\omega^*)$. An individual with private signal t forms a posterior belief about the true quality $P(\omega|t)$. Miller et al. design payment schemes to obtain truthful opinions based on probability scoring rules. Jurca and Faltings [6] show how to minimize the payments needed to offset the potential gain from lying. In both cases, the authors set rewards that depend explicitly on the prior distribution. In a similar setting, Prelec [7] suggests an alternative approach by delegating to participants the estimation of the distribution. Individuals are asked for both their private opinion and their beliefs about the posterior distribution of opinions given their information. In Prelec's mechanism, the rewards do not depend on any distribution directly, but indirectly through the reports of the participants.

In the mechanisms described so far, the goal is to enforce truthful reports from each participant through a Nash implementation. This is achieved by choosing payment schemes in function of the distributions, either provided by the mechanism designer or by the participants. This poses a number of practical difficulties: the mechanism designer usually does not know the parameters of the model, and asking individuals to report a distribution may be unnatural and infeasible with many outcomes. Fortunately, to obtain a sample of opinions that is representative of the population, we need not know the opinion of a specific individual.

Jurca and Faltings [8] consider the related problem of obtaining the distribution of opinions in an online setting. In contrast to previous work, their payment schemes does not depend on any prior. However, their mechanism is limited to binary opinions, such as yes/no answers, and is only correct asymptotically as the number of participants grows to infinity. In particular, it cannot be used to obtain a true sample of opinions.

Besides, in all cases, the authors consider a bayesian model with common prior. However, in many situations of interest, information is asymmetric and no general assumption can be made about the knowledge of individuals. For example, when rating an hotel, people who often travel in rural areas will form a different belief about the distribution of hotel quality than those who frequently visit large cities. Some individuals may be more informed than others, for example a frequent business traveler staying over an extended period can hold more accurate beliefs than occasional travelers with short stays. In general, when beliefs depend on information that is common knowledge nor part of the private signal being reported, mechanisms that assume a common prior are no longer incentive-compatible.

To construct robust mechanisms that do not rely on any particular assumption about the knowledge of the agents, it is common to look for dominant-strategy implementations. While a dominant-strategy implementation cannot be achieved in the present setting, we propose mechanisms to obtain independent samples of

opinions representative of the population, based on a Nash implementation that does not rely on any particular knowledge structure. In particular, no common prior is needed, and there may be asymmetric information. Our mechanisms provide incentives through a payment scheme that depends only on the reports of opinions of the individuals being surveyed. When at least one participant may be trusted, our mechanisms ensure that at all Nash equilibria correspond to true samples of opinions.

The paper is organized as follows. We present the problem and the model in Section 2. In Section 3, we propose an intermediary mechanism to elicit random values from given distributions. Those results are used in Section 4, which presents our main survey mechanisms. We conclude in Section 5.

2 Model

We consider a large population of individuals, each of whom owns an opinion regarding a given question (e.g., what is the quality of this hotel? what will be the price of a barrel of oil in 10 years?). We assume opinions can be expressed as real values in some closed interval \mathcal{I} , for example a scale between 0 (worst hotel quality) and 10 (best hotel quality). F represents the distribution of opinions across the population: for each opinion value x , the quantity $F(x)$ denotes the proportion of the population with an opinion less than or equal to x . Formally, we may consider that the population forms a continuum of individuals in the interval \mathcal{I} distributed according to F (hereafter referred to as the *population distribution*). We assume that F is absolutely continuous (i.e., admits a density function), and that the density is positive on \mathcal{I} . The objective of the surveyor is to obtain n independent samples of opinions, which may be used for example to estimate the population distribution or to perform statistical analysis, such as hypothesis testing, goodness-of-fit, etc.

The process of surveying the population is accomplished by a survey mechanism. Formally, a survey mechanism is a tuple (\mathcal{I}, n, Π) . \mathcal{I} is the interval of possible values of opinions, n is the number of agents being surveyed, and $\Pi : \mathcal{I}^n \mapsto \mathbb{R}^n$ is the vector of payments. The mechanism is interpreted as follows:

- Step 1.** The surveyor selects n individuals at random from the population, referred to as “agent 1, . . . , agent n ”.
- Step 2.** Each agent i reports an opinion $r_i \in \mathcal{I}$.
- Step 3.** Each agent i gets a payment $\Pi_i(r_1, \dots, r_n)$.

Each individual knows whether she is being surveyed, but does not know the identities of the other agents being surveyed. As agents are selected at random, their opinions are (ex-ante) identically and independently drawn from F . We assume that agents are rational and seek to maximize their expected payment. The population distribution is, a priori, not known.

We will be interested in mechanisms that satisfy certain properties described below. For a mechanism (\mathcal{I}, n, Π) :

BUDGET-BALANCE. The mechanism is *budget-balanced* when it generates no profit nor loss: for all possible reports $r_1, \dots, r_n \in \mathcal{I}$,

$$\sum_{1 \leq i \leq n} \Pi_i(r_1, \dots, r_n) = 0.$$

ANONYMITY. The mechanism is *anonymous* when payments do not depend on the ordering of the agents: for all possible reports $r_1, \dots, r_n \in \mathcal{I}$, all agent i , and all permutations σ of $\{1, \dots, n\}$,

$$\Pi_i(r_1, \dots, r_n) = \Pi_{\sigma(i)}(r_{\sigma^{-1}(1)}, \dots, r_{\sigma^{-1}(n)}).$$

The surveyor's objective is to obtain samples of opinions that are representative of the population, and is captured by the following two properties:

ACCURACY. The mechanism is *accurate* when each agent reporting an opinion drawn (ex-ante) according to the population distribution is a Nash equilibrium.

STRONG ACCURACY. The mechanism is *strongly accurate* when each agent reporting an opinion drawn (ex-ante) according to the population distribution constitute the *only* Nash equilibria.

Note that reporting one's true opinion is an accurate strategy, because the opinion of an agent selected at random from the population is ex-ante distributed according to F . However, depending on the information available to each individual, there are accurate strategies that are not truthful: for example, each agent reporting the opinion of her neighbor would still lead to accuracy. This is not limiting, as the surveyor is not interested in the opinion of a particular individual, but only in reports of opinions representative of the population.

Although we do not consider an implementation in dominant strategies, our results hold independently of the knowledge structure of the population. In addition to her own opinion, each individual may have some knowledge about the population, about the knowledge of the population, about the knowledge of the knowledge of the population, etc. For example, individuals may be ignorant and know nothing about the population distribution. Or individuals may be omniscient and know the opinion of each individual in the population. Alternatively, there may be asymmetric information: some individuals may be ignorant and others may know exactly the population distribution. There may be a common prior, or different priors conditional on the history of each member of the population. There may be publicly available information, such as the mean opinion, etc. For simplicity, the reader may consider a complete information setting in which opinions all individuals of the population are common knowledge, however our results are much more general.

3 Generating Random Values

In this section, we present a mechanism to elicit random values drawn from any given distributions. The results of this section will be used to prove properties of our survey mechanisms.

3.1 Mechanism Description

We consider a group of n agents $1, \dots, n$. For all agent i , let F_i be a cumulative distribution on a closed interval \mathcal{I} , absolutely continuous with positive density. The distributions F_1, \dots, F_n are common knowledge. We define the following *random generator mechanism*:

Step 1. Each agent i is asked to report a value randomly drawn from F_i .

Step 2. Each agent i is rewarded a payment given by

$$\begin{aligned} \Pi_i(r_1, \dots, r_n) = & \\ & \frac{1}{n-1} (|\{j \mid F_i(r_i) < F_j(r_j)\}| - |\{j \mid F_i(r_i) > F_j(r_j)\}|) \\ & + 2F_i(r_i) - \frac{2}{n-1} \sum_{j \neq i} F_j(r_j) . \end{aligned}$$

This mechanism creates incentives for each agent i to report a random value drawn from F_i , as shown in the next theorem.

Theorem 1. *The random generator mechanism satisfies the following properties:*

1. *The mechanism is budget-balanced.*
2. *If $F_1 = \dots = F_n$, the mechanism is anonymous.*
3. *The payments take values in the range $[-1, 1]$.*
4. *There exists a unique Nash equilibrium, corresponding to each agent i reporting a random number drawn according to F_i .*

Proof. Let $\mathbb{1}_B$ be the function that equals 1 if the boolean statement B is true, and 0 otherwise.

Items 1., 2., and 3. are easily shown, the proof is omitted due to space constraints.

Item 4. We proceed in two steps. We begin by showing that each agent i choosing a value r_i at random from F_i is a Nash equilibrium, then we show the equilibrium is unique.

Let $\mathcal{I} = [a, b]$, and consider any particular agent i . Assume that any other agent $j \neq i$ chooses to report a value r_j distributed according to F_j . Let r_k be the report of any agent k .

Writing the payment of i as

$$\Pi_i(r_1, \dots, r_n) = \frac{2}{n-1} \sum_{j \neq i} \left[\frac{1}{2} \mathbb{1}_{F_i(r_i) < F_j(r_j)} - \frac{1}{2} \mathbb{1}_{F_i(r_i) > F_j(r_j)} + F_i(r_i) - F_j(r_j) \right] ,$$

we get the expected payment for agent i , given her report r_i :

$$\begin{aligned} \mathbb{E}_{r_j \sim F_j, j \neq i} [\Pi_i(r_1, \dots, r_n)] &= \frac{2}{n-1} \sum_{j \neq i} \left[\int_a^{F_j^{-1}(F_i(r_i))} \left(-\frac{1}{2}\right) f_j(r_j) dr_j \right. \\ &\quad + \int_{F_j^{-1}(F_i(r_i))}^b \left(\frac{1}{2}\right) f_j(r_j) dr_j \\ &\quad \left. + \int_a^b (F_i(r_i) - F_j(r_j)) f_j(r_j) dr_j \right] \\ &= \frac{1}{n-1} \sum_{j \neq i} \left[-\frac{F_i(r_i)}{2} + \frac{1 - F_i(r_i)}{2} + F_i(r_i) \right. \\ &\quad \left. - \int_a^b F_j(r_j) f_j(r_j) dr_j \right] \\ &= 0 . \end{aligned}$$

Therefore the expected payment of agent i is null for any report r_i . In particular, a randomized value distributed according to F_i is a best response. By symmetry, each agent i choosing a value at random drawn from F_i is a Nash equilibrium.

We now prove that the Nash equilibrium is unique. For all $1 \leq i \leq n$, let G_i be (cumulative) distributions such that each agent i choosing to report a random value distributed according to G_i is a Nash equilibrium (with the convention that pure strategies correspond to point mass distributions).

Agent i 's expected payment is

$$\mathbb{E}_{r_j \sim G_j, j \neq i} [\Pi_i(r_1, \dots, r_n)] = \frac{1}{n-1} \sum_{j \neq i} \int_a^b H_j(F_i(r_i)) dG_i(r_i) \tag{1}$$

under the Riemann-Stieltjes integral, with

$$\begin{aligned} H_j(\alpha) &= \int_a^{F_j^{-1}(\alpha)} \left(-\frac{1}{2}\right) dG_j(r_j) + \int_{F_j^{-1}(\alpha)}^b \frac{1}{2} dG_j(r_j) \\ &\quad + \int_a^b [\alpha - F_j(r_j)] dG_j(r_j) . \end{aligned}$$

After simplification and rearranging the terms,

$$H_j(\alpha) = \left[\alpha - G_j(F_j^{-1}(\alpha)) \right] - \left[-\frac{1}{2} + \int_a^b F_j(r_j) dG_j(r_j) \right] . \tag{2}$$

By Hewitt’s theorem [9], we may integrate by parts the second term and apply the change of variable $y = F_j(x_j)$:

$$\begin{aligned} -\frac{1}{2} + \int_a^b F_j(r_j) dG_j(r_j) &= -\frac{1}{2} + [F_j(r_j)G_j(r_j)]_a^b + \int_a^b G_j(r_j)f_j(r_j) dr_j \\ &= \frac{1}{2} - \int_a^b G_j(r_j)f_j(r_j) dr_j \\ &= \frac{1}{2} - \int_0^1 G_j(F_j^{-1}(x)) dx \\ &= \int_0^1 [x - G_j(F_j^{-1}(x))] dx , \end{aligned}$$

where we observed that $\int_0^1 x dx$. We replace the last term of (2) and get

$$H_j(\alpha) = \Gamma_j(\alpha) - \int_0^1 \Gamma_j(y) dy , \tag{3}$$

where we defined $\Gamma_j(y) = y - G_j(F_j^{-1}(y))$. Let

$$\zeta_i = \frac{1}{n-1} \sum_{j \neq i} \Gamma_j . \tag{4}$$

By putting together (1), (3) and (4), we get

$$\mathbb{E}_{r_j \sim G_j, j \neq i} [H_i(r_1, \dots, r_n)] = \int_a^b \left[\zeta_i(F_i(r_i)) - \int_0^1 \zeta_i(y) dy \right] dG_i(r_i) . \tag{5}$$

Suppose by contradiction that there exists i such that $\zeta_i \neq 0$. Then we show that there exists some possible report r_i^* such that agent i choosing r_i^* makes a positive expected payment.

We first prove that $\int_0^1 \zeta_i(y) dy < \sup \zeta_i$. If the inequality is false, then $\zeta_i = \sup \zeta_i$ almost everywhere, however since $\zeta_i \neq 0$ and $\zeta_i(0) = \zeta_i(1) = 0$, we can choose $y < 1$ such that $\zeta_i(y) < \sup \zeta_i$. As G_j and F_j^{-1} are nondecreasing, $\Gamma_j(y + \epsilon) < \sup \zeta_i$ for $\epsilon > 0$ small enough, so that ζ_i does not almost everywhere equal $\sup \zeta_i$.

Since $\int_0^1 \zeta_i(y) dy < \sup \zeta_i$, there exists y^* such that $\zeta_i(y^*) - \int_0^1 \zeta_i(y) dy > 0$, and so by taking $r_i^* = F_i^{-1}(y^*)$, we find that agent i choosing the pure strategy r_i^* would make a positive expected payment according to (5).

Since i plays a Nash equilibrium, i ’s strategy is a best response and her expected payment is at least that obtained by choosing the pure strategy r_i^* and so is strictly positive. Therefore, if the Nash equilibrium is such that $\zeta_i \neq 0$ for some i , then i ’s payment is strictly positive, otherwise $\zeta_i = 0$ and i ’s expected profit is null. So the expected profit of every agent is non-negative, and if there exists at least one agent i such that $\zeta_i \neq 0$, agent’s i profit is strictly positive, which is impossible as the mechanism is budget-balanced. Hence for all i , $\zeta_i = 0$, which implies $y = G_j(F_j^{-1}(y))$: the only possible Nash equilibrium corresponds to $G_j = F_j$, for all j .

3.2 Graphical Interpretation

The *lazy hiker race* gives an intuitive interpretation of our mechanism, and may be described as follows. A group of n hikers starts a march on a mountain with 1 mile high. Each hiker has a designated trail, which is common knowledge. All the trails share a common starting/ending point. Hikers are able to keep track of the distance they cover along their own trail, but cannot observe the progression of others. After 10 hours, the march stops and hikers are ranked in decreasing order of altitude.

Hikers want to win the race, and are strong enough to climb to the top within the time limit. But they are also lazy and prefer to win by making as little effort as possible. When there are two hikers, the winner gets the maximum satisfaction (+1) when he wins by being just above the other hiker. He gets the worse satisfaction (−1) when he wins by being at the top, while the other hiker remains at the bottom. The loser’s satisfaction is the opposite of the winner’s. Satisfaction is linear in the difference of altitude, so that the winner is indifferent between winning and losing when the difference of altitude between the hikers is 0.5 miles. Figure 1 illustrates the case of two hikers. If there are more than two hikers, the satisfaction of a hiker equals the average satisfaction when he compares himself to each other hiker.




Fig. 1. Profile of trails for two hikers A and B

Hikers strategize about the distance they should cover so as to maximize their average satisfaction. We observe that, if $F_i(d_i)$ denotes the altitude reached by hiker i after walking a distance d_i , the payment $\Pi_i(d_1, \dots, d_n)$ of the random generator mechanism described previously corresponds exactly to the satisfaction of hiker i when hikers $1, \dots, n$ cover the respective distances d_1, \dots, d_n (where we take $\mathcal{I} = [0, 1]$).

To simplify our argument, let’s take the case of two hikers. We first note that there is no pure Nash equilibrium: if the loser knows where the winner is, he will

¹ For simplification d_i is normalized and equals to the proportion of the total distance covered.

change his strategy to place himself slightly above the winner. Therefore hikers should cover a random distance. When one hiker chooses to cover a distance so that his altitude is uniformly random, each hiker gets a null satisfaction on average, no matter what the other hiker decides to do. If, however, one hiker will likely stop at a low altitude, the other hiker would get a likely positive satisfaction by stopping at a medium altitude. More generally, when a hiker makes frequent stops at some altitudes, the other can choose a location so as to get a positive expected satisfaction. Therefore any choices of random distances that result in nonuniform distribution of altitudes cannot lead to a Nash equilibrium. A similar argument applies to groups of any size. Note that uniform distributions of altitudes are obtained only when each hiker i covers a distance d_i chosen at random according to the distribution F_i . For a given distribution F with density $f = F'$, one can verify that the trail with profile given by

$$x(y) = \int_0^y \sqrt{\frac{L}{f(F^{-1}(h))^2} - 1} \, dh$$

will generate the mixed-Nash equilibrium strategy with distribution F , where L is the desired length of the trail, with $L > \max f^2$. Figure 2 shows some density functions and their associated trails.

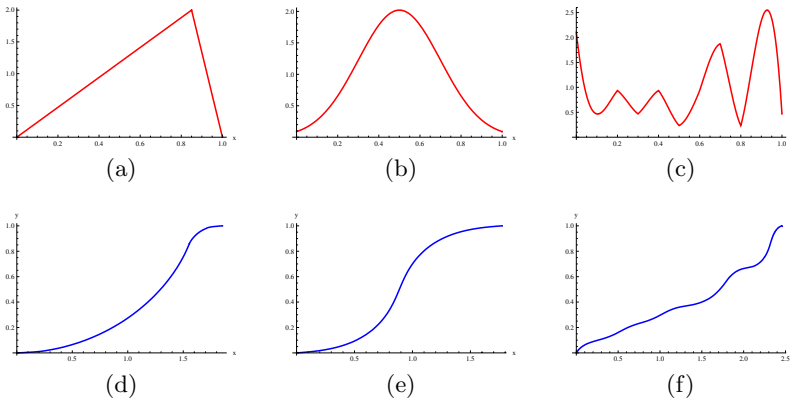


Fig. 2. The equilibrium strategy for trails with profiles (d), (e) and (f) is to choose a random distance with respective densities (a), (b) and (c)

4 Mechanisms for Truthful Surveys

We now describe our survey mechanisms. Given a random sample of k opinions x_1, \dots, x_k , let $\tilde{F}^{x_1, \dots, x_k}(x)$ be a statistical estimator of the proportion of the population having an opinion less than or equal to x . In practice, it is common to use a probabilistic model with parameterized densities. The maximum-likelihood

parameters may for example be obtained through the Expectation-Maximization algorithm [10]. The statistical estimator is said to be *unbiased* when, for all x ,

$$\mathbb{E}_{X_1, \dots, X_k \sim F} [\tilde{F}^{X_1, \dots, X_k}(x)] = F(x)$$

for all population distribution F . For example, the empirical distribution is an unbiased estimator:

$$\tilde{F}^{x_1, \dots, x_k}(x) = \frac{1}{k} \sum_{1 \leq i \leq k} \mathbb{1}_{x_i < x} .$$

with $\mathbb{1}_{x_i < x} = 1$ if $x_i < x$ and $\mathbb{1}_{x_i < x} = 0$ otherwise.

Let $\mathcal{G}_1, \dots, \mathcal{G}_k$ be a partition of the n agents into k groups, $k \geq 2$. Let $\mathcal{G}(i)$ be the group that includes agent i , and let $\mathcal{S}_i = \{1, \dots, n\} \setminus \mathcal{G}(i)$ be the set of agents that doesn't include the group containing i . Our *basic survey mechanism* uses the random generator mechanism of the previous section to incentivize each agent i to reveal an opinion that corresponds to a statistical estimate of the distribution population. For a given interval of possible opinions \mathcal{I} and a number of agents n , the payments of our survey mechanism are defined by

$$\begin{aligned} \Pi_i(r_1, \dots, r_n) = & \frac{1}{|\mathcal{G}(i)| - 1} [|\{j \mid r_i < r_j\}| - |\{j \mid r_i > r_j\}|] \\ & + 2\tilde{F}_i(r_i) - \frac{2}{|\mathcal{G}(i)| - 1} \sum_{j \in \mathcal{G}(i), j \neq i} \tilde{F}_j(r_j) \end{aligned}$$

where $\tilde{F}_i(x) = \tilde{F}^{\{r_j\}_{j \in \mathcal{S}_i}}(x)$ is an unbiased statistical estimator of $F(x)$ given the reports of agents in \mathcal{S}_i . By linearity of expectation, and noting that the payments are linear in the estimators, the following can be derived from Theorem 1.

Theorem 2. *The basic survey mechanism is budget-balanced, anonymous, and accurate.*

Payments take values in the interval $[-3, 3]$. Payments near the interval bounds occur only with estimators of high variance. As the variance decreases, payments become restricted to the interval $[-1, 1]$. One may also offset/rescale the payments, to get for example payments in the interval $[0, 1]$ so as to provide participation incentives and strict individual rationality.

The mechanism easily adapts to the case of sequential elicitation, often desired in online surveys. The surveyor should form groups of 2 or 3 people, progressively as new reports come in, and reward individuals of a group as soon as the group is finalized. Distribution estimates should be computed from reports of previous groups only, with the exception of the distribution estimates used for rewarding the first group, which could take as input reports of the second group.

As opposed to the work of Miller et al. [5] and Jurca and Faltings [6], truthful revelation is a non-strict equilibrium of our mechanism. This limitation is due to the lack of common prior: it is easily shown that, with our general knowledge structures, no survey mechanism may implement truthful reporting as a strict Nash equilibrium. Indeed, the expected payment for an agent who reports her

true opinion must be maximized under all possible distributions of opinions, and therefore must be constant for all possible reports. We observe by the same argument that any survey mechanism that implements reporting one’s true opinion as a Nash equilibrium is also accurate, in the sense of Section 2.

Note that, as accuracy only requires that truthful reporting is a Nash equilibrium, a trivial mechanism that assigns a zero payoff to all agents also meets the accuracy criterium: all strategies constitute a Nash equilibrium. However, our mechanism is not trivial, as it admits only restricted Nash equilibria. Indeed it can be shown that the only Nash equilibria correspond to each group reporting the same aggregate distribution. As agents do not know, *ex-ante*, which group they belong to, the only strategies that always result in a Nash equilibrium correspond to all agents reporting an opinion drawn from the same distribution. Therefore it would require significant coordination among the agents to play a Nash equilibrium that is not truthful.

Besides, the Nash equilibria of our mechanism are not unique, and all methods for eliciting subjective information suffer from the multiplicity of Nash equilibria, since rewards can only be a function of information submitted by the agents. However, unlike other methods such as Miller et al. [5] in which non-truthful Nash-equilibria may lead to higher revenue for all agents, in our mechanism all Nash equilibria lead to a null expected payment for all agents.

If there are trusted individuals, we can ensure uniqueness of the Nash equilibria that correspond to accurate samples, so that the surveyor is guaranteed to obtain true random samples. Let \mathcal{T} be a group of trusted individuals who provide their true opinion. \mathcal{T} may not be empty but can be of any positive size, larger groups are generally preferred as they reduce the variance of individual payments. For simplicity we assume that trusted agents form a separate group from the n surveyed agents. The payments of our *trusted-survey mechanism* are defined as follows:

$$\begin{aligned} \Pi_i(r_1, \dots, r_n) = & \frac{1}{n-1} [|\{j \mid r_i < r_j\}| - |\{j \mid r_i > r_j\}|] \\ & + 2\tilde{F}(r_i) - \frac{2}{n-1} \sum_{j \neq i} \tilde{F}(r_j), \end{aligned}$$

where $\tilde{F}(x) = \tilde{F}^{\{r_j\}_{j \in \mathcal{T}}}(x)$ is an unbiased estimator of $F(x)$ given by the reports of trusted agents in \mathcal{T} .

Our next theorem claims that the trusted-survey mechanism is guaranteed to elicit true random samples of opinions. As for Theorem 2, the proof follows from Theorem 1.

Theorem 3. *The trusted-survey mechanism is budget-balanced, anonymous, and strongly accurate.*

5 Conclusion

We have investigated the problem of incentivizing individuals to elicit samples of opinions that are representative of a population. We have proposed a nontrivial

budget-balanced, anonymous mechanism for which reporting a true sample of opinion, in particular reporting one's true opinion, is a Nash equilibrium. When some opinions can be trusted, we propose a variation of our mechanism which guarantees that the only Nash equilibria correspond to providing true samples. Although we use a Nash implementation as opposed to a dominant strategy implementation—impossible to achieve in our setting—our results do not depend on the knowledge structure of the population, in particular we do not make use of a common prior.

We believe an important avenue for future work is that of empirical studies. Our analysis has focused on theoretical considerations. However, it is not clear how individuals would behave in practice. Experiments studying and comparing our approach with those whose payments depend on a common prior, either provided by the surveyor as in Miller et al. [5], or provided by the agents as in Prelec [7], or simply assumed by all agents as in Jurca and Faltings [8], would need to be performed to help assess the validity of each method, and their potential applicability to practical contexts. In particular, more work would be needed to understand the limitations raised by the common prior assumption and those implied by the weakness of the Nash equilibria in our mechanisms.

References

1. Montgomery, D.: Design and Analysis of Experiments. Wiley, Chichester (1984)
2. Kish, L.: Survey sampling. Wiley, Chichester (1995)
3. Winkler, R., Muñoz, J., Cervera, J., Bernardo, J., Blattenberger, G., Kadane, J., Lindley, D., Murphy, A., Oliver, R., Ríos-Insua, D.: Scoring rules and the evaluation of probabilities. *TEST* 5(1), 1–60 (1996)
4. Lambert, N., Pennock, D., Shoham, Y.: Eliciting Properties of Probability Distributions. In: Proceedings of the 9th ACM Conference on Electronic Commerce, pp. 129–138 (2008)
5. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science* 51(9), 1359–1373 (2005)
6. Jurca, R., Faltings, B.: Minimum payments that reward honest reputation feedback. In: Proceedings of the 7th ACM Conference on Electronic Commerce, pp. 190–199 (2006)
7. Prelec, D.: A Bayesian Truth Serum for Subjective Data. *Science* 306(5695), 462–466 (2004)
8. Jurca, R., Faltings, B.: Incentives for Expressing Opinions in Online Polls. In: Proceedings of the 9th ACM Conference on Electronic Commerce, pp. 119–128 (2008)
9. Hewitt, E.: Integration by Parts for Stieltjes Integrals. *The American Mathematical Monthly* 67(5), 419–423 (1960)
10. Dempster, A., Laird, N., Rubin, D., et al.: Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society* 39(1), 1–38 (1977)

Correlated Equilibrium of Bertrand Competition

John Wu

School of Business, East China Normal University,
200241 Shanghai, China
jswoo@sina.com

Abstract. This paper explores the relation between equilibrium coarsenings and equilibrium refinements via Bertrand competition example and similar situations, it shows that the typical equilibrium coarsening — a unique correlated equilibrium — is equivalent to the unique Nash equilibrium itself, is also equivalent to the equilibrium refinement, for the standard n -firms Bertrand competition model with linear demand and symmetric, linear costs in the most special and simplest case, and compares some wonderful and remarkable differences of the existence, uniqueness, stability, connectivity, and strategic property of Nash equilibrium and correlated equilibrium between Cournot and Bertrand model. We also propose some open questions.

Keywords: Equilibrium coarsenings, equilibrium refinements, strategic correlation principle, positive correlated equilibrium, negative correlated equilibrium, duality gap.

1 Introduction

Van Damme(2002), Aumann and Dreze(2008), etc. call three basic solution concepts — objective correlated equilibrium(henceforth correlated equilibrium), subjective correlated equilibrium (henceforth subjective equilibrium) and rationalizability — coarsenings of Nash equilibrium(henceforth equilibrium coarsenings), in order to contrast with the well known refinements of Nash equilibrium(henceforth equilibrium refinements)¹.

Before these classifications have been made, Aumann himself research both subjects and hope the world will enjoy on both too. He rejected the idea that correlated equilibrium represent the truth or that subgame perfect equilibrium represent the truth solely, and said one should not do this or that exclusively, but one should develop both kind of concepts and see where they lead².

These classifications and Aumann's hope remind us to realize that equilibrium coarsenings are at least the same important as well as equilibrium refinements, and to pursue the relation between them simultaneously³.

¹ See any “principle” textbook of game theory for equilibrium refinements and Section 4 for a brief discussion.

² See Van Damme(1998)'s interview with Robert Aumann.

³ See Section 4 for this relation and Subsection 6.g for a conjecture and its various meanings.

The pursuit of simultaneously uniqueness issue of the set of correlated equilibrium and the set of Nash equilibrium is started by many authors. In the case of two firms, i.e., Cournot and Bertrand duopoly, the result of the uniqueness of the correlated equilibrium is known and follows from two different directions, rationalizability and supermodular games. On the one hand, the fact that there is a unique correlated equilibrium for the Cournot duopoly is deduced from Bernheim's work (1984) because correlated equilibrium are independently rationalizable for a two-player game (Brandenburger and Dekel 1987). On the other hand, when there are two firms, the supermodular game, Milgrom and Roberts (1990) prove that if the game has a unique pure strategy Nash equilibrium, this equilibrium is the unique rationalizable strategy profile and hence the unique correlated equilibrium. However, when the number of firms is more than 2, Cournot and Bertrand games are not supermodular and have a continuum of rationalizable strategies.

Liu asserts that the unique Nash equilibrium of the homogeneous good n -firms Cournot model with linear demand and linear costs coincides with the unique correlated equilibrium. Neyman(1997) calls this is a surprising result⁴ and generalizes it. Neyman considers the class of strategic games that have convex strategy sets, smooth concave potential functions and bounded payoffs and shows that for this class of games the set of correlated equilibrium coincides with the set of mixtures of pure strategies that maximize the potential. In games with continuous action space, he also obtains a sufficient condition of uniqueness of correlated equilibrium if, in addition, the potential function is strictly concave. A recent paper by Ui (2008) further generalizes the condition to a weaker one suffices for the uniqueness of a correlated equilibrium. Bergemann and Morris(2008) extend it to belief-free incomplete information games and show that if a game has a smooth concave potential for every payoff type profile and also has an ex post equilibrium, then the ex post equilibrium forms a unique incomplete information correlated equilibrium.

Yi (1997) also extends Liu's results in two directions by introducing symmetric product differentiation and by allowing for convex cost functions. Yi explores conditions on demand and cost functions under which the unique pure-strategy Nash equilibrium of the Cournot model with symmetric product differentiation (which includes the case of no differentiation as a special case) is also the unique correlated equilibrium, and establishes the equivalence of the pure strategy Nash equilibrium set and the correlated equilibrium set for linear demand and weakly convex cost functions.

Viossat (2005)⁵ and Calvó-Armengol(2006) also get the result of the set of correlated equilibrium and Nash equilibrium of competitive games coincide and are reduced to one single point, but both of them does not mention the differences between correlated equilibrium of Cournot model and Bertrand model⁶.

Many people think that the unique Nash equilibrium must always be the unique correlated equilibrium since any Nash equilibrium is correlated equilibrium. Sekiguchi (2005) shows it is true for any finitely repeated game with imperfect monitoring, if all stage game correlated equilibrium are equilibrium minimaxing and have the same

⁴ I thank the referee for pointing out it's not totally surprising that games with unique Nash equilibrium have also unique correlated equilibrium now, because this is not generally the case; see below for more in-depth information.

⁵ I thank the referee for pointing out Viossat's work.

⁶ See Section 5 for general differences and Subsection 6.c for the sign of covariance.

payoff vector, then any correlated equilibrium of the repeated game must be a repetition of stage game correlated equilibrium and therefore must have the same payoff vector. He also give a surprising counterexample that a stage game has a unique Nash equilibrium and has not a unique correlated equilibrium when anyone of two conditions is relaxed, so the finitely repeated game may have an equilibrium outcome that is quite different from its stage game equilibrium.

It is the very pursuing process and the attractive property of Bertrand model that constitutes the macro and micro background of this paper; we will to explore the relation between equilibrium coarsenings and equilibrium refinements via Bertrand competition example and its related topics.

This work is just a beginning and a by-product of the small part of the large project⁷, it shows that the set of typical equilibrium coarsening —(the unique) correlated equilibrium — is equivalent to the set of Nash equilibrium itself, is also equivalent to the set of equilibrium refinement in the most special and simplest case and similar situations — Where there are a lot of interesting properties, i.e., there is consensus and the consensus is very clear on which refinement is appropriate because the equilibrium uniqueness and can greatly simplify or reduce information complexity, communication complexity, computational complexity and strategy complexity in correlated mechanism design or correlated market design, and can goes a long way only carrying a little rationality, and so on.

The plan of the paper is as follows: After summarizing preliminaries in Section 2, we devote Section 3 to a careful conceptual discussion and constructive proof of the main result. Section 4 briefly recalls the literature and lists a kind of feasible classifications of equilibrium coarsenings. Section 5 is devoted to general discussion about some wonderful and remarkable differences of the unique Nash equilibrium and correlated equilibrium between Cournot and Bertrand model, and Section 6 to listing some open questions.

2 Preliminaries

2.1 Bertrand Model

We begin by considering static games of complete information. The Bertrand model, in which several firms compete to supply a homogeneous good with a given market demand of the good, is well presented in a strategic game form. Let $N=\{1, \dots, i, \dots, n\}$ be a finite set of firms. For each firm $i \in N$, let S^i be a set of price possibilities of firm i . Let S be the Cartesian product of all S^i , i.e., $S=\prod_{i \in N} S^i$, an n -tuple price distribution is $p=(p^i)_{i \in N} \in S$. For each $i \in N$ and $p \in S$, $p^{-i}=\{p^1, \dots, p^{i-1}, p^{i+1}, \dots, p^n\}$ denotes the price distribution supplied by every firm but firm i ; thus $p^{-i} \in S^{-i}=\prod_{j \in N \setminus \{i\}} S^j$ and $p=\{p^{-i}, p^i\}$. Firm i has the cost function $C_i(P^i)$ and the demand function $D_i(P^i)$, P^i is price of firm i . Firm i 's profits are $\pi^i=p^i D_i(P^i) - C_i(D_i(P^i))$; i.e., for each firm $i \in N$, $\pi^i:S \rightarrow R$ is the payoff function of firm i . A Bertrand-Nash equilibrium of the above model is a combination, $p^* \in S$, such that for every i in N , $p^{*i} \in \text{Argmax}_{p^i \in S^i} \pi(p^{*-i}, p^i)$.

⁷ See section 4 for the project.

The model that is studied in the paper is a standard Bertrand model with linear demand and symmetric, linear cost functions. The demand function $D_i(P^i)$ is linear, i.e.,

$$Q^i = D_i(P^1, \dots, P^n) = \begin{cases} D(P^i) = (a - P^i) / b > 0 & P^i < \min\{a, P^j\} & i, j \in N, i \neq j \\ D(P^i) / m = (a - P^i) / mb > 0 & P^i = P^j \equiv P < a & i \in N, j \in M, i \neq j \\ 0 & P^i > P^j & i, j \in N, i \neq j \\ 0 & P^i > a & i \in N \end{cases}$$

where $a > 0, b > 0; m = \#\{j \in N: P^j = P^i\}, 1 < m \leq n, M = \{1, \dots, i, \dots, m\}$. Each firm i has a linear cost function $C_i(D_i(P^i))$, i.e., $C_i(D_i(P^i)) = c_i D_i(P^i)$, where c_i is the constant marginal cost. Therefore, the profit of firm i is given by

$$\pi^i = P^i D_i(P^i) - C_i(D_i(P^i)) = P^i D_i(P^i) - c_i D_i(P^i) = (P^i - c_i) D_i(P^i). \tag{1}$$

It is well known that in this model there exists a unique Bertrand-Nash equilibrium point, $P^* = \{P^1, \dots, P^{*i}, \dots, P^{*n}\}$, where the price is $P^{*i} = c_i$, and all the firms produce a positive amount is the symmetric one, i.e., $Q^* = D(P) / n = (a - c) / nb > 0, m = n$, with the above symmetric price.

2.2 Correlated Equilibrium

A correlated equilibrium (Aumann 1974) is a Nash equilibrium of a game where each player gets a private or public signal from a correlation mechanism before the beginning of the original game.

For a finite game (N and S are finite), a probability vector $F = F(p) p \in S$ on S is a correlated equilibrium distribution if for all $i \in N$, the following incentive compatibility constraint inequality is satisfied:

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) [\pi^i(p^{-i}, p^i) - \pi^i(p^{-i}, \xi^i(p^i))] \geq 0. \tag{2}$$

3 The Main Result

In this section we will to explore the relation between equilibrium coarsenings and equilibrium refinements in Bertrand competition example, together with some similar basic propositions.

For the class of n -firms Bertrand models with linear demand and linear, symmetric cost functions, we find a following equivalent property between their correlated and Nash equilibrium.

Theorem. An n -firms Bertrand competition with linear demand and symmetric, linear cost functions has a unique correlated equilibrium, which is the unique Nash equilibrium.

Proof. Assume that there exists a correlated equilibrium p^i , which differs from the unique Nash equilibrium P^{*i} ; i.e., there exists a probability distribution F on S such that $F(S \setminus \{p^*\}) > 0$ and satisfying incentive constraint (2),

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) \left\{ \begin{array}{l} [P^i(a - P^i) / nb - c_i(a - P^i) / nb] \\ - [\xi^i(P^i)(a - \xi^i(P^i)) / nb - c_i(a - \xi^i(P^i)) / nb] \end{array} \right\} \geq 0.$$

i.e.,

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) \{ [(P^i - c_i)(a - P^i) / nb] - [(\xi^i(P^i) - c_i)(a - \xi^i(P^i)) / nb] \} \geq 0.$$

$\forall i \in N$, where the function ξ^i represents an arbitrary deviation of firm i , $\xi^i : S^i \rightarrow S^i$, which must not be profitable.

Let ζ be an n -tuple of deviations, i.e., $\zeta = \{\zeta^1, \dots, \zeta^i, \dots, \zeta^n\}$, and $\zeta^i : S^i \rightarrow S^i$.

The above inequalities can be rewritten as

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) \{ (P^i - \xi^i(P^i)) [a - c_i - p^i - \xi^i(P^i)] / nb \} \geq 0, \tag{3}$$

$\forall i \in N$, We consider an n -tuple of deviations ζ from $p \in S$ to the Nash equilibrium point P^* , i.e.,

$$\zeta(P) = \{\zeta^1(P^1), \dots, \zeta^i(P^i), \dots, \zeta^n(P^n)\} = \{P^{*1}, \dots, P^{*i}, \dots, P^{*n}\},$$

where for each firm i ,

$$\xi^i(P^i) = P^{*i} = c_i. \tag{4}$$

Then (3) can be rewritten by using (4)

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) [(P^i - c_i)(a - P^i) / nb] \geq 0 \tag{5}$$

$\forall i \in N$. There are three possible cases for the inequalities (5), i.e., for each firm i :

(1) If

$$\xi^i(P^i) = P^{*i} = c_i > P^i,$$

Then

$$(a - P^i) / nb > 0$$

And

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) [(P^i - c_i)(a - P^i) / nb] < 0,$$

This deviation must not be profitable; what's more, by (5), we have,

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) [(P^i - c_i)(a - P^i) / nb] \geq 0,$$

A contradiction.

(2) If

$$\xi^i(P^i) = P^{*i} = c_i < P^i,$$

Then

$$(a - P^i) / nb > 0.$$

And

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) [(P^i - c_i)(a - P^i) / nb] > 0,$$

Although this deviation satisfies (5)

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) [(P^i - c_i)(a - P^i) / nb] \geq 0,$$

But

$$(a - P^i) / nb > 0,$$

A contradiction.

(3) If

$$\xi^i(P^i) = P^{*i} = c_i = P^i,$$

Then

$$(a - P^i) / nb > 0$$

And

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) [(P^i - c_i)(a - P^i) / nb] = 0,$$

This deviation satisfies both

$$(a - P^i) / nb > 0$$

And

$$\sum_{p^{-i} \in S^{-i}} F(p^{-i}, p^i) [(P^i - c_i)(a - P^i) / nb] \geq 0.$$

Since the first deviation case must not be profitable and the second deviation case must not be impossible, so only the third deviation case $\xi^i(P^i) = P^{*i} = c_i = P^i$ holds in the above mapping, i.e., the original correlated equilibrium strategy must equal the Nash equilibrium strategy almost surely,

$$P^i = P^{*i}. \tag{6}$$

$$\forall i \in N. \tag{Q.E.D.}$$

This result is similar to Milgrom and Roberts(1990)' earlier one, but our method is totally different from theirs. Their result is an indirect and implicit one based on the solution of the supermodular game. First, our proof is a direct one, which uses Aumann's inequality directly, and follows Liu (1996) closely in notation, with some changes to accommodate Bertrand competition. Second, our proof is more compact in notation as well as generally applicable, which can be used in the analysis of correlated equilibrium with arbitrary number of players, i.e., not only the oligopolistic competition of "small" games with few players — duopoly is the special case of oligopoly competition only with two players, but also the monopolistic competition of large games with many players. Third, our proof is constructive and computable, which can be easily extended to construct and analyze the Bertrand competition in network or combinatorial markets, and can be used in the design of the algorithm of correlated equilibrium or correlated mechanism.

Corollary 1. The set of typical equilibrium coarsening — (the unique) correlated equilibrium — is equivalent to the set of Nash equilibrium itself, is also equivalent to the set of equilibrium refinement, for n-firms Bertrand competition model with linear demand and symmetric, linear costs.

Proof. The first half is obvious by the above Theorem. Because there is only one game, its subgame is the original game itself; hence, the set of Nash equilibrium is

also equivalent to the set of equilibrium refinement. So, the second half is obvious too. Q.E.D.

Corollary 2. The set of typical equilibrium coarsening — (the unique) correlated equilibrium — is equivalent to the set of Nash equilibrium itself, is also equivalent to the set of equilibrium refinement, for n-firms Cournot competition model with linear demand and asymmetric, linear costs.

Proof. The first half is obvious by the above Theorem or Liu(1996)'s equilibrium uniqueness theorem. Because there is only one game, its subgame is the original game itself; hence, the set of Nash equilibrium is also equivalent to the set of equilibrium refinement. So, the second half is obvious too. Q.E.D.

Corollary 3. The set of typical equilibrium coarsening — (the unique) correlated equilibrium — is equivalent to the set of Nash equilibrium itself, is also equivalent to the set of equilibrium refinement, for belief-free incomplete information games with supermodularity.

Proof. The first half is obvious by Bergemann and Morris(2008)'s equilibrium uniqueness theorem. Because there is only one game, its subgame is the original game itself; hence, the set of Nash equilibrium is also equivalent to the set of equilibrium refinement. So, the second half is obvious too. Q.E.D.

Corollary 4. The set of typical equilibrium coarsening — (the unique) correlated equilibrium — is equivalent to the set of Nash equilibrium itself, is also equivalent to the set of equilibrium refinement, for any finitely repeated game with imperfect monitoring which all stage game correlated equilibrium are equilibrium minimaxing and have the same payoff vector.

Proof. The first half is obvious by Sekiguchi (2005)'s equilibrium uniqueness theorem. Because there is only one game, its subgame is the original game itself; hence, the set of Nash equilibrium is also equivalent to the set of equilibrium refinement. So, the second half is obvious too. Q.E.D.

Remark. The result of Corollary 4 is not hold when anyone of two conditions is relaxed.

4 Literature

In this section we briefly recall the literature and list a feasible one of equilibrium coarsenings from various logically reasonable classifications.

Historically, equilibrium refinements once succeeded in getting rid of many equilibria, but there is no consensus on which refinement is appropriate and came a cost: the sharper the equilibrium refinement, the more the information complexity, the communication complexity, the computational complexity and the strategy complexity, and the higher the degree of rationality needed on the part of players(Maskin1999). On the contrary, the duller the equilibrium coarsening, the less the information complexity, the communication complexity, the computational complexity and the strategy complexity, and the lower the degree of rationality needed on the part of players.

Although the area of equilibrium refinements has not died out, but it is no longer as active as before because too much effort was put into this area, it has reached a certain

maturity and has become quite esoteric at some point (Hart2007), whereas equilibrium coarsenings are well studied since 1970s, the literature on these subjects is too large to survey it here, we can only suggest a possible rough framework in statistical game-theoretic language, e.g.,

- I. correlated equilibrium
 - a. positive correlated objective equilibrium
 - b. zero correlated objective equilibrium
 - c. negative correlated objective equilibrium
- II. subjective equilibrium
 - a. positive correlated subjective equilibrium
 - b. zero correlated subjective equilibrium
 - c. negative correlated subjective equilibrium
- III. rationalizability
 - a. positive correlated rationalizability
 - b. zero correlated rationalizability
 - c. negative correlated rationalizability

Note that there exists a small conceptual puzzle in zero correlated objective equilibrium which is old in statistics and new in game theory, we may call it Gul's paradox because it was pointed out by Frank Gul(1998), similarly in zero correlated subjective equilibrium and zero correlated rationalizability. In order to avoid this kind of conceptual puzzle, zero correlated objective equilibrium, subjective equilibrium and rationalizability here, which are neatly permutation of all the equilibrium coarsening concepts, just mean independent objective equilibrium, subjective equilibrium and rationalizability.

Moreover, we must make effort to find which equilibrium coarsening is interesting because it is not all equilibrium coarsenings are meaningful. In the rest sections, we restrict our attention to positive correlated equilibrium and negative correlated equilibrium.

Recently, Young(2007) even proposes a new concept of coarse correlated equilibrium or the "coarse" notion of correlated equilibrium which maybe be called it equilibrium coarsenings' coarsenings, he also conjectures that coarse correlated equilibrium may prove useful in describing the behavior of experimental subjects which has never been investigated systematically.

Since equilibrium coarsenings is becoming a more and more important concept of game theory, it is time to summarize their rapid developments and to explore the relation between equilibrium coarsenings and equilibrium refinements now. This constitutes a large project which remains us to accomplish in the future, and requires us to dig deep into or renew our toolkits to carry out this mission in the new era of algorithmic game theory or algorithmic economics.

5 Discussion

There are some wonderful and remarkable differences of existence, uniqueness, stability, connectivity, and strategic property of Nash equilibrium and correlated equilibrium between Cournot and Bertrand model:

- a. The existence of Nash equilibrium and correlated equilibrium of Cournot model with linear demand and asymmetric, linear cost functions, according to Liu (1996); whereas the existence of Nash equilibrium and correlated equilibrium of Bertrand model with linear demand and symmetric, linear cost functions, according to this paper.
- b. The uniqueness of Nash equilibrium and correlated equilibrium of Cournot model and Bertrand model can be described as the same conditions as 5.a.
- c. The stability of Nash equilibrium and correlated equilibrium of Cournot model is evolutionary stable strategies, according to Qin and Stuart (1997), whereas the stability of the unique Nash equilibrium and correlated equilibrium of Bertrand model is not evolutionary stable strategies.
- d. The connectivity, such as pairwise stable network (Jackson and Wolinsky 1996), of Nash equilibrium and correlated equilibrium of Cournot model, if links have a negligible cost, one firm gains with each link that it adds; if link costs are small enough, then the complete network is the unique pairwise stable network, according to Jackson (2006), whereas the pairwise stable network of Nash equilibrium and correlated equilibrium of Bertrand model has no such links.
- e. The strategic property of Nash equilibrium and correlated equilibrium of Cournot model is the firms' outputs are strategic substitutes: each firm's best response function slopes downwards; according to Bulow, Geanakoplos and Klemperer (1985), whereas the strategy property stability of the Nash equilibrium and correlated equilibrium of Bertrand model is the firms' outputs are strategic complements: each firm's best response function, assuming they exist, slopes upward.

6 Open Questions

Our work suggests a number of interesting open questions because there are so many useful aspects of equilibrium coarsenings which we need to know.

- a. The fundamental difference of the unique Nash equilibrium and correlated equilibrium between Cournot and Bertrand model maybe due to strategic riskiness valuation based on real option consideration of the players (Aumann and Serrano 2007, Foster and Hart 2007, 2008).
- b. When the number of players of strategic competition is two and the power of the two players is not equal, there is a duality gap between the two players, according to linear program method. The set of correlated equilibrium may converge towards the set of Nash equilibrium while the duality gap become smaller and smaller. The set of correlated equilibrium and the set of Nash equilibrium both degenerate to a singleton while the duality gap equivalent to zero — which the set of correlated equilibrium and the set of Nash equilibrium both degenerate to a singleton (Nau, Gomez Canovas and Hansen 2004). If this claim is true, is the duality gap method equivalent to Myerson (1997)'s dual reduction method?
- c. The covariance of correlated equilibrium (Chwe 2006) is related to the equilibrium coarsenings. The sign of covariance of correlated equilibrium of Bertrand model is positive, whereas the sign of covariance of correlated equilibrium of Cournot model is negative.

- d. Bertrand and Cournot correlated equilibrium represent two types of most basic correlated equilibrium: positive correlated equilibrium and negative correlated equilibrium. The sign of covariance of correlated equilibrium is positive represent positive correlated equilibrium; the sign of covariance of correlated equilibrium is negative represent negative correlated equilibrium.
- e. positive correlated equilibrium and negative correlated equilibrium are the key concept of strategic correlation principle of strategic competition, measured by strategic positive correlation coefficient, strategic negative correlation coefficient, correspond to strategic complements and strategic substitutes — another pair of the traditional key concept of strategic competition.
- f. The indicators of positive correlated equilibrium and negative correlated equilibrium, such as strategic positive correlation coefficient and strategic negative correlation coefficient, can be partial or complete identified by nonparametric or semiparametric estimation methods in statistical game theory, and it is also important in network game setting; similarly to other indicators, e.g. variance, volatility, signal to noise ratio (Fudenberg and Levine2008), covariance (Chwe2006), likelihood ratio (Abreu, Milgrom and Pearce 1991, Aoyagi2002, Zheng2008), hazard rate(Mahdian, McAfee and Pennock2008), etc.
- g. Equilibrium coarsenings and equilibrium refinements may be viewed as results of information set's coarsenings and information set's refinements in some condition. Moreover, they may be also viewed as mutually converse algorithm in the measure of information complexity, communication complexity, computational complexity and strategy complexity. This analogy make game theory more like Newton and Leibniz's calculus, i.e., coarsenings may like integral, whereas refinements may like differential, which was pointed out by Von Neumann and Morgenstern in their great masterpiece over sixty years ago. If so, then we are going into a new golden age of mathematics, science and philosophy.

Acknowledgments. I am grateful to Robert Aumann for his bringing me into the exciting area of game theory, and six referees, Lin Zhou for their useful comments and suggestions, and Elchnan Ben-Porath, Chengzhong Qin, Myrna Wooders for their time and patience in reading an early draft. Conversations with Ken Binmore, Colin Camerer, Xiaotie Deng, Drew Fudenberg, Yinyu Ye is particularly helpful, their intellectual generosity is greatly appreciated.

References

1. Abreu, D., Milgrom, P., Pearce, D.G.: Information and Timing in Repeated Partnerships. *Econometrica* 59(6), 1713–1733 (1991)
2. Aoyagi, M.: Collusion in dynamic Bertrand oligopoly with correlated private signals and communication. *Journal of Economic Theory* 102, 229–248 (2002)
3. Aumann, R.J.: Subjectivity and correlation in randomized strategies. *Journal of Mathematic Economics* 1, 67–96 (1974)
4. Aumann, R.J., Dreze, J.H.: Rational Expectations in Games. *American Economic Review* 98(1), 72–86 (2008); See also: When All is Said and Done, How Should You Play and What Should You Expect? Discussion paper no.387, Center for the Study of Rationality of The Hebrew University of Jerusalem (2005)

5. Aumann, R.J., Serrano, R.: An economic Index of Riskiness. Discussion paper no.446, Center for the Study of Rationality of The Hebrew University of Jerusalem (2007)
6. Bergemann, D., Morris, S.: Belief Free Incomplete Information Games. In: Papadimitriou, Zhang (eds.) Proceeding of Workshop on Internet and Network Economics. Springer, Heidelberg (2008)
7. Bernheim, D.: Rationalizable strategic behavior. *Econometrica* 52, 1007–1028 (1984)
8. Brandenburger, A., Dekel, E.: Rationalizability and correlated equilibrium. *Econometrica* 55, 1391–1402 (1987)
9. Bulow, J., Geanakoplos, J.D., Klemperer, P.D.: Multimarket oligopoly: Strategic Substitutes and Complements. *Journal of Political Economy* 93(3), 488–511 (1985)
10. Calvó-Armengol, A.: The Set of Correlated Equilibria of 2×2 Games, <http://selene.uab.es/acalvo>
11. Chwe, M.S.-Y.: Incentive Compatibility Implies Signed Covariance (2006), <http://www.chwe.net/michael/i.pdf>
12. Foster, D.P., Hart, S.: An Operational Measure of Riskiness, Discussion paper no.454, Center for the Study of Rationality of The Hebrew University of Jerusalem (2007)
13. Foster, D.P., Hart, S.: A Reserve-based Axiomatization Of the Measure of Riskiness, Discussion paper, Center for the Study of Rationality of The Hebrew University of Jerusalem (2008)
14. Fudenberg, D., Levine, D.: Repeated Games with Frequent Signals. *Quarterly Journal of Economics* (to appear, 2008)
15. Gul, F.: A Comment on Aumann's Bayesian View. *Econometrica* 66, 923–927 (1998)
16. Hart, S.: Five Questions on Game Theory. In: Hendricks, V.F., Hansen, P.G. (eds.) *Game Theory: 5 Questions*, pp. 97–107. Automatic Press (2007)
17. Jackson, M.O.: The Economics of Social Networks. In: Blundell, R., Newey, W., Persson, T. (eds.) *Advances in Economics and Econometrics, Theory and Applications*. Cambridge University Press, Cambridge (2006)
18. Jackson, M.O., Wolinsky, A.: A strategic Model of Social and Economic networks. *Journal of Economic Theory* 71, 44–74 (1996)
19. Liu, L.: Correlated equilibrium of Cournot oligopoly competition. *Journal of Economic Theory* 68, 544–548 (1996)
20. Mahdian, M., McAfee, R.P., Pennock, D.: The Secretary Problem with a Hazard Rate Condition. In: Papadimitriou, Zhang (eds.) *Proceeding of Workshop on Internet and Network Economics*. Springer, Heidelberg (2008)
21. Maskin, E.S.: *Introduction to Recent Developments in Game Theory*. Edward Elgar Publishing (1999)
22. Milgrom, P., Roberts, J.: Rationalizability, learning and equilibrium in games with strategic complementarities. *Econometrica* 58, 1255–1278 (1990)
23. Myerson, R.B.: Dual Reduction and Elementary Games. *Games and Economic Behavior* 21, 183–202 (1997)
24. Nau, R.F., Gomez Canovas, S., Hansen, P.: On the Geometry of Nash Equilibria and Correlated Equilibria. *International Journal of Game Theory* 32, 443–453 (2004)
25. Neyman, A.: Correlated Equilibrium and Potential Games. *International Journal of Game Theory* 26, 223–227 (1997)
26. Pearce, D.G.: Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica* 52(4), 1029–1050 (1984)
27. Qin, C.-Z., Stuart, C.: Are Cournot and Bertrand equilibria evolutionary stable strategies? *Journal Evolutionary Economics* 7, 41–47 (1997)

28. Sekiguchi, T.: Uniqueness of equilibrium payoffs in finitely repeated game with imperfect monitoring. *The Japanese Economic Review* 56(3), 317–331 (2005)
29. Van Damme, E.: Strategic Equilibrium. In: Aumann, R., Hart, S. (eds.) *Handbook of Game Theory*, ch. 41, vol. III. North Holland, Amsterdam (2002)
30. Van Damme, E.: On the State of the Art in Game Theory: An Interview with Robert Aumann. *Games and Economic Behavior* 24, 181–210 (1998)
31. Viossat, Y.: Openness of the set of games with a unique correlated equilibrium. *cahier du laboratoire d'économétrie* 2005-28, Ecole polytechnique, France (revised in, 2006)
32. Ui, T.: Correlated Equilibrium and Concave Games. *International Journal of Game Theory* 37(1), 1–13 (2008)
33. Yi, S.: On the Existence of a Unique Correlated Equilibrium in Cournot Oligopoly. *Economics Letters* 54, 235–239 (1997)
34. Young, H.P.: The Possible and the Impossible in Multi-Agent Learning. *Artificial Intelligence* 171, 429–433 (2007)
35. Zheng, B.: Approximate efficiency in repeated games with correlated private signal. *Games and Economic Behavior* 63(1), 406–416 (2008)

Diffusion of Innovations on Random Networks: Understanding the Chasm

Marc Lelarge

INRIA-ENS

Paris, France

marc.lelarge@ens.fr

Abstract. We analyze diffusion models on sparse random networks with neighborhood effects. We show how large cascades can be triggered by small initial shocks and compute critical parameters: contagion threshold for a random network, phase transition in the size of the cascade.

1 Introduction

In Crossing the Chasm [12], Moore begins with the diffusion of innovations theory from Everett Rogers [16], and argues there is a chasm between the early adopters of the product (the technology enthusiasts and visionaries) and the early majority (the pragmatists). According to Moore, the marketer should focus on one group of customers at a time, using each group as a base for marketing to the next group. The most difficult step is making the transition between visionaries (early adopters) and pragmatists (early majority). This is the chasm that he refers to.

In this paper, we analyze a simple model of diffusion with neighborhood effects on random networks and we show that it can explain this chasm. Most of the epidemic models [14], [15] consider a transmission mechanism which is independent of the local condition faced by the agents concerned. But if there is a factor of persuasion or coordination involved, relative considerations tend to be important in understanding whether some new behavior or belief is adopted [17].

We begin by discussing one of the most basic game-theoretic diffusion models proposed by Morris [13]. Consider a graph G in which the nodes are the individuals in the population and there is an edge (i, j) if i and j can interact with each other. Each node has a choice between two possible behaviors labeled A and B . On each edge (i, j) , there is an incentive for i and j to have their behaviors match, which is modeled as the following coordination game parameterized by a real number $q \in (0, 1)$: if i and j choose A (resp. B), they each receive a payoff of q (resp. $(1 - q)$); if they choose opposite strategies, then they receive a payoff of 0. Then the total payoff of a player is the sum of the payoffs with each of her neighbors. If the degree of node i is d_i and S_i^B is her number of neighbors playing B , then the payoff to i from choosing A is $q(d_i - S_i^B)$ while the payoff from choosing B is $(1 - q)S_i^B$. Hence i should adopt B if $S_i^B > qd_i$ and A if $S_i^B \leq qd_i$. A number of qualitative insights can be derived from a diffusion model

even at this level of simplicity. Specifically, consider a network where all nodes initially play A . If a small number of nodes are forced to adopt strategy B (the seed) and we apply best-response updates to other nodes in the network, then these nodes will be repeatedly applying the following rule: switch to B if enough of your neighbors have already adopted B . There can be a cascading sequence of nodes switching to B such that a network-wide equilibrium is reached in the limit. This equilibrium may involve uniformity with all nodes adopting B or it may involve coexistence, with the nodes partitioned into a set adopting B and a set sticking to A . Morris [13] considers the case of infinite regular graph G and provides graph-theoretic characterizations for when these different types of equilibria arise.

Our work allows us to study rigorously an extension of this model, the symmetric threshold model, when the underlying network is a random network with given vertex degrees. We are able to characterize the relation between the network and the individual behavior. In particular, we compute the contagion threshold of the random network and validate a heuristic result of Watts [18]. We also show that there is a phase transition for the set of adopters at a critical value of the size of the initial seed. To the best of our knowledge, this result is new and our work is the first rigorous analysis of a general threshold model on a random network. Although random graphs are not considered to be highly realistic models of most real-world networks, they are often used as first approximation and are a natural first choice for a sparse interaction network in the absence of any known geometry of the problem.

In [4], the influence maximization problem is defined as follows: given a social network, find a small set of 'target' individuals so as to maximize the number of customers who will eventually purchase the product following the effect of word-of-mouth. Hardness results have been obtained in [8], [3] and there is a large literature on this topic. However, in most practical cases, the structure of the underlying network is not known and then one has to rely on distributional assumptions (like distribution of the degrees). Our model allows to answer the probabilistic version of the influence maximization problem, when the exact topology of the social network is not known.

The rest of the paper is organized as follows. In Section [2], we describe our model. Section [3] contains the main results in particular, the contagion threshold is computed and the phase transition phenomena is explained. Section [4] contains technical details and we conclude in Section [5].

2 Model

2.1 The Configuration Model

In this section, we define our random graph model which is standard in the literature on random graphs [2]. Let $n \in \mathbb{N}$ and let $(d_i)_1^n = (d_i^{(n)})_1^n$ be a sequence of non-negative integers such that $\sum_{i=1}^n d_i$ is even. We define a random multigraph with given degree sequence $(d_i)_1^n$, denoted by $G^*(n, (d_i)_1^n)$ by the configuration

model [2]: take a set of d_i half-edges for each vertex i and combine the half-edge into pairs by a uniformly random matching of the set of all half-edges. Conditioned on the multigraph $G^*(n, (d_i)_1^n)$ being a simple graph, we obtain a uniformly distributed random graph with the given degree sequence, which we denote by $G(n, (d_i)_1^n)$.

We will let $n \rightarrow \infty$ and assume that we are given $(d_i)_1^n$ satisfying the following regularity conditions [11]:

Condition 1. For each n , $(d_i)_1^n = (d_i^{(n)})_1^n$ is a sequence of non-negative integers such that $\sum_{i=1}^n d_i$ is even and, for some probability distribution $(p_r)_{r=0}^\infty$ independent of n ,

- (i) $\#\{i : d_i = r\}/n \rightarrow p_r$ for every $r \geq 0$ as $n \rightarrow \infty$;
- (ii) $\lambda := \sum_r r p_r \in (0, \infty)$;
- (iii) $\sum_{i=1}^n d_i/n \rightarrow \lambda$ as $n \rightarrow \infty$;
- (iv) $\sum_i d_i^2 = O(n)$.

In words, (p_r) describes the distribution of the degrees, λ is the average mean degree in the graph, condition (iii) ensures that the number of edges divided by n tends to the average degree divided by 2. The technical condition (iv) is required to transfer the results from $G^*(n, (d_i)_1^n)$ to $G(n, (d_i)_1^n)$ [5].

The results of this work can be applied to some other random graphs models too by conditioning on the vertex degrees. For example, for the Erdős-Rényi graph $G(n, p)$ with $np \rightarrow \lambda \in (0, \infty)$, the assumptions hold with p_r the distribution of a Poisson random variable with mean λ .

We consider asymptotics as $n \rightarrow \infty$ and say that an event holds w.h.p. (with high probability) if it holds with probability tending to 1 as $n \rightarrow \infty$.

2.2 Symmetric Threshold Model

The contagion model of [13] is the simplest model for cascading behavior in a social network: people switch to a new behavior when a certain threshold fraction of neighbors have already switched. Our symmetric threshold model generalizes this model by allowing the threshold fraction be a random variable with distribution depending on the degree of the node and which are independent among nodes. This is to account for our lack of knowledge of the exact threshold value of each individual. Formally, we define for each $d \in \mathbb{N}$, a sequence of i.i.d. random variables in \mathbb{N} denoted by $(K(d), K_i(d))_{i=1}^\infty$. The threshold associated to node i is $K_i(d_i)$ where d_i is the degree of node i .

Now the progressive dynamics of the behavior operates as follows: some set of nodes S starts out adopting the new behavior B ; all other nodes start out adopting A . We will say that a node is active if it is following B . Time operates in discrete steps $t = 1, 2, 3, \dots$. At a given time t , any inactive node i becomes active if its fraction of active neighbors exceeds its threshold $K_i(d_i) + 1$. This in turn may cause others nodes to become active leading to potentially cascading adoption of behavior B . We will suppose that $K_i(1) = 0$ for all i , so that any leaf of the network is active as soon as its parent becomes active.

It is easy to see that the final set of active nodes (after n time steps if the network is of size n) only depends on the initial set S (and not on the order of the activations) and can be obtained as follows: set $X_i = 1$ for all i in the set of initial adopters. Then as long as there exists i such that $\sum_{j \sim i} X_j > K_i(d_i)$, set $X_i = 1$. When this algorithm finishes, the final state of node i is represented by X_i : $X_i = 1$ if node i is active and $X_i = 0$ otherwise. It is easily seen that the linear threshold model [9] is covered by our framework (see [10] for a proof).

3 Main Results

3.1 Contagion Threshold of a Random Graph

We consider the simple contagion model studied by Morris in [13] on a random graph, i.e. $K_i(d) = qd$ for all i . We define the contagion threshold of the graph to be the maximum q for which a single individual can trigger a global cascade, i.e. activate a strictly positive fraction of the total population, w.h.p. This notion is the natural extension of the contagion threshold defined in [13] for regular graphs.

Proposition 1. *The contagion threshold q_c is given by*

$$q_c = \sup \left\{ q : \sum_{1 \leq s < 1/q} s(s-1)p_s > \lambda \right\}.$$

This result is in accordance with the heuristic result of [18] (see in particular the cascade condition Eq. 5 in [18]) and is proved in Section 4. Figure 1 gives the contagion threshold as a function of λ , the mean degree of the graph.

Note that q is related to the quality of the new technology: the lower q is, the better the quality of the new technology is. In particular if $q < 1/2$, then technology B is better than technology A . Hence q_c can be interpreted as the minimal quality for technology B to get a non-negligible adoption with a finite initial seed of adopters.




Fig. 1. q_c for the contagion model on a Poisson random graph (green dashed) and on a Power-law random graph (red) as a function of λ

3.2 Phase Transition in the Contagion Model

We now still consider the contagion model but for $q > q_c$. In this case, in order to trigger a large cascade, the set of initial adopters must be a non-negligible fraction of the total population. For simplicity, we assume that each node of the network is part of the initial set of adopters with probability α independently of everything else. In particular, the fraction of initial adopters is α and we now compute the final proportion of active nodes: $\Phi(\alpha) = \lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n X_i}{n}$.

We need to introduce some notation first. For integers $\ell \geq 0$ and $0 \leq r \leq \ell$ let $b_{\ell r}$ denote the binomial probabilities $b_{\ell r}(p) := \binom{\ell}{r} p^r (1-p)^{\ell-r}$. We denote by D a random variable with distribution $\mathbb{P}(D = r) = p_r$. For $0 \leq p \leq 1$ we let D_p be the thinning of D obtained by taking D points and then randomly and independently keeping each of them with probability p : $\mathbb{P}(D_p = r) = \sum_{\ell=r}^{\infty} p_{\ell} b_{\ell r}(p)$. We now define $h(p) = \mathbb{E}[D_p \mathbf{1}(D_p \geq (1-q)D)]$.

The following proposition shows that the map $\alpha \mapsto \Phi(\alpha)$ exhibits point of discontinuity.

Proposition 2. *Consider a random graph such that $p_1 > 0$ and let \tilde{p} be the largest local maximum point of $\psi(p) = h(p)/p^2$ in $(0, 1)$. Then there is a phase transition at $\alpha_c = 1 - \frac{\lambda}{\psi(\tilde{p})}$: the function $\Phi(\cdot)$ is discontinuous at α_c .*

Figure 2 shows an example of such a phase transition in the case of Poisson random graphs.




Fig. 2. Function $\Phi(\alpha)$ for the contagion model on a Poisson random graph with parameter $\lambda = 6$ and $q = 0.3$

Returning to the (probabilistic) influence maximization problem, our derivation of the function $\Phi(\alpha)$ is of crucial importance. In particular, the fact that this function is highly non-linear seems not to have been taken into account so far and will have a big impact on the optimal strategy. In the case where the marketer knows the degree of each individual (but not the underlying social network), our derivation of $\Phi(\alpha)$ will allow her to target her effort, by choosing the variable α .

3.3 Dynamic of the Epidemic

In previous section, α was related to the amount spent by the marketer and q corresponded to the quality of the new technology. We now consider that α is actually fixed and corresponds to the fraction of technology enthusiasts in the population. The rest of the population consists of pragmatists. Then the marketer’s effort allows to increase the perceived quality by decreasing the value of q . It is easy to see that the phase transition described in previous section translates in a phase transition in the parameter q . Moreover, let consider the simple following dynamic of the epidemic: the edges of the active nodes become active(meaning that the end-point of the edge actually notices that his neighbor is active) at rate 1 (see [10] for more details) . Then Figure 3 shows the case where the real quality of the technology is $q = 0.3$. Without any marketing, a small fraction of the pragmatists adopt the new technology but with marketing, the diffusion is able to ‘cross the chasm’ and a large fraction of the population adopt the new technology.




Fig. 3. Dynamic of the epidemic for Poisson random graph with $\lambda = 6$, $\alpha = 0.06$ and for $q = 0.29$ and $q = 0.3$

4 Exact Asymptotics

In this section, we state the theorem which is the corner stone of our work (see [10] for a proof). Recall that D_p is the thinning of D (defined in Section 3.2). We define the functions

$$h(p) := \mathbb{E}[D_p \mathbf{1}(D_p \geq D - K(D))], \tag{1}$$

$$h_1(p) := \mathbb{P}(D_p \geq D - K(D)). \tag{2}$$

Theorem 1. Consider the graph $G(n, (d_i)_1^n)$ satisfying Condition 1 where each node is part of the initial set of adopters with probability α independently of everything else. Let $\hat{p} := \max\{p \in [0, 1] : (1 - \alpha)h(p) = \lambda p^2\}$.

- (i) If $(1 - \alpha)h(p) < \lambda p^2$ for all $p \in (0, 1]$, which is equivalent to $\hat{p} = 0$, then w.h.p. $\Phi(\alpha) = 1$.

(ii) If $(1 - \alpha)h(p) \geq \lambda p^2$ for some $p \in (0, 1]$, which is equivalent to $\hat{p} \in (0, 1]$, and further \hat{p} is not a local maximum point of $(1 - \alpha)h(p) - \lambda p^2$, then w.h.p. $\Phi(\alpha) = 1 - (1 - \alpha)h_1(\hat{p})$.

The proof of this Theorem consists in an extension of the work of Janson and Luczak [7] where the k -core problem is studied. Our model is related to the bootstrap percolation which is more or less the opposite of taking the k -core: with our notation, it consists in taking $K_i(d_i) = k$ a fix constant. For regular graphs (i.e. $d_i = d$ for all i), this process has been studied. Theorem 1 of [1] or Theorem 5.1 of [6] correspond exactly to our Theorem in the particular case of a d -regular graph, with fixed threshold.

Proof. of Proposition 1: Following the heuristic in [18], we introduce the following threshold: $K_i(d) = (d + 1)\mathbb{1}(d \geq 1/q)$. In words, a node i becomes active if one of his neighbor is active and $d_i < 1/q$. Clearly the nodes that become active in this model need to have only one active neighbor in the original contagion model with parameter q . For any node i , let C_i denotes the final set of active nodes when starting with only i as active node. Clearly, if $j \in C_i$, then we have $C_i = C_j$. Now if we prove that $\Phi(0+) = \lim_{\alpha \rightarrow 0} \Phi(\alpha) - \alpha > 0$, then for any $\Phi(0+) > \alpha > 0$, at least one of the $n\alpha$ nodes say i in the initial set has activated at least $\frac{\Phi(0+)}{\alpha}n$ nodes. Hence we have $\#C_i/n \geq \frac{\Phi(0+)}{\alpha}$ and any point in C_i will activate at least the set C_i in the original contagion model. We now prove that for $q < q_c$, we have $\Phi(0+) > 0$. This will implies that the contagion threshold is larger than q_c . We have

$$\begin{aligned} h(p) &= \mathbb{E}[D_p \mathbb{1}(D \geq 1/q) + D \mathbb{1}(D < 1/q, D_p = D)] \\ &= \sum_{s \geq 1/q} spp_s + \sum_{s < 1/q} sp_s p^s \\ &= p \left(\sum_{s \geq 1/q} sp_s + \sum_{s < 1/q} sp_s p^{s-1} \right). \end{aligned}$$

Let $f(p) = \lambda p - \frac{h(p)}{p}$. The condition $\Phi(0+) > 0$ is equivalent to for $\epsilon > 0$ small enough $f(1 - \epsilon) > 0$. We have $f(1 - \epsilon) = \epsilon \left(-\lambda + \sum_{s < 1/q} s(s - 1)p_s \right) + o(\epsilon)$, which is the condition of the proposition. The proof that for $q > q_c$ a single active node cannot activate a positive fraction of the population is similar and omitted.

Proof. of Proposition 2: Let $f(p, \alpha) = \lambda p^2 - (1 - \alpha)h(p)$. Note that $f(0, \alpha) = 0$ and $f(1, \alpha) = \alpha\lambda$. Then we have $f(p, \alpha) \sim -(1 - \alpha)p_1 p < 0$ as $p \rightarrow 0$ and $\alpha\lambda - f(p, \alpha) \sim (1 - p) \left((1 - \alpha) \left(\sum_{s < 1/q} s(s - 1)p_s - \lambda \right) - 2\alpha\lambda \right) < 0$ as $p \rightarrow 1$ and the result follows easily.

5 Conclusion

We proposed a simple model of diffusion with neighborhood effects which allows to explain the 'chasm'. We should emphasize that the random graph model

considered eliminates a lot of the network structure from the problem (only the degree distributions are preserved). We expect that other local effects like clustering will have a significant impact on the diffusion. However our work shows that neighborhood effects 'alone' can explain the 'chasm' and we think that these effects will actually 'add up'. These issues are left for future research.

References

1. Balogh, J., Pittel, B.G.: Bootstrap percolation on the random regular graph. *Random Structures Algorithms* 30(1-2), 257–286 (2007)
2. Bollobás, B.: *Random graphs*, 2nd edn. Cambridge Studies in Advanced Mathematics, vol. 73. Cambridge University Press, Cambridge (2001)
3. Chen, N.: On the approximability of influence in social networks. In: *SODA 2008* (2008)
4. Domingos, P., Richardson, M.: Mining the network value of customers. In: *KDD 2001* (2001)
5. Janson, S.: The probability that a random multigraph is simple, [arxiv.org:math/0609802](http://arxiv.org/math/0609802) (2006)
6. Janson, S.: On percolation in random graphs with given vertex degrees, arxiv.org:0804.1656 (2008)
7. Janson, S., Luczak, M.J.: A simple solution to the k -core problem. *Random Structures Algorithms* 30(1-2), 50–62 (2007)
8. Kempe, D., Kleinberg, J., Tardos, É.: Maximizing the spread of influence through a social network. In: *KDD 2003* (2003)
9. Kleinberg, J.: Cascading behavior in networks: Algorithmic and economic issues. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
10. Lelarge, M.: Diffusion and cascading behavior in random networks (in preparation)
11. Molloy, M., Reed, B.: A critical point for random graphs with a given degree sequence. *Random Structures Algorithms* 6(2-3), 161–179 (1995)
12. Moore, G.: *Crossing the Chasm*. HarperBusiness, New York (2002)
13. Morris, S.: Contagion. *Rev. Econom. Stud.* 67(1), 57–78 (2000)
14. Newman, M.E.J.: Spread of epidemic disease on networks. *Phys. Rev. E* 66(1), 016128 (2002)
15. Pastor-Satorras, R., Vespignani, A.: *Evolution and Structure of the Internet: A Statistical Physics Approach*. Cambridge University Press, New York (2004)
16. Rogers, E.M.: *Diffusion of Innovations*, 5th edn. Free Press (August 2003)
17. Vega-Redondo, F.: *Complex Social Networks*. Cambridge University Press, Cambridge (2007)
18. Watts, D.J.: A simple model of global cascades on random networks. *Proc. Natl. Acad. Sci. USA* 99(9), 5766–5771 (2002) (electronic)

An Efficient PTAS for Two-Strategy Anonymous Games

Constantinos Daskalakis*

Microsoft Research

Abstract. We present a novel polynomial time approximation scheme for two-strategy anonymous games, in which the players' utility functions, although potentially different, do not differentiate among the identities of the other players. Our algorithm computes an ϵ -approximate Nash equilibrium of an n -player 2-strategy anonymous game in time $\text{poly}(n) \cdot (1/\epsilon)^{O(1/\epsilon^2)}$, which significantly improves upon the running time $n^{O(1/\epsilon^2)}$ required by the algorithm of Daskalakis & Papadimitriou, 2007. The improved running time is based on a new structural understanding of approximate Nash equilibria: We show that, for any ϵ , there exists an ϵ -approximate Nash equilibrium in which either only $O(1/\epsilon^3)$ players randomize, or all players who randomize use the same mixed strategy. To show this result we employ tools from the literature on Stein's Method.

1 Introduction

It has been recently established that computing a Nash equilibrium is an intractable problem [19,11,6,14], even in the case of two-player games [7]. In view of this hardness result, research has been directed towards the computation of approximate Nash equilibria, which are states of the game in which no player has more than some small ϵ incentive to change her strategy. But, despite much research in this direction [23,22,12,18,13,5,28], only constant ϵ 's can be achieved in polynomial time. Yet, an approximate Nash equilibrium in which the players have regret equal to a significant fraction of their payoffs is not an attractive solution concept; after all, there is no reason to expect a player to keep her strategy if she can significantly improve by changing to a different one. On the contrary, if ϵ were arbitrarily small, it could be that the cost of switching one's strategy is larger than the regret ϵ that she suffers. Hence, approximate equilibria with arbitrarily close approximation could be credible solutions concepts. The following question then emerges: *Is there a Polynomial Time Approximation Scheme for approximate Nash equilibria?*

The question remains open for general games, but there are special classes known to be tractable. It is well-known, for example, that zero-sum games are solvable exactly in polynomial time by Linear Programming [25,9]. This tractability result has been extended to a generalization of zero-sum games,

* This research was done while the author was a student at UC Berkeley. Supported by a Microsoft Research Fellowship, NSF grant CCF - 0635319 and a MICRO grant.

called *two-player low-rank games*, in which the sum of the players' payoff tables has fixed rank; in this case there is a PTAS for approximate Nash equilibria. It has also been shown that symmetric multi-player games with (about logarithmically) few strategies per player can be solved exactly in polynomial time by a reduction to the theory of real closed fields [26]. In congestion games, we can compute in polynomial time a pure Nash equilibrium, if the game is a symmetric network congestion game [17], and an approximate pure Nash equilibrium, if the congestion game is symmetric (but not necessarily network) and the utilities are somehow "continuous" [8].

In this paper, we consider another important class of games, called *anonymous*. These are games in which each player's utility function does not differentiate among the identities of the other players. That is, the payoff of a player depends on the strategy that she chooses and only the *number* of other players choosing each strategy. Anonymous games comprise a broad and well studied class of games (see, e.g., [3,4,20,24] for recent work on this subject by economists) which are of special interest to the Algorithmic Game Theory community, as they capture important aspects of auctions and markets, as well as of Internet congestion.

But, what do we know about computing Nash equilibria in anonymous games? It was recently established that there is a PTAS for the case of a constant number of strategies per player [15,16]. The running time of the algorithm given in [16] is $n^{O(f(s,1/\epsilon))}$, where ϵ is the desired approximation, s the number of strategies available to the players, and f some function which is polynomial in $1/\epsilon$, but superpolynomial in s . Hence, although theoretically efficient for any fixed ϵ and s , the algorithm is highly non-practical. Even for the simpler case of two-strategy anonymous games the running time achieved by [15] is $n^{O(1/\epsilon^2)}$.

In this paper, we present a more efficient algorithm for 2-strategy anonymous games, which runs in time $\text{poly}(n) \cdot (1/\epsilon)^{O(1/\epsilon^2)}$. The improved running time is due to a novel understanding of certain structural properties of approximate Nash equilibria. In particular, we show that, for any integer k , there exists an ϵ -approximate Nash equilibrium, with $\epsilon = O(1/k)$, in which

- (a) either at most $k^3 = O((1/\epsilon)^3)$ players use randomized strategies, and their strategies are integer multiples of $1/k^2$; ¹
- (b) or all players who randomize choose the same mixed strategy which is also an integer multiple of $\frac{1}{kn}$.

To derive the above characterization, we study mixed strategy profiles in the proximity of a Nash equilibrium. We establish that there always exists a nearby mixed strategy profile which is of one of the types (a) or (b) described above and satisfies the Nash equilibrium conditions to within an additive ϵ , thus corresponding to an ϵ -approximate equilibrium. Given this structural result (see Theorem 1), an ϵ -approximate equilibrium can be found by dynamic programming (see Theorem 2).

¹ Note that, since every player has 2 strategies, a mixed strategy is a number in $[0, 1]$.

We feel that a more sophisticated analysis can establish similar structural properties for approximate Nash equilibria in multi-strategy anonymous games, extending our efficient PTAS to anonymous games with any fixed number of strategies.

Overview of Techniques: Let us track down the effect in the Nash equilibrium resulting by replacing a mixed Nash equilibrium $(p_1, \dots, p_n) \in [0, 1]^n$ by another strategy profile $(q_1, \dots, q_n) \in [0, 1]^n$, where the probabilities p_i and q_i correspond to the mixed strategy of player i in the two strategy profiles. It is not hard to see that the approximation achieved by the strategy profile (q_1, \dots, q_n) can be, loosely speaking, bounded by the total variation distance between the distribution of the sum of n Bernoulli random variables with expectations p_1, \dots, p_n and that of another sum of Bernoulli random variables with expectations q_1, \dots, q_n . Hence, to establish the structural property described above it is sufficient to show that given any set of probability values (p_1, \dots, p_n) there is another set (q_1, \dots, q_n) which satisfies either Property (a) or Property (b) and is such that the total variation distance between the two sums of Bernoulli random variables with expectations $\{p_i\}_i$ and $\{q_i\}_i$ respectively is at most ϵ .

To give some insight into the construction of the set $\{q_i\}_i$, let us consider the following scenarios for an integer k :

- (i) at least k^3 of the p_i 's fall in the set $[1/k, 1 - 1/k]$ and the others are either 0 or 1;
- (ii) at most k^3 of the p_i 's fall in the set $[1/k, 1 - 1/k]$ and the others are either 0 or 1;

Let us consider Case (i) first. It is reasonable to expect, by the Central Limit Theorem, or finitary versions thereof, that the sum of at least k^3 Bernoulli random variables with expectations from the set $[1/k, 1 - 1/k]$ is close in total variation distance to a Normal distribution, with the appropriate mean and variance, and, hence, to a Binomial distribution which approximates that Normal distribution. In Section 6.1 we show that this is indeed the case, even if we only allow the probability of the Binomial distribution to be an integer multiple of $1/kn$. Hence, the sum of the original Bernoulli random variables with expectations from the set $[1/k, 1 - 1/k]$ can be approximated by another set of Bernoulli random variables which all share the same mean, which, moreover, is an integer multiple of $1/kn$. Hence, an approximate equilibrium satisfying Property (b) above can be defined.

In the Case (ii), approximating by a Normal distribution is not tight enough to give overall total variation distance of $O(1/k)$. We resort instead to the following structural result shown in 1.5: Given any set of Bernoulli random variables with expectations p_1, \dots, p_n , there is a way to round the probabilities to multiples of $1/k^2$, for any k , so that the distribution of the sum of these n variables is affected by an additive $O(1/k)$ in total variational distance. Hence, an approximate equilibrium satisfying Property (a) above can be defined (see Section 6.2).

It remains to argue that the Cases (i) and (ii) are general enough. For this, we describe an iterative procedure which alters the values of those p_i 's falling outside the set $\{0, 1\} \cup [1/k, 1 - 1/k]$ in such a way that, in the end of this

procedure, all p_i 's do fall in the set $\{0, 1\} \cup [1/k, 1 - 1/k]$ and the distribution of the sum of Bernoulli random variables with expectations p_i does not change by more than $O(1/k)$ in total variation distance. Let us, e.g., consider the p_i 's falling in the set $(0, 1/k)$ and round some of them to 0 and some of them to $1/k$ so that their sum is approximated to within $O(1/k)$. By the law of rare events (the precise result we use is from the literature on Poisson approximations, see, e.g., [1]), the sum of Bernoulli's before and after the rounding is distributed like two Poisson distributions with means equal to the sum of p_i 's before and after the rounding respectively. Since these means are within $O(1/k)$, the two Poisson distributions are within $O(1/k)$ in total variation distance and so are the sums of the Bernoulli's before and after the rounding. For details see Section 5.

2 Definitions and Notation

A game has n players, $1, \dots, n$, and t strategies, $1, 2, \dots, t$, available to them, so that each player gets some payoff for every selection of strategies by her and the other players. The game is called *anonymous*, if the payoff of each player depends on her strategy and only the *number*, but not the identities, of the other players who choose each of the t strategies.

In this paper, we study two-strategy anonymous games. In these games, the payoff function of each player i is specified by giving $u_1^i, u_2^i : \{0, 1, \dots, n - 1\} \rightarrow [0, 1]$ ² so that $u_s^i(m)$ is the payoff of i , if she chooses strategy s and m of the other players choose strategy 2. Hence, the game is *succinctly representable* [26], in the sense that its representation requires $2n^2$ numbers, as opposed to the (exponential in the number of players) nt^n numbers required for general games. Arguably, succinct games are the only multiplayer games that are computationally meaningful (see [26] for an extensive discussion of this point).

A *mixed strategy profile* is a set of n probability values $p_1, p_2, \dots, p_n \in [0, 1]$, corresponding to the probability with which each player chooses strategy 2. A mixed strategy profile is an ϵ -Nash equilibrium if, for all $i \in [n]$, the following hold

$$\begin{aligned}
 E_{\{p_j\}_{j \neq i}} u_1^i(x) > E_{\{p_j\}_{j \neq i}} u_2^i(x) + \epsilon &\Rightarrow p_i = 0, \\
 E_{\{p_j\}_{j \neq i}} u_2^i(x) > E_{\{p_j\}_{j \neq i}} u_1^i(x) + \epsilon &\Rightarrow p_i = 1,
 \end{aligned}$$

where for the purposes of the expectation x is drawn from $\{0, \dots, n - 1\}$ by tossing $n - 1$ independent coins with probabilities $\{p_j\}_{j \neq i}$. That is, a mixed strategy profile is an ϵ -Nash equilibrium if every player is only randomizing among strategies which, when played against the mixed strategies of the other players, achieve expected payoff within (additive) ϵ from the expected payoff achieved by the best strategy.

The notion of ϵ -Nash equilibrium is closely related to the notion of ϵ -approximate Nash equilibrium, defined as any mixed strategy profile in which no player can improve his expected payoff by more than ϵ by changing to a different mixed strategy. It is easy to see that any ϵ -Nash equilibrium is also an ϵ -approximate

² In the literature on Nash approximation, utilities are usually normalized in this way so that the approximation error is additive.

Nash equilibrium, but the opposite implication is not true in general. In this paper, we present algorithms for computing (the stronger notion of) ϵ -Nash equilibria.

Anonymous games can be extended to ones in which there is also a finite number of *types* of players, and utilities depend on how many players of each type play each of the available strategies. Our algorithm can be easily generalized to this framework, with the number of types multiplying the exponent of the running time.

Let us conclude this section with a few more definitions. We define the *total variation distance* between two distributions \mathbb{P} and \mathbb{Q} supported on a finite set \mathcal{A} as follows

$$\|\mathbb{P} ; \mathbb{Q}\| := \frac{1}{2} \sum_{\alpha \in \mathcal{A}} |\mathbb{P}(\alpha) - \mathbb{Q}(\alpha)|.$$

Similarly, if X and Y are two random variables ranging over a finite set, their total variation distance, denoted $\|X ; Y\|$, is defined as the total variation distance between their distributions.

We also define the *Translated Poisson distribution* as follows.

Definition 1 ([27]). *We say that an integer random variable Y has a translated Poisson distribution with parameters μ and σ^2 and write $\mathcal{L}(Y) = TP(\mu, \sigma^2)$ if $\mathcal{L}(Y - \lfloor \mu - \sigma^2 \rfloor) = \text{Poisson}(\sigma^2 + \{\mu - \sigma^2\})$, where $\{\mu - \sigma^2\}$ represents the fractional part of $\mu - \sigma^2$.*

Finally, for a positive integer ℓ , we denote $[\ell] := \{1, \dots, \ell\}$.

3 Statement of Results

We show the following probabilistic lemma, whose proof is given in Sections 4-6.

Theorem 1. *Let $\{p_i\}_{i=1}^n$ be arbitrary probability values, $p_i \in [0, 1]$ for all $i = 1, \dots, n$, let $\{X_i\}_{i=1}^n$ be independent indicator random variables such that X_i has expectation $\mathcal{E}[X_i] = p_i$, and let k be a positive integer. Then there exists another set of probabilities $\{q_i\}_{i=1}^n$, $q_i \in [0, 1]$, $i = 1, \dots, n$, which satisfy the following properties:*

1. *if $\{Y_i\}_{i=1}^n$ are independent indicator random variables such that Y_i has expectation $\mathcal{E}[Y_i] = q_i$, then,*

$$\left\| \sum_i X_i ; \sum_i Y_i \right\| = O(1/k), \tag{1}$$

and, for all $j = 1, \dots, n$,

$$\left\| \sum_{i \neq j} X_i ; \sum_{i \neq j} Y_i \right\| = O(1/k). \tag{2}$$

2. the set $\{q_i\}_{i=1}^n$ is such that:

(a) if $p_i = 0$ then $q_i = 0$;

(b) one of the following is true:

- i. there exists $S \subseteq [n]$ and some value q which is an integer multiple of $\frac{1}{kn}$, such that, for all $i \notin S$, $q_i \in \{0, 1\}$, and, for all $i \in S$, $q_i = q$;
- ii. or, there exists $S \subset [n]$, $|S| < k^3$ such that, for all $i \notin S$, $q_i \in \{0, 1\}$, and, for all $i \in S$, q_i is an integer multiple of $\frac{1}{k^2}$.

Our main result (Theorem 2 below) is based on the following observation: if we replace a strategy profile $(p_i)_{i=1}^n$ that is a Nash equilibrium by the nearby strategy profile $(q_i)_{i=1}^n$ specified by Theorem 1, then the change in each player's utility is bounded by the total variation distance between the number of players playing their second strategy in the two strategy profiles. It follows that the search for approximate Nash equilibria can be restricted to the strategy profiles of the form 2(b)i or 2(b)ii specified in Theorem 1, and this search can be done efficiently with dynamic programming. Due to space limitations, the proof of the following theorem is given in the full version of the paper [10].

Theorem 2. *For any $\epsilon < 1$, an ϵ -Nash equilibrium of a two-strategy anonymous game with n players can be computed in time $\text{poly}(n) \cdot U \cdot (1/\epsilon)^{O(1/\epsilon^2)}$, where U is the number of bits required to represent a payoff value of the game.*

4 Overview of the Proof of Theorem 1

We employ a *hybrid argument*. In particular, we define first a set of probability values $\{p'_i\}_{i \in [n]}$ and a corresponding set of independent Bernoulli random variables $\{Z_i\}_{i \in [n]}$ with expectations $\mathcal{E}[Z_i] = p'_i$, for all $i \in [n]$, such that

$$\left\| \sum_i X_i ; \sum_i Z_i \right\| = O(1/k), \tag{3}$$

and, moreover, for all $j = 1, \dots, n$, $\left\| \sum_{i \neq j} X_i ; \sum_{i \neq j} Z_i \right\| = O(1/k)$. (4)

The set of probability values $\{p'_i\}_{i \in [n]}$ does not necessarily satisfy Property 2(b)i or 2(b)ii in the statement of Theorem 1, but will allow us to define the set of probabilities $\{q_i\}_{i \in [n]}$ which *does* satisfy Property 2(b)i or 2(b)ii and, moreover,

$$\left\| \sum_i Z_i ; \sum_i Y_i \right\| = O(1/k), \tag{5}$$

and, for all $j = 1, \dots, n$, $\left\| \sum_{i \neq j} Z_i ; \sum_{i \neq j} Y_i \right\| = O(1/k)$. (6)

By the triangle inequality, (3) and (5) imply (1), and (4) and (6) imply (2). Let us call *Stage 1* the process of determining the p'_i 's and *Stage 2* the process of determining the q_i 's. The two stages are described briefly below, and in detail in Sections 5 and 6 respectively.

Stage 1: The goal of this stage is to eliminate any probability value p_i falling in the set $\mathcal{T}_k := (0, \frac{1}{k}) \cup (1 - \frac{1}{k}, 1)$, that is, any p_i that is either too small, but non-zero, or too large, but not one. To remove the small p_i 's we round some of them to 0 and some of them to $1/k$ in such a way that their sum changes by at most $1/k$. Similarly, we round some of the large p_i 's to $1 - 1/k$ and some of them to 1, so that their sum changes by at most $1/k$. Finally, we leave the p_i 's falling outside \mathcal{T}_k unchanged. Our work from [15] implies that the set of probability values $\{p'_i\}_i$ thus defined satisfies (3) and (4). See details in Section 5.

Stage 2: The definition of the set $\{q_i\}_i$ depends on the number m of p'_i 's which are different than 0 and 1. The case $m \geq k^3$ corresponds to the Case 2(b)i in the statement of Theorem 1, and the case $m < k^3$ corresponds to the Case 2(b)ii. In both cases we set $q_i = p'_i$, if $p'_i \in \{0, 1\}$; and here is how we round the p'_i 's from the index set $\mathcal{M} := \{i \mid p'_i \notin \{0, 1\}\}$:

- Case $m \geq k^3$: Using results from the literature on Stein's method, we show that the sum of Bernoulli random variables with expectations p'_i , $i \in \mathcal{M}$, can be approximated by a Binomial distribution $B(m', q)$, where $m' \leq m$ and q is an integer multiple of $\frac{1}{kn}$. In particular, we show that an appropriate choice of m' and q implies (5) and (6), if we set m' of the q_i 's from the index set \mathcal{M} equal to q and the remaining equal to 0.
- Case $m < k^3$: The Binomial approximation may not be tight enough for small values of m . To remedy this, we follow our rounding scheme from [15]. That is, we delicately round the p'_i 's to nearby multiples of $\frac{1}{k^2}$ so that their sum $\sum_{i \in \mathcal{M}} p'_i$ is approximated to within $1/k^2$ by the sum $\sum_{i \in \mathcal{M}} q_i$. Our results from [15] imply that (5) and (6) hold in this case.

5 Details of Stage 1

We describe the definition of the set $\{p'_i\}_i$. For concreteness, let

$$\mathcal{L} := \{i \mid i \in [n] \wedge p_i \in (0, 1/k)\} \quad \text{and} \quad \mathcal{H} := \{i \mid i \in [n] \wedge p_i \in (1 - 1/k, 1)\}.$$

We set $p'_i = p_i$, for all $i \in [n] \setminus \mathcal{L} \cup \mathcal{H}$; that is, we leave the probabilities p_i falling outside the set \mathcal{T}_k unchanged. It follows that

$$\left\| \sum_{i \in [n] \setminus \mathcal{L} \cup \mathcal{H}} X_i ; \sum_{i \in [n] \setminus \mathcal{L} \cup \mathcal{H}} Z_i \right\| = 0, \tag{7}$$

and, for all $j \in [n] \setminus \mathcal{L} \cup \mathcal{H}$,

$$\left\| \sum_{i \in [n] \setminus (\mathcal{L} \cup \mathcal{H}) \setminus \{j\}} X_i ; \sum_{i \in [n] \setminus (\mathcal{L} \cup \mathcal{H}) \setminus \{j\}} Z_i \right\| = 0. \tag{8}$$

To round the probabilities p_i , $i \in \mathcal{L}$, we use the following procedure:

1. Let $S_{\mathcal{L}} := \sum_{i \in \mathcal{L}} p_i$; $m = \lfloor \frac{S_{\mathcal{L}}}{1/k} \rfloor$.
2. Let $\mathcal{L}' \subseteq \mathcal{L}$ be an arbitrary subset of \mathcal{L} with cardinality $|\mathcal{L}'| = m$.
3. Set $p'_i = \frac{1}{k}$, for all $i \in \mathcal{L}'$, and $p'_i = 0$, for all $i \in \mathcal{L} \setminus \mathcal{L}'$.

An application of Lemma 3.9 from [15] with $\alpha = 1$ implies immediately that

$$\left\| \sum_{i \in \mathcal{L}} X_i ; \sum_{i \in \mathcal{L}} Z_i \right\| \leq \frac{3}{k}, \tag{9}$$

and, for all $j \in \mathcal{L}$,

$$\left\| \sum_{i \in \mathcal{L} \setminus \{j\}} X_i ; \sum_{i \in \mathcal{L} \setminus \{j\}} Z_i \right\| \leq \frac{6}{k}, \tag{10}$$

where we used that $\left| \sum_{i \in \mathcal{L}} p_i - \sum_{i \in \mathcal{L}} p'_i \right| \leq \frac{1}{k}$ and, for all $j \in \mathcal{L}$, $\left| \sum_{i \in \mathcal{L} \setminus \{j\}} p_i - \sum_{i \in \mathcal{L} \setminus \{j\}} p'_i \right| \leq \frac{2}{k}$.

We follow a similar rounding scheme for the probabilities p_i , $i \in \mathcal{H}$; that is, we round some to $1 - 1/k$ and some to 1 in such a way that their sum is preserved to within $1/k$. As a result, we get (to see this, repeat the argument we employed above to the variables $1 - X_i$ and $1 - Z_i$, $i \in \mathcal{H}$)

$$\left\| \sum_{i \in \mathcal{H}} X_i ; \sum_{i \in \mathcal{H}} Z_i \right\| \leq \frac{3}{k}, \tag{11}$$

and, for all $j \in \mathcal{H}$,

$$\left\| \sum_{i \in \mathcal{H} \setminus \{j\}} X_i ; \sum_{i \in \mathcal{H} \setminus \{j\}} Z_i \right\| \leq \frac{6}{k}. \tag{12}$$

Using (7), (8), (9), (10), (11), (12) and the coupling lemma we get (3) and (4).

6 Details of Stage 2

Recall that $\mathcal{M} := \{i \mid p'_i \notin \{0, 1\}\}$ and $m := |\mathcal{M}|$. The definition of the probability values $\{q_i\}_i$ will depend on whether $m \geq k^3$ or $m < k^3$. In particular, the case $m \geq k^3$ will correspond to the Case 2(b)i in the statement of Theorem 1, and the case $m < k^3$ will correspond to the Case 2(b)ii. In both cases we set $q_i = p'_i$, for all $i \in [n] \setminus \mathcal{M}$. It follows that

$$\left\| \sum_{i \in [n] \setminus \mathcal{M}} Z_i ; \sum_{i \in [n] \setminus \mathcal{M}} Y_i \right\| = 0, \tag{13}$$

and, for all $j \in [n] \setminus \mathcal{M}$,

$$\left\| \sum_{i \in [n] \setminus \mathcal{M} \setminus \{j\}} Z_i ; \sum_{i \in [n] \setminus \mathcal{M} \setminus \{j\}} Y_i \right\| = 0. \tag{14}$$

6.1 The Case $m \geq k^3$

We show that the random variable $\sum_{i \in \mathcal{M}} Z_i$ is within total variation distance $O(1/k)$ from a Binomial distribution $B(m', q)$ with

$$m' := \left\lceil \frac{\left(\sum_{i \in \mathcal{M}} p'_i\right)^2}{\sum_{i \in \mathcal{M}} p_i'^2} \right\rceil \quad \text{and} \quad q := \frac{\ell^*}{kn},$$

where $\ell^* \in \{0, \dots, kn\}$ satisfies $\frac{\sum_{i \in \mathcal{M}} p'_i}{m'} \in [\frac{\ell^*}{kn}, \frac{\ell^*+1}{kn}]$.

In particular, let us choose an arbitrary subset $\mathcal{M}' \subseteq \mathcal{M}$ with cardinality m' (note that $m' \leq m$ by the Cauchy-Schwarz inequality) and let us set $q_i = q$, for all $i \in \mathcal{M}'$, and $q_i = 0$, for all $i \in \mathcal{M} \setminus \mathcal{M}'$. We shall now compare the distributions of the random variables $\sum_{i \in \mathcal{M}} Z_i$ and $\sum_{i \in \mathcal{M}} Y_i$. For this let us set

$$\begin{aligned} \mu &:= \mathcal{E} \left[\sum_{i \in \mathcal{M}} Z_i \right] \quad \text{and} \quad \mu' := \mathcal{E} \left[\sum_{i \in \mathcal{M}} Y_i \right], \\ \sigma^2 &:= \text{Var} \left[\sum_{i \in \mathcal{M}} Z_i \right] \quad \text{and} \quad \sigma'^2 := \text{Var} \left[\sum_{i \in \mathcal{M}} Y_i \right]. \end{aligned}$$

The following lemma compares the values $\mu, \mu', \sigma, \sigma'$.

Lemma 1. *The following hold*

$$|\mu - \mu'| \leq \frac{1}{k}, \tag{15}$$

$$|\sigma'^2 - \sigma^2| \leq 1 + \frac{3}{k}, \tag{16}$$

$$\mu \geq k^2, \tag{17}$$

$$\sigma^2 \geq k^2 \left(1 - \frac{1}{k}\right). \tag{18}$$

The proof of Lemma 1 is given in the full version [10]. To compare $\sum_{i \in \mathcal{M}} Z_i$ and $\sum_{i \in \mathcal{M}} Y_i$ we approximate both by Translated Poisson distributions. To do this, we make use of the following theorem, due to Röllin [27].

Theorem 3 ([27]). *Let J_1, \dots, J_n be a sequence of independent random indicators with $\mathcal{E}[J_i] = p_i$. Then*

$$\left\| \sum_{i=1}^n J_i ; TP(\mu, \sigma^2) \right\| \leq \frac{\sqrt{\sum_{i=1}^n p_i^3(1-p_i)} + 2}{\sum_{i=1}^n p_i(1-p_i)},$$

where $\mu = \sum_{i=1}^n p_i$ and $\sigma^2 = \sum_{i=1}^n p_i(1-p_i)$.

Theorem 3 implies that

$$\begin{aligned} \left\| \sum_{i \in \mathcal{M}} Z_i ; TP(\mu, \sigma^2) \right\| &\leq \frac{\sqrt{\sum_{i \in \mathcal{M}} p_i'^3(1-p'_i)} + 2}{\sum_{i \in \mathcal{M}} p_i'(1-p'_i)} \leq \frac{\sqrt{\sum_{i \in \mathcal{M}} p_i'(1-p'_i)} + 2}{\sum_{i \in \mathcal{M}} p_i'(1-p'_i)} \\ &\leq \frac{1}{\sqrt{\sum_{i \in \mathcal{M}} p_i'(1-p'_i)}} + \frac{2}{\sum_{i \in \mathcal{M}} p_i'(1-p'_i)} = \frac{1}{\sigma} + \frac{2}{\sigma^2} \\ &\leq \frac{1}{k\sqrt{1-1/k}} + \frac{2}{k^2(1-\frac{1}{k})} = O(1/k). \quad (\text{using (18)}) \end{aligned}$$

Similarly,

$$\begin{aligned} \left\| \sum_{i \in \mathcal{M}} Y_i ; TP(\mu', \sigma'^2) \right\| &\leq \frac{1}{\sigma'} + \frac{2}{\sigma'^2} \\ &\leq \frac{1}{k \sqrt{1 - \frac{1}{k} - \frac{1}{k^2} - \frac{3}{k^3}}} + \frac{2}{k^2 \left(1 - \frac{1}{k} - \frac{1}{k^2} - \frac{3}{k^3}\right)} = O(1/k). \end{aligned} \quad (\text{using (I6), (I8)})$$

By the triangle inequality we then have that

$$\begin{aligned} \left\| \sum_{i \in \mathcal{M}} Z_i ; \sum_{i \in \mathcal{M}} Y_i \right\| &\leq \left\| \sum_{i \in \mathcal{M}} Z_i ; TP(\mu, \sigma^2) \right\| + \left\| \sum_{i \in \mathcal{M}} Y_i ; TP(\mu', \sigma'^2) \right\| + \|TP(\mu, \sigma^2) ; TP(\mu', \sigma'^2)\| \\ &= O(1/k) + \|TP(\mu, \sigma^2) ; TP(\mu', \sigma'^2)\|. \end{aligned} \quad (19)$$

It remains to bound the total variation distance between the two Translated Poisson distributions. We make use of the following lemma.

Lemma 2 (I2). *Let $\mu_1, \mu_2 \in \mathbb{R}$ and $\sigma_1^2, \sigma_2^2 \in \mathbb{R}_+ \setminus \{0\}$ be such that $\lfloor \mu_1 - \sigma_1^2 \rfloor \leq \lfloor \mu_2 - \sigma_2^2 \rfloor$. Then*

$$\|TP(\mu_1, \sigma_1^2) - TP(\mu_2, \sigma_2^2)\| \leq \frac{|\mu_1 - \mu_2|}{\sigma_1} + \frac{|\sigma_1^2 - \sigma_2^2| + 1}{\sigma_1^2}.$$

Lemma I2 implies

$$\begin{aligned} \|TP(\mu, \sigma^2) ; TP(\mu', \sigma'^2)\| &\leq \frac{|\mu - \mu'|}{\min(\sigma, \sigma')} + \frac{|\sigma^2 - \sigma'^2| + 1}{\min(\sigma^2, \sigma'^2)} \\ &\leq \frac{1/k}{k \sqrt{1 - \frac{1}{k} - \frac{1}{k^2} - \frac{3}{k^3}}} + \frac{2 + 3/k}{k^2 \left(1 - \frac{1}{k} - \frac{1}{k^2} - \frac{3}{k^3}\right)} \quad (\text{using Lemma I1}) \\ &= O(1/k^2). \end{aligned} \quad (20)$$

Using (I9) and (20) we get

$$\left\| \sum_{i \in \mathcal{M}} Z_i ; \sum_{i \in \mathcal{M}} Y_i \right\| = O(1/k). \quad (21)$$

We claim that the following is also true (see proof in the full version I10).

Lemma 3. *For all $j \in \mathcal{M}$:*

$$\left\| \sum_{i \in \mathcal{M} \setminus \{j\}} Z_i ; \sum_{i \in \mathcal{M} \setminus \{j\}} Y_i \right\| = O(1/k). \quad (22)$$

6.2 The Case $m < k^3$

Theorem 3.1 of [15] implies that there exists a set of probability values $\{q_i\}_{i \in \mathcal{M}}$, such that

– q_i is an integer multiple of $\frac{1}{k^2}$, for all $i \in \mathcal{M}$;

$$- \left\| \sum_{i \in \mathcal{M}} Z_i ; \sum_{i \in \mathcal{M}} Y_i \right\| = O(1/k); \tag{23}$$

$$- \text{ and, for all } j \in \mathcal{M}, \left\| \sum_{i \in \mathcal{M} \setminus \{j\}} Z_i ; \sum_{i \in \mathcal{M} \setminus \{j\}} Y_i \right\| = O(1/k). \tag{24}$$

6.3 Concluding Stage 2

In the case $m \geq k^3$ considered in Section 6.1, a set of probability values $\{q_i\}_i$ was defined which satisfied Property 2(b)i in the statement of Theorem 1. In the case $m < k^3$ considered in Section 6.2, the resulting set $\{q_i\}_i$ satisfied Property 2(b)ii. Moreover, in both cases, the following were satisfied

$$\left\| \sum_{i \in \mathcal{M}} Z_i ; \sum_{i \in \mathcal{M}} Y_i \right\| = O(1/k) \tag{25}$$

$$\left\| \sum_{i \in [n] \setminus \mathcal{M}} Z_i ; \sum_{i \in [n] \setminus \mathcal{M}} Y_i \right\| = 0, \tag{26}$$

and,

$$\text{for all } j \in \mathcal{M}, \left\| \sum_{i \in \mathcal{M} \setminus \{j\}} Z_i ; \sum_{i \in \mathcal{M} \setminus \{j\}} Y_i \right\| = O(1/k), \tag{27}$$

$$\text{for all } j \in [n] \setminus \mathcal{M}, \left\| \sum_{i \in [n] \setminus \mathcal{M} \setminus \{j\}} Z_i ; \sum_{i \in [n] \setminus \mathcal{M} \setminus \{j\}} Y_i \right\| = 0. \tag{28}$$

Using (25), (26), (27), (28) and the coupling lemma, we get (5) and (6).

Acknowledgment. We thank Christos Papadimitriou for helpful discussions.

References

1. Barbour, A.D., Holst, L., Janson, S.: Poisson Approximation. Oxford University Press, New York (1992)
2. Barbour, A.D., Lindvall, T.: Translated Poisson Approximation for Markov Chains. Journal of Theoretical Probability 19(3) (July 2006)
3. Blonski, M.: Anonymous Games with Binary Actions. Games and Economic Behavior 28(2), 171–180 (1999)
4. Blonski, M.: The women of Cairo: Equilibria in large anonymous games. Journal of Mathematical Economics 41(3), 253–264 (2005)
5. Bosse, H., Byrka, J., Markakis, E.: New Algorithms for Approximate Nash Equilibria in Bimatrix Games. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 17–29. Springer, Heidelberg (2007)

6. Chen, X., Deng, X.: 3-NASH is PPAD-Complete. *Electronic Colloquium in Computational Complexity*, TR05-134 (2005)
7. Chen, X., Deng, X.: Settling the Complexity of Two-Player Nash Equilibrium. In: *FOCS* (2006)
8. Chien, S., Sinclair, A.: Convergence to Approximate Nash Equilibria in Congestion Games. In: *SODA* (2007)
9. Dantzig, G.B.: *Linear Programming and Extensions*. Princeton University Press, Princeton (1963)
10. Daskalakis, C.: An Efficient PTAS for Two-Strategy Anonymous Games. *ArXiv Report* (2008)
11. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The Complexity of Computing a Nash Equilibrium. In: *STOC* (2006)
12. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: A Note on Approximate Nash Equilibria. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) *WINE 2006*. LNCS, vol. 4286, pp. 297–306. Springer, Heidelberg (2006)
13. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: Progress in Approximate Nash Equilibria. In: *EC* (2007)
14. Daskalakis, C., Papadimitriou, C.H.: Three-Player Games Are Hard. *Electronic Colloquium in Computational Complexity*, TR05-139 (2005)
15. Daskalakis, C., Papadimitriou, C.H.: Computing Equilibria in Anonymous Games. In: *FOCS* (2007)
16. Daskalakis, C., Papadimitriou, C.H.: Discretized Multinomial Distributions and Nash Equilibria in Anonymous Games. In: *FOCS* (2008)
17. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The Complexity of Pure Nash Equilibria. In: *STOC* (2004)
18. Feder, T., Nazerzadeh, H., Saberi, A.: Approximating Nash Equilibria Using Small-Support Strategies. In: *EC* (2007)
19. Goldberg, P.W., Papadimitriou, C.H.: Reducibility Among Equilibrium Problems. In: *STOC* (2006)
20. Kalai, E.: Partially-Specified Large Games. In: Deng, X., Ye, Y. (eds.) *WINE 2005*. LNCS, vol. 3828, pp. 3–13. Springer, Heidelberg (2005)
21. Kannan, R., Theobald, T.: Games of Fixed Rank: A Hierarchy of Bimatrix Games. In: *SODA* (2007)
22. Kontogiannis, S.C., Panagopoulou, P.N., Spirakis, P.G.: Polynomial Algorithms for Approximating Nash Equilibria of Bimatrix Games. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) *WINE 2006*. LNCS, vol. 4286, pp. 286–296. Springer, Heidelberg (2006)
23. Lipton, R., Markakis, E., Mehta, A.: Playing Large Games Using Simple Strategies. In: *EC* (2003)
24. Milchtaich, I.: Congestion Games with Player-Specific Pay off Functions. *Games and Economic Behavior* 13, 111–124
25. von Neumann, J., Morgenstern, O.: *Theory of Games and Economic Behavior*. Princeton University Press, Princeton (1944)
26. Papadimitriou, C.H., Roughgarden, T.: Computing Equilibria in Multi-Player Games. In: *SODA* (2005)
27. Röllin, A.: Translated Poisson Approximation Using Exchangeable Pair Couplings. *ArXiv Report* (2006)
28. Tsaknakis, H., Spirakis, P.G.: An Optimization Approach for Approximate Nash Equilibria. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 42–56. Springer, Heidelberg (2007)

Equilibria of Graphical Games with Symmetries (Extended Abstract)

Felix Brandt¹, Felix Fischer¹, and Markus Holzer²

¹ Institut für Informatik, Universität München, 80538 München, Germany
{brandtf,fischerf}@tcs.ifi.lmu.de

² Institut für Informatik, Technische Universität München, 85748 Garching, Germany
holzer@in.tum.de

Abstract. We study graphical games where the payoff function of each player satisfies one of four types of symmetry in the actions of his neighbors. We establish that deciding the existence of a pure Nash equilibrium is NP-hard in general for all four types. Using a characterization of games with pure equilibria in terms of even cycles in the neighborhood graph, as well as a connection to a generalized satisfiability problem, we identify tractable subclasses of the games satisfying the most restrictive type of symmetry. Hardness for a different subclass is obtained via a satisfiability problem that remains NP-hard in the presence of a matching, a result that may be of independent interest. Finally, games with symmetries of two of the four types are shown to possess a symmetric *mixed* equilibrium which can be computed in polynomial time. We thus obtain a class of games where the pure equilibrium problem is computationally harder than the mixed equilibrium problem, unless $P=NP$.

1 Introduction

The idea underlying *graphical games* [10] is that in games with a large number of players, the payoff of any particular player will often depend only on the actions of a small number of other players in a local neighborhood. More formally, a graphical game is given by a (directed or undirected) graph on the set of players of a normal-form game, such that the payoff of each player depends only on the actions of his neighbors in this graph. If neighborhoods are bounded, graphical games can be represented using space polynomial in the number of players. *Symmetric games* constitute another natural and well-studied class of games, characterized by the fact that players can not, or need not, distinguish between other players. In this paper, we consider graphical games where the payoff function of each player is symmetric in the actions of his neighbors. For instance, consider a setting where each player is faced with the decision of producing one of two types of complementary goods within a regional neighborhood. Players are not only producers but also consumers and thus happier when both products are available within their neighborhood. We will see in Section 3.3 that deciding the existence of a pure Nash equilibrium, *i.e.*, a profile of mutual best responses, in such a setting is highly nontrivial yet computationally tractable.

Related Work: The computational problem of finding Nash equilibria in graphical games with degree bounded by 3 has recently been shown equivalent to the same problem for general n -player games with $n \geq 4$ [8], and thus complete for the complexity class PPAD [4]. It is not surprising that the structure of the neighborhood graph greatly influences the complexity of the equilibrium problem. PPAD-hardness holds even if the underlying graph has constant pathwidth, but becomes tractable for graphs of degree 2, *i.e.*, for *paths* [5]. All known algorithms for the more general case of *trees* have exponential worst-case running time even on trees with bounded degree and pathwidth 2, but equilibria satisfying various fairness criteria can be computed in polynomial time if additionally there are only two actions per player and the best response policy, a data structure representing *all* Nash equilibria of a game, has polynomial size [6].

A different line of research has investigated the problem of deciding the existence of *pure* Nash equilibria, *i.e.*, equilibria where the support of each strategy contains only a single action. Unlike Nash equilibria in *mixed* strategies, *i.e.*, probabilistic combinations of actions, pure equilibria are not guaranteed to exist. If they exist, however, pure equilibria have two distinct advantages over mixed ones. For one, requiring randomization in order to reach a stable outcome has been criticized on various grounds. In multi-player games, where action probabilities in equilibrium can be irrational numbers, randomization is particularly questionable. Secondly, pure equilibria as computational objects are usually much smaller in size than mixed ones. The pure equilibrium problem is NP-complete for graphical games on directed graphs with outdegree bounded by 2 and with only two actions for each player and two different payoffs, and tractable for graphs with bounded treewidth [9, 7].

Brandt et al. [1] analyze four classes of symmetric games, and show that the pure equilibrium problem is tractable if the number of actions is a constant, and complete for NP or PLS, respectively, if the number of actions grows in the number of players. One of the classes, in which all players have identical payoff functions, is guaranteed to possess a symmetric equilibrium, *i.e.*, one where all players play the same strategy. This equilibrium is not necessarily pure, but can be found efficiently if the number of actions is not too large compared to the number of players [14]. A larger class, allowing different payoff functions for different players, admits an approximation by a factor depending on the Lipschitz constant of the payoff function and on the square of the number of actions, and a polynomial-time approximation scheme for the case of two actions [3].

These results fuel hope that tractability results can be obtained for larger classes of games satisfying some kind of symmetry. In this regard, Daskalakis and Papadimitriou [2] consider games on a d -dimensional undirected torus or grid with payoff functions that are identical for all players and symmetric in the actions of the players in the neighborhood, a condition that will be called *symmetry* in this paper. In particular, they show that deciding the existence of a pure Nash equilibrium in such a game is NL-complete when $d = 1$ and NEXP-complete for $d \geq 2$. In this paper, we investigate the pure equilibrium problem in graphical games satisfying the kinds of symmetries considered by Brandt et al.

[1]. Our work can thus be seen as a refinement of the work of Gottlob et al. [9] and of Daskalakis and Papadimitriou [2].

Paper Structure and Results: We begin by introducing the necessary game-theoretic concepts. In Section 3, we then investigate the computational complexity of the pure equilibrium problem in graphical games satisfying four different types of symmetries. The question of tractable classes of graphical games is answered mostly in the negative. For three of the four symmetry classes, deciding the existence of a pure equilibrium is NP-hard already for the case of two actions, two payoffs, and neighborhoods of size two. Assuming the most restricted type of symmetry, the problem becomes NP-hard when there are three different payoffs, or neighborhoods of size four. On the other hand, we use interesting connections of the latter class to even cycles in directed graphs and to generalized satisfiability to identify tractable classes of games. As a corollary, we exhibit a satisfiability problem that remains NP-hard in the presence of a matching. We present this result, which may be of independent interest, in Section 4. Finally, in Section 5, we show that mixed equilibria in games with two of the above symmetry types can be found in polynomial time if the number of actions grows only slowly in the neighborhood size. Quite interestingly, there exists a class of games where deciding the existence of a pure equilibrium problem is likely to be harder than finding a mixed equilibrium. We assume the reader to be familiar with the complexity classes P, NP, and #P, and the notion of polynomial-time reducibility [e.g., 13].

2 Preliminaries

An accepted way to model situations of strategic interaction is by means of a normal-form game [e.g., 11].

Definition 1 (normal-form game). A game in normal-form is a tuple $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ where N is a set of players and for each player $i \in N$, A_i is a nonempty set of actions available to i , and $p_i : (\times_{i \in N} A_i) \rightarrow \mathbb{R}$ is a function mapping each action profile of the game, i.e., combination of actions, to a real-valued payoff for i .

A vector $s \in \times_{i \in N} A_i$ of actions is also called a profile of *pure strategies*. This concept can be generalized to (*mixed*) *strategy profiles* $s \in S = \times_{i \in N} S_i$, by letting players randomize over their actions. Here, we have S_i denote the set of probability distributions over player i 's actions, or (*mixed*) *strategies* available to player i . We further write $n = |N|$ for the number of players in a game, s_i for the i th strategy in profile s , and s_C for the vector of strategies for all players in a subset $C \subseteq N$.

A graphical game is given by a graph on the set of players, such that the payoff of a player only depends only on his own action, and on the actions of his neighbors in the graph. In the following definition, the underlying graph is directed, corresponding to a neighborhood relation that is not necessarily symmetric.

Definition 2 (graphical game). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game, $\nu : N \rightarrow 2^N$. Γ is a graphical game with neighborhood ν if for all $i \in N$ and $s, s' \in A^N$, $p_i(s) = p_i(s')$ whenever $s_{\hat{\nu}(i)} = s'_{\hat{\nu}(i)}$, where $\hat{\nu}(i) = \nu(i) \cup \{i\}$.

A game Γ has k -bounded neighborhoods if there exists $\nu : N \rightarrow 2^N$ such that Γ is a graphical game with neighborhood ν and for all $i \in N$, $|\nu(i)| \leq k$.

We assume throughout the paper that graphical games are encoded by listing the payoffs of each player as a function of the actions of his neighbors.

Symmetry as a property of a mathematical object refers to its invariance under a certain type of transformation. Symmetries of games usually mean invariance of the payoffs under automorphisms of the set of action profiles induced by some group of permutations of the set of players. *Anonymous* games as considered by Daskalakis and Papadimitriou [3], for example, require the set of available actions to be the same for all players, and the payoff of a particular player to remain the same under any permutation of the elements of an action profile. This imposes constraints on individual payoff functions only and can therefore directly be applied to graphical games as well. In general, however, it does not make much sense from a computational point of view to consider symmetries of the payoff functions without requiring the neighborhood graph to be “symmetric” in an appropriate way as well. Consider, for example, the class of all graphical games whose payoff functions are invariant under automorphisms in the automorphism group of the neighborhood graph. While this class of games is very natural, it does not impose meaningful computational restrictions. Indeed, it is not too hard to see that any graphical game can be encoded by a game in the above class that has a neighborhood graph with a trivial automorphism group. Hardness results for both pure and mixed equilibria thus carry over immediately.

In general, different types of restrictions on the neighborhood structure will be required for different kinds of symmetries of the payoff functions. In this paper, we take a slightly different approach. We consider properties found in anonymous and symmetric games, and study graphical games that possess these properties. A characteristic feature of symmetries in games is the inability to distinguish between other players. Following Daskalakis and Papadimitriou [3], the most general class of games with this property will be called *anonymous*. Four different classes of games are obtained by considering two additional characteristics: *identical payoff functions* for all players¹ and the ability to *distinguish oneself* from the other players. The games obtained by adding the former property will be called *symmetric*, and presence of the latter will be indicated by the prefix “*self*.” For the obvious reason, we will henceforth talk about games where the set of actions is the same for all players and write $A = A_1 = \dots = A_n$ and $k = |A|$, respectively, to denote this set and its cardinality.

An intuitive way to describe anonymous games is in terms of equivalence classes of the aforementioned automorphism group, using a notion introduced by Parikh [15] in the context of context-free languages. Given a set A of actions, the

¹ We assume the set of actions and the payoff function to be the same for *all* players rather than just those with intersecting neighborhoods. This is only done for ease of exposition.

commutative image of an action profile $s \in A^N$ is given by $\#(s) = (\#(a, s))_{a \in A}$ where $\#(a, s) = |\{i \in N \mid s_i = a\}|$. In other words, $\#(a, s)$ denotes the number of players playing action a in action profile s , and $\#(s)$ is the vector of these numbers for all the different actions. This definition naturally extends to action profiles for subsets of players.

Definition 3 (symmetries). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a graphical game with neighborhood ν , A a set of actions such that for all $i \in N$, $A_i = A$. Γ is called

- anonymous if for all $i \in N$ and all $s, s' \in A^N$, $p_i(s) = p_i(s')$ whenever $s_i = s'_i$ and for all $a \in A$, $\#(a, s_{\nu(i)}) = \#(a, s'_{\nu(i)})$;
- symmetric if for all $i, j \in N$ and all $s, s' \in A^N$, $|\nu(i)| = |\nu(j)|$ and $p_i(s) = p_j(s')$ whenever $s_i = s'_j$ and for all $a \in A$, $\#(a, s_{\nu(i)}) = \#(a, s'_{\nu(j)})$;
- self-anonymous if for all $i \in N$ and all $s, s' \in A^N$, $p_i(s) = p_i(s')$ whenever for all $a \in A$, $\#(a, s_{\hat{\nu}(i)}) = \#(a, s'_{\hat{\nu}(i)})$; and
- self-symmetric if for all $i, j \in N$ and all $s, s' \in A^N$, $|\nu(i)| = |\nu(j)|$ and $p_i(s) = p_j(s')$ whenever for all $a \in A$, $\#(a, s_{\hat{\nu}(i)}) = \#(a, s'_{\hat{\nu}(j)})$.

It should be noted that a graphical game in one of the four classes does not necessarily belong to the corresponding class of general normal-form games as defined by Brandt et al. [1], unless the neighborhood of every player contains all other players. When talking about self-anonymous and self-symmetric games with two actions, we write $\mathbf{p}_i(m) = p_i(s)$ where $\#(1, s_{\hat{\nu}(i)}) = m$ for the payoff of player i when m players in his neighborhood, including i himself, play action 1, and $\mathbf{p}_i = (\mathbf{p}_i(m))_{0 \leq m \leq |\hat{\nu}(i)|}$ for the vector of payoffs for the possible values of m .

One of the best-known solution concepts for strategic games is Nash equilibrium [12]. In Nash equilibrium, no player is able to increase his payoff by unilaterally changing his strategy.

Definition 4 (Nash equilibrium). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game. A strategy profile $s \in S$ of Γ is called Nash equilibrium if for each player $i \in N$ and each strategy $s'_i \in S_i$, $p_i(s) \geq p_i((s_{N \setminus \{i\}}, s'_i))$. A Nash equilibrium is called pure if it is a pure strategy profile.

3 Complexity of the Pure Equilibrium Problem

For graphical games with neighborhoods of size one, symmetries do not impose any restrictions. The pure equilibrium problem for such games can be decided in polynomial time [e.g., 7]. On the other hand, the game used by Schoenebeck and Vadhan [18] to show NP-completeness of the pure equilibrium problem in general graphical games with neighborhoods of size two is anonymous. We thus have the following initial result.

Theorem 1 (Schoenebeck and Vadhan [18]). Deciding whether a graphical game has a pure Nash equilibrium is NP-complete, even if every player has only two neighbors, two actions, and two different payoffs, and when restricted to anonymous games.

3.1 Symmetry and Self-symmetry

We now turn to more restrictive kinds of symmetry. The following theorem concerns games where the utility functions of all players are identical. The proof of this theorem is similar to a construction used by Schoenebeck and Vadhan [18] where each gate of a Boolean circuit corresponds to a player in a graphical game, and two additional players play a game with or without a pure equilibrium, depending on the output of the circuit. The main difficulty is to model these two steps using only a single payoff function. The proof of the theorem is given in the full version of this paper.

Theorem 2. *Deciding whether a graphical game has a pure Nash equilibrium is NP-complete, even if every player has only two actions, and when restricted to symmetric games with two different payoffs or to self-symmetric games with three different payoffs.*

3.2 Self-anonymity and Two Different Payoffs

Since self-symmetric games form a subset of self-anonymous games, Theorem 2 also implies NP-hardness of the self-anonymous case. However, the result is not tight in that three different payoffs are required for hardness. A natural question is what happens for self-anonymous games with only two different payoffs. In this section we will prove a tight result for the most restricted version of self-anonymity, *i.e.*, the case with only two different payoff *functions*.

The problem with anonymity and the construction used in the proof of Theorem 2 is that two different payoffs are not enough to make a player care about his own action no matter which actions are played by his neighbors. With four different values for $\#(1, s_{\hat{v}(i)})$, there will either be an equilibrium where all players play the same action, or a situation where a player is indifferent between both of his actions. When we want to use games to compute a function, such indifference is clearly undesirable. The key idea that will enable us to prove the following theorem is to isolate pure equilibria that are themselves *symmetric* in the actions of a subset of the players, *i.e.*, equilibria in which these players all play the same action. To enforce that two particular players play the same action in every equilibrium, we will add two additional players, each of which observes the other as well as one of the original players. Depending on the actions of the original players, the new players will either play a game with a unique pure equilibrium, or a game that is prototypical both for self-anonymous games and for games without pure equilibria, namely Matching Pennies. We proceed with the statement of the theorem, a detailed proof is deferred to the full version of this paper.

Theorem 3. *Deciding whether a graphical game has a pure Nash equilibrium is NP-complete, even if every player has only two neighbors, two actions, and two different payoffs, and when restricted to self-anonymous games with two different payoff functions.*

3.3 Self-symmetry and Two Different Payoffs

Let us return to self-symmetric graphical games. Self-symmetric games as studied by Brandt et al. [11] always possess a pure Nash equilibrium due to the fact that they are common-payoff games. This is not the case for self-symmetric graphical games, even when there are only two different payoffs. In particular, there exists a seven-player game in the latter class that does not have a pure equilibrium, and in which each player has exactly two actions and two neighbors. It will be instructive to view a graphical game as a hypergraph, with each vertex corresponding to a player and each edge to the set of players in the neighborhood of one particular player *including* the player himself. Corresponding to the set of games with m -neighborhood is the set of $(m + 1)$ -uniform hypergraphs that possess a matching in the sense of Seymour [19], *i.e.*, a bijection from the set of vertices to the set of edges that maps every vertex to an edge containing it. Then, a self-symmetric game with two actions and payoffs $\mathbf{p}_i = (0, 1, 1, 0)$ for all $i \in N$ has a pure Nash equilibrium if and only if the corresponding hypergraph is vertex two-colorable. Given a two-coloring, every player observes either one or two players in his neighborhood, including himself, who play action 1, and thus obtains the maximum payoff of 1. If on the other hand there is no two-coloring, then there is at least one player for every action profile who plays the same action as all of his neighbors and can deviate to obtain a higher payoff. Figure 1 shows the neighborhood of a graphical game with seven players and two neighbors for each player. This graph induces the 3-uniform square hypergraph corresponding to the lines of the Fano plane, which in turn cannot be two-colored [*e.g.*, 19]. We leave it to the avid reader to verify that there is no game with the above properties and less than seven players.

An interesting property of the neighborhood graph on the left of Figure 1 is that it does not have any cycles of even length. We will begin our investigation of the pure equilibrium problem in self-symmetric games by generalizing this observation to games with arbitrary neighborhoods and $\mathbf{p}_i = (0, 1, 1, \dots, 1, 0)$ for all $i \in N$. The following lemma characterizes games with pure equilibria in the above subclass in terms of cycles in the neighborhood graph. Seymour [19] provides a similar characterization of the minimal uniform square hypergraphs that do not have a two-coloring.

Lemma 1. *Let Γ be a self-symmetric graphical game with two actions per player and payoffs p_i such that for all $i \in N$, $\mathbf{p}_i = (0, 1, 1, \dots, 1, 0)$. Then, Γ has a pure Nash equilibrium if and only if for all $i \in N$, there exists $j \in N$ reachable from i such that j lies on a cycle of even length.*

Proof. For the implication from left to right, assume that there exists a pure equilibrium, *i.e.*, a two-coloring $c : N \rightarrow \{0, 1\}$ of the neighborhood graph such that the neighborhood of every player contains some player playing action 0 and some player playing action 1. Now consider an arbitrary player $v_1 \in N$. Using the above property of c , we can construct a path $v_1, v_2, \dots, v_{|N|+1}$, $v_i \in N$, such that for all i , $1 \leq i \leq |N|$, $c(v_i) = 1 - c(v_{i+1})$. By the pigeonhole principle, there must exist i, j , $1 \leq i < j \leq |N| + 1$, such that $v_i = v_j$ and for all j' , $i < j' < j$, $v_{j'} \neq v_i$. Then, v_i, v_{i+1}, \dots, v_j is a cycle of even length reachable from v_1 .




Fig. 1. Neighborhood graph of a graphical game with seven players (left), corresponding to the three-uniform square hypergraph given by the lines of the Fano plane (right). A directed edge from vertex i to vertex j of the neighborhood graph denotes that $j \in \nu(i)$.

For the implication from right to left, let $N' \subseteq N$ be a set of players such that for every $i \in N$ there exists a directed path to some $j \in N'$, and such that N' induces a set of vertex-disjoint cycles of even length. We construct a two-coloring $c : N \rightarrow \{0, 1\}$, corresponding to an assignment of actions to players, as follows. First color the members of N' such that for all $i \in N'$ and $j \in \nu(i) \cap N'$, $c(i) = 1 - c(j)$. While there are uncolored vertices left, find $i, j \in N$ such that $j \in \nu(i)$, i is uncolored, and j is colored. Such a pair of vertices must always exist, since for every member of N there is a directed path to some member of N' , and thus to a vertex that has already been colored. Color i such that $c(i) = 1 - c(j)$. It is now easily verified that at any given time, and for all $i \in N$ that have already been colored, there exist $j, j' \in \hat{\nu}(i)$ with $c(j) = 0$ and $c(j') = 1$. If all vertices have been colored, then every neighborhood will contain at least one player playing action 0, and at least one player playing action 1. The corresponding action profile is a pure Nash equilibrium. \square

Thomassen [20] has shown that for every k , there exists a directed graph without even cycles where every vertex has outdegree k . Together with Lemma 1, this means that the pure equilibrium problem for the considered class of games is nontrivial.

Corollary 1. *For every $m \in \mathbb{N}$, $m > 0$, there exist self-symmetric graphical games Γ and Γ' with two actions where for all $i \in N$, $|\nu(i)| = m$ and $\mathbf{p}_i = (0, 1, 1, \dots, 1, 0)$, such that Γ has a pure Nash equilibrium and Γ' does not.*

We are now ready to identify several classes of graphical games where the existence of a pure equilibrium can be decided in polynomial time.

Theorem 4. *Let Γ be a self-symmetric graphical game with payoffs p_i . The pure equilibrium problem for Γ can be decided in polynomial time if one of the following properties holds:*

- (i) for all $i \in N$, $\mathbf{p}_i(0) \geq \mathbf{p}_i(1)$ or for all $i \in N$, $\mathbf{p}_i(|\hat{\nu}(i)|) \geq \mathbf{p}_i(|\hat{\nu}(i)| - 1)$;
- (ii) for all $i \in N$ and all m , $1 \leq m \leq |\nu(i)|$, $\mathbf{p}_i(m-1) > \mathbf{p}_i(m)$ and $\mathbf{p}_i(m+1) > \mathbf{p}_i(m)$, or $\mathbf{p}_i(m-1) < \mathbf{p}_i(m)$ and $\mathbf{p}_i(m+1) < \mathbf{p}_i(m)$;
- (iii) for all $i \in N$ and all m , $1 \leq m < |\nu(i)|$, $\mathbf{p}_i(m) = \mathbf{p}_i(m+1)$.

Proof. It is easy to see that a game Γ satisfying [\(i\)](#) possesses a pure equilibrium s in which $\#(0, s) = 0$ or $\#(1, s) = 1$.

For a game Γ satisfying [\(ii\)](#), we observe that in every equilibrium s , $p_i(s) = 1$ for all $i \in N$. The pure equilibrium problem for Γ thus corresponds to a variant of generalized satisfiability, with clauses induced by neighborhoods of Γ . The constraints associated with this particular variant require that the number of variables in each clause set to true is odd, and can be written as a system of linear equations over $GF(2)$. Tractability of the pure equilibrium problem for Γ then follows from Theorem 2.1 of Schaefer [\[17\]](#).

Finally, a game satisfying [\(iii\)](#) but not [\(i\)](#) can be transformed into a best response equivalent one that satisfies the conditions of Lemma [1](#). We further claim that we can check in polynomial time whether for every $i \in N$, there exists $j \in N$ on a cycle of even length and reachable from i . For a particular $i \in N$, this problem is equivalent to checking whether the subgraph induced by the vertices reachable from i contains an even cycle. The latter problem has long been open, but was recently shown to be solvable in polynomial time [\[16\]](#). \square

It is readily appreciated that every self-symmetric graphical game Γ with two different payoffs and neighborhoods of size two or three can be transformed into a game Γ' with the same set of players and the same neighborhoods, such that Γ and Γ' have the same set of pure equilibria and Γ' satisfies one of the conditions of Theorem [4](#). We thus have the following.

Corollary 2. *The problem of deciding whether a self-symmetric graphical game with two different payoffs and three-bounded neighborhood has a pure equilibrium is in P.*

3.4 Self-symmetry and Larger Neighborhoods

The question that remains is whether the pure equilibrium problem can be solved in polynomial time for all self-symmetric graphical games with two payoffs, or whether there is some bound on the neighborhood size where this problem again becomes hard. We will show in this section that the latter is true, and that the correct bound is indeed four, as suggested by Corollary [2](#).

We will essentially use the same tools as in Section [3.2](#) but will extract the necessary complexity from only a single payoff function. The additional insight required for this extraction will be that “constant” players, *i.e.*, players who play the same action in every pure equilibrium of a game, can be used to prune a larger payoff table and effectively obtain different payoff functions for smaller neighborhoods that can then be used to proceed with the original proof. Constructing such players will prove a rather difficult task in its own right. A detailed proof is again deferred to the full version of the paper.

Theorem 5. *Deciding whether a self-symmetric graphical game with two different payoffs has a pure Nash equilibrium is NP-complete, even if every player has exactly four neighbors.*

Observing that in the constructions used in the proofs of Theorems 2, 3, and 5 there is a one-to-one correspondence between satisfying assignments of a Boolean circuit and pure equilibria of a game, we have that counting the number of pure equilibria in the respective games is as hard as computing the permanent of a matrix.

Corollary 3. *For graphical games with neighborhoods of size two, counting the number of pure Nash equilibria is #P-hard, even when restricted to symmetric games with two different payoffs, to self-anonymous games with two different payoffs and two different payoff functions, or to self-symmetric games with three different payoffs. The same holds for self-symmetric graphical games with neighborhoods of size four and two different payoffs.*

4 Interlude: Generalized Satisfiability in the Presence of a Matching

The analysis at the end of the previous section allows us to derive a corollary that may be of independent interest. Schaefer [17] completely characterizes which variants of the generalized satisfiability problem are in P and which are NP-complete. Some of the variants become tractable if there exists a matching, *i.e.*, a bijection from variables to clauses that maps every variable to a clause it appears in. For not-all-equal 3SAT, this holds by equivalence with two-colorability of three-uniform hypergraphs and from the work of Robertson et al. [16]. On the other hand, the proof of Theorem 5 identifies a variant that is NP-complete and remains so in the presence of a matching. We thus have the following.

Corollary 4. *Generalized satisfiability is NP-complete, even if there exists a matching and all clauses have size five.*

We leave a complete characterization for future work. While the proof techniques developed in this paper will certainly be useful in this respect, it should be noted that the equivalence between generalized satisfiability and the pure equilibrium problem covered by Theorem 5 may fail to hold for instances of the latter where $p_i(s) = p_i(s') = 0$ for s, s' such that $\#(1, s_{\hat{v}(i)}) = \#(1, s'_{\hat{v}(i)}) + 1$. For example, it would not be possible to show hardness of one-in-three 3SAT [17] using the same approach.

5 Mixed Equilibria

Let us now briefly look at the problem of finding a mixed equilibrium. The following theorem states that this problem is tractable in symmetric graphical games if the number of actions grows slowly in the neighborhood size. The proof relies on the fact that such games always have a symmetric equilibrium and is given in the full version of the paper.

Theorem 6. *Let $\Gamma = (N, A^N, (p_i)_{i \in N})$ be a symmetric graphical game such that for all $i \in N$, $|A| = O(\log |\nu(i)| / \log \log |\nu(i)|)$. Then, a Nash equilibrium of Γ can be computed in polynomial time.*

This result applies in particular to the case where both the number of actions and the neighborhood size are bounded. Since the pure equilibrium problem in symmetric graphical games is NP-complete even in the case of two actions, we have identified a class of games where computing a mixed equilibrium is computationally easier than deciding the existence of a pure one, unless $P=NP$. A different class of games with the same property is implicit in Theorem 3.4 of Daskalakis and Papadimitriou [2]. On the other hand, existence of a symmetric equilibrium does not in general extend to games that are not anonymous or in which players have different payoff functions.

6 Open Problems

In this paper we have mainly considered neighborhoods of constant size. The construction used in the proof of Theorem 5 can be generalized to arbitrary neighborhoods of even size, but it is unclear what happens for odd-sized neighborhoods. The extreme case when the neighborhood of every player consists of all other players yields ordinary symmetric games, and it is known that the pure equilibrium problem in these games is in P when the number of actions is bounded [1]. It is an open problem at which neighborhood size the transition between membership in P and NP-hardness occurs. Another open question concerns the complexity of the mixed equilibrium problem in anonymous graphical games. A promising direction for proving hardness would be to make the construction of Goldberg and Papadimitriou [8] anonymous. Finally, as suggested in Section 4, it would be interesting to study the complexity of generalized satisfiability problems in the presence of matchings.

Acknowledgements. This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/3-2. We are grateful to Michael Tautschnig for providing us with a customized version of his ABSolver software.

References

1. Brandt, F., Fischer, F., Holzer, M.: Symmetries and the complexity of pure Nash equilibrium. *Journal of Computer and System Sciences* (to appear, 2008)
2. Daskalakis, C., Papadimitriou, C.H.: The complexity of games on highly regular graphs. In: Brodal, G.S., Leonardi, S. (eds.) *ESA 2005*. LNCS, vol. 3669, pp. 71–82. Springer, Heidelberg (2005)
3. Daskalakis, C., Papadimitriou, C.H.: Computing equilibria in anonymous games. In: *Proc. of 48th FOCS Symposium*, pp. 83–93. IEEE Press, Los Alamitos (2007)
4. Daskalakis, C., Goldberg, P., Papadimitriou, C.: The complexity of computing a Nash equilibrium. In: *Proc. of 38th STOC*, pp. 71–78. ACM Press, New York (2006)

5. Elkind, E., Goldberg, L.A., Goldberg, P.W.: Equilibria in graphical games on trees revisited. In: Proc. of 7th ACM-EC Conference, pp. 100–109. ACM Press, New York (2006)
6. Elkind, E., Goldberg, L.A., Goldberg, P.W.: Computing good Nash equilibria in graphical games. In: Proc. of 8th ACM-EC Conference, pp. 162–171. ACM Press, New York (2007)
7. Fischer, F., Holzer, M., Katzenbeisser, S.: The influence of neighbourhood and choice on the complexity of finding pure Nash equilibria. *Information Processing Letters* 99(6), 239–245 (2006)
8. Goldberg, P.W., Papadimitriou, C.H.: Reducibility among equilibrium problems. In: Proc. of 38th STOC, pp. 61–70. ACM Press, New York (2006)
9. Gottlob, G., Greco, G., Scarcello, F.: Pure Nash equilibria: Hard and easy games. *Journal of Artificial Intelligence Research* 24, 195–220 (2005)
10. Kearns, M.J., Littman, M.L., Singh, S.P.: Graphical models for game theory. In: Proc. of 17th UAI Conference, pp. 253–260. Morgan Kaufmann, San Francisco (2001)
11. Luce, R.D., Raiffa, H.: *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc., Chichester (1957)
12. Nash, J.F.: Non-cooperative games. *Annals of Mathematics* 54(2), 286–295 (1951)
13. Papadimitriou, C.H.: *Computational Complexity*. Addison-Wesley, Reading (1994)
14. Papadimitriou, C.H., Roughgarden, T.: Computing equilibria in multi-player games. In: Proc. of 16th SODA, pp. 82–91. SIAM, Philadelphia (2005)
15. Parikh, R.: On context-free languages. *Journal of the ACM* 13(4), 570–581 (1966)
16. Robertson, N., Seymour, P.D., Thomas, R.: Permanents, Pfaffian orientations, and even directed circuits. *Annals of Mathematics* 150, 929–975 (1999)
17. Schaefer, T.J.: The complexity of satisfiability problems. In: Proc. of 10th STOC, pp. 216–226. ACM Press, New York (1978)
18. Schoenebeck, G., Vadhan, S.: The computational complexity of Nash equilibria in concisely represented games. In: Proc. of 7th ACM-EC Conference, pp. 270–279. ACM Press, New York (2006)
19. Seymour, P.D.: On the two-colouring of hypergraphs. *The Quarterly Journal of Mathematics* 25, 303–312 (1974)
20. Thomassen, C.: Even cycles in directed graphs. *European Journal of Combinatorics* 6, 85–89 (1985)

Equilibrium Points in Fear of Correlated Threats^{*}

Spyros C. Kontogiannis^{1,2} and Paul G. Spirakis²

¹ Computer Science Dept., University of Ioannina, Greece
kontog@cs.uoi.gr

² R.A. Computer Technology Institute, N. Kazantzaki Str., Patras University
Campus, 26500 Patras, Greece
{kontog,spirakis}@cti.gr

Abstract. The present work considers the following computational problem: Given any finite game in normal form G and the corresponding infinitely repeated game G^∞ , determine in *polynomial time* (wrt the representation of G) a profile of strategies for the players in G^∞ that is an equilibrium point wrt the limit-of-means payoff. The problem has been solved for two players [10], based mainly on the *implementability* of the threats for this case. Nevertheless, [4] demonstrated that the traditional notion of threats is a computationally hard problem for games with at least 3 players (see also [8]). Our results are the following: (i) We propose an alternative notion of *correlated threats*, which is polynomial time computable (and therefore credible). Our correlated threats are also more severe than the traditional notion of threats, but not overwhelming for any individual player. (ii) When for the underlying game G there is a *correlated strategy* with payoff vector *strictly larger* than the correlated threats vector, we efficiently compute a polynomial-size (wrt the description of G) equilibrium point for G^∞ , for any *constant* number of players. (iii) Otherwise, we demonstrate the construction of an equilibrium point for an arbitrary number of players and up to 2 concurrently positive payoff coordinates in any payoff vector of G . This completely resolves the cases of 3 players, and provides a direction towards handling the cases of more than 3 players. It is mentioned that our construction is *not* a Nash equilibrium point, because the correlated threats we use are implemented via, not only full synchrony (as in [10]), but also coordination of the other players' actions. But this seems to be a fair trade-off between efficiency of the construction and players' coordination, in particular because it only affects the punishments (which are anticipated never to be used).

1 Introduction

Consider an arbitrary game in normal form, $G = \langle N, (S_p)_{p \in N}, (U_p : \times_{q \in N} S_q \mapsto \mathbb{R})_{p \in N} \rangle$. For any *mixed profile of strategies* $\mathbf{x} \in \times_{p \in N} \Delta(S_p)$, we consider the

^{*} This work has been partially supported by the Programmes of the European Union (i) ERC/IDEAS under contract number ERC-StG-2008-210743 (RIMACO), and (ii) ICT under contract number ICT-2008-215270 (FRONTS).

¹ With respect to.

expected payoffs of the players $U_p(\mathbf{x}) \equiv \mathbb{E}_{\mathbf{s} \sim \mathbf{x}}\{U_p(\mathbf{s})\}$, ie, the expected values of the utility functions U_p when the outcome \mathbf{s} of the game is randomly chosen according to the product distribution induced by \mathbf{x} . Similarly, for any *correlated strategy* $\sigma \in \Delta(\times_{p \in N} S_p)$, the expected payoff to player p is $U_p(\sigma) \equiv \mathbb{E}_{\mathbf{s} \sim \sigma}\{U_p(\mathbf{s})\}$. Given G , we may consider an infinitely repeated game G^∞ , in each round $t \geq 1$ of whose a new realization of G takes place, where the players have complete knowledge of the *history pattern* h^{t-1} of previous realizations of G . The *minmax value* for $p \in N$ determines the minimum payoff that p would accept in a realization of G , against uncoordinated strategies of the opponents, conditioning on his/her selfish behavior. This value, aka the **threat value** for p , is given by:

$$\theta_p(G) \equiv \min_{\mathbf{x}_{-p} \in \times_{q \neq p} \Delta(S_q)} \max_{\mathbf{x}_p \in \Delta(S_p)} U_p(\mathbf{x}_{-p}, \mathbf{x}_p) \tag{1}$$

The vector $\theta(G) = (\theta_p(G))_{p \in N}$, called the **threat point** of the game, is a key concept for the standard Folk Theorem arguments, as it represents the worst possible uncoordinated punishment that can be inflicted during a round t of G^∞ on a player, for deviating (in a previous round) from some agreed behavioral plan which assures a payoff vector within the individually rational region of G .

Although the Folk Theorem asserts that finding equilibria of G^∞ should be easier than finding equilibria of G itself, it was proved in [4] that even for a 3-player, win-lose stage game G , it is **NP**-hard to approximate the threat value of a player to within $\frac{3}{n^2}$ (n is the maximum number of actions for a player), let alone the computation of those strategies for the players that actually implement this threat against a given player p . In the same paper it was also proved that for any $k \geq 2$, finding a Nash equilibrium of an infinitely repeated $(k + 1)$ -player game is **PPAD**-complete, via a reduction from the well known **PPAD**-complete problem [5,6] of finding a Nash equilibrium in an arbitrary k -player normal form game. But then, the question is *how much credible a threat can be when it is not efficiently computable by any of the players?* On the other hand, [10] demonstrated how to efficiently compute Nash equilibrium points in an infinitely repeated 2-player game (wrt to limit-of-means criterion). The catch is that in this case both the threat point and the other player’s aggressive strategy are polynomial time computable. In this work we aim at tackling the intractability of the threat point, in order to efficiently compute equilibrium points of infinitely repeated games among at least 3 players. To this direction, we suggest an alternative notion of *correlated threats*. In particular, we consider the following **correlated threat values** of a normal form game G :

$$\forall p \in N, \varphi_p(G) \equiv \min_{\sigma_{-p} \in \Delta(\times_{q \neq p} S_q)} \max_{\mathbf{x}_p \in \Delta(S_p)} U_p(\sigma_{-p}, \mathbf{x}_p) \tag{2}$$

$\varphi(G) = (\varphi_p(G))_{p \in N}$ is the **correlated threat point** of G . Observe that for any stage game G , it holds that $\varphi(G) \leq \theta(G)$. The only difference with the classical notion of threats is that now the other players may cooperate against the player who must be punished, by adopting a *correlated strategy* $\sigma_{-p} \in \Delta(\times_{q \neq p} S_q)$ rather than a product distribution $\mathbf{x}_{-p} \in \times_{q \neq p} \Delta(S_q)$, against player p .

On the positive side, we prove that the correlated threat point constitutes a credible threat for the players, since both this point and the punishment plans that implement these threats are polynomial-time computable. Additionally, we demonstrate how to exploit this kind of threat to construct an equilibrium point for the infinitely repeated game, despite the alleged difficulties claimed in [10] for more than 2 players. On the negative side, the players will have to coordinate their actions *against* any player that violates the designated behavioral plan which would assure individually rational payoffs for them. But this tradeoff (of sacrificing players' independence for the sake of efficiency of an individually rational payoff point) is only artificial: It is only the punishments that demand coordination, and at equilibrium there will be no need to implement these punishments. During protocol-abiding play, each player follows an independent and uncoordinated (but in total synchrony) plan of actions. Nevertheless, the punishments have to be convincing for each player, in order to constitute a credible threat for them. It is also mentioned that the correlated threat point is not more powerful than it should: It respects each player's individual power in the game, by respecting his/her payoff function. It can be seen as a "*one for all and all (others) against any defector*" scheme that enforces compliance with the agreed behavioral plan, so long as this plan assures payoffs within the individually rational region of the stage game G . Finally, by means of communication complexity, our notion of equilibrium is not far from the classical Nash equilibrium, since during protocol-abiding rounds each player behaves independently of the other players, as would be the case for a classical Nash equilibrium of G^∞ .

Observe that "simple" ways to solve the problem via central enforcement of a (polynomial-time computable) correlated equilibrium for G have actually big deficiencies: (i) They make all players subordinates of a correlation device in each round. (ii) They induce an infinite amount of *private* communication: In each round the correlation device must secretly tell everyone how to play. (iii) They just achieve single points of "common behavior", in contrast to the Folk Theorems rationale according to which *any* feasible payoff point strictly above the (correlated) threats point should be enforceable in G^∞ . In contrast, our approach minimizes the possible use of the correlation device. It is only used as a credible threat, like the heavy charges of extra pages in proceedings (never to be paid indeed) that repel the authors from exceeding the desired page limit. And it does not require private communication in any round (other than public monitoring of the other players' behavior, which is necessary in any case).

An interesting line of research is to see how the necessary correlation for realizing the threats could be implemented via signals embodied in the players' actions (eg, as in [13]). The present paper does not deal with such issues. It mainly focuses on the efficient construction and representation of an equilibrium point, leaving the implementation issues for future study. It is finally mentioned that our results provide *efficient* constructions of equilibrium points that are *not* Nash equilibrium points, since they demand some sort of (not only synchrony but also) coordination of the players via correlated strategies, in order to punish potential defectors from the desired behavioral plan. Therefore, our result is

not in contradiction with that of [4] concerning the **PPAD**–completeness of computing a Nash equilibrium point of (eg) an infinitely repeated 3–player game.

1.1 Definitions and Notation

Mathematical Notation. For any $k \in \mathbb{N}$, let $[k] \equiv \{1, 2, \dots, k\}$. We denote by $A \in F^{m \times n}$ a $m \times n$ matrix whose elements have values in some set F . A $k \times 1$ matrix is also considered to be a **k -vector**. Vectors are denoted by (bold latin, or greek) small letters (eg, $\mathbf{x}, \mathbf{y}, \sigma$). We denote by $\mathbf{1}_k$ ($\mathbf{0}_k$) the k -vector having 1s (0s) in all its coordinates. For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we denote the *component-wise comparison* by $\mathbf{x} \geq \mathbf{y}$: $\forall i \in [n], x_i \geq y_i$. For $A \in \mathbb{R}^{m \times n}$, $A[\star, j]$ is its j -th column (as an $m \times 1$ vector), $A[i, \star]$ is the i -th row (as a transposed $1 \times n$ vector) and $A[i, j]$ is the (i, j) -th element. For $\mathbf{v} \in \mathbb{R}^k$, the value of its i -th coordinate is given by $\mathbf{v}[i]$, or alternatively, by v_i . For any finite k -element set S , $\Delta(S) \equiv \{\mathbf{z} \in \mathbb{R}^k : \mathbf{z} \geq \mathbf{0}; (\mathbf{1}_k)^T \mathbf{z} = 1\}$ is the set of probability distributions over the elements of S . For any $\sigma \in \Delta(S)$, and any real function $U : S \mapsto \mathbb{R}$, $\mathbb{E}_{s \sim \sigma}\{U(s)\}$ is the expected value of U wrt the probability measure σ .

Game Theoretic Notation. Let $G = \langle [k], (S_p)_{p \in [k]}, (U_p)_{p \in [k]} \rangle$ denote an arbitrary k -player normal form game, in which, $\forall p \in [k]$, the **action set** S_p contains all possible actions that this player may take, and the **payoff** $U_p : \times_{p \in [k]} S_p \mapsto \mathbb{R}$ expresses the preferences of p over all possible action profiles \mathbf{s} from the **action space** $S \equiv \times_{p \in [k]} S_p$. Given G , a k -player **infinitely repeated game** G^∞ is an extensive form game which consists of an infinite number of realizations of G (G is usually called the **stage game**, or **constituent game**). G captures the idea that a player will have to take into account the impact of his/her current action on the future actions of other players (sometimes called the **reputation** of this player). In particular, G^∞ is considered to be played in **rounds**. In each round $t \geq 1$, the k players simultaneously choose a strategies profile $\mathbf{x}^{(t)} \in \times_{p \in [k]} \Delta(S_p)$ and get a payoff vector $\mathbf{U}^{(t)} = \mathbf{U}(\mathbf{x}^{(t)}) = (\mathbb{E}_{\mathbf{s} \sim \mathbf{x}^{(t)}}\{U_p(\mathbf{s})\})_{p \in [k]}$. Thus, $\mathbf{U}^{(t)}[p] = \mathbb{E}_{\mathbf{s} \sim \mathbf{x}^{(t)}}\{U_p(\mathbf{s})\}, \forall p \in [k]$. $S^* \equiv \{\emptyset\} \cup (\cup_{t=1}^\infty S^t)$ denotes the collection of **history patterns** $h^t \in S^t$, ie, sequences of action profiles of the G for the first t rounds of G^∞ , for all $t \geq 0$. For player $p \in [k]$, a strategy (wrt G^∞) is any mapping $M_p : S^* \mapsto \Delta(S_p)$ of history patterns to mixed strategies of this player. We regard this as follows: At step $t \geq 1$ player p chooses $\mu_p^{(t)} = M_p(h^{(t-1)})$ as his/her own mixed strategy for the current realization of G , that determines his/her reaction to the history pattern $h^{(t-1)}$ that has appeared so far in G^∞ . We extend this notion of strategy for infinitely repeated games, in order to signal abnormal situations (eg, violation of the prescribed protocol by some player) that have to be handled centrally. Therefore, we consider in this work that a strategy is a mapping $M_p : S^* \mapsto \Delta(S_p) \cup [k]$, ie, each player either chooses a probability distribution according to which he/she will choose an action, or indicates a player (eg, the one of minimum id) that violated the prescribed protocol in the past (eg, during the previous round). In the latter case, the action of p is determined (possibly for a sequence of rounds) by a proper correlation device, ie, which uses a *fixed and publicly known* probability distribution over the action space.

An important feature of an infinitely repeated game is how each player’s preferences are modeled. The most important preference measures are: (i) **Limit-of-Means:** Player $p \in [k]$ gets the average payoff over T periods, as $T \rightarrow \infty$: $U_p(M) = \lim_{T \rightarrow \infty} \frac{1}{T} \cdot \sum_{t=1}^T \mathbb{E}_{\mathbf{s}^{(1)} \sim M(h^{(0)}); \dots; \mathbf{s}^{(t)} \sim M(h^{(t-1)})} \{U_p(\mathbf{s}^{(t)})\}$. (ii) **Overtaking:** Player $p \in [k]$ prefers $M = (M_q)_{q \in [k]}$ from $M' = (M_{-p}, M'_p)$ iff:

$$\lim_{T \rightarrow \infty} \sum_{t=1}^T \left(\mathbb{E}_{\mathbf{s}^{(1)} \sim M(h^{(0)}); \dots; \mathbf{s}^{(t)} \sim M(h^{(t-1)})} \{U_p(\mathbf{s}^{(t)})\} - \mathbb{E}_{\mathbf{s}^{(1)} \sim M'(h^{(0)}); \dots; \mathbf{s}^{(t)} \sim M'(h^{(t-1)})} \{U_p(\mathbf{s}^{(t)})\} \right) > 0$$

(iii) **Discounting:** The valuation of the game diminishes with time depending on some discount parameter $\delta \in (0, 1)$: $U_p(M) = (1 - \delta) \cdot \sum_{t=1}^{\infty} \delta^{t-1} \cdot \mathbb{E}_{\mathbf{s}^{(1)} \sim M(h^{(0)}); \dots; \mathbf{s}^{(t)} \sim M(h^{(t-1)})} \{U_p(\mathbf{s}^{(t)})\}$. In this work we focus our interest on the case of the limit-of-means criterion. Nevertheless, it is mentioned that if an infinite sequence $(\mathbf{v}^{(t)})_{t \geq 1}$ of payoff vectors is preferred to another sequence $(\mathbf{u}^{(t)})_{t \geq 1}$ according to the limit-of-means criterion, then there is a proper discounting factor (close enough to 1) such that $(\mathbf{v}^{(t)})_{t \geq 1}$ is preferred to $(\mathbf{u}^{(t)})_{t \geq 1}$ according to the discounting criterion as well. We say that a strategies profile $M = (M_p)_{p \in [k]}$ is an **equilibrium point** of G^∞ iff for any player $p \in [k]$ and any strategy M'_p the profile M is not inferior to the profile $M' = (M_{-p}, M'_p)$, wrt the limit-of-means criterion. It is obvious that when the players adopt an equilibrium play for the stage game G in every round (but not necessarily the same NE in all rounds) of G^∞ , then such a profile is an equilibrium for G^∞ as well. The presence of additional equilibrium points is because the threat of retaliation is real, since one will play the game infinitely often against the same set of opponents. Every expected payoffs vector \mathbf{u} for the stage game whose coordinates are rational numbers is called a **feasible payoff vector**. If in \mathbf{u} each player p gets payoff (strictly) greater than $\varphi_p(G)$, then \mathbf{u} is called **(strictly) enforceable** or **(strictly) individually rational** payoff vector².

It is well known that any feasible, strictly enforceable payoff point (wrt to classical threats $\theta_p(G)$) of G can be induced as a Nash equilibrium point of G^∞ . This is known as the “Folk Theorem” in the related literature, and was originally proved for the limit-of-means criterion (cf. [2]) and consequently for the discounting criterion, at least for non-degenerate games [7]. Indeed, in many occasions it is true that the optimal method of playing a repeated game is *not* to repeatedly play a Nash profile of the underlying stage game, but to cooperate and play a Pareto optimal strategy (eg, the “cooperate–cooperate” profile in the repeated prisoners’ dilemma case). An analogous fact holds also for the correlated threats. Any feasible, strictly enforceable payoff point can be implemented as an equilibrium point of the infinitely repeated game: Either the players cooperate, or all the players turn against a single defector. This can be interpreted as a “social norm” and one essential part of infinitely repeated games is *punishing*

² The enforceability is defined in the literature wrt the original notion $\theta_p(G)$ of threat values. We use the same term here also wrt our correlated threat values $\varphi_p(G)$, because the meaning is in complete analogy with the classical definition.

players who deviate from such a desirable cooperative behavior. The punishment may be something like playing a strategy which leads to a reduced payoff for a deviating player for some steps, or even for the rest of the game (sometimes called a **trigger** strategy).

Wrt representation issues, our assumption is that all the involved numbers in the description of the (stage) game are rational numbers, of no more than P bits each. Let also $\forall p \in [k]$, $n_p = |S_p|$ be the number of alternative actions that player p has to choose from in each round. The running time of an *efficient* algorithm for computing an equilibrium point of G^∞ must be polynomial in the representation size of G , $size(G) = P \cdot k \cdot \prod_{p \in [k]} n_p$. A polynomial-sized number is one whose representation is bounded by some polynomial function $poly(size(G))$ on the size of G . Multiplication, division, addition or subtraction of polynomial-sized rational numbers, as well as the solution of a polynomial-sized system of linear equations or a linear program, are also of polynomial size [12].

As said before, strategies in infinitely repeated games can be *infinitely large* sequences, mapping the interaction history (so far) to either a probability distribution for selecting the action for the next round, or the index of a player. Hence there is a problem with representing a strategy for each player in such a game. For our discussion to be meaningful, we must consider some sort of a *finite representation* of profiles, when computing equilibrium points of G^∞ . We shall consider here the representation via finite-state automata, as well as their counting-node extensions (cf. [10][11]). In such a description, each node (belonging to a player p) is labeled by a mixed strategy $\sigma_p \in \Delta(S_p)$, determining how player p will choose his/her next action, along with a (square) label indicating the number of rounds that p will stay in this node, before considering migration to another node. If the label is an action, it then implies that p deterministically chooses this action so long as he/she stays at the present node. The outgoing arcs from a node indicate the transition to another node, and they also labels that express the proper condition that should be satisfied in order to follow these arcs. If the label of an arc is a pure profile $\mathbf{x} \in S$, then this implies that in the origin-node it must have been the case that this profile appeared in all the last rounds during which p stayed at this node. Alternatively, the label of an arc may be “*”, indicating the default transition, or some other, polynomial-time computable logical expression that must be satisfied in order to follow the arc. There is though a significant difference with the state-machine strategies defined in [10]. In our case, apart from the players’ finite-state automata, describing their anticipated behavior, we also allow the existence of some additional, globally visible states, which actually represent correlation devices (ie, publicly known probability distributions over S) that implement the correlated threats. But it is only up to the players to decide whether or not to abide with the recommendation of any such state, for a (fixed) number of rounds.

1.2 Contribution and Roadmap

The present work considers the following computational problem: Given any finite game in normal form $G = \langle [k], (S_p)_{p \in [k]}, U_p : \times_{p \in [k]} S_p \mapsto \mathbb{R} \rangle$ and the

corresponding infinitely repeated game G^∞ , determine in *polynomial time* (wrt the representation $size(G)$ of the stage game) a profile of state-automata of size $poly(size(G))$ for the players in G^∞ that is an equilibrium point wrt the limit-of-means criterion. Our main findings are the following: In Section 2 we prove that our alternative notion of *correlated* threats is polynomial time computable (and therefore credible). Although more severe than the traditional notion of threats, the correlated threats do not compromise the actual power of each individual player. In Section 3 we study the case when the underlying stage game G has a non-empty strictly individually rational region (denoted by $sirr(G)$). We provide in polynomial time a polynomial-size equilibrium point for G^∞ , for any constant number of players, that induces an arbitrary feasible point of polynomial-size representation in $sirr(G)$. Indeed, we manage to choose a feasible and enforceable point that is also quality assuring wrt the individual players’ levels of satisfaction. In Section 4 we focus on the case where $sirr(G) = \emptyset$. We demonstrate the efficient construction of an equilibrium point for an arbitrary number of players and up to 2 concurrently positive payoff coordinates in any payoff vector of the stage game. Along with the previous result, this completely resolves the cases of 3 players, and provides a direction towards handling the cases of more players. Due to lack of space some proofs are deferred to the full version of the paper [9].

2 The Correlated Threat Point

Fix any normal form k -player game $G = \langle [k], (S_p)_{p \in [k]}, (U_p)_{p \in [k]} \rangle$ with rational payoff functions $U_p : S \mapsto \mathbb{Q}$ for the players. Let $Z = \{ \mathbf{z} \in \mathbb{Q}^k : \exists \mathbf{s} \in S \text{ s.t. } \forall p \in [k], U_p(\mathbf{s}) = \mathbf{z}[p] \}$ be the set of all the payoff vectors that the k players may get at an actions profile $\mathbf{s} \in S$ of G . We also denote by $conv(Z)$ the convex hull of this point set in \mathbb{R}^k : $conv(Z) = \{ \sum_{\mathbf{s} \in S} \lambda_{\mathbf{s}} \cdot \mathbf{U}(\mathbf{s}) \in \mathbb{R}^k : \sum_{\mathbf{s} \in S} \lambda_{\mathbf{s}} = 1; \forall \mathbf{s} \in S, \lambda_{\mathbf{s}} \geq 0 \}$. Observe that any *correlated strategy* $\sigma \in \Delta(S)$ corresponds to a vector of (expected) payoffs $\mathbf{U}(\sigma) = \mathbb{E}_{\mathbf{s} \sim \sigma} \{ \mathbf{U}(\mathbf{s}) \} \in conv(Z)$, and vice versa. Our first goal is to determine in polynomial time the minimum acceptable payoff for player p , as well as the player p ’s defensive strategy (ensuring it) and the other players’ aggressive (correlated) strategies (enforcing it). This minimum will be exactly p ’s correlated threat value $\varphi_p(G)$.

Theorem 1. *For any fixed constant natural number $k \geq 2$ and any finite game in normal form $G = \langle [k], (S_p)_{p \in [k]}, (U_p)_{p \in [k]} \rangle$, the correlated threat values, the defensive strategies (of each player) and the aggressive strategies (of the other players against each player) are polynomial time computable wrt $size(G)$.*

Proof. For each player $p \in [k]$, consider the $n_p \times (\prod_{q \neq p} n_q)$ matrix defined as follows: $\forall (s_p, \mathbf{s}_{-p}) \in S_p \times S_{-p}, P_p[s_p, \mathbf{s}_{-p}] = U_p(s_p, \mathbf{s}_{-p})$. Observe that P_p contains all possible payoffs that player p may receive, in all possible outcomes of the stage game G . The rows of P_p correspond to actions of player p , while the columns correspond to *joint* actions of all the other players, ie, elements $\mathbf{s}_{-p} \in S_{-p}$. We consider the zero-sum bimatrix game $\langle P_p, P_{-p} \rangle$, and we define the following primal-dual pair of LPs determining the *maximin strategy* $\mathbf{d}_p \in$

$\Delta(S_p)$ (and payoff $\varphi_p(G) = V_p$) for player p , and the *maximin* (due to zero sum property) correlated strategy $\mathbf{a}_p \in \Delta(S_{-p})$ of the other players against p :

$$(V_p, \mathbf{d}_p) \in \arg \max \{ \bar{V}_p : \forall \mathbf{s}_{-p} \in S_{-p}, \bar{\mathbf{d}}^T \cdot P_p[\star, \mathbf{s}_{-p}] \geq \bar{V}_p; \bar{\mathbf{d}}_p \in \Delta(S_p) \} \quad (3)$$

$$(V_p, \mathbf{a}_p) \in \arg \min \{ \bar{V}_p : \forall s_p \in S_p, P_p[s_p, \star] \cdot \bar{\mathbf{a}} \leq \bar{V}_p; \bar{\mathbf{a}}_p \in \Delta(S_{-p}) \} \quad (4)$$

It is well known that these maximin-play strategies $(\mathbf{d}_p, \mathbf{a}_p)$ comprise a Nash equilibrium, in the zero-sum (bimatrix) game $\langle P_p, P_{-p} \rangle$, while V_p is the minimum payoff that player p may accept in any equilibrium play (either in $\langle P_p, -P_p \rangle$ or G), even if his/her opponents join forces against him/her. Recall that \mathbf{a}_p is a *correlated* strategy for all the other players (assumed to conspire against p). By duality theorem of linear programming, we also know that the maximin solution in this primal–dual pair equals the minmax solution for p on one hand, and all the other players against p on the other hand, which is the solution to our notion of correlated threats (cf. Equation 2). We call \mathbf{d}_p the **defensive strategy** of p and $V_p = \varphi_p(G)$ is his/her correlated threat value, while $\mathbf{a}_p \in \Delta(S_{-p})$ is the **aggressive (correlated) strategy** of the other players against p , that pushes p 's payoff down to $\varphi_p(G)$. So long as the number of players, k is *constant*, we can efficiently compute both the defensive and the aggressive strategies for each player $p \in [k]$, as well as their correlated threat values: Since the defensive and aggressive strategies, as well as the correlated threat values, are solutions to proper linear programs of polynomial size, they are also of polynomial-size. \square

From now on we shall consider wlog³ for the correlated threat point that $\varphi(G) = \mathbf{0}$. If this is not the case, we simply consider a stage game G' properly shifted utility functions by $\varphi(G)$, so that it has $\varphi(G') = \mathbf{0}$. Clearly the two games have exactly the same set of (even approximate) Nash equilibria. Our next step is, given the tractability of the correlated threat point, to construct in polynomial time an equilibrium point for the infinitely repeated game G^∞ . We distinguish two main cases, depending on whether the **strictly individually rational region** of G , $\text{sirr}(Z) = \text{conv}(Z) \cap \{ \mathbf{z} \in \mathbb{R}^k : \mathbf{z} > \varphi(G) \}$, is empty or not. If $\text{sirr}(G) \neq \emptyset$ then we say that there exist payoff points of G guaranteeing a **mutual advantage** (to comply with the desired plan) for the players, following the terminology of [10]. In such a case we implement the rationale of the Folk Theorem, by locating such a (feasible) payoff point in polynomial time and enforcing it as the result of the players' behavioral plan in G^∞ . If $\text{sirr}(Z) = \emptyset$, then we provide an alternative construction of equilibrium points based on the stability of subsets of players.

3 The Mutual Advantage Case

In our quest for a polynomial-time computable equilibrium point of G^∞ , a crucial task is to discover whether $\text{sirr}(Z) \neq \emptyset$. In that case, as we shall see shortly, any feasible, strictly individually rational payoff point of G may be interpreted into an equilibrium of polynomial–size description, for G^∞ .

³ Without loss of generality.

Lemma 1. *For any fixed integer $k \geq 2$ and game $G = \langle [k], (S_p)_{p \in [k]}, (U_p)_{p \in [k]} \rangle$, we can determine in time $\text{poly}(\text{size}(G))$ whether $\text{sirr}(G) \neq \emptyset$.*

Proof. See full version of the paper. \square

Remark: The payoff vector that is constructed in the proof of Lemma 1 is also some sort of quality assuring social norm, in the sense that it maximizes in $\text{sirr}(G)$ the minimum payoff assigned to any of the players.

The following theorem is the main result of the mutual-advantage case:

Theorem 2. *For any constant $k \geq 2$, and game $G = \langle [k], (S_p)_{p \in [k]}, (U_p)_{p \in [k]} \rangle$ such that $\text{sirr}(G) \neq \emptyset$, there is a strategies profile $M = (M_p)_{p \in [k]}$ for the players that is an equilibrium of G^∞ , whose description size is $\text{poly}(\text{size}(G))$.*

Proof (sketch). In order to prove this statement, we proceed with the description of such a profile for the infinitely repeated game. Initially we describe a behavioral plan for the players that assures an average payoff vector equal to the one point of $\mathbf{z}^* = \sum_{i=1}^k \hat{\lambda}_i \hat{\mathbf{z}}_i$ of $\text{sirr}(G)$ we have already located in the proof of Lemma 1. This was the outcome of some LP that used as input a subset of k vertices $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k \in Z$. Consequently we describe how we guarantee the compliance of all the players to this behavioral plan, via an implementation of the correlated threats devices.

Assuring a Behavioral Plan with Payoff Vector in $\text{sirr}(Z)$. The payoff vector $\mathbf{z}^* = \sum_{i=1}^k \hat{\lambda}_i \hat{\mathbf{z}}_i$ indicated in Lemma 1 is a rational payoff vector, since it is the outcome of some LP with rational coefficients, over a set of k pure strategies (also rational) payoff points $\hat{\mathbf{z}}_1, \dots, \hat{\mathbf{z}}_k \in Z$. Additionally, the vector $\hat{\lambda}$ indicates how \mathbf{z}^* is expressed as a *convex combination* of these payoff points, and is also rational. Finally, both these vectors are of size $\text{poly}(\text{size}(G))$, as the solutions of an LP. Let $\mathbf{x}_i = \mathbf{s}(\hat{\mathbf{z}}_i) \in S$ be the pure strategies profile providing the payoff vector $\hat{\mathbf{z}}_i$, for each $i \in [k]$. It then holds that: $\forall i \in [k], \hat{\lambda}_i = \frac{\gamma_i}{\Gamma_i} = \frac{\gamma_i \prod_{j \neq i} \Gamma_j}{\prod_{j \in [k]} \Gamma_j} = \frac{\xi_i}{\Xi}$, where $\forall i \in [k], 0 \leq \gamma_i \leq \Gamma_i$, $\xi_i = \gamma_i \prod_{j \neq i} \Gamma_j$ and $\Xi = \prod_{j \in [k]} \Gamma_j$ are natural numbers of size $\text{poly}(\text{size}(G))$. Because $\sum_{i=1}^k \hat{\lambda}_i = 1 \Leftrightarrow \sum_{i=1}^k \xi_i = \Xi$, we can construct a Ξ -state machine M_p for each player $p \in [k]$, so that M_p starts with exactly ξ_1 rounds of $\mathbf{x}_1[p]$, followed by ξ_2 rounds of $\mathbf{x}_2[p]$, followed by ξ_3 rounds of $\mathbf{x}_3[p]$, and so on. When all these machines operate in *full synchrony*, it holds that during each phase of Ξ rounds the profile \mathbf{x}_i appears exactly ξ_i times, for all $i \in [k]$. The payoff vector for a whole phase of Ξ consecutive rounds is therefore $\sum_{i=1}^k \xi_i \cdot \hat{\mathbf{z}}_i = \Xi \cdot \mathbf{z}^* \geq \Xi \cdot \hat{\mathbf{c}} > 0$, since we have assumed non-emptiness of $\text{sirr}(Z)$. I.e, the average payoff vector per round is exactly $\mathbf{z}^* \geq \hat{\mathbf{c}} \cdot \mathbf{1}$.

Inducing an Equilibrium via Cooperative Punishments. In order to enforce the compliance of all the players to the previously mentioned behavioral plan, which assures the average payoff vector \mathbf{z}^* as an equilibrium point of the repeated game, we must assure that the other players discourage each player from disobeying the prescribed protocol. To this direction, we shall use the aggressive

(correlated) strategy of the other players against the *defector with the smallest index* (in order for all players to attack the same defector). In particular, assume that at some round $t \geq 1$, player $q \in [k]$ is the defecting player with the smallest index. Then, starting from round $t + 1$, the other players exit their own finite state machines and give up control of their choices for a sufficient number of rounds to an independent correlation device that uses the mixed profile \mathbf{a}_q to determine their profiles in each of these rounds. We call this the **punishment phase** of player q , and its duration is such that the positive gain that player q had at step t is completely absorbed by his/her cumulative loss during the punishment phase. See more details in the full version of the paper. \square

4 Handling Games with No Mutual Advantage

We now explore what happens if $\text{sirr}(Z) = \emptyset$. This situation is easily recognizable by the fact that the point $\mathbf{z}^* = \sum_{i=1}^k \hat{\lambda}_i \cdot \hat{\mathbf{z}}_i$ determined in the (full) proof of Lemma 1 has minimum payoff $\hat{\zeta} = 0$. The next lemma calculates the maximum number $\mu(G)$ of players that may concurrently get a positive payoff, at some point in $\text{conv}(Z)$.

Lemma 2. *For any constant $k \geq 2$ and any game $G = \langle [k], (S_p)_{p \in [k]}, (U_p)_{p \in [k]} \rangle$, $\mu(G)$ is computable in time $\text{poly}(\text{size}(G))$.*

Proof. See the full version of the paper. \square

The next theorem gathers some particular cases we have handled, that may work as a guide for coping with the general no-mutual-advantage case. Along with the mutual-advantage result, this completely resolves the case of 3 players.

Theorem 3. *For any constant $k \geq 2$ and any game $G = \langle [k], (S_p)_{p \in [k]}, (U_p)_{p \in [k]} \rangle$ with $\text{sirr}(G) = \emptyset$, there is an efficiently computable equilibrium point for G^∞ , when $\mu(G) \leq 2$.*

Proof (sketch). Let $\mu = \mu(G)$. Observe that if $\mu = 0$ then the defensive profile \mathbf{d} , which assures a non-negative payoff for every player, is a NE point of G (and therefore, its infinite repetition is an equilibrium of G^∞). But what happens in the general (no-mutual-advantage) case? Here we demonstrate here how to construct an equilibrium point of G^∞ for $\mu = 1$. In the full version of the paper we provide a similar (but more involved) argument for $\mu = 2$. Our conjecture is that an analogous argument also works for the more general cases of $\mu \in \{3, \dots, k-1\}$.

The Case of $\mu = 1$: For the defensive profile $\mathbf{d} \in \times_{p \in [k]} \Delta(S_p)$ let $\forall p \in [k]$, $s_p^{(1)} \in \arg \max_{s_p \in S_p} \{U_p(s_p, \mathbf{d}_{-p})\}$; $v_p^{(1)} = U_p(s_p^{(1)}, \mathbf{d}_{-p})$ be a pure best response, and the corresponding (maximum possible) payoff, of player p against the defensive profile \mathbf{d}_{-p} of the other players. If $\mathbf{v}^{(1)} \equiv \left(v_p^{(1)} \right)_{p \in [k]} = \mathbf{U}(\mathbf{d})$, then the defensive profile \mathbf{d} is NE of G , since it already assures a maximum possible payoff for every player against the defensive profile. So we assume (wlog) that $v_1^{(1)} > U_1(\mathbf{d}) \geq 0$ (since \mathbf{d} is the profile of defensive strategies). We shall prove that

the profile $(s_1^{(1)}, \mathbf{d}_{-1})$ is NE for G . Suppose not. Then, let $\forall p \in [k] \setminus \{1\}, s_p^{(2)} \in \arg \max_{s_p \in S_p} \{U_p(s_1^{(1)}, s_p, \mathbf{d}_{-\{1,p\}})\}$; $v_p^{(2)} = U_p(s_1^{(1)}, s_p^{(2)}, \mathbf{d}_{-\{1,p\}})$. It must be the case for some $p \in [k] \setminus \{1\}$ (say, for $p = 2$) that $v_2^{(2)} > U_2(s_1^{(1)}, \mathbf{d}_{-1}) \geq 0$. Observe that the profile $(s_1^{(1)}, \mathbf{d}_{-1})$ produces a payoff vector $\mathbf{U}(s_1^{(1)}, \mathbf{d}_{-1}) = (\bar{v}_1, \bar{\mathbf{u}}_{-1}) \in \text{conv}(Z)$ where $\bar{v}_1 > 0$ and $\bar{\mathbf{u}}_{-1} = \mathbf{0}$ (since $\mu = 1$ and the other players adopt their defensive strategies). On the other hand, the profile $(s_1^{(1)}, s_2^{(2)}, \mathbf{d}_{-\{1,2\}})$ produces a payoff vector $\mathbf{U}(s_1^{(1)}, s_2^{(2)}, \mathbf{d}_{-\{1,2\}}) = (\hat{u}_1, \hat{v}_2, \hat{\mathbf{u}}_{-\{1,2\}}) \in \text{conv}(Z)$ for which we know that $\hat{v}_2 > 0$ and $\hat{\mathbf{u}}_{-\{1,2\}} = \mathbf{0}$. Consider the points of the following line segment in $\text{conv}(Z)$: $\forall \lambda \in [0, 1], \mathbf{z}(\lambda) = \lambda(\bar{v}_1, \bar{\mathbf{u}}_{-1}) + (1 - \lambda)(\hat{u}_1, \hat{v}_2, \hat{\mathbf{u}}_{-\{1,2\}}) = (\lambda\bar{v}_1 + (1 - \lambda)\hat{u}_1, \lambda\bar{u}_2 + (1 - \lambda)\hat{v}_2, \lambda\bar{\mathbf{u}}_{-\{1,2\}} + (1 - \lambda)\hat{\mathbf{u}}_{-\{1,2\}})$, for which we observe that: (i) $\forall \lambda \in [0, 1] \lambda\bar{\mathbf{u}}_{-\{1,2\}} + (1 - \lambda)\hat{\mathbf{u}}_{-\{1,2\}} = \mathbf{0}$, (ii) $\forall \lambda \in [0, 1), \lambda\bar{u}_2 + (1 - \lambda)\hat{v}_2 = (1 - \lambda)\hat{v}_2 > 0$, and (iii) $\forall \lambda \in \left(\frac{|\hat{u}_1|}{\bar{v}_1 + |\hat{u}_1|}, 1\right], \lambda\bar{v}_1 + (1 - \lambda)\hat{u}_1 \geq \lambda\bar{v}_1 - (1 - \lambda)|\hat{u}_1| > 0$. That is, we have proved that for any value $\lambda \in \left(\frac{|\hat{u}_1|}{\bar{v}_1 + |\hat{u}_1|}, 1\right)$ it holds that $\mathbf{z}(\lambda) \in \text{conv}(Z)$ has *two strictly positive* payoffs (all the other are zero payoffs), which is a contradiction to $\mu = 1$. Therefore we conclude that $\mathbf{v}^{(2)} = \mathbf{U}(s_1^{(1)}, \mathbf{d}_{-1})$ and thus the profile $(s_1^{(1)}, \mathbf{d}_{-1})$ is NE point of G .

The Case of $\mu = 2$: The details of this construction can be found in the full version of the paper. The main idea is as follows: We locate a point in $\text{conv}(Z)$ corresponding to a mixed profile of the players, in which exactly two players deviate from their defensive strategies, and get positive payoffs, while the remaining $k - 2$ players get (exactly) zero payoffs, which are *provably* the best possible payoffs they can get all over $\text{conv}(Z)$. This guarantees that their defensive strategies are weakly dominant strategies⁴ for these $k - 2$ players.

We then focus on the first two players, who may (even concurrently) get positive payoffs at some points in $\text{conv}(Z)$. Indeed, since $\mathbf{d}_{-\{1,2\}}$ is a profile of weakly dominant strategies for the remaining players, we shall take for granted that all these players adopt this profile. That is, we consider the bimatrix produced by the *expected* payoffs of the first two players, given that all the other players comply with the profile $\mathbf{d}_{-\{1,2\}}$. In particular, consider the $n_1 \times n_2$ payoff matrices $\forall p \in \{1, 2\}, Q_p = [\mathbb{E}_{\mathbf{s}_{-\{1,2\}} \sim \mathbf{d}_{-\{1,2\}}} \{U_p(s_1, s_2, \mathbf{s}_{-\{1,2\}})\}]_{(s_1, s_2) \in S_1 \times S_2}$, which are polynomial-time computable due to the constant number k of players that we consider. Observe also that these two payoff matrices are rational and of size $\text{poly}(\text{size}(G))$, given the rationality of the defensive strategies and of the payoff values of the players in G . For the normal form (bimatrix) game $\Gamma = \langle Q_1, Q_2 \rangle$ it

⁴ A strategy is **weakly dominant** if, regardless of what the other players do, the strategy assures for this player a payoff at least as high as any other strategy, and additionally, it assures a strictly higher payoff for some profile of other players' strategies. If a strategy of a player is always strictly better than all other strategies of this player, for all profiles of other players' strategies, then it is called a **strictly dominant** strategy.

is true that the correlated threats vector $\varphi(I)$ is non-negative: Both players still have the option of adopting their defensive strategies wrt the original game. For the infinitely repeated game I^∞ between players 1 and 2 we determine (inductively) an equilibrium point against correlated threats. It now holds that so long as the other players keep playing according to the profile $\mathbf{d}_{-\{1,2\}}$, the first two players will have no intention to unilaterally defect from the prescribed equilibrium (between them) in I^∞ , in fear of one another's retaliation. On the other hand, due to the weak dominance of the strategies in $\mathbf{d}_{-\{1,2\}}$, no other player is willing to change strategy, no matter how the first two players behave, for the whole duration of G^∞ , because there is no payoff point (even correlated) in $\text{conv}(Z)$ that assures them a strictly positive payoff, and their defensive strategies already assure a zero payoff to each of them. This completes the proof of Theorem 3. \square

Acknowledgements. We thank Leonidas Palios for helpful discussions, especially concerning the geometric part of our arguments.

References

1. Abraham, I., Dolev, D., Gonen, R., Halpern, J.: Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In: Proc. of 25th ACM Symp. on Princ. of Distr. Comp. (PODC 2006) (2006)
2. Aumann, R.J., Shapley, L.S.: Long-term competition – a game theoretic analysis. In: Megiddo, N. (ed.) Essays in Game Theory, pp. 1–15 (1994)
3. Bárány, I.: Fair distribution protocols or how the players replace fortune. *Mathematics of Operations Research* 17(2), 327–340 (1992)
4. Borgs, C., Chayes, J., Immorlica, N., Kalai, A.T., Mirrokni, V., Papadimitriou, C.: The myth of the folk theorem. In: Proc. of 40th ACM Symp. on Th. of Comp. (STOC 2008) (2008)
5. Chen, X., Deng, X.: Settling the complexity of 2-player nash equilibrium. In: Proc. of 47th IEEE Symp. on Found. of Comp. Sci (FOCS 2006). IEEE Comp. Soc. Press, Los Alamitos (2006)
6. Daskalakis, C., Goldberg, P., Papadimitriou, C.: The complexity of computing a nash equilibrium. In: Proc. of 38th ACM Symp. on Th. of Comp (STOC 2006). Assoc. of Comp. Mach. ACM (2006)
7. Fudenberg, D., Maskin, E.: Folk theorems for repeated games with discounting and incomplete information. *Econometrica* 54, 533–554 (1986)
8. Hansen, K.A., Hansen, T.D., Miltersen, P.B., Sørensen, T.B.: Approximability and parameterized complexity of minmax values. In: Proc. of 4th W. on Internet and Net Econ. (WINE 2008) (2008)
9. Kontogiannis, S., Spirakis, P.: Equilibrium points in fear of correlated threats. In: Proc. of 4th W. on Internet and Net. Econ. (WINE 2008) (2008), <http://www.cs.uoi.gr/~kontog/pubs/repeated-nash.pdf>
10. Littman, M.L., Stone, P.: A polynomial-time nash equilibrium algorithm for repeated games. *Decision Support Systems* 39(1), 55–66 (2005)
11. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press, Cambridge (1994)
12. Schrijver, A.: *Theory of Linear and Integer Programming*. Wiley–Interscience, New York (1986)

Performance Evaluation of a Descent Algorithm for Bi-matrix Games

Haralampos Tsaknakis¹, Paul G. Spirakis^{1,2}, and Dimitrios Kanoulas²

¹ Research Academic Computer Technology Institute (RACTI), Greece
{tsaknak,spirakis}@cti.gr

² Dept. of Computer Eng. and Informatics, Patras University, Patras, Greece
kanoulas@ceid.upatras.gr

Abstract. In this paper we present an implementation and performance evaluation of a descent algorithm that was proposed in [1] for the computation of approximate Nash equilibria of non-cooperative bi-matrix games. This algorithm, which achieves the best polynomially computable ϵ -approximate equilibria till now, is applied here to several problem instances designed so as to avoid the existence of easy solutions. Its performance is analyzed in terms of quality of approximation and speed of convergence. The results demonstrate significantly better performance than the theoretical worst case bounds, both for the quality of approximation and for the speed of convergence. This motivates further investigation into the intrinsic characteristics of descent algorithms applied to bi-matrix games. We discuss these issues and provide some insights about possible variations and extensions of the algorithmic concept that could lead to further understanding of the complexity of computing equilibria. We also prove here a new significantly better bound on the number of loops required for convergence of the descent algorithm.

1 Introduction and Definitions

The problem of computing approximate Nash equilibria in polynomial time has received attention due to the intractability results for the problem of finding exact Nash equilibria ([4]), even for two-player games. The two-player game is intractable in the sense that it is PPAD-complete. Simple polynomial time algorithms have been presented in the past two years for finding ϵ -approximate equilibria for $\epsilon = \frac{3}{4}$ and $\epsilon = \frac{1}{2}$ ([2], [6]). Furthermore, some more complicated polynomial time algorithms have recently been presented, achieving approximations $\epsilon = 0.38$, $\epsilon = 0.36$ and $\epsilon = 0.3393$ ([5], [7], [1]). The last result achieves the best constant approximation reported in the literature so far and is based on an optimization approach applied to the problem of finding approximate equilibria for bi-matrix games. In particular, it is based on a descent algorithm aiming at minimizing the value of a regret function representing the approximation of equilibria. This work is concerned with the implementation and evaluation of the performance of this approach as well as the presentation of a new complexity result on the convergence rate.

Let R, C denote the $m \times n$ row and column players' payoff matrices respectively, for m, n any positive integers. We assume, without loss of generality, that both payoff matrices are positively normalized, i.e. the maximum entry is 1 and the minimum entry is 0 for each matrix.

Let us denote by e_k the k -dimensional column vector having all its entries equal to 1 (for positive integer k) and let $\Delta_k = \{u : u \in R^k, u \geq 0, e_k^\tau u = 1\}$ denote the k -dimensional standard simplex (superscript τ denotes transpose). Also, let $(1, k)$ denote the set of indices from 1 to k . We will need the following additional definitions for any vector $u \in R^k$:

- $supp(u) = \{i \in (1, k) : u_i \neq 0\}$ (the set of indices in $(1, k)$ for which u is non-zero).
- $suppmax(u) = \{i \in (1, k) : u_i \geq u_j \ \forall j \in (1, k)\}$ (the set of indices in $(1, k)$ for which the maximum value of u is attained).
- $max(u) = \{u_i : i \in (1, k), u_i \geq u_j \ \forall j \in (1, k)\}$ (the maximum value of u).
- $max_S(u) = \{u_i, i \in S : u_i \geq u_j \ \forall j \in S\}$ (the maximum value of u in the index set $S \subset (1, k)$).

The problem of finding an ϵ -approximate Nash equilibrium of the game (R, C) , for some $\epsilon \geq 0$, is to compute a pair of strategies $\hat{x} \in \Delta_m$ and $\hat{y} \in \Delta_n$ such that the following relationships hold: $x^\tau R \hat{y} \leq \hat{x}^\tau R \hat{y} + \epsilon \ \forall x \in \Delta_m$ and $\hat{x}^\tau C y \leq \hat{x}^\tau C \hat{y} + \epsilon \ \forall y \in \Delta_n$.

Following the analysis in [1], we define the following function, mapping $\Delta_m \times \Delta_n$ into $[0, 1]$: $f(x, y) = max\{f_R(x, y), f_C(x, y)\}$, where, $f_R(x, y) = max(Ry) - x^\tau Ry$ and $f_C(x, y) = max(C^\tau x) - x^\tau Cy$. For any pair of strategies $(x, y) \in \Delta_m \times \Delta_n$, this function measures the maximum distance between the players' payoffs (achieved by that pair of strategies) and their respective best response payoffs. We call this a **regret function**.

The descent approach attempts to minimize the regret function through an iterative process moving along feasible descent directions in the space of strategies for both players simultaneously. The descent directions are computed by solving linear programs. A descent algorithm produces a monotonically decreasing regret function which always terminates with a **stationary point** in the space of strategies. It was proven in [1] that at any stationary point we obtain an ϵ -approximate Nash equilibrium, where ϵ is ≤ 0.3393 (for positively normalized payoff matrices).

2 Algorithm Description

In this section we give the main points underlying the descent algorithm as derived in [1]. From any given point $(x, y) \in \Delta_m \times \Delta_n$, we consider motions away from it along feasible directions of the form $(1 - \epsilon)[x^\tau, y^\tau] + \epsilon[(x')^\tau, (y')^\tau]$, where, $(x', y') \in \Delta_m \times \Delta_n$ is another pair of strategies and $\epsilon : 0 \leq \epsilon \leq 1$ (vectors in brackets denote $m + n$ -dimensional vectors). The computation of a feasible direction $(x', y') \in \Delta_m \times \Delta_n$ that is also a descent direction for the regret function at a point (x, y) involves the solution of an appropriate linear program which is formulated as follows, for two distinct cases (A) and (B):

(A) If $f_R(x, y) \neq f_C(x, y)$, then:

(a1) If $f_R(x, y) > f_C(x, y)$, keep y fixed and solve the following LP with respect to x' :

$$\min_{x'} \{ \max(Ry) - (x')^\tau Ry \}$$

subject to: $\max(C^\tau x') - (x')^\tau Cy \leq \max(Ry) - (x')^\tau Ry$ and $x' \in \Delta_m$.

A minimizer x' defines a descent direction (x', y) .

(a2) If $f_R(x, y) < f_C(x, y)$, keep x fixed and solve the following LP with respect to y' :

$$\min_{y'} \{ \max(C^\tau x) - x^\tau Cy' \}$$

subject to: $\max(Ry') - x^\tau Ry' \leq \max(C^\tau x) - x^\tau Cy'$ and $y' \in \Delta_n$.

A minimizer y' defines a descent direction (x, y') .

(B) If $f_R(x, y) = f_C(x, y)$, then solve the following $(m + n)$ -dimensional LP in mini-max form:

$$\min_{(x', y')} \max_{(w, z, \rho)} [\rho w^\tau, (1 - \rho) z^\tau] G(x, y) \begin{bmatrix} y' \\ x' \end{bmatrix}$$

where:

(i) The maximum is taken with respect to dual variables w, z, ρ such that:

$$w \in \Delta_m \text{ and } \text{supp}(w) \subset S_R(y) \equiv \text{suppmax}(Ry)$$

$$z \in \Delta_n \text{ and } \text{supp}(z) \subset S_C(x) \equiv \text{suppmax}(C^\tau x)$$

$$\rho \in [0, 1]$$

(ii) the minimum is taken with respect to $(x', y') \in \Delta_m \times \Delta_n$, and

(iii) the matrix $G(x, y)$ is the following $(m + n) \times (m + n)$ matrix:

$$G(x, y) = \begin{bmatrix} R - e_m x^\tau R & -e_m y^\tau R^\tau + e_m e_m^\tau x^\tau Ry \\ -e_n x^\tau C + e_n e_n^\tau x^\tau Cy & C^\tau - e_n y^\tau C^\tau \end{bmatrix}$$

The descent direction is specified by a minimizer (x', y') of the above problem.

It should be pointed out that in case (A) a solution of the corresponding LP always leads to a point where the values of the two components of the regret function are equal, i.e. $f_R(x', y) = f_C(x', y)$ (or $f_R(x, y') = f_C(x, y')$) and the regret function at this point is strictly smaller than the previous one. We make this statement precise in the following Lemma:

Lemma 1. *At a given pair of strategies (x, y) , if $f_R(x, y) > f_C(x, y)$ and x' is a minimizer of the LP in (a1) above, then:*

$$f(x', y) = f_R(x', y) = f_C(x', y) \leq \frac{f_R(x, y)}{1 + f_R(x, y) - f_C(x, y)}$$

(A similar statement holds if $f_R(x, y) < f_C(x, y)$ as in case (a2) by interchanging the roles of the row and column players).

Proof. For fixed y , the equality of the two regrets (for the row and column players) follows from the fact that at an optimal solution of the LP in (a1) above, not all constraints can be inactive since the upper bound of these constraints (which is the regret of the row player and the objective function to be minimized) can always be made equal to 0 by a choice of a best response strategy x' such that $\text{supp}(x') \subset \text{suppmax}(Ry)$.

Now, choose any $x_1 \in \Delta_m$ such that: $\text{supp}(x_1) \subset \text{suppmax}(Ry)$. Consider the following function:

$$g(\epsilon) = f_C((1 - \epsilon)x + \epsilon x_1, y) - f_R((1 - \epsilon)x + \epsilon x_1, y) \text{ for } 0 \leq \epsilon \leq 1$$

It is easy to verify that $g(\epsilon)$ is continuous and convex in ϵ and that it satisfies: $g(0) = -(f_R(x, y) - f_C(x, y)) < 0$ and $0 \leq g(1) = f_C(x_1, y) \leq 1$. So, there is an ϵ' such that $g(\epsilon') = 0$. By convexity, we should have $g(\epsilon') \leq (1 - \epsilon')g(0) + \epsilon'g(1)$ which, in view of the above relationships, implies: $(1 - \epsilon') \leq \frac{1}{1 + f_R(x, y) - f_C(x, y)}$. Furthermore, choosing $x' = (1 - \epsilon')x + \epsilon'x_1$, we have $f_C(x', y) = f_R(x', y) = (1 - \epsilon')f_R(x, y)$. Finally, the assertion of the Lemma follows from the last two relationships. \square

From the above Lemma it can also be observed that if we take any arbitrary strategy y for the column player and a strategy x for the row player that is a best response to it, then, the resulting value of the regret function in the first step will always be $\leq \frac{1}{2}$. So, for the computations in the main step (B) of the algorithm we can always start at a point with a regret function value that is $\leq \frac{1}{2}$ and for which the two components of the regret are equal, i.e. $f_R(x, y) = f_C(x, y)$. We call the overall process of computing a descent direction a **Descent Direction** step.

After obtaining a descent direction, say (x', y') , from any given point (x, y) , we move on to compute the minimum of the function $f(x + \epsilon(x' - x), y + \epsilon(y' - y))$ with respect to the scalar parameter ϵ , producing thus a new pair of strategies with smaller regret. We call this process a **Line Search** step.

For the line search computations, we exploit the fact that the above function is piece-wise quadratic with respect to ϵ (for any x, y, x', y') and the total number of switches from one quadratic to another is less than $m + n$. For the purposes of the implementation described here, we have considered stepsizes equal to the switching point ϵ^* that is closest to 0. We provide explicit estimates of this stepsize in the proof of Lemma 2 below.

Lemma 2. *Let (x, y) be a pair of strategies such that $f_R(x, y) = f_C(x, y)$ and let $(x', y') \in \Delta_m \times \Delta_n$ be a solution of the LP defined in (B) above. Also, let $V(x, y)$ denote the optimal value of the LP. Then, if $V(x, y) - f(x, y) < 0$, there is an $\epsilon^* > 0$ such that $f(x + \epsilon(x' - x), y + \epsilon(y' - y)) - f(x, y) < 0$ for all $\epsilon \in (0, \epsilon^*)$.*

Proof. The proof is based on the construction of the above difference for any $\epsilon : 0 \leq \epsilon \leq 1$ and on using the properties of a solution of the mini-max LP in (B). We also provide a new explicit bound on the decrease of the function f at each step of the algorithm (compared to the one given in [11]), which is useful for a more refined analysis of its convergence properties.

At first, it can be easily verified (by setting $x' = x, y' = y$) that we always have $V(x, y) - f(x, y) \leq 0$. Let $\overline{S_R}(y)$ and $\overline{S_C}(x)$ be the complements of the index sets $S_R(y)$ and $S_C(x)$ with respect to the index sets $(1, m)$ and $(1, n)$ respectively (we have defined $S_R(y) \equiv \text{suppmax}(Ry)$ and $S_C(x) \equiv \text{suppmax}(C^\tau x)$). Let $\epsilon^* = \min\{\epsilon_1^*, \epsilon_2^*, 1\}$, where:

$$\epsilon_1^* = \min_i \left[\frac{\max(Ry) - (Ry)_i}{\max(Ry) - (Ry)_i + (Ry')_i - \max_{S_R(y)}(Ry')} \right]$$

(where the minimum is taken over i such that: $i \in \overline{S_R}(y)$ and $(Ry')_i - \max_{S_R(y)}(Ry') \geq 0$)

$$\epsilon_2^* = \min_j \left[\frac{\max(C^\tau x) - (C^\tau x)_j}{\max(C^\tau x) - (C^\tau x)_j + (C^\tau x')_j - \max_{S_C(x)}(C^\tau x')} \right]$$

(where the minimum is taken over j such that: $j \in \overline{S_C}(x)$ and $(C^\tau x')_j - \max_{S_C(x)}(C^\tau x') \geq 0$).

Also, let

$$\Delta = \min \left[\min_{i \in \overline{S_R}(y)} [\max(Ry) - (Ry)_i], \min_{j \in \overline{S_C}(x)} [\max(C^\tau x) - (C^\tau x)_j] \right]$$

It is clear that we always have $\Delta > 0$ and that $\epsilon^* \geq \frac{\Delta}{1+\Delta}$.

Following the derivations in [11] and after several manipulations that were performed taking into account that the algorithm minimizes the difference with respect to $\epsilon \in (0, \epsilon^*]$ at every step, we finally get the following relationship for the new value of f (we drop the indices for notational simplicity) that is obtained at every step:

$$f_{new} - f \leq -\min \left[\frac{|V - f|^2}{4(1 - V)}, \left(\frac{\Delta}{1 + \Delta} \right)^2 f + \frac{\Delta}{(1 + \Delta)^2} |V - f| \right]$$

From the above relationship it is clear that there is always a decrease of the value of f unless $V - f = 0$. In the latter case we have a stationary point. Notice that a stationary point always exists as a limit point of the sequence of values of f produced by the descent algorithm. Such a sequence has always a limit since it is monotonically decreasing and bounded from below by 0. □

The descent algorithm essentially consists of an iterative loop containing the two basic steps as defined above. The execution of these steps continue until the stationarity condition $V(x, y) - f(x, y) = 0$ is satisfied. We call the common value of $V(x, y)$ and $f(x, y)$ at stationarity the **value** of the stationary point.

Schematically, the basic steps of the algorithm can be described as follows:

0. Start at an arbitrary pair of strategies
1. Apply **Descent Direction**
2. Apply **Line Search**
3. Stop if termination condition is satisfied. If not, return to step 1.

3 Convergence Rate

We provide an improved complexity result of the descent algorithm compared to the result presented in [1]. This is based on an estimate of a tighter bound on the number of loops required for convergence to a stationary point.

Fix $\delta > 0$ and let (x, y) be the current pair of strategies obtained during the descent procedure for which $f_R(x, y) = f_C(x, y)$. We define the following index sets:

$$S_R(y, \delta) = \{i \in (1, m) : (Ry)_i \geq \max(Ry) - \delta\}$$

$$S_C(x, \delta) = \{j \in (1, n) : (C^T x)_j \geq \max(C^T x) - \delta\}$$

Using the above index sets throughout the algorithm (for the computation of descent directions and all quantities associated with them), we define a δ -stationary point as a point that satisfies the relationship $V(x, y) - f(x, y) \geq -\delta$. Let $\overline{S}_R(y, \delta)$ and $\overline{S}_C(x, \delta)$ be the complements of these sets with respect to the index sets $(1, m)$ and $(1, n)$ respectively. Then, we have $\max(Ry) - (Ry)_i > \delta$ for all $i \in \overline{S}_R(y, \delta)$ and $\max(C^T x) - (C^T x)_j > \delta$ for all $j \in \overline{S}_C(x, \delta)$. Therefore, $\delta < \Delta$, where Δ is as defined in the proof of Lemma 2 above. The convergence rate result is expressed in Lemma 3 below.

Lemma 3. *The descent algorithm converges to a δ -stationary point in $O\left(\frac{1}{\delta} \log\left(\frac{1}{\delta}\right)\right)$ loops.*

Proof. Let us denote by b the value of f at a limit point of the descent algorithm. We should have $V \leq b \leq f$ at every loop of the algorithm (b cannot be more than 0.3393 and the initial value of f is $\leq \frac{1}{2}$). Using the last expression in the proof of Lemma 2 for the new value of f as a function of the previous value, it can be verified by direct calculation, that at every loop we obtain either a descent of the form $f_{new} - b \leq (f - b) - \frac{(f-b)^2}{4(1-b)}$, or a descent of the form $f_{new} - b \leq \left(1 - \frac{\Delta}{1+\Delta}\right)(f - b) - \frac{\Delta^2}{(1+\Delta)^2}b$. So we have two types of positive sequences $s_k, k = 1, 2, \dots$ starting from values $< \frac{1}{2}$ and converging to 0 according to the relationships $s_{k+1} \leq s_k - \frac{s_k^2}{4(1-b)}$ and $s_{k+1} \leq \left(1 - \frac{\Delta}{1+\Delta}\right)s_k - \frac{\Delta^2}{(1+\Delta)^2}b$. It can be verified that the first sequence requires at most $O\left(\frac{1}{\delta}\right)$ steps to reach a point s_k such that $s_k \leq \delta$ and the second sequence requires at most $O\left(\frac{1}{\Delta} \log\left(\frac{1}{\delta}\right)\right)$ steps to reach such a point. The result follows from these facts and from $\delta < \Delta$. \square

4 Evaluation Scenarios

For the evaluation of performance we have applied the algorithm to a large number of bi-matrix game problem instances with sizes ranging from 10×10 to 100×100 . The games that were generated were basically of two kinds: Those consisting of randomly generated real-valued payoff matrices R, C and those consisting of 0-1 (win-loose) matrices. The example matrices were generated so as to avoid the existence of easy solutions such as pure strategy equilibria, 2×2 equilibria and uniform distribution equilibria. Furthermore, the instances were selected

so as to avoid being close to constant sum games (for which every stationary point is a Nash equilibrium). All example matrices were positively normalized before running the algorithm. The 0-1 matrices R, C that were generated were of two kinds:

(a) Randomly generated matrices, under the constraint that no small support equilibria exist (pure and 2×2) and the games are not close to constant sum games.

(b) Specifically designed matrices satisfying the above constraints and also containing arbitrary cycles of consecutively interchanging assignments of ordered pairs $(0, 1)$ and $(1, 0)$ to the entries of $(R_{i,j}, C_{i,j})$ horizontally and perpendicularly (3). In particular, according to this pattern, we start from an arbitrary entry (i, j) and assign either $(0, 1)$ or $(1, 0)$ to it. Then, keeping i (or j) fixed, pick an arbitrary column (or row) and assign the opposite pair to it, i.e. $(1, 0)$ or $(0, 1)$. The next step is to move in the perpendicular direction, i.e. to keep j (or i) fixed and continue this process an arbitrary number of times to finally close the cycle. We can have several such cycles in the specification of the matrices R, C . It is expected that the presence of such cycles makes the problem of finding equilibria more difficult. Indeed, uniform distributions of either small support or large support (close to the dimension of the matrices) are not likely to be close to equilibria under such patterns of payoff assignments.

For each example, we used several different starting points in order to check the response of the algorithm to the variation of the starting distributions. In particular, we used the uniform distributions, as one starting point, but also several arbitrarily chosen pure strategies for both players. The algorithm was implemented in C and the CPLEX Program was used as an LP solver. The parameter δ was fixed to $\delta = 0.001$ for all instances.

5 Evaluation of the Results

The results obtained from all the runs that were performed generally indicate that the descent algorithm is a highly efficient and practical algorithm that converges fast to stationary points providing high quality approximate Nash equilibria. In fact, for almost all cases, the approximations achieved are less than the precision parameter δ , far better than the theoretical worst case approximation (0.3393). This experimental conclusion is typical for all categories of instances (as described above) upon which the algorithm was applied.

The worst approximation obtained across all experimental instances that we considered was 0.015 (notably, this happened for a small game). The reduction factor of the regret function f from the start to the end appears to be much larger than the one dictated by the theoretical worst case approximation. The curves showing the reduction of the regret function f with respect to the number of loops appear to exhibit a more or less similar pattern across all instances of all categories of experiments. Furthermore, for instances with relatively large

approximation (i.e. large values of f at stationarity but still less than the above-mentioned 0.015), it was sufficient to choose a different starting point (for example, an arbitrary pure strategy pair) for the algorithm to converge to an exact equilibrium (for the same instance).

Overall, it appears that if one moves along paths determined by the descent algorithm, it is very likely to hit a Nash equilibrium (or a stationary point close to a Nash equilibrium) along the way.

These experimental results indicate that the stationary points of bi-matrix games are rather unstable to the operation of the descent process and that this instability tends to increase rapidly with their value f . It is also possible that the stationary points with larger values are less probable (in fact significantly less probable) than the ones with smaller values. It is conjectured that with an appropriate definitions of stability of a stationary point, it may be possible to formulate a rigorous approach that could potentially lead to better approximation guarantees than the one currently available, including the possibility of obtaining a PTAS for the equilibrium problem. More specifically, it is conjectured that a modification of the descent algorithm to include restarts from new strategies obtained by small perturbations around a stationary point, could provide an effective way to bypass stationary points with high values.

Significant further insight into the problem will certainly be achieved if we can find harder instances. In this respect, it is worth investigating the existence and construction of instances for which the descent process could be more or less easily trapped into relatively high stationary points.

In regard to convergence rate, the experiments indicate that the number of loops required for convergence differ across starting points and categories but typically it was between 4 and 20 with a median value of 10. Also, the supports of the resulting (approximate equilibria) distributions were relatively large, typically between $n/3$ and $n/2$ (n being the size of the game).

Overall, the number of loops required for convergence to a δ -stationary point was found in all experiments to be much smaller than the worst case bound theoretically predicted in [11]. This motivated a closer look into the convergence properties of the descent algorithm which resulted in a new bound of the order of $O\left(\frac{1}{\delta} \log\left(\frac{1}{\delta}\right)\right)$ (formulated and proved here in Lemma 3), a significant improvement over the previous $O\left(\frac{1}{\delta^2}\right)$.

However, it appears that a better complexity bound of the descent algorithm is possible, in terms of speed of convergence to a stationary point. A possible way to obtain improved convergence results is to investigate the maximum number of small steps that the algorithm can go through as a function of the size of the game. It appears from the experiments that for each small step there is an increase of the size of the support (which often occurs in large chunks rather than one at a time), so, the total number of such small steps in a row cannot be more than a fraction of n . Also, it was observed that the large steps were too often large enough to enforce very fast convergence to a stationary point. Actually, only a small number of large steps were typically needed for convergence.

6 Discussion

The experimental results were surprising, particularly in regard to the quality of approximation to Nash equilibria. It seems that it is quite hard to create hard instances for the descent algorithm. We believe that the results motivate further investigation into the complexity of finding Nash equilibria along the following lines: (a) Study of the issue of stability of stationary points as a function of their value, (b) Investigation of ways to bypass stationary points via small perturbations around them, (c) Creation of hard instances for the descent algorithm, i.e. instances for which it is possible to get stuck to stationary points with large values, and (d) Some further investigation of the convergence rate of the algorithm for an improved bound.

References

1. Tsaknakis, H., Spirakis, P.G.: An Optimization Approach for Approximate Nash Equilibria. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 42–56. Springer, Heidelberg (2007)
2. Kontogiannis, S., Panagopoulou, P., Spirakis, P.: Polynomial algorithms for approximating nash equilibria in bimatrix games. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 286–296. Springer, Heidelberg (2006)
3. Kontogiannis, S., Spirakis, P.G.: Efficient algorithms for constant well supported approximate equilibria of bimatrix games. In: Arge, L., Cachin, C., Jurdziński, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596. Springer, Heidelberg (2007)
4. Chen, X., Deng, X.: Settling the complexity of 2-player nash equilibrium. In: Proc. of the 47th IEEE Symp. on Found. of Comp. Sci (FOCS 2006), pp. 261–272. IEEE Comp. Soc. Press, Los Alamitos (2006)
5. Daskalakis, C., Mehta, A., Papadimitriou, C.: Progress in approximate nash equilibrium. In: Proc. Of the 8th ACM Conf. on Electr. Commerce (EC 2007) (2007)
6. Daskalakis, C., Mehta, A., Papadimitriou, C.: A note on approximate nash equilibria. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 297–306. Springer, Heidelberg (2006)
7. Boss, H., Byrka, J., Markakis, E.: New Algorithms for Approximate Nash Equilibria in Bimatrix Games. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 17–29. Springer, Heidelberg (2007)

Worst-Case Nash Equilibria in Restricted Routing*

Pinyan Lu and Changyuan Yu

Institute for Theoretical Computer Science, Tsinghua University
Beijing, 100084, P.R. China
{lpy, yucy05}@mails.tsinghua.edu.cn

Abstract. We study a restricted related model of the network routing problem. There are m parallel links with possibly different speeds, between a source and a sink. And there are n users, and each user i has a traffic of weight w_i to assign to one of the links from a subset of all the links, named his/her allowable set. We analyze the *Price of Anarchy* (denoted by PoA) of the system, which is the ratio of the maximum delay in the worst-case Nash equilibrium and in an optimal solution. In order to better understand this model, we introduce a parameter λ for the system, and define an instance to be λ -good if for every user, there exist a link with speed at least $\frac{s_{max}}{\lambda}$ in his/her allowable set. In this paper, we prove that for λ -good instances, the Price of Anarchy is $\Theta\left(\min\left\{\frac{\log \lambda m}{\log \log \lambda m}, m\right\}\right)$. We also show an important application of our result in coordination mechanism design for task scheduling game. We propose a new coordination mechanism, *Group-Makespan*, for unrelated selfish task scheduling game. Our new mechanism ensures the existence of pure Nash equilibrium and its PoA is $O\left(\frac{\log^2 m}{\log \log m}\right)$. This result improves the best known result of $O(\log^2 m)$ by Azar, Jain and Mirrokni in [2].

1 Introduction

Network routing is one of the most important problems in the network management. In most networks, especially in a large-scale network like internet, it is unlikely that there is a centralized controller who can coordinate the behavior of all the users in the network. In such situations, every user in the network decides how to rout his/her traffic, aware of the congestion caused by other users. Users only care about the delay they suffer, and their selfish behavior often leads the whole network to a suboptimal state. Recently, researchers start to investigate the performance degradation due to the lack of the coordination for the users.

In the model first studied by Koutsoupias and Papadimitriou [9], there are m identical parallel links from the same origin to the same destination. There

* Supported by the National Natural Science Foundation of China Grant 60553001 and the National Basic Research Program of China Grant 2007CB807900, 2007CB807901.

are n users, and each with a traffic of weight w_i . We assume that the traffic of each user can not be split and as a result each user chooses exactly one link. After all the users choose their links, the delay of a link is equal to the total weight of the traffics on it, and the delay a user suffers is equal to the delay of the link he chooses. The performance of the system we consider here is the maximum delay of all the links. We are mainly interested in stable states, where no user can decrease his delay by unilaterally changing his choice. In game theory, such a state is also called a *Nash equilibrium*. In order to measure the performance degradation, they compared the performance of Nash equilibrium with the optimal solution when there is centralized coordination. In particular, we analyze the *Price of Anarchy* (PoA for short) of the system, which is defined to be the performance ratio between the worst-case Nash equilibrium and an optimal solution. In [9], Koutsoupias and Papadimitriou showed that the PoA of that system is at most $2 - 1/m$.

Since then, a lot of research works have been done along this line. There are mainly two generalized models of this problem which are well studied. One model is routing with related links, where different links may have different speeds and the delay of a link is equal to the total weight on this link over its speed. In this uniform related model, Czumaj and Vöcking proved that the PoA is $\Theta\left(\frac{\log m}{\log \log m}\right)$ [4]. The other model is routing with restricted links, where each user i is only allowed to choose links from a subset S_i of all the links. However the links are still identical in the sense that the speed of each link is the same. In this restricted model, Awerbuch et al. proved that the PoA is also $\Theta\left(\frac{\log m}{\log \log m}\right)$ [1].

In light of these results, one may conjecture that the common extension of these two models, where the links are both related and restricted, also has a PoA of $\Theta\left(\frac{\log m}{\log \log m}\right)$. In fact, this model was studied by Gairing et al. in [7], and they showed that the PoA of this problem can be as large as $m - 1$. However, in their bad instance demonstrating the lower bound of $m - 1$, some users can only use extremely slow links (with speed less than $\frac{s_{\max}}{(m-1)!}$, where s_{\max} is the largest speed). This is a little artificial and unlikely to appear in the real world. So in order to better understand this model, we introduce a property called λ -goodness for the system. An instance is called λ -good if and only if every user can at least use a link with speed no less than $\frac{s_{\max}}{\lambda}$. Now in our notation, the result in [7] says that the PoA can be as large as $m - 1$ when the system is only $(m - 1)!$ -good. So what is the exact relation between the PoA and the λ -goodness of a system? In this paper, we answer this question completely by giving a tight bound for the PoA of a λ -good system in term of λ .

Theorem 1. *For λ -good instances, the price of anarchy is $\Theta\left(\min\left\{\frac{\log \lambda m}{\log \log \lambda m}, m\right\}\right)$.*

In the proof of Czumaj and Vöcking for related links, they essentially used the property of uniform related, which means that each link has a fixed speed and all the users can choose it. And in the proof of Awerbuch et al. for restricted links, they essentially used the property of identical, which means that all the links have the same speed. In our extended model, namely restricted related

links, none of the two properties hold and as a result none of their technique can be adopted to analyze the PoA of the new model directly. In this paper, we use a new proof approach. We calculate the delay of links interval by interval, obtain some recursive relations between them based on the property of Nash equilibrium, and finally we are able to derive a bound of the maximum delay in the system.

Our result also has an important application in task scheduling game with coordination mechanism. Task scheduling can be viewed as another model for routing problem by treating the links as machines, the traffics as tasks, the delay of a user as the completion time of his/her task, and the delay of the system as the makespan of the system. Then we have scheduling with identical machines, related machines, and restricted machines corresponding to the above three models of routing problems. Further more, we also have a more general model, called scheduling with unrelated machines, in which each machine may have different speeds for different tasks. An instance of scheduling unrelated machines is denoted by a matrix $t = (t_{ij})$, where t_{ij} denotes the processing time that machine j needs for task i . In this language, when each machines uses the *Makespan* policy, i.e. to process its tasks in such a parallel way that all of them are completed at the same time, the task scheduling game is essentially the same as the routing problem. However, as observed by Christodoulou, Koutsoupias and Nanavati in [3], the scheduling policies of the machines may affect the choices of the users, and hence the PoA of the system. So they considered the problem of designing a set of local scheduling policies such that the PoA of the system is small. Such a set of scheduling policies are called coordination mechanism, and the PoA of the system with a coordination mechanism is also called the PoA of this mechanism.

Using our main result, we propose a new coordination mechanism, named *Group-Makespan* mechanism, for scheduling unrelated machines. This Group-Makespan mechanism ensures the existence of a pure Nash equilibrium and its PoA is $O(\frac{\log^2 m}{\log \log m})$, improving the best known result $O(\log^2 m)$ by Azar, Jain and Mirrokni in [2].

2 Preliminaries and Notations

In this section, we define our problem formally. There are m independent links from certain origin to destination, and n independent users. We use $[m]$ and $[n]$ to denote the link set $\{1, \dots, m\}$ and user set $\{1, \dots, n\}$ respectively. Each link $j \in [m]$ has a speed s_j and w.l.o.g, we assume $s_1 \geq s_2 \geq \dots \geq s_m$. Each user $i \in [n]$ has a traffic of weight w_i , which can only be assigned to a link from a set $S_i \subseteq [m]$. We use $\langle w, s, \mathcal{S} \rangle$ to denote an instance of the problem, where $w = (w_1, \dots, w_n)$, $s = (s_1, \dots, s_m)$ and $\mathcal{S} = \{S_1, \dots, S_n\}$ denote the weights, speeds and allowable link sets. We introduce the property of λ -goodness for a instance $\langle w, s, \mathcal{S} \rangle$.

Definition 1 (λ -Goodness). An instance $\langle w, s, \mathcal{S} \rangle$ is λ -good if and only if the following condition holds: for any user $i \in [n]$, there exists a machine $j \in S_i$ such that the speed s_j is at least s_1/λ .

We consider pure strategies for users, and each user’s strategy is to decide which link to assign his/her traffic. We use $a = (a_1, \dots, a_n) \in S_1 \times \dots \times S_n$ to denote a combination of all users’ strategies, where user i selects a link $a_i \in S_i$. We also use a_{-i} to denote the strategies of all the other users except user i . In a state a , the delay of link j , denoted by l_j^a , is the total weights on it over its speed, and the delay of the system, denoted by l^a , is the maximum delay over all the links. That is $l_j^a = \frac{1}{s_j} \sum_{i:a_i=j} w_i$, $l^a = \max_j l_j^a$.

We consider the optimum when there is centralized coordination, that is, the minimal delay of the system over all the possible states. We use opt to denote the optimum as well as an optimal solution.

We assume the users are all non-cooperative and each one wishes to minimize his/her own cost, without any regard to the performance of the system. The cost of user i in a state a is the delay of link a_i and we use c_i^a to denote it. We have $c_i^a = l_{a_i}^a$.

Now we define the Nash equilibria of the system formally.

Definition 2 (Nash Equilibrium). A state a is called a Nash Equilibrium (NE for short) of the system if and only if no user can decrease his/her cost by unilaterally changing a link. That is, for any user $i \in [n]$, any strategy $a'_i \in S_i$ and $a' = (a_{-i}, a'_i)$, we have $c_i^a \leq c_i^{a'}$.

For any instance of the problem, pure Nash Equilibrium always exists. The proof of this fact is using a quite common method with an elegant potential function, which is pointed out in several places(see [5] for example).

Theorem 2 (Existence of Nash Equilibrium). For $\lambda \geq 1$ and any λ -good instance $\langle w, s, \mathcal{S} \rangle$, there exists a Nash Equilibrium state a of it.

To compare the performance of Nash Equilibrium with the optimum, we give the definition of *Price of Anarchy*.

Definition 3 (Price of Anarchy). For instance of restricted routing problem, the Price of Anarchy (PoA for short) is defined as the performance ratio between the worst-case Nash equilibrium and the optimal solution. That is

$$PoA = \max_{\substack{a \in S_1 \times \dots \times S_n \\ a \text{ is a NE}}} \frac{l^a}{opt}.$$

And for any family of instances, its Price of Anarchy is defined to be the largest PoA among all its possible instances.

3 PoA of λ -Good Restricted Routing

In this section, we prove our main result Theorem 1. If $\lambda > (m-1)!$, Gairing et al. gave a tight bound $\Theta(m)$ [7]. So in this section, we always assume $\lambda \leq (m-1)!$

and prove that the PoA of the family of λ -good instances is $\Theta\left(\frac{\log \lambda m}{\log \log \lambda m}\right)$. We only give the proof for the upper bound and omit the tight example here.

Theorem 3 (Upper Bound). *Given any λ -good instance $\langle w, s, \mathcal{S} \rangle$ and a state $a \in S_1 \times \dots \times S_n$ which is a Nash equilibrium, delay of the system l^a is at most $opt \cdot O\left(\frac{\log \lambda m}{\log \log \lambda m}\right)$.*

For notational simplicity, we scale the speeds and weights such that $s_1 = 1$ and $opt = 1$. We also define several notations used in the proof. For any $k \in R^+$ and $j \in [m]$, let $W_j^k = \max\{l_j^a - k, 0\} \cdot s_j$ and $W^k = \sum_{j \in [m]} W_j^k$. Especially, we use $W_j = W_j^0$ to denote the total weight assigned to link j , and $W = W^0$ to denote the total weight of all the users. Fix an optimal solution opt , let O_j be the set of users assigned to link j in opt . We also define O_j^k to be the set of users who choose link j in opt and have costs at least k , that is, $O_j^k = \{i \in O_j, c_i^a \geq k\}$.

Our proof of the upper bound theorem comes from the following lemmas. In Lemma 1, we give a initial condition of W^k and this is the only point we use the condition that the instance is λ -good. Then Lemma 2 and Lemma 3 give recursive relations between W^k s, which basically says that W^k should increase significantly when k become small. So we can bound the total weight W from below in terms of makespan l^a and λ . On the other hand, the total weigh is bounded from above by m . Putting things together, we can bound l^a .

Lemma 1. *For any λ -good instance and any Nash equilibrium a , we have $W^{l^a-2} \geq \frac{1}{\lambda}$.*

Proof. Consider a link whose delay achieves l^a , say link j^* . Let i be a user on link j^* , and let link $j \in S_i$ has the maximum speed in S_i . Now if $j = j^*$, we have $l_j^a = l^a$. If $j \neq j^*$, since a is a Nash equilibrium, i cannot decrease his/her cost by changing from link j^* to link j . We have $l^a = c_i^a \leq l_j^a + \frac{w_i}{s_j}$.

As in the optimal solution, task i can only be assigned to a link from S_i , whose speed is at most s_j , we have $w_i/s_j \leq opt = 1$. Therefore, we have $l_j^a \geq l^a - 1$. So no matter whether $j = j^*$ or not, we have $l_j^a \geq l^a - 1$, hence

$$W^{l^a-2} \geq W_j^{l^a-2} \geq 1 \cdot s_j \geq \frac{1}{\lambda},$$

where the last inequality is because the instance is λ -good. □

Lemma 2. *For any Nash equilibrium a and $0 \leq k \leq l^a - 2$, we have $W^k \geq \frac{l^a}{l^a - (k+2)} W^{k+2}$.*

Proof. Firstly, we want to prove that $W_j^k \geq \sum_{i \in O_j^{k+2}} w_i$. If O_j^{k+2} is empty, then we are done. Otherwise, for any task $i \in O_j^{k+2}$, $c_i^a \geq k + 2$, by the definition of Nash equilibrium, we have

$$k + 2 \leq c_i^a \leq l_j^a + w_i/s_j \leq l_j^a + 1.$$

The last inequality is because that the task i is assigned to link j in opt . Therefore, $l_j^a \geq k + 1$ and $W_j^k \geq 1 \cdot s_j \geq \sum_{i \in O_j} w_i \geq \sum_{i \in O_j^{k+2}} w_i$. Noticing that $\bigcup_j O_j^k = \{i : c_i^a \geq k\}$, we can bound W^k as follows:

$$W^k = \sum_{j \in [m]} W_j^k \geq \sum_{j \in [m]} \sum_{i \in O_j^{k+2}} w_i = \sum_{i: c_i^a \geq k+2} w_i = \sum_{j: l_j^a \geq k+2} W_j \tag{1}$$

By the definition of W_j and W_j^{k+2} , for any j , $l_j^a > k + 2$, we have:

$$W_j = \frac{l_j^a}{l_j^a - (k + 2)} W_j^{k+2} \geq \frac{l^a}{l^a - (k + 2)} W_j^{k+2} \tag{2}$$

The last inequality is because the function $f(x) = \frac{x}{x - (k+2)}$ is monotone decreasing when $x > k + 2$ and for all j , we have $l_j^a \leq l^a$.

So from (1) and (2), we have:

$$W^k \geq \frac{l^a}{l^a - (k + 2)} \sum_{j: l_j^a > k+2} W_j^{k+2} = \frac{l^a}{l^a - (k + 2)} W^{k+2}$$

From lemma 1 and lemma 2, we have recursive relation about W^k and an initial condition. These ensures us to prove an upper bound on l^a , which is $O(\log \lambda m)$. There is a little gap between our expected bound. The reason is that in the above estimation in (2), we bounded all the l_j^a from above by l^a . This is a little weak since there cannot be too many links with large l_j^a . The following lemma uses a more careful estimation, and explores a recursive relation between W^k, W^{k+2} , and W^{k+4} , which helps us to obtain a better bound on l^a .

Lemma 3. *For any λ -good instance and any Nash equilibrium a , we have $W^k \geq \frac{k+6}{4}(W^{k+2} - 2W^{k+4})$.*

Proof. First, we omit some links in the summation of the last term in (1), then

$$W^k \geq \sum_{j: l_j^a > k+2} W_j \geq \sum_{j: k+6 \geq l_j^a > k+2} W_j.$$

Now, the estimation occurred in (2) can be more tight: for any link j such that $k + 6 \geq l_j^a > k + 2$, we have

$$W_j = \frac{l_j^a}{l_j^a - (k + 2)} W_j^{k+2} \geq \frac{k + 6}{k + 6 - (k + 2)} W_j^{k+2} = \frac{k + 6}{4} W_j^{k+2}$$

So, we can bound W^k as

$$W^k \geq \frac{k + 6}{4} \sum_{j: k+6 \geq l_j^a > k+2} W_j^{k+2} = \frac{k + 6}{4} \left(W^{k+2} - \sum_{j: l_j^a > k+2} W_j^{k+2} \right) \tag{3}$$

For $\forall j, l_j^a > k + 6$, we have

$$W_j^{k+2} = (l_j^a - (k + 2)) \cdot s_j \quad \text{and} \quad W_j^{k+6} = (l_j^a - (k + 6)) \cdot s_j,$$

Hence $2W_j^{k+4} = W_j^{k+2} + W_j^{k+6}$. Using this equality, we can bound the negative term in (3) as follows:

$$\sum_{j:l_j^a > k+6} W_j^{k+2} \leq \sum_{j:l_j^a > k+6} 2W_j^{k+4} \leq 2W^{k+4}.$$

Substituting this into (3), and we finish the proof. □

Putting things together, we have the proof of Theorem 3.

Proof of Theorem 3: Let $k_0 = \lfloor \frac{l^a}{6} \rfloor$. For any $k \geq l^a - 2k_0 \geq \frac{2l^a}{3}$, we have:

$$\begin{aligned} W^k &\geq \frac{k + 6}{4} \left(W^{k+2} - 2W^{k+4} \right) \\ &\geq \frac{k + 6}{4} \left(W^{k+2} - 2 \cdot \frac{l^a - (k + 4)}{l^a} W^{k+2} \right) \\ &= \frac{2(k + 4) - l^a}{4l^a} \cdot (k + 6)W^{k+2} \\ &\geq \frac{2(\frac{2l^a}{3} + 4) - l^a}{4l^a} \cdot (\frac{2l^a}{3} + 6)W^{k+2} \\ &\geq \frac{l^a}{18} W^{k+2} \end{aligned}$$

The first inequality is by lemma 3 and the second inequality is by lemma 2.

So using this recursive relation and lemma 1, we have:

$$W^{l^a - 2k_0} \geq W^{l^a - 2} \cdot \left(\frac{l^a}{18} \right)^{\frac{l^a}{6}} \geq \frac{1}{\lambda} \cdot \left(\frac{l^a}{18} \right)^{\frac{l^a}{6}}.$$

Since $\forall j, s_j \leq s_1 = 1$, and $opt = 1$, we have $W \leq opt \cdot \sum_j s_j \leq m$. By $W \geq W^{l^a - 2k_0}$, we have $\left(\frac{l^a}{18} \right)^{\frac{l^a}{6}} \leq \lambda m$. Since the solution to the equation $x^x = y$ is $x = \Theta\left(\frac{\log y}{\log \log y}\right)$, we can obtain that l^a is at most $O\left(\frac{\log \lambda m}{\log \log \lambda m}\right)$. □

4 An Application in Coordination Mechanism

In this section, we see an application in coordination mechanism design for selfish task scheduling game. We give the high level ideas of our new mechanism. This new mechanism is inspired by the mechanism *Split & Shortest* in [2]. Given an instance t for scheduling with unrelated machines, we can define $t_i = \min_{j \in [m]} t_{ij}$ as the weight of task i , and define the speed s_{ij} of a machine

j with respect to a task i as $s_{ij} = t_i/t_{ij}$, namely the minimum running time of task i on all the machines over the running time of task i on machine j . In our Group-Makespan mechanism, every machine simulates $\log m$ sub machines and submachine k of machine j only run those tasks i for which machine j has speed $s_{ij} \in [2^{-k}, 2^{-k+1})$. We artificially delay a task so that the k -th sub machines of different machines all have fixed speed 2^{-k} . Each machine simulates its sub machines by round-robin, and for each submachine we use the Makespan scheduling policy. In the submachine level, each submachine has a fixed speed, and a task can only be assigned to some of the sub machines. So it becomes a problem of scheduling with restricted related machines. Further more, all the instance obtained in this way have a very good structure, namely it is 1-good. Therefore in the submachine level, the PoA is bounded by $\Theta(\frac{\log m}{\log \log m})$. Since each machine has to simulate $\log m$ machines all the time, this may loss a factor of at most $\log m$.

We give the theorem as following, and omit the formal definition of the Group-Makespan mechanism and the proof of this theorem due to the space limitation. Readers may see [2] for idea of the submachine and related analysis.

Theorem 4. *The Group-Makespan mechanism for scheduling m unrelated machines ensures the existence of pure Nash equilibria, and the PoA of the task scheduling game with this mechanism is $O(\frac{\log^2 m}{\log \log m})$.*

References

1. Awerbuch, B., Azar, Y., Richter, Y., Tsur, D.: Tradeoffs in worst-case equilibria. *Theoretical Computer Science* 361, 200–209 (2006)
2. Azar, Y., Jain, K., Mirrokni, V.: (Almost) optimal coordination mechanisms for unrelated machine scheduling. In: *Proceedings of SODA 2008*, pp. 323–332 (2008)
3. Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination mechanisms. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) *ICALP 2004*. LNCS, vol. 3142, pp. 345–357. Springer, Heidelberg (2004)
4. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. In: *Proceedings of SODA 2002*, pp. 413–420 (2002)
5. Fotakis, D., Kontogiannis, S., Koutsoupias, E., Mavronicolas, M., Spirakis, P.: The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) *ICALP 2002*. LNCS, vol. 2380, pp. 123–134. Springer, Heidelberg (2002)
6. Gairing, M., Lucking, T., Mavronicolas, M., Monien, B.: Computing nash equilibria for scheduling on restricted parallel links. In: *Proceedings of STOC 2004*, pp. 613–622 (2004)
7. Gairing, M., Lucking, T., Mavronicolas, M., Monien, B.: The Price of Anarchy for Restricted Parallel Links. *Parallel Processing Letters* 16(1), 117–132 (2006)
8. Immorlica, N., Li, L., Mirrokni, V., Schulz, A.: Coordination mechanisms for selfish scheduling. In: Deng, X., Ye, Y. (eds.) *WINE 2005*. LNCS, vol. 3828, pp. 55–69. Springer, Heidelberg (2005)
9. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: *Proceedings of STACS 1999*, pp. 404–413 (1999)
10. Mavronicolas, M., Spirakis, P.: The Price of Selfish Routing. *Algorithmica* 48(1), 91–126 (2007)

Stackelberg Routing in Arbitrary Networks^{*}

Vincenzo Bonifaci^{1,2,**}, Tobias Harks^{3,***}, and Guido Schäfer^{3,†}

¹ Università degli Studi dell’Aquila, Italy

² Sapienza Università di Roma, Italy

bonifaci@dis.uniroma1.it

³ Technische Universität Berlin, Germany

{harks,schaefer}@math.tu-berlin.de

Abstract. We investigate the impact of *Stackelberg routing* to reduce the price of anarchy in network routing games. In this setting, an α fraction of the entire demand is first routed centrally according to a predefined *Stackelberg strategy* and the remaining demand is then routed selfishly by (nonatomic) players. Although several advances have been made recently in proving that Stackelberg routing can in fact significantly reduce the price of anarchy for certain network topologies, the central question of whether this holds true in general is still open. We answer this question negatively. We prove that the price of anarchy achievable via Stackelberg routing can be unbounded even for single-commodity networks. In light of this negative result, we consider bicriteria bounds. We develop an efficiently computable Stackelberg strategy that induces a flow whose cost is at most the cost of an optimal flow with respect to demands scaled by a factor of $1 + \sqrt{1 - \alpha}$. Finally, we analyze the effectiveness of an easy-to-implement Stackelberg strategy, called SCALE. We prove bounds for a general class of latency functions that includes polynomial latency functions as a special case. Our analysis is based on an approach which is simple, yet powerful enough to obtain (almost) tight bounds for SCALE in general networks.

1 Introduction

Over the past years, the impact of the behavior of selfish, uncoordinated users in congested networks has been investigated intensively in the theoretical computer science literature. In this context, *network routing games* have proved to be an appropriate means of modeling selfish behavior in networks. The basic idea is to model the interaction between the selfish network users as a *noncooperative game*. We are given a directed graph with latency functions on the arcs and a set of origin-destination pairs, called *commodities*. Every commodity has a *demand* associated with it, which specifies the amount of flow that needs to be sent from the respective origin to the destination.

^{*} This work was supported by the European Regional Development Fund (ERDF).

^{**} Research supported by European Commission project ARRIVAL FP6-021235-2. This work was done while the author was a postdoctoral fellow at TU Berlin.

^{***} Research supported by the Federal Ministry of Education and Research (BMBF grant 03MOPAI1).

[†] Research supported by the DFG Research Center MATHEON “Mathematics for key technologies”.

We assume that every demand represents a large population of players, each controlling an infinitesimal small amount of flow of the entire demand (such players are also called *nonatomic*). The latency that a player experiences to traverse an arc is given by a (non-decreasing) function of the total flow on that arc. We assume that every player acts selfishly and routes his flow along a minimum-latency path from its origin to the destination; this corresponds to a common solution concept for noncooperative games, that of a *Nash equilibrium* (here *Nash* or *Wardrop flow*). In a Nash flow no player can improve his own latency by unilaterally switching to another path.

It is well known that Nash equilibria can be *inefficient* in the sense that they need not achieve socially desirable objectives [217]. In the context of network routing games, a Nash flow in general does not minimize the total cost; or said differently, selfish behavior may cause a performance degradation in the network. Koutsoupias and Papadimitriou [13] initiated the investigation of the efficiency loss caused by selfish behavior. They introduced a measure to quantify the inefficiency of Nash equilibria which they termed the *price of anarchy*. The price of anarchy is defined as the worst-case ratio of the cost of a Nash equilibrium over the cost of a system optimum. In recent years, considerable progress has been made in quantifying the degradation in network performance caused by the selfish behavior of noncooperative network users. In a seminal work, Roughgarden and Tardos [21] showed that the price of anarchy for network routing games with nonatomic players and linear latency functions is $4/3$; in particular, this bound holds independently of the underlying network topology. The case of more general families of latency functions has been studied by Roughgarden [16] and Correa, Schulz, and Stier-Moses [3]. (For an overview of these results, we refer to the book by Roughgarden [19].) Despite these bounds for specific classes of latency functions, it is known that the price of anarchy for general latency functions is unbounded even on simple parallel-arc networks [21].

Due to this large efficiency loss, researchers have proposed different approaches to reduce the price of anarchy in network routing games. One of the most prominent approaches is the use of *Stackelberg routing* [12][18]. In this setting, it is assumed that a fraction $\alpha \in [0, 1]$ of the entire demand is controlled by a central authority, termed *Stackelberg leader*, while the remaining demand is controlled by the selfish nonatomic players, also called the *followers*. In a *Stackelberg game*, the Stackelberg leader first routes the centrally controlled flow according to a predetermined policy, called the *Stackelberg strategy*, and then the remaining demand is routed by the selfish followers. The aim is to devise Stackelberg strategies so as to minimize the price of anarchy of the resulting combined flow.

Although Roughgarden [18] showed that computing the *best* Stackelberg strategy, i.e., one that minimizes the price of anarchy of the induced flow, is NP-hard even for parallel-arc networks and linear latency functions, several advances have been made recently in proving that Stackelberg routing can indeed significantly reduce the price of anarchy in network routing games. As an example, Roughgarden [18] showed that for parallel-arc networks Stackelberg strategies exist that reduce the price of anarchy to $1/\alpha$, *independently* of the latency functions. That is, even if the Stackelberg leader controls only a small constant fraction of the overall demand, the price of anarchy reduces to a constant (while it is unbounded in the absence of any centralized control). More

recently, Swamy [23] obtained a similar result for single-commodity, series-parallel networks and Fotakis [8] for parallel-arc networks and unsplittable flows. Despite these positive results, a central question regarding the effectiveness of Stackelberg routing is still open: Does there always exist a Stackelberg strategy such that the price of anarchy is bounded? This question has been posed explicitly by Roughgarden [17, Open Problem 4].

Besides these efforts, researchers have also tried to characterize the effectiveness of easy-to-implement Stackelberg strategies for specific classes of latency functions. One of the simplest Stackelberg strategies is SCALE (see also [18]), which simply computes an optimal flow for the entire demand and then scales this flow by α . The currently best known bound for the price of anarchy induced by SCALE on multi-commodity networks and linear latency functions is due to Karakostas and Kolliopoulos [11]. More recently, Swamy [23] derived the first general bounds for polynomial latency functions.

Our Results. We investigate the impact of Stackelberg routing to reduce the price of anarchy in network routing games with nonatomic players. Our contribution is threefold:

1. We show that there are single-commodity networks for which every Stackelberg strategy induces a price of anarchy of at least $\Omega(n)$, where n is the number of nodes of the network. The result holds independently of the fraction $\alpha \in (0, 1)$ of the centrally controlled demand. This settles the open question raised by Roughgarden [17].
2. In light of this negative result, we investigate the effectiveness of Stackelberg routing strategies compared to an optimum flow for a larger demand; i.e., we consider bicriteria bounds. We develop an efficiently computable Stackelberg strategy inducing a flow whose cost is at most the cost of an optimal flow with respect to demands increased by a factor of $1 + \sqrt{1 - \alpha}$.
3. We give upper bounds on the efficiency of SCALE for a general class of latency functions which, among others, contains polynomial latency functions with non-negative coefficients. We also derive the first tight lower bounds for SCALE. Our bound is tight for concave latency functions; for higher degree polynomials our bounds are almost tight (though there remains a small gap for small values of α).

Significance and Techniques. Our first result settles an important open question regarding the applicability of Stackelberg routing in general networks. While most existing results show that the performance degradation due to the absence of central control is *independent* of the underlying network topology, our result shows that the network topology matters in the context of Stackelberg routing. Our negative result also carries over to the unsplittable flow setting. However, due to lack of space, we omit the details from this extended abstract.

One important application of Stackelberg routing is the routing of Internet traffic within the domain of an Internet service provider, see also Sharma and Williamson [22]. Here, the Internet service provider centrally controls a fraction of the overall traffic traversing its domain. In this setting, our second result provides the Internet service provider with an efficient algorithm to route the centrally controlled traffic. The performance of this routing algorithm is characterized by a smooth trade-off curve that scales between the absence of centralized control (doubling the demands is sufficient)

and completely centralized control (no scaling is necessary). Additionally, our result has a nice interpretation for the class of (practical relevant) $M/M/1$ -latency functions that model arc-capacities: In order to beat the cost of an optimal flow, it is sufficient to scale all arc capacities by $1 + \sqrt{1 - \alpha}$. Our bound is a natural generalization of the bicriteria bound by Roughgarden and Tardos [21] (see Correa et al. [4] for other related results).

We introduce a general approach, which we term λ -*approach*, to prove upper bounds on the price of anarchy of Stackelberg strategies for specific classes of latency functions. This approach is simple, yet powerful enough to obtain (almost) tight bounds for SCALE in general networks. For polynomial latency functions, our approach yields upper bounds that significantly improve the bounds by Swamy [23]. For linear latency functions, we derive an upper bound that coincides with a previous bound of Karakostas and Kolliopoulos in [11]. Their analysis is based on a (rather involved) machinery presented in [15]. However, our analysis is much simpler; in particular, we do not rely on the machinery in [15]. Moreover, we show that this bound also holds for concave latency functions. We present a generalized Braess instance that shows that for the linear case our bound is tight; a similar instance can be used to show that for higher degree polynomials our bounds are almost tight, leaving only a small gap for small values of α . We are confident that our λ -approach will prove useful to derive upper bounds on the price of anarchy also in other settings. For instance, the λ -approach can be applied to prove upper bounds when flows are unsplittable; details will be given in the full version of the paper. So far, such upper bounds for general networks are only known for linear latency functions (see Fotakis [8]).

Related Work. The idea of using Stackelberg strategies to improve the performance of a system was first proposed by Korilis, Lazar, and Orda [12]. The authors identified necessary and sufficient conditions for the existence of Stackelberg strategies that induce a system optimum; their model differs from the one discussed here. Roughgarden [18] first formulated the problem and model considered here. He also proposed some natural Stackelberg strategies such as SCALE and Largest-Latency-First (LLF). For parallel-arc networks he showed that the price of anarchy for LLF is bounded by $4/(3 + \alpha)$ and $1/\alpha$ for linear and arbitrary latency functions, respectively. Both bounds are tight. He also showed that for certain types of Stackelberg strategies, which he termed *weak* strategies (see Section 2 for a definition), the price of anarchy for multi-commodity networks can be unbounded [18]. However, this did not rule out the existence of effective Stackelberg strategies in general. Moreover, he also proved that it is NP-hard to compute the best Stackelberg strategy. Kumar and Marathe [14] investigated approximation schemes to compute the best Stackelberg strategy. The authors gave a PTAS for the case of parallel-arc networks.

Karakostas and Kolliopoulos [11] proved upper bounds on the price of anarchy for SCALE and LLF. Their bounds hold for arbitrary multi-commodity networks and linear latency functions. Their analysis is based on a result obtained by Perakis [15] to bound the price of anarchy for network routing games with asymmetric and non-separable latency functions. Furthermore, Karakostas and Kolliopoulos [11] showed that their analysis for SCALE is almost tight. More recently, Swamy [23] obtained upper bounds on the price of anarchy for SCALE and LLF for polynomial latency functions. Swamy

also proved a bound of $1 + 1/\alpha$ for single-commodity, series-parallel networks with arbitrary latency functions. Fotakis [8] studied LLF and a randomized version of SCALE for the case of unsplittable flows. He proved upper and lower bounds on the price of anarchy for linear latency functions. For parallel-arc networks, Fotakis proved that LLF still achieves an upper bound of $1/\alpha$ for arbitrary latency functions in this case.

Correa and Stier-Moses [5] proved, besides some other results, that the use of *opt-restricted strategies*, i.e., strategies in which the Stackelberg leader sends no more flow on every edge than the system optimum, does not increase the price of anarchy. Sharma and Williamson [22] considered the problem of determining the smallest value of α such that the price of anarchy can be improved. They obtained results for parallel-arc networks and linear latency functions. Kaporis and Spirakis [10] studied a related question of finding the minimum demand that the Stackelberg leader needs to control in order to enforce an optimal flow.

2 Model

In a network routing game we are given a directed network $G = (V, A)$ and k origin-destination pairs $(s_1, t_1), \dots, (s_k, t_k)$ called *commodities*. For every commodity $i \in [k]$, a demand $r_i > 0$ is given that specifies the amount of flow with origin s_i and destination t_i . Let \mathcal{P}_i be the set of all paths from s_i to t_i in G and let $\mathcal{P} = \cup_i \mathcal{P}_i$. A *flow* is a function $f : \mathcal{P} \rightarrow \mathbb{R}_+$. The flow f is *feasible* (with respect to r) if for all i , $\sum_{P \in \mathcal{P}_i} f_P = r_i$. For a given flow f , we define the flow on an arc $a \in A$ as $f_a = \sum_{P \ni a} f_P$. Moreover, each arc $a \in A$ has an associated variable *latency* denoted by $\ell_a(\cdot)$. For each $a \in A$ the latency function ℓ_a is assumed to be nonnegative, nondecreasing and differentiable. If not indicated otherwise, we also assume that ℓ_a is defined on $[0, \infty)$ and that $x\ell_a(x)$ is a convex function of x . Such functions are called *standard* [16]. The latency of a path P with respect to a flow f is defined as the sum of the latencies of the arcs in the path, denoted by $\ell_P(f) = \sum_{a \in P} \ell_a(f_a)$. The triple (G, r, ℓ) is called an *instance*. The *cost* of a flow f is $C(f) = \sum_{P \in \mathcal{P}} f_P \ell_P(f)$. Equivalently, $C(f) = \sum_{a \in A} f_a \ell_a(f_a)$. The feasible flow of minimum cost is called *optimal* and denoted by o . A feasible flow f is a *Nash flow*, or *selfish flow*, if for every $i \in [k]$ and $P, P' \in \mathcal{P}_i$ with $f_P > 0$, $\ell_P(f) \leq \ell_{P'}(f)$. In particular, if f is a Nash flow, all s_i - t_i paths to which f assigns a positive amount of flow have equal latency. It is well-known that if f_1 and f_2 are Nash flows for the same instance, then $C(f_1) = C(f_2)$, see e.g. [21].

In a Stackelberg network game we are given, in addition to G , r and ℓ , a parameter $\alpha \in (0, 1)$. A (*strong*) *Stackelberg strategy* is a flow g feasible with respect to $r' = (\alpha_1 r_1, \dots, \alpha_k r_k)$, for some $\alpha_1, \dots, \alpha_k \in [0, 1]$ such that $\sum_{i=1}^k \alpha_i r_i = \alpha \sum_{i=1}^k r_i$. If $\alpha_i = \alpha$ for all i , g is called a *weak Stackelberg strategy*. Thus, both strong and weak strategies route a fraction α of the overall traffic, but a strong strategy can choose how much flow of each commodity is centrally controlled. For single-commodity networks the two definitions coincide. A Stackelberg strategy g is called *opt-restricted* if $g_a \leq o_a$ for all $a \in A$. Given a Stackelberg strategy g , let $\tilde{\ell}_a(x) = \ell_a(g_a + x)$ for all $a \in A$ and let $\tilde{r} = r - r'$. Then a flow h is *induced by* g if it is a Nash flow for the instance $(G, \tilde{r}, \tilde{\ell})$.

The Nash flow h can be characterized by the following *variational inequality* [6]: h is a Nash flow induced by g if and only if for all flows x feasible with respect to \tilde{r} ,

$$\sum_{a \in A} h_a \ell_a(g_a + h_a) \leq \sum_{a \in A} x_a \ell_a(g_a + h_a). \tag{1}$$

We will mainly be concerned with the cost of the combined induced flow $g + h$, given by $C(g + h) = \sum_{a \in A} (g_a + h_a) \ell_a(g_a + h_a)$. In particular, we are interested in bounding the ratio $C(g + h)/C(o)$, called the *price of anarchy*.

Due to lack of space, we omit some of the proofs from this extended abstract; details will be given in the full version of the paper.

3 Limits of Stackelberg Routing

In this section, we prove that there does not exist a Stackelberg strategy that induces a price of anarchy bounded by a function of α only. More precisely, we show that for any fixed $\alpha \in (0, 1)$, the ratio between the cost of the flow induced by any Stackelberg strategy and the optimum can be arbitrarily large, even in single-commodity networks.

Theorem 1. *Let $M > 0$ and $\alpha \in (0, 1)$. Then, there exists a single-commodity instance $\mathcal{I} = (G, r, \ell, \alpha)$ such that, if g is any Stackelberg strategy for \mathcal{I} inducing a Nash flow h , and o is an optimal flow for the instance (G, r, ℓ) , then $C(g + h) \geq M \cdot C(o)$.*

To prove the theorem we use the instance $G_k = (V_k, A_k)$ depicted in Figure 1. For a positive integer k , the graph G_k has $4k + 4$ nodes. There is a single commodity (s, t) , with unit demand. Define $r_0 := (1 - \alpha)/2$ and $r_1 := (1 + \alpha)/2k$. Note that the total demand is equal to $r_0 + kr_1$. Every arc is of one of five different types $\{A, B, C, D, E\}$ as indicated in Figure 1. The latency of an arc is determined by its type. Type B arcs have constant latency 1, and type C arcs have constant latency 0. Arcs of type A have the following latency function:

$$\ell_0(x) = \begin{cases} 0, & \text{if } x \leq r_0 \\ 1 - \frac{r_0 + r_1 - x}{r_1}, & \text{if } x > r_0. \end{cases}$$

Although $\ell_0(x)$ is not differentiable in r_0 , it can be approximated with arbitrarily small error by standard functions.




Fig. 1. The graph G_k , used in the proof of Theorem 1. Arcs are labeled with their type.

For fixed $L, \tau > 0$, let $u_{L,\tau}(x)$ be any standard function satisfying $u_{L,\tau}(L) = 0$ and $u_{L,\tau}(L + \tau) = M/\tau$. Type D arcs have latency $u_{r_0, \delta/3k^3}(x)$, and type E arcs have latency $u_{r_1, \delta/3k^3}(x)$. We will fix the constant δ later in the proof.

Lemma 1. $C(o) \leq 1$.

Proof. Let P_0 be the path $(s, s_0, p_1, q_1, p_2, \dots, p_k, q_k, t_0, t)$, and for $i \in [k]$, let P_i be the path (s, s_i, t_i, t) . Consider the feasible flow f such that $f_{P_0} = r_0$ and $f_{P_i} = r_1$ for $i \in [k]$. The latency induced by f is 0 on arcs of type A, C, D, E and 1 on arcs of type B. So $C(o) \leq C(f) = k \cdot r_1 = (1 + \alpha)/2 \leq 1$. \square

The following lemma will allow us to focus on the case where the combined flow on type D and E arcs does not exceed a certain threshold value.

Lemma 2. For any Stackelberg strategy g inducing a Nash flow h , the following hold:

- (i) If a is a type D arc and $g_a + h_a \geq r_0 + \delta/3k^3$, then $C(g + h) \geq M \cdot C(o)$.
- (ii) If a is a type E arc and $g_a + h_a \geq r_1 + \delta/3k^3$, then $C(g + h) \geq M \cdot C(o)$.

Proof. We prove statement (i); the proof for (ii) is similar. We have $C(g + h) \geq (g_a + h_a) \cdot \ell_a(g_a + h_a) = (g_a + h_a) \cdot u_{r_0, \delta/3k^3}(g_a + h_a) \geq (r_0 + \delta/3k^3) \cdot M/(\delta/3k^3) \geq M$. The proof follows from Lemma 1. \square

For the remainder of the proof we assume that there is no arc satisfying the conditions of Lemma 2; otherwise the theorem follows immediately.

Lemma 3. For any Stackelberg strategy g inducing a Nash flow h , the following hold:

- (i) For any arc $a = (q_{i-1}, p_i)$, $i \in [k]$, $g_a + h_a \geq r_0 - \delta/k$.
- (ii) For any arc $a = (s, s_i)$, $i \in [k]$, $g_a + h_a \geq r_1 - \delta/k$.

We are now ready to conclude the proof of Theorem 1. \square

Proof (Theorem 1). For any $i \in [k]$, consider the i th block in the graph (Figure 2). Let g_i, h_i be the Stackelberg and selfish flow on the arc (s_i, t_i) , respectively. We have two cases:

1. $h_i = 0$: in this case, using Lemma 3, the flow on arc (p_i, q_i) is at least $r_0 - \delta/k + r_1 - \delta/k - g_i$. The latency on that same arc is thus at least $\ell_0(r_0 + r_1 - 2\delta/k - g_i)$.
2. $h_i > 0$: in this case, the Nash flow on path $P'_i = (s, s_i, t_i, t)$ is strictly positive. Consider the path $P''_i = (s, s_i, p_i, q_i, t_i, t)$. By definition of a Nash flow, we get $\ell_{P''_i}(g + h) \geq \ell_{P'_i}(g + h)$. Notice that the two paths P'_i, P''_i share all their nonzero-latency




Fig. 2. The i th block of the graph G_k

arcs except for (s_i, t_i) (only present in P'_i) and (p_i, q_i) (only present in P''_i). Thus $\ell_{P'_i}(g+h) \geq \ell_{P''_i}(g+h)$ implies $\ell_{(p_i, q_i)}(g+h) \geq \ell_{(s_i, t_i)}(g+h) = 1$. As a consequence, $\ell_{(p_i, q_i)}(g+h) \geq 1 = \ell_0(r_0 + r_1) \geq \ell_0(r_0 + r_1 - 2\delta/k - g_i)$ since g_i and δ/k are non-negative.

In both cases, $\ell_{(p_i, q_i)}(g+h) \geq \ell_0(r_0 + r_1 - 2\delta/k - g_i) \geq 1 - \frac{g_i + 2\delta/k}{r_1}$.

The latency $\ell_{P_0}(g+h)$ on the path $P_0 = (s, s_0, p_1, q_1, \dots, p_k, q_k, t_0, t)$ is at least

$$\sum_{i=1}^k \ell_{(p_i, q_i)}(g+h) \geq \sum_{i=1}^k \left(1 - \frac{g_i + 2\delta/k}{r_1}\right) \geq k - \frac{\alpha}{r_1} - \frac{2\delta}{r_1} = \left(\frac{1 - \alpha - 4\delta}{1 + \alpha}\right) k.$$

The last inequality is a consequence of the fact that the total Stackelberg flow is α , so $\sum_i g_i \leq \alpha$.

Choosing $\delta < (1 - \alpha)/4$, we can conclude that $\ell_{P_0}(g+h) = \Omega(k)$. Together with Lemma 1 and Lemma 3, this gives

$$C(g+h) \geq (r_0 - \delta/k) \cdot \ell_{P_0}(g+h) \geq \left(\frac{1}{2} \cdot (1 - \alpha) - \delta\right) \cdot \Omega(k) = \Omega(k) \cdot C(o).$$

Thus, $C(g+h)/C(o)$ can be made arbitrarily large by picking a sufficiently large k . \square

Remark 1. Suppose the Stackelberg leader (e.g., a navigation systems provider) is solely interested in minimizing the travel time of his players (customers), i.e., $C_1(g+h) = \sum_{a \in A} g_a \ell_a(g_a + h_a)$. Our result also implies that even the ratio $C_1(g+h)/C(o)$ can be unbounded, independent of the Stackelberg strategy g .

4 A Bicriteria Bound for General Latency Functions

As we have seen in the previous section, no Stackelberg strategy controlling a constant fraction of the traffic can reduce the price of anarchy to a constant, even if we consider single-commodity networks. In light of this negative result, we therefore compare the cost of a Stackelberg strategy on an instance $\mathcal{I} = (G, r, \ell, \alpha)$ to the cost of an optimal flow for the instance $\mathcal{I}^\beta = (G, \beta r, \ell)$ in which the demand vector has been scaled up by a factor $\beta > 1$.

We propose the following simple Stackelberg strategy, which we term *Augmented SCALE (ASCALE)*:

1. Compute an optimal flow o^β for the instance \mathcal{I}^β .
2. Define the Stackelberg flow by $g := \frac{\alpha}{\beta} o^\beta$.

We prove that the resulting flow induced by the Stackelberg strategy ASCALE satisfies $C(g+h) \leq C(o^\beta)$ if we choose $\beta = 1 + \sqrt{1 - \alpha}$. This result can be seen as a generalization of the result by Roughgarden and Tardos that the cost of a Nash flow is always less than or equal to the cost of the optimal flow for an instance in which demands have been doubled [21]. Our bound gives a smooth transition from absence of centralized control (where doubling the demands is sufficient) to completely centralized control (where no augmentation is necessary).

Theorem 2. If g is the ASCALE strategy, $C(g+h) \leq \frac{1}{\beta-1} \cdot (1 - \frac{\alpha}{\beta}) \cdot C(o^\beta)$. Furthermore, this bound is tight.

Corollary 1. Let $\beta = 1 + \sqrt{1 - \alpha}$. If g is the ASCALE strategy, then $C(g+h) \leq C(o^\beta)$.

For a given instance $\mathcal{I} = (G, r, \ell, \alpha)$, the SCALE strategy is defined as $g = \alpha o$, where o is an optimal flow for (G, r, ℓ) . The next theorem shows that our result for ASCALE has a consequence for the SCALE strategy as well.

Theorem 3. Let $g = \alpha o$ be the SCALE strategy for instance $\mathcal{I} = (G, r, \ell, \alpha)$. Define a modified instance $\hat{\mathcal{I}} = (G, r, \hat{\ell}, \alpha)$ with latency functions $\hat{\ell}_a(x) = \ell_a(x/\beta)/\beta$ for every arc a , where $\beta = 1 + \sqrt{1 - \alpha}$, and let $\hat{C}(\cdot)$ denote the cost of a flow with respect $\hat{\ell}$. Let \hat{h} be the Nash flow induced by $\hat{g} = g$ in $\hat{\mathcal{I}}$. Then, $\hat{C}(\hat{g} + \hat{h}) \leq C(o)$.

5 Bounds for Specific Classes of Latency Functions

In this section, we first present a general approach, which we call λ -approach, to analyze the price of anarchy of opt-restricted Stackelberg strategies. We then use the λ -approach to derive bounds on the price of anarchy of the SCALE strategy for a general class of latency functions, including polynomial latency functions with nonnegative coefficients.

λ -Approach. We start by proving an upper bound on the cost of the combined flow induced by an opt-restricted Stackelberg strategy.

Lemma 4. For any opt-restricted strategy g , $C(g+h) \leq \sum_{a \in A} o_a \ell_a(g_a + h_a)$.

Proof. The proof follows immediately by applying the variational inequality (II) with $x = o - g$. \square

For any latency function ℓ_a and nonnegative numbers g_a, λ , we define the following nonnegative value:

$$\omega(\ell_a; g_a, \lambda) := \sup_{o_a, h_a \geq 0} \frac{o_a}{g_a + h_a} \cdot \frac{\ell_a(g_a + h_a) - \lambda \ell_a(o_a)}{\ell_a(g_a + h_a)}. \quad (2)$$

(We assume by convention $0/0 = 0$.) In order to bound the price of anarchy, we use the variational inequality (Lemma 4) and bound the cost of the induced flow on every arc by some λ -fraction of the optimal cost plus some ω -fraction of the cost of the induced flow itself:

$$C(g+h) \leq \sum_{a \in A} \lambda \cdot o_a \ell_a(o_a) + \omega(\ell_a; g_a, \lambda) \cdot (g_a + h_a) \ell_a(g_a + h_a). \quad (3)$$

Now, the idea is to determine a λ that provides the tightest bound possible. Choosing $\lambda = 1$, the above approach resembles the one that was previously used by Correa, Schulz, and Stier-Moses [3] to bound the price of anarchy of network routing games; however, optimizing over the parameter λ provides an additional means to obtain better bounds. The idea of introducing the scaling parameter λ was first introduced in

the context of bounding the price of anarchy in atomic splittable network games (see Harks [9]).

For a given opt-restricted strategy g we further define $\omega(g, \lambda) = \max_{a \in A} \omega(\ell_a; g_a, \lambda)$. Before we state the main theorem, we need one additional definition. Given an opt-restricted strategy g , the *feasible λ -region* is defined as $\Lambda(g) := \{\lambda \in \mathbb{R}_+ \mid \omega(g, \lambda) < 1\}$. Notice that every $\lambda \in \Lambda(g)$ induces a bound on the price of anarchy.

Theorem 4. *Let $\lambda \in \Lambda(g)$. Then $C(g+h) \leq \frac{\lambda}{1-\omega(g,\lambda)}C(o)$.*

Proof. The proof follows immediately from [3], Lemma 4 and the definition of ω . \square

Bounds for SCALE. In the following, we will analyze the SCALE strategy defined by $g = \alpha o$. Let \mathcal{L}_d , $d \geq 1$, be a class of continuous, nondecreasing, and standard latency functions satisfying $\ell(cz) \geq c^d \ell(z)$ for all $c \in [0, 1]$. \mathcal{L}_d contains, among others, polynomials with nonnegative coefficients and degree at most d . This characterization has been used before by Correa et al. [3].

Lemma 5. *Assume $\lambda \in [0, 1]$ and latency functions in \mathcal{L}_d . Then, we have*

$$\omega(\alpha o, \lambda) \leq \max \left\{ \frac{1}{\alpha}(1-\lambda), \frac{d}{d+1} \cdot \frac{1}{((d+1)\lambda)^{1/d}} \right\}.$$

Proof. By the definition of $\omega = \omega(\ell_a; \alpha o_a, \lambda)$:

$$\omega = \sup_{o_a, h_a \geq 0} \frac{o_a}{\alpha o_a + h_a} \cdot \frac{\ell_a(\alpha o_a + h_a) - \lambda \ell_a(o_a)}{\ell_a(\alpha o_a + h_a)}.$$

We consider two cases: (i) $\alpha o_a + h_a \geq o_a$. Define $\mu := \frac{o_a}{\alpha o_a + h_a} \in [0, 1]$. We have

$$\begin{aligned} \omega &= \sup_{o_a, h_a \geq 0, \mu \in [0, 1]} \mu \cdot \frac{\ell_a(\alpha o_a + h_a) - \lambda \ell_a(\mu(\alpha o_a + h_a))}{\ell_a(\alpha o_a + h_a)} \\ &\leq \max_{\mu \in [0, 1]} \mu (1 - \lambda \mu^d) = \frac{d}{d+1} \cdot \frac{1}{((d+1)\lambda)^{1/d}}. \end{aligned}$$

where the last inequality follows from the definition of \mathcal{L}_d . The second case (ii) $\alpha o_a + h_a \leq o_a$ leads to

$$\begin{aligned} \omega &\leq \sup_{o_a, h_a \geq 0} \frac{o_a}{\alpha o_a + h_a} \cdot \frac{\ell_a(\alpha o_a + h_a) - \lambda \ell_a(\alpha o_a + h_a)}{\ell_a(\alpha o_a + h_a)} \\ &\leq \sup_{o_a, h_a \geq 0} \frac{o_a}{\alpha o_a + h_a} (1 - \lambda) \leq \frac{1}{\alpha}(1 - \lambda), \end{aligned}$$

where the first inequality is valid since latencies are nondecreasing. \square

Lemma 6. *There is a unique $\lambda \in (0, 1)$, call it λ_d , such that $\frac{1}{\alpha}(1 - \lambda) = \frac{d}{d+1} \cdot \frac{1}{((d+1)\lambda)^{1/d}}$. Then: $\lambda_d = z_d^d / (d+1)$, where $z_d \geq 1$ is the unique solution to the equation $z^{d+1} - (d+1)z + \alpha d = 0$.*

Proof. Substituting $\lambda = z_d^d / (d+1)$ in the starting equation and rewriting yields $z^{d+1} - (d+1)z + \alpha d = 0$. To verify that this equation has indeed exactly one solution larger than 1, use for example Descartes' rule of signs. \square

Theorem 5. *The price of anarchy of the SCALE strategy for latency functions in the class \mathcal{L}_d is at most*

$$\frac{(d+1)z_d - \alpha d}{(d+1)z_d - d},$$

where $z_d \geq 1$ is the unique solution of the equation $z^{d+1} - (d+1)z + \alpha d = 0$.

Proof. We will use Theorem 4 with $\lambda = \lambda_d$. However, in order to apply the theorem, we first need to upper bound $\omega(\alpha, \lambda_d)$. Using Lemma 5 and Lemma 6, we know that

$$\omega(\alpha, \lambda_d) \leq \frac{d}{d+1} \cdot ((d+1)\lambda_d)^{-1/d} = \frac{d}{d+1} \cdot z_d^{-1} < 1.$$

This implies $\lambda_d \in \Lambda(\alpha)$ and we can invoke Theorem 4 to obtain a bound on the price of anarchy given by

$$\frac{\lambda_d}{1 - \omega(\alpha, \lambda_d)} \leq \frac{z_d^d / (d+1)}{1 - \frac{d}{d+1} z_d^{-1}} = \frac{z_d^{d+1}}{(d+1)z_d - d} = \frac{(d+1)z_d - \alpha d}{(d+1)z_d - d}. \quad \square$$

The bound thus obtained gives an improvement with respect to the previously best bounds obtained by Swamy [23].

For the class of \mathcal{L}_1 latency functions, which, in particular, contains continuous, non-decreasing, standard, and concave latencies, the above theorem reads as stated in Corollary 2 below. The same bound has been proven by Karakostas and Kolliopoulos [11] for the special case of affine latencies.

Corollary 2. *The price of anarchy of the SCALE strategy for latency functions in \mathcal{L}_1 is at most $((1 + \sqrt{1 - \alpha})^2) / (2(1 + \sqrt{1 - \alpha}) - 1)$.*

A lower bound for polynomial latency functions of degree d can be obtained by considering generalized Braess graphs [120] (details omitted).

Theorem 6. *Let $n \geq 2$ be an integer and let $c = (1 - (n-1)\alpha/n)^d$. Then, the price of anarchy of the SCALE strategy for latency functions in the class \mathcal{L}_d is at least $(nc^{1+1/d} + (n-1)\alpha c) / ((n-1)c + n^{-d})$.*

Note that the theorem does not fix n , so it is possible to optimize n based on α . For functions in \mathcal{L}_1 the stated lower bound pointwise matches the upper bound of Corollary 2 for infinitely many values of α . More precisely, the upper bound is matched for all values of α such that $1/\sqrt{1-\alpha}$ is an integer. To the best of our knowledge, this is the first tight bound for values of $\alpha \neq 0, 1$.

References

1. Babaioff, M., Kleinberg, R., Papadimitriou, C.H.: Congestion games with malicious players. In: Proc. of the 8th ACM Conf. on Electronic Commerce, pp. 103–112 (2007)
2. Braess, D.: Über ein Paradoxon der Verkehrsplanung. *Unternehmensforschung* 11, 258–268 (1968)

3. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: Selfish routing in capacitated networks. *Mathematics of Operations Research* 29, 961–976 (2004)
4. Correa, J.R., Schulz, A.S., Stier-Moses, N.E.: On the inefficiency of equilibria in congestion games. In: *Proc. of the 11th Int. Conf. on Integer Programming and Combinatorial Optimization*, pp. 167–181. Springer, Heidelberg (2005)
5. Correa, J.R., Stier-Moses, N.E.: Stackelberg routing in atomic network games. Technical report, Columbia Business School (February 2007)
6. Dafermos, S., Sparrow, F.: The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards, Series B* 73, 91–118 (1969)
7. Dubey, P.: Inefficiency of Nash equilibria. *Mathematics of Operations Research* 11, 1–8 (1986)
8. Fotakis, D.: Stackelberg strategies for atomic congestion games. In: *Proc. of the 15th European Symp. on Algorithms*, pp. 299–310. Springer, Heidelberg (2007)
9. Harks, T.: Stackelberg strategies and collusion in network games with splittable flow. In: *Proc. of the 6th Workshop on Approximation and Online Algorithms (WAOA)*. Springer, Heidelberg (2008)
10. Kaporis, A., Spirakis, P.: The price of optimum in Stackelberg games on arbitrary single commodity networks and latency functions. In: *Proc. of the 18th ACM Symp. on Parallelism in Algorithms and Architectures*, pp. 19–28. ACM Press, New York (2006)
11. Karakostas, G., Kolliopoulos, S.G.: Stackelberg strategies for selfish routing in general multicommodity networks. *Algorithmica* (to appear, 2007)
12. Korilis, Y.A., Lazar, A.A., Orda, A.: Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Transactions on Networking* 5(1), 161–173 (1997)
13. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: Meines, C., Tison, S. (eds.) *Proc. of the 16th Symp. on Theoretical Aspects of Computer Science*, pp. 404–413. Springer, Heidelberg (1999)
14. Kumar, V.S.A., Marathe, M.V.: Improved results for Stackelberg scheduling strategies. In: *Proc. of the 33rd Int. Colloquium of Automata, Languages and Programming*, pp. 776–787. Springer, Heidelberg (2002)
15. Perakis, G.: The price of anarchy when costs are non-separable and asymmetric. In: *Proc. of the 10th Int. Conference on Integer Programming and Combinatorial Optimization*, pp. 46–58. Springer, Heidelberg (2004)
16. Roughgarden, T.: The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences* 67, 341–364 (2002)
17. Roughgarden, T.: *Selfish Routing*. PhD thesis, Cornell University (2002)
18. Roughgarden, T.: Stackelberg scheduling strategies. *SIAM Journal on Computing* 33(2), 332–350 (2004)
19. Roughgarden, T.: *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge (2005)
20. Roughgarden, T.: On the severity of Braess’s paradox: Designing networks for selfish users is hard. *Journal of Computer and System Sciences* 72(5), 922–953 (2006)
21. Roughgarden, T., Tardos, E.: How bad is selfish routing? *Journal of the ACM* 49(2), 236–259 (2002)
22. Sharma, Y., Williamson, D.: Stackelberg thresholds in network routing games or the value of altruism. In: *Proc. of the 8th ACM Conf. on Electronic Commerce*, pp. 93–102 (2007)
23. Swamy, C.: The effectiveness of Stackelberg strategies and tolls for network congestion games. In: *Proc. of the 18th ACM-SIAM Symp. on Discrete Algorithms*, pp. 1133–1142. SIAM, Philadelphia (2007)

Computational Aspects of a 2-Player Stackelberg Shortest Paths Tree Game^{*}

Davide Bilò¹, Luciano Gualà², Guido Proietti^{3,4}, and Peter Widmayer¹

¹ Institut für Theoretische Informatik, ETH, 8092 Zürich, Switzerland

² Dipartimento di Matematica, Università di Tor Vergata, 00133 Roma, Italy

³ Dipartimento di Informatica, Università di L'Aquila, 67010 L'Aquila, Italy

⁴ Istituto di Analisi dei Sistemi ed Informatica, CNR, 00185 Roma, Italy
{guala,proietti}@di.univaq.it, {dbilo,widmayer}@inf.ethz.ch

Abstract. Let a communication network be modelled by a directed graph $G = (V, E)$ of n nodes and m edges. We consider a one-round two-player network pricing game, the *Stackelberg Shortest Paths Tree* (STACKSPT) game. This is played on G , by assuming that edges in E are partitioned into two sets: a set E_F of edges with a fixed positive real weight, and a set E_P of edges that should be priced by one of the two players (the *leader*). Given a distinguished node $r \in V$, the STACKSPT game is then as follows: the leader prices the edges in E_P in such a way that he will maximize his *revenue*, knowing that the other player (the *follower*) will build a shortest paths tree of G rooted at r , say $S(r)$, by running a publicly available algorithm. Quite naturally, for each edge selected in the solution, the leader's revenue is assumed to be equal to the *loaded price* of an edge, namely the product of the edge price times the number of paths from r in $S(r)$ that use it. First, we show that the problem of maximizing the leader's revenue is NP-hard as soon as $|E_P| = \Theta(n)$. Then, in search of an effective method for solving the problem when the size of E_P is constant, we focus on the basic case in which $|E_P| = 2$, and we provide an efficient $O(n^2 \log n)$ time algorithm. Afterwards, we generalize the approach to the case $|E_P| = k$, and we show that it can be solved in polynomial time whenever $k = O(1)$.

Keywords: Communication Networks, Shortest Paths Tree, Stackelberg Games, Network Pricing Games.

1 Introduction

From a game theoretic point of view, the Internet is usually seen as an economic system in which a multitude of firms (i.e., the final users) compete on a wide range of applications involving the allocation of the network resources. Quite naturally, the firms are assumed to share the very same market quote, and then the economic model which usually captures the strategic side of the system is that of a *Nash game*, in which all firms simultaneously compete against each other. However, as soon as a firm has a leadership position (which often happens

^{*} Part of this work has been developed while the third author was visiting ETH.

in an oligopoly of service providers like the Internet tendentially is), it could in practice manipulate the market to its own advantage. This setting requires a different economic model, so as to be able to discriminate between the *leader* (i.e., the dominating subject), and the *followers* (i.e., the residual players).

Leader-follower games were introduced to model heterogeneous markets by von Stackelberg as early as 1934 [11]. In the basic formulation, the game is played in a single round by only two players: the leader which moves first, and the follower which observes the leader's move and then makes its own move, after which the game is over. In our setting, the follower computes a solution by optimizing an objective (public) function, while the leader has its own objective, based on the solution used by the follower. Therefore, the strategic aspect of the game consists of the fact that the leader can anticipate the response of the follower and take it into account when optimizing his own objective.

1.1 Stackelberg Network Games

Despite their importance in economics, leader-follower games have not yet been extensively studied from a computational complexity point of view. Instead, the so called *network games* have been mainly studied under the Nash model which served as a paradigm in most of the developments of *algorithmic game theory* [8]. In these games, classic combinatorial optimization problems on networks (e.g., network design problems, flow networks, etc.) are revised in a strategic setting in which users own physical network components, or selfishly decide which part of the network must be used for their personal purposes. In this context, the potentially large size of an instance of a network game revealed itself as a boosting motivation for focusing on its computational aspects. However, network games can be easily regarded as Stackelberg games, as soon as a situation arises in which a subset of the players has a dominant position, by controlling a higher-level decision phase in which part of the game instance is set (e.g., cost of a subset of network arcs, routing of a substantial amount of the flow, etc.). In particular, Stackelberg network games in which the leader(s) fixes the costs of a subset of links in the network, which are of interest for this paper, are widely known as *network pricing games* (NPG).

The probably oldest NPG is that in which we are given a directed graph $G = (V, E)$, where E is partitioned into a set E_F of edges with a fixed positive real weight and a set E_P of priceable edges, and two nodes $s, t \in V$; then, the follower wants to travel along a *shortest path* in G between s and t , after the leader has fixed the costs for edges in E_P (see [10] for a survey). The problem for the leader of fixing the costs so as to maximize its revenue (given by the sum of the costs of edges in E_P used by the follower), has been shown to be strongly NP-hard, as well as $O(\log |E_P|)$ -approximable [9]. For the case of multiple followers (each with a specific source-destination pair), Labbé *et al.* [6] derived a bilevel LP formulation of the problem (and proved NP-hardness), while Grigoriev *et al.* [5] gave algorithms for a restricted shortest path problem on parallel edges.

Another basic NPG game, which quite surprisingly was considered only very recently, is that in which the follower wants to eventually use a *minimum*

spanning tree of G (now considered as undirected). For this game, in [3] the authors proved the APX-hardness already when the number of possible weights for the edges in E_F is 2, and gave a logarithmic-approximation algorithm.

1.2 Our Problem

In this paper, we focus on a NPG concerning one of the most popular network topologies, namely the *single-source shortest paths tree*. Our setting is the following: We assume the existence of a leader which can set the costs for using a certain set of links in a network modelled by a directed graph. The follower needs to send a message from a given, fixed source node to each of the other nodes. This models the situation in which the follower realizes a broadcast by a sequence of *unicasts*, the predominant protocol in the Internet. More formally, our game, that we call STACKSPT, can be described as a bilevel optimization problem (all the paths are assumed to be directed):

- *Instance*: A directed graph $G = (V, E = E_F \cup E_P, E_F \cap E_P = \emptyset)$, a function $w : e \in E_F \mapsto \mathbb{R}^+$, and a source node $r \in V$;
- *Leader feasible solution*: A pricing $p : e \in E_P \mapsto \mathbb{R}^+$;
 - *Follower feasible solution*: A spanning arborescence T of G rooted at r (after edges in E_P have been priced);
 - *Follower objective*: Minimize the sum of all the path lengths in T from r to any node in G ;
- *Leader objective*: For a given optimal solution w.r.t. the follower objective, say $S(r) = (V, E_S)$, maximize the revenue

$$\rho(S(r)) = \sum_{e \in E_P \cap E_S} p(e) \cdot ||e||$$

where $||e||$ is the number of paths in $S(r)$ emanating from r and using e .

For the follower’s objective we assume that in the graph G , edges in E_F are weighted w.r.t. $w(\cdot)$, while edges in E_P are weighted w.r.t. $p(\cdot)$. Moreover, to make the problem bounded, we assume that in $G - E_P = (V, E_F)$, for each $v \in V$ there exists a path from r to v . Finally, when multiple optimal solutions are available for the follower, we adopt the standard tie-breaking rule which enforces to select an optimal solution maximizing the leader’s revenue.¹ We will see later (see Lemma 3) how this tie-breaking rule can be actually implemented. Moreover, in the paper we adopt the following notation:

- $\pi(u, v)$: a shortest path in $G = (V, E)$ between u and v ;
- $e_i = (u_i, v_i)$: an edge in E_P directed from u_i to v_i ;
- $\pi_{i_1 \dots i_h \hat{j}_1 \dots \hat{j}_k}(u, v)$: for a subset $e_{i_1}, \dots, e_{i_h}, e_{j_1}, \dots, e_{j_k}$ of edges of E_P , a shortest path in G between u and v constrained to contain edges in e_{i_1}, \dots, e_{i_h} , and to avoid edges in e_{j_1}, \dots, e_{j_k} ;
- $|\pi_{i_1 \dots i_h \hat{j}_1 \dots \hat{j}_k}(u, v)|$: the length of the path.

¹ While this rule may appear unnatural in our competitive setting, it simply avoids the undesired technicality that an optimal solution for the leader can only be reached within an arbitrarily small subtracted term.

1.3 Our Results

First of all, notice that if $|E_P| = 1$, then our problem can be solved efficiently by using a classic sensitivity analysis algorithm. More precisely, it suffices to compute, for each $v \in V$, the length of a shortest path between r and v in the graph deprived of the priceable edge, and the length of a shortest path between r and v in the graph in which the weight of the priceable edge is set to 0. Then, we consider all the nodes for which these two lengths differ, and we sort in non-increasing order the difference between the two lengths. These values, once multiplied by their position in the sorted sequence, correspond exactly to the revenue associated with each given threshold price. Since finding the length of all the replacement paths can be done in $O(m + n \log n)$ time [7], it is not hard to see that this is an upper bound for solving the problem.

On the other hand, for $|E_P| > 1$ our problem can be faced by using the very general results given in [2], where the authors show how to handle NPGs in which the followers optimize a polynomial-time network optimization problem. More precisely, it can be shown that our problem can be rephrased as an $(n - 1)$ -follower game, for which the authors provide, for any $\epsilon > 0$, a $H_{(n-1)|E_P|}(1 + \epsilon)$ -approximation algorithm, where H_i denotes the i -th harmonic number (see Theorem 2 in [2]).

In this paper, we contribute to the knowledge on the STACKSPT game by analyzing it under several respects. More precisely, firstly we analyze the complexity of the game, and we show that finding an optimal pricing for the leader's edges is NP-hard, as soon as $|E_P| = \Theta(n)$. Then, we turn our attention to the development of an efficient algorithm for the case in which $|E_P|$ is fixed. To this aim, we analyze in detail the basic case $|E_P| = 2$, for which we provide an efficient $O(n^2 \log n)$ time algorithm. This result forms the basis for the case $|E_P| = k$, for which we are then able to design an $O\left(\frac{4^{k^2}}{k!} n^k (k^3 + m + n \log n)\right)$ time algorithm.

2 NP-Hardness of the STACKSPT Game

In this section we prove the following:

Theorem 1. *STACKSPT in directed graphs is NP-hard.*

Proof. The reduction is from 3-SAT, and is an extension of the hardness result proved in [9] for the NPG in which the follower wants to compute a shortest path between two given nodes. Let k be the number of clauses in the formula. Figure 1 shows the gadget for the i -th clause, and how the gadgets are linked in order to obtain the graph corresponding to the entire formula. In particular, there is a shortcut of cost $1/2$ connecting the head of an edge $e_{i,p}$ with the tail of an edge $e_{j,q}$, with $j > i$ and $p, q \in \{1, 2, 3\}$, iff in the corresponding clauses, the variable $x_{i,p}$ is the negation of $x_{j,q}$. Moreover, we add a set of $h - 1$ vertices that can be reached from the last gadget with edges of cost 0. The edges that have to be priced are $\{e_1, \dots, e_k\} \cup \bigcup_{i=1}^k \{e_{i,1}, e_{i,2}, e_{i,3}\}$.




Fig. 1. The gadget for clause i , and the network for the formula $(x_1 \vee x_2 \vee \bar{x}_3) \wedge (\bar{x}_2 \vee x_3 \vee x_4) \wedge (\bar{x}_1 \vee \bar{x}_3 \vee \bar{x}_4)$. Priceable edges are in bold, the weight of backward edges from t_i -nodes in the network is 6, while all other unspecified weights are equal to 0.

A slightly different version of the following lemma was already proven in [9]:

Lemma 1. *The formula is satisfiable if and only if there is a pricing for the corresponding graph such that the revenue obtained from the path from r to v_1 is strictly greater than $2k - 1/2$. Moreover, if the formula is satisfiable then the revenue that can be achieved from the path is exactly $2k$.*

Proof. Given a truth assignment for the formula, we can define a pricing inducing a revenue of value $2k$ as follows. We set the price of $e_{i,j}$ to 1 if $e_{i,j}$ corresponds to a true literal, a sufficiently high value otherwise. All the other priceable edges are set to 1. On the other hand, consider a pricing inducing a revenue of value strictly greater than $2k - 1/2$. First, notice that there is fixed cost path of length $2k$ from r to v_1 . Therefore the shortest path P from r to v_1 selected by the follower cannot use any edge having a fixed cost greater or equal to $1/2$. As a consequence, P contains $2k$ priceable edges each having a price in the interval $(1/2, 1]$. Now, putting the literals corresponding to the selected priceable edges to true provides a valid truth assignment: each clause has a true literal, and the path does not contain two edges corresponding to a literal and its negation, since in that case the shortest path would take the shortcut of cost $1/2$. \square

Now, let $h > 48k^2$. Then we claim that the 3-SAT instance has a positive answer if and only if the corresponding STACKSPT instance admits a pricing yielding a total revenue strictly greater than $24k^2 + (2k - 1/2)h$. Indeed, if the formula is satisfiable, from Lemma 1 we can price the edges in such a way that the total revenue is at least $2kh > 24k^2 + (2k - 1/2)h$ (we can guarantee a revenue of $2k$ for each v_j). On the other hand, if the formula does not admit a truth assignment, then the total revenue is upper bounded by $24k^2 + (2k - 1/2)h$. To see this, notice that the revenue obtained from each path from r to any v_j is at most $2k - 1/2$ (Lemma 1). Moreover, a path from r to any other vertex can bring a revenue that is at most $3k$, since each vertex belonging to a gadget can be reached with a path of non-priceable edges of total cost upper bounded by $3k$. Since there are $8k$ of such vertices, the claim follows. \square

3 The Case $|E_P| = 2$

Let us focus on the case in which $E_P = \{e_1, e_2\}$. For the sake of generality, let us assume that for $p(e_1) = p(e_2) = 0$, a SPT of G rooted at r contains both edges.

Moreover, we also assume that $e_2 = (u_2, v_2)$ descends from $e_1 = (u_1, v_1)$ in such a tree, since the case in which the two edges are unrelated is just a particularization of our analysis. From our assumptions, the aforementioned SPT will induce a partition of the nodes of G into three sets: those using both edges, those using only e_1 , and finally those not using them. In the following, we analyze in detail the set of nodes of the first type, since the set of nodes of the second type can be studied in a very similar and simpler way, while nodes in the third set are just not affected by the pricing strategy.

Let then $v \in V$ be a node whose path from the root in the SPT passes through u_1, v_1, u_2 and v_2 , in this order (it is easy to see that such order will then be maintained regardless of the values $p(e_1), p(e_2)$). Then, we need to compute the partition of the set of points $(p(e_1), p(e_2)) \in \mathbb{R}^+ \times \mathbb{R}^+$ into four regions, say $v[\{0, 1\}^2]$, in which the first (resp., second) component of the argument is set to 1 if e_1 (resp., e_2) is used in the (shortest) path in $S(r)$ between r and v . The four regions are then defined by the following set of points:

$$\begin{aligned} v[1, \cdot] &:= \{(p(e_1), p(e_2)) \text{ s.t. } |\pi_1(r, v)| < |\pi_{\hat{1}}(r, v)|\}; \\ v[\cdot, 1] &:= \{(p(e_1), p(e_2)) \text{ s.t. } |\pi_2(r, v)| < |\pi_{\hat{2}}(r, v)|\}; \\ v[1, 1] &:= v[1, \cdot] \cap v[\cdot, 1]; \quad v[0, 0] := \mathbb{R}^+ \times \mathbb{R}^+ \setminus (v[1, \cdot] \cup v[\cdot, 1]). \end{aligned}$$

As far as $v[1, \cdot]$ is concerned, it can be exploited as a function of whether e_2 either belongs or not to $\pi_1(r, v), \pi_{\hat{1}}(r, v)$ (in the formula we omit the (r, v) argument for the sake of readability):

$$\begin{aligned} v[1, \cdot] &= \{(p(e_1), p(e_2)) \text{ s.t. } \min\{|\pi_{12}|, |\pi_{1\hat{2}}|\} < \min\{|\pi_{\hat{1}\hat{2}}|, |\pi_{2\hat{1}}|\}\} \\ &= \{(p(e_1), p(e_2)) \text{ s.t. } (|\pi_{12}| < |\pi_{2\hat{1}}|) \wedge (|\pi_{12}| < |\pi_{\hat{1}\hat{2}}|)\} \cup \quad (1) \\ &\quad \cup \{(p(e_1), p(e_2)) \text{ s.t. } (|\pi_{1\hat{2}}| < |\pi_{2\hat{1}}|) \wedge (|\pi_{1\hat{2}}| < |\pi_{\hat{1}\hat{2}}|)\} \end{aligned}$$

The above can be rewritten as the union of the solutions of the following two systems of inequalities:

$$(S_1) := \left\{ \begin{array}{l} |\pi_{12}| < |\pi_{\hat{1}\hat{2}}| \\ |\pi_{12}| < |\pi_{2\hat{1}}| \end{array} \right. \quad (S_2) := \left\{ \begin{array}{l} |\pi_{1\hat{2}}| < |\pi_{\hat{1}\hat{2}}| \\ |\pi_{1\hat{2}}| < |\pi_{2\hat{1}}|. \end{array} \right.$$

By noticing that π_{12} and $\pi_{2\hat{1}}$ both contains e_2 , (S_1) can be rewritten as:

$$\left\{ \begin{array}{l} |\pi_{\hat{1}\hat{2}}(r, u_1)| + p(e_1) + |\pi_{\hat{1}\hat{2}}(v_1, u_2)| + p(e_2) + |\pi_{\hat{1}\hat{2}}(v_2, v)| < |\pi_{\hat{1}\hat{2}}(r, v)| \\ |\pi_{\hat{1}\hat{2}}(r, u_1)| + p(e_1) + |\pi_{\hat{1}\hat{2}}(v_1, u_2)| + |\pi_{\hat{1}\hat{2}}(v_2, v)| < |\pi_{\hat{1}}(r, u_2)| + |\pi_{\hat{1}}(v_2, v)| \end{array} \right.$$

and from our assumptions on the relative positions of u_1, v_1, u_2, v_2 w.r.t. r , it follows that $|\pi_{\hat{1}\hat{2}}(v_2, v)| = |\pi_{\hat{1}}(v_2, v)| = |\pi(v_2, v)|$, and $|\pi_{\hat{1}\hat{2}}(r, u_1)| = |\pi(r, u_1)|$, and $|\pi_{\hat{1}\hat{2}}(v_1, u_2)| = |\pi(v_1, u_2)|$, and then (see Figure 2 for checking out the lines defining the half-planes associated with the inequalities)

$$(S_1) := \left\{ \begin{array}{l} p(e_1) + p(e_2) < |\pi_{\hat{1}\hat{2}}(r, v)| - |\pi(r, u_1)| - |\pi(v_1, u_2)| - |\pi(v_2, v)| \quad (\text{line } \ell_6) \\ p(e_1) < |\pi_{\hat{1}}(r, u_2)| - |\pi(r, u_1)| - |\pi(v_1, u_2)|. \quad (\text{line } \ell_2) \end{array} \right.$$

Similarly, for (S_2) we obtain

$$(S_2) := \begin{cases} p(e_1) < |\pi_{\hat{1}\hat{2}}(r, v)| - |\pi(r, u_1)| - |\pi_2(v_1, v)| & \text{(line } \ell_4) \\ p(e_1) - p(e_2) < |\pi_{\hat{1}}(r, u_2)| + |\pi(v_2, v)| - |\pi(r, u_1)| - |\pi_2(v_1, v)|. & \text{(line } \ell_3) \end{cases}$$

Simmetrically, for $v[\cdot, 1]$ we get the following systems of inequalities:

$$(S'_1) := \begin{cases} p(e_1) + p(e_2) < |\pi_{\hat{1}\hat{2}}(r, v)| - |\pi(r, u_1)| - |\pi(v_1, u_2)| - |\pi(v_2, v)| & \text{(line } \ell_6) \\ p(e_2) < |\pi_2(v_1, v)| - |\pi(v_2, v)| - |\pi(v_1, u_2)| & \text{(line } \ell_1) \end{cases}$$

$$(S'_2) := \begin{cases} p(e_2) < |\pi_{\hat{1}\hat{2}}(r, v)| - |\pi(v_2, v)| - |\pi_{\hat{1}}(r, u_2)| & \text{(line } \ell_5) \\ p(e_1) - p(e_2) > |\pi_{\hat{1}}(r, u_2)| + |\pi(v_2, v)| - |\pi(r, u_1)| - |\pi_2(v_1, v)|. & \text{(line } \ell_3) \end{cases}$$

The partition of the plane induced by $v[1, \cdot]$ and $v[\cdot, 1]$ depends on the mutual positions of the half-planes obtained from the constraints in the system. In particular, if $|\pi_{\hat{1}\hat{2}}(r, v)| - |\pi(r, u_1)| - |\pi_2(v_1, v)| > |\pi_{\hat{1}}(r, u_2)| - |\pi(r, u_1)| - |\pi(v_1, u_2)|$ (i.e., ℓ_4 is to the right of ℓ_2), which implies $|\pi_{\hat{1}\hat{2}}(r, v)| - |\pi(v_2, v)| - |\pi_{\hat{1}}(r, u_2)| > |\pi_2(v_1, v)| - |\pi(v_2, v)| - |\pi(v_1, u_2)|$ (i.e., ℓ_5 is above ℓ_1), then we have the partition on the left side of Figure 2 otherwise we have the partition on the right side (notice that sloped lines have angular coefficient ± 1).




Fig. 2. The two possible partitions of the plane

Now, observe that w.r.t. any given node v , the maximum revenue is obtained by selecting the maximum in either of the points belonging to the associated solid lines, depending on the plane partition. However, when all the nodes have to be considered simultaneously for selecting an optimal solution, say $(p^*(e_1), p^*(e_2))$, all the plane partitions have to be overlaid. Such an overlaying will generate a new partition of the plane into closed regions holding the following property:

Lemma 2. *For each region \mathcal{R} , the relative maximum is placed onto some vertex of its boundary.*

Proof. Indeed, \mathcal{R} is convex and is homogenous w.r.t. each node, meaning that for any node $v \in V$, \mathcal{R} is covered by either of $v[0, 0], v[0, 1], v[1, 0], v[1, 1]$. Then, from the monotonicity of the revenue function, it is easy to see that for each of these cases, the maximum for node v is retained onto a vertex of the boundary of \mathcal{R} , from which the claim follows. \square

Thus, to find $(p^*(e_1), p^*(e_2))$, we check all the points of the plane in which the boundary lines of the various partitions intersect. Since the partition of the plane for every vertex is determined by a constant number of lines, we have $O(n)$ lines in total. As a consequence, the $O(n)$ lines can generate $O(n^2)$ intersection points. Observe now that all these intersection points can be found in $O(n^2)$ time, after an $O(m + n \log n)$ preprocessing time. Indeed, to compute each node partition, it suffices to find all the distances from the root in the four possible configurations where $p(e_1), p(e_2) \in \{0, \infty\}$. Thus, in order to solve our problem, it remains to check efficiently each of these points. The following lemma (which holds independently of the size of E_P) provides a way to do that:

Lemma 3. *For a given pricing $p : E_P \rightarrow \mathbb{R}^+$, a shortest path tree maximizing the total revenue can be computed in $O(m + n \log n)$ time.*

Proof. Consider the graph where all priceable edges are priced according to $p(\cdot)$. By a trivial modification of Dijkstra’s algorithm, it is possible to compute in $O(m + n \log n)$ time the subgraph H of G containing an edge $f \in E(G)$ iff f is an edge of some shortest path tree of G rooted at r . Now we show how to compute in $O(m)$ time a shortest path tree $S^*(r)$ of H rooted at r which gives the maximum revenue for the leader w.r.t. p . We define a new weight function $\phi(\cdot)$ as follows. For each edge $e \in E(H)$, $\phi(e)$ is defined as $-p(e)$ if $e \in E_P$, 0 otherwise. Then we compute an SPT of H rooted at r w.r.t. $\phi(\cdot)$, say $S^*(r)$. This can be done in $O(m)$ time since H is acyclic [4]. Now the claim follows from the fact that for any SPT $S(r)$ of H we have that $\rho(S(r)) = \sum_{e \in E(S(r)) \cap E_P} p(e) \|e\| = - \sum_{e \in E(S(r))} \phi(e) \|e\|$. \square

Hence, from the above discussion, it follows that the case $|E_P| = 2$ can be handled in $O(mn^2 + n^3 \log n)$ time. As we will see, in next section we will extend this result to the case in which $|E_P| = k$. In the remaining of this section we show that the running time for the case $|E_P| = 2$ can be actually reduced to $O(n^2 \log n)$ after a careful geometrical analysis and the use of geometric data structures. Although we leave open the problem of generalizing this faster algorithm to the case $|E_P| = k$, we believe this result is of independent interest, due to the importance of the case $|E_P| = 2$, which models the atomic situation in which there is an agent owning more than one edge in the network.

3.1 A Faster Algorithm

In the sequel, we denote each of the possible lines as depicted in Figure 2 and we only refer to intersections between lines associated with distinct nodes. We start by observing that since the abscissa of points A and C is independent of

node v , then all ℓ_2 -segments lie on the same line. Thus, the closed rectangles in the first partition and the pentagons in the second partition are all adjacent on three sides (the two axis and line ℓ_2). Moreover, as all the lines are vertical, horizontal or with angular coefficient ± 1 , it follows that

Fact 1. *Each intersection point is on the boundary of at most 8 regions.*

For each region \mathcal{R} of the partitioning of the plane obtained by overlaying all the partitions associated with the nodes, let $\mathcal{R}(|e_1|, |e_2|)$ denote the (constant) load of each interior point of the region, namely the number of nodes using e_1 and the number of nodes using e_2 .

Let us divide the plane in two zones, the part we will call *zone 1*, which is on the left side of line ℓ_2 (line ℓ_2 is included), and the part we will call *zone 2*, which is on the right side of line ℓ_2 (line ℓ_2 is included). Since all ℓ_2 -segments lay on the same line, we have that all B -points of the first partition are in zone 2, all A -points of the first partition lay on ℓ_2 , while all the C -points of the second partition are in zone 1.

Let (x, y) be any point in zone 2 which is internal to a region, say \mathcal{R} . Partition zone 2 into three subzones as shown in Figure 3a by drawing the lines $p(e_1) = x, p(e_2) = y$, and the line with angular coefficient 1 passing through (x, y) . Let B_{e_1}, B_{e_2} , and B_{not} be the natural partitioning of the B -points (multiple occurrences are allowed as the B -points of some vertices might be the same) induced by the three subzones. From Figure 3a it is not hard to see that $\mathcal{R}(|e_1|, |e_2|) = (|B_{e_1}|, |B_{e_2}|)$.




Fig. 3. (a) Partitioning of zone 2 into three subzones w.r.t. point (x, y) : the three dotted lines emanating from some B -point are the lines determining the plane subdivision of some vertex that generates that B -point; (b) Partitioning of zone 1 into three subzones w.r.t. point (x, y) : the three dotted lines emanating from some C -point are the lines determining the plane subdivision of some vertex that generates that C -point, while the dashed line emanating from some A -point is the line determining the plane subdivision of some vertex that generates that A -point; (c) How to compute the load of all the regions whose boundaries contain the point (x, y) (the circle has a very tiny radius)

² Observe that there is no B -point which is on some of the three lines emanating from (x, y) (see Figure 3a), as (x, y) is internal to region \mathcal{R} and because of the structure of the first partition (see Figure 2).

Let (x, y) be any point in zone 1 which is internal to a region, say \mathcal{R} . Partition zone 1 into three subzones as shown in Figure 3b by drawing the lines $p(e_1) = x, p(e_2) = y$, and the line with angular coefficient -1 passing through (x, y) . Let P_{both}, P_{e_1} , and P_{not} be the natural partitioning of the A -points and C -points (again, multiple occurrences are allowed) induced by the three subzones. From Figure 3b it is not hard to see that $\mathcal{R}(\|e_1\|, \|e_2\|) = (|P_{\text{both}}| + |P_{e_1}|, |P_{\text{both}}|)$. Moreover, observe that at most a new intersection point belonging to zone 1 is generated if we overlap the partitions associated with any two nodes. We can state the following

Theorem 2. STACKSPT with $|E_P| = 2$ can be solved in $O(n^2 \log n)$ time.

Proof. First of all, recall that the set X of $O(n^2)$ intersection points (which clearly includes also the A, B, C -points) can be found in $O(n^2)$ time, after an $O(m + n \log n)$ preprocessing time.

For all the intersection points $(x, y) \in X$, we compute the best revenue the leader can get if its edges are priced $p(e_1) = x, p(e_2) = y$. From Lemma 2, we have that the maximum revenue is obtained on some of these points. To compute the best revenue for (x, y) it is enough to compute the load of the regions whose boundaries contain the point (x, y) , and check which load gives the best revenue. As from Fact 1 (x, y) belongs to the boundaries of at most 8 regions, and as (x, y) can be generated by lines which are vertical, horizontal or with angular coefficient ± 1 , it is enough to compute the load for the 8 points which are close to (x, y) as shown in Figure 3c. This guarantees that all the regions whose boundary contains (x, y) are taken into account. Indeed, using standard data structures known for range counting problems (see for example [1]), after a preprocessing phase requiring $O(n \log n)$ time, we can compute in $O(\log n)$ time the number of B -points belonging to all the three regions in Figure 3a, as well as the number of A -points and C -points belonging to all the three regions in Figure 3b, where (x, y) is an input parameter. Since it is possible to check in constant time if (x, y) is in zone 1 or in zone 2, we have that the best revenue of every point can be computed in $O(\log n)$ time. As we have $O(n^2)$ points, the claim follows. \square

4 The Case $|E_P| > 2$

In this section we extend the basic result we obtained for $|E_P| = 2$ to the general case in which the leader owns k edges. As a consequence of the generalization, the runtime of the algorithm we propose is $O\left(\frac{4^{k^2}}{k!} n^k (k^3 + m + n \log n)\right)$.

Before providing the description of the algorithm, we need to introduce some notation which was not needed for the case $|E_P| = 2$. Given a vector b , we will denote by b_i the i -th component of the vector, while (b_{-i}, b'_i) denotes the vector which is equal to b except for its i -th component whose value is b'_i . Finally, by

³ Observe that neither A -points nor C -points can be on some of the three lines emanating from (x, y) (see Figure 3b), as (x, y) is internal to region \mathcal{R} and because of the structure of the second partition (see Figure 2).

(\cdot_{-i}, \bar{b}_i) we will denote the set of all vectors for which the i -th component is fixed to be \bar{b}_i .

Along the same line of the previous section, we focus our attention on the more general case in which, given a vertex $v \in V$, we have that a shortest path in G from r to v passes through all the k edges, once their weight is set to 0. For vertices using only a subset of edges, the analysis will follow by properly considering the corresponding set of vector components. Now, for a given $i, 1 \leq i \leq k$, we define $v[\cdot_{-i}, 1] = \{ (p(e_1), \dots, p(e_k)) \text{ s.t. } |\pi_i(r, v)| < |\pi_{\bar{i}}(r, v)| \}$. Obviously $v[\cdot_{-i}, 0] = (\mathbb{R}^+)^k \setminus v[\cdot_{-i}, 1]$. Therefore, for any given vector $b = (b_1, \dots, b_k) \in \{0, 1\}^k$, we have that $v[b_1, \dots, b_k] = \bigcap_{i=1}^k v[\cdot_{-i}, b_i]$.

Let $\pi_b(r, v)$ denotes a path between r and v which includes e_i iff $b_i = 1$. Expanding the definition of $v[\cdot_{-i}, 1]$ leads to the following collection of systems of linear inequalities

$$\begin{aligned} v[\cdot_{-i}, 1] &= \{ (p(e_1), \dots, p(e_k)) \text{ s.t. } |\pi_i(r, v)| < |\pi_{\bar{i}}(r, v)| \} \\ &= \bigcup_{b' \in \{0, 1\}^k} \{ (p(e_1), \dots, p(e_k)) \text{ s.t. } |\pi_{(b'_{-i}, 1)}(r, v)| < |\pi_{\bar{i}}(r, v)| \} \\ &= \bigcup_{b' \in \{0, 1\}^k} \left\{ (p(e_1), \dots, p(e_k)) \text{ s.t. } \bigwedge_{b'' \in \{0, 1\}^k} |\pi_{(b'_{-i}, 1)}(r, v)| < |\pi_{(b''_{-i}, 0)}(r, v)| \right\}. \end{aligned}$$

Hence, $v[b_1, \dots, b_k]$ is determined by a system of linear inequalities. Let \mathcal{H} denote the set of all the hyperplanes corresponding to some linear inequality of the system. More formally

$$\mathcal{H} = \bigcup_{v \in V \setminus \{r\}} \bigcup_{i=1}^k \left\{ |\pi_{b'}(r, v)| = |\pi_{b''}(r, v)| \text{ s.t. } b', b'' \in \{0, 1\}^k, b'_i = 1, b''_i = 0 \right\}. \quad (2)$$

We are now ready to give our main result for the case in which $|E_P| = k$:

Theorem 3. STACKSPT can be solved in $O\left(\frac{4^{k^2}}{k!} n^k (k^3 + m + n \log n)\right)$ time.

Proof. First of all, notice that since the coefficients of each hyperplane can be computed in $O(m + n \log n)$ time by means of two constrained shortest-path computations, and given that by definition $|\mathcal{H}| \leq n4^k$, we have that \mathcal{H} can be found in $O(n4^k(m + n \log n))$ time.

Then, the algorithm first computes a set X of possible pricings for the leader’s edges. Consider the subdivision of the k -dimensional Euclidean space into regions obtained from overlaying all the hyperplanes in \mathcal{H} . The set X computed by the algorithm contains all the extremal points for all the regions resulting from the subdivision specified above. Clearly, inside every region the load of the leader’s edges is constant. Obviously, every region is convex because its boundary is determined by a system of hyperplanes in \mathcal{H} . Therefore, an extremal point of some of these regions gives an optimal pricing of the edges. This proves that X contains an optimal pricing for the edges.

Next, the algorithm searches for a point in X that maximizes the leader's revenue. From Lemma 3 this can be done in $O(|X|(m + n \log n))$. Therefore, the running time of the algorithm is $O(|X|(m + n \log n))$ plus the time needed to compute X . In what follows we prove that it is possible to compute X in $O\left(\frac{4^{k^2}}{k!}k^3n^k\right)$ time, thus proving the claim. Any extremal point is determined by the intersection of k hyperplanes in \mathcal{H} . Thus, we have

$$|X| \leq \binom{|\mathcal{H}|}{k} \leq \binom{n4^k}{k} \leq \frac{n^k 4^{k^2}}{k!}.$$

Since each point in X can be found in $O(k^3)$ time by solving the corresponding $k \times k$ system of hyperplanes [4], we have that X can be computed in $O\left(\frac{4^{k^2}}{k!}k^3n^k\right)$ time. Since from Lemma 3 each point can be checked in $O(m + n \log n)$ time, the claim follows. \square

References

1. Agarwal, P.K., Erickson, J.: Geometric range searching and its relatives. In: Chazelle, B., Goodman, J.E., Pollack, R. (eds.) *Advances in Discrete and Computational Geometry*. Contemporary Mathematics, vol. 23, pp. 1–56. American Mathematical Society Press, Providence (1999)
2. Briest, P., Hoefer, M., Krysta, P.: Stackelberg network pricing games. In: 25th Annual Symposium on Theoretical Aspects of Computer Science (STACS), pp. 133–142 (2008), <http://drops.dagstuhl.de/opus/volltexte/2008/1340>
3. Cardinal, J., Demaine, E.D., Fiorini, S., Joret, G., Langerman, S., Newman, I., Weimann, O.: The Stackelberg minimum spanning tree game. In: Dehne, F., Sack, J.-R., Zeh, N. (eds.) *WADS 2007*. LNCS, vol. 4619, pp. 64–76. Springer, Heidelberg (2007)
4. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: *Introduction to Algorithms*. MIT Press and McGraw-Hill (2001)
5. Grigoriev, A., van Hoesel, S., van der Kraaij, A., Uetz, M., Bouhtou, M.: Pricing network edges to cross a river. In: Persiano, G., Solis-Oba, R. (eds.) *WAOA 2004*. LNCS, vol. 3351, pp. 140–153. Springer, Heidelberg (2005)
6. Labbé, M., Marcotte, P., Savard, G.: A bilevel model of taxation and its application to optimal highway pricing. *Management Science* 44(12), 608–622 (1998)
7. Malik, K., Mittal, A.K., Gupta, S.K.: The k most vital arcs in the shortest path problem. *Oper. Res. Letters* 8, 223–227 (1989)
8. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.): *Algorithmic Game Theory*. Cambridge University Press, New York (2007)
9. Roch, S., Savard, G., Marcotte, P.: An approximation algorithm for Stackelberg network pricing. *Networks* 46(1), 57–67 (2005)
10. van Hoesel, S.: An overview of Stackelberg pricing in networks. *European Journal of Operational Research* 189(3), 1393–1402 (2008)
11. von Stackelberg, H.: *Marktform und Gleichgewicht (Market and Equilibrium)*. Verlag von Julius Springer, Vienna, Austria (1934)

Local Two-Stage Myopic Dynamics for Network Formation Games

Esteban Arcaute¹, Ramesh Johari², and Shie Mannor³

¹ Institute for Computational and Mathematical Engineering, Stanford University
arcaute@stanford.edu

² Department of Management Science and Engineering, Stanford University
rjohari@stanford.edu

³ Department of Electrical and Computer Engineering, McGill University
shie@ece.mcgill.ca

Abstract. *Network formation games* capture two conflicting objectives of self-interested nodes in a network. On one hand, such a node wishes to be able to reach all other nodes in the network; on the other hand, it wishes to minimize its cost of participation. We focus on myopic dynamics in a class of such games inspired by transportation and communication models. A key property of the dynamics we study is that they are *local*: nodes can only deviate to form links with others in a restricted neighborhood. Despite this locality, we find that our dynamics converge to efficient or nearly efficient outcomes in a range of settings of interest.

1 Introduction

Viewing modern data-networks, such as the Internet or ad-hoc networks, as a federation of selfish, independent actors leads to a range of interesting game theoretic questions. The goal of a selfish agent in a data network is two-fold. On one hand, it wishes to be able to reach all other agents in the network. On the other hand, it wishes to minimize its cost for participating in the network. The recent literature on *network formation games* (NFGs) provides a natural model with which to study this tradeoff between cost and connectivity; NFGs have been suggested as models for many domains, from trade networks to peer-to-peer networks (see [1] for a comprehensive review).

In the model we consider, nodes derive utility from connectivity to each other, and incur a cost comprised of three terms: (1) traffic related costs; (2) costs to maintain links to other nodes; and (3) payments made to other nodes. The payments are a natural mechanism for users to compensate each other for the costs required to establish links; we assume these payments are bilaterally negotiated through *contracts*. Because link formation is bilateral in NFGs, the equilibrium concept of interest for NFGs is *pairwise Nash stable equilibrium* [2], as opposed to Nash equilibrium in traditional game theoretic models. Roughly speaking, a graph is pairwise Nash stable if it is a Nash equilibrium, and if it is *pairwise stable* as first introduced in [3]: no node can profitably remove edges adjacent to it, and no two nodes can deviate jointly by altering their contract to improve both their payoffs.

This paper studies dynamic processes of network formation in the setting described above; in particular, we study whether realistic dynamics exist that also naturally select

efficient equilibria. The most basic dynamic network formation process is *best response dynamics*, where at each round a profitable deviation is undertaken by one or a pair of nodes at a time. An outcome is a fixed point of best response dynamics if and only if it is a pairwise stable equilibrium. Although best response dynamics are attractively simple, they may fail to converge; further, since any pairwise Nash stable equilibrium is a fixed point, best response dynamics can lead to inefficient equilibria.

Our main insights are that, in a model we previously introduced in [4], there exist a natural class of dynamics that also select efficient equilibria. When there is no cost to maintain links and some equilibria with redundant edges exist, our dynamics can select those non-tree equilibria; however, in this case the resulting equilibria may be inefficient.

The dynamics we study are a generalization of local-best response dynamics. First, nodes only deviate with other nodes in a local neighborhood; this is consistent with the observation that in modern data networks, nodes are typically only aware of the local topology of the network around them. Second, we allow nodes to deviate *twice* in succession, using the first deviation to improve its bargaining power in the second step.

Formally, the dynamics proceed in *rounds*. Each round is divided into two stages. During the first stage, an exogenously designated node u considers all possible unilateral deviations. During the second stage, the same node considers all possible unilateral and bilateral deviations with nodes in its local neighborhood. The key assumption is that u seeks to maximize its prospective payoff *at the end of the round*, while nodes involved in bilateral deviations choose their actions in order to maximize prospective payoff *at the end of the stage*, considering the current network when making a decision. As such, one can think of our dynamics as constrained local-best response dynamics with *one-step look-ahead* for the active node. It is straightforward to show in the models we consider that any fixed point of these dynamics must be pairwise Nash stable.

The one-step look-ahead feature allows the deviating node more flexibility than best-response dynamics—and thus in principle encourage more selfish behavior. The main benefit of the proposed dynamics is that they serve as a local and decentralized equilibrium selection mechanism. In particular, we show that for the utility model from [4], where potentially inefficient pairwise Nash stable equilibria exist, these dynamics in fact converge to desirable pairwise Nash stable equilibria instead. Also, when there is no link maintenance cost and equilibria with redundant edges exist, we prove that it is possible for the most inefficient equilibria selected to be less efficient than those selected when there is a positive link-maintenance cost. Thus we prove that there is a possible efficiency loss when the link maintenance cost vanishes.

In the NFG model we consider, a node incurs a traffic related cost that depends on the total volume of traffic routed through it. In [4], it was shown that, for a very restricted class of global dynamics, the network converges to a desirable pairwise Nash stable equilibrium under certain conditions. Our paper recovers this convergence result and extends it to the case when the link maintenance cost is nil using significantly more general and natural, local dynamics; at the same time, the proofs are simplified via an ordinal potential function argument (see [2]).

Related Work

Our work touches on several related threads of the literature. Most closely related is the work on network formation games in economics (see [1] for a survey). In the context of communication networks [5] considered a static network formation game related to the model in this paper; by contrast, the focus of the current paper is on the dynamics of such formation processes. The work of Jackson and Watts also considers dynamics for network formation games [6, 7], but for a utility model that is unrelated to ours. While in their dynamics only a unilateral or a bilateral deviation may occur in a given round, our dynamics are designed so that each round consists of two stages, thus allowing a unilateral deviation to be *followed* by a bilateral deviation. This latter property allows our dynamics to *select* desirable equilibria. Despite the dissimilarities in cost structure and dynamics, random activation of nodes in the dynamics is needed for obtaining convergence results both in our setting and in that of [7].

In [4], we considered essentially the same utility model and static game as in this paper. There are several important differences that we summarize below.

- In [4] we used a variant of pairwise stability as the solution concept. Here we use pairwise Nash stability, which allows us to define our dynamics in a more natural framework, and also allows us to provide a complete characterization of equilibrium outcomes.
- In [4] we considered a restricted version of the dynamics we introduce in this paper. Each round is again composed of two stages. In [4], at the beginning of each round, an edge is sampled from the set of all possible edges, the active node is selected among the two end-points of the sampled edge. Two main differences can be pointed out:
 - During the first stage, the active node is allowed to unilaterally deviate *only* with respect to the sampled edge. This is in sharp contrast to the dynamics of this paper where the active node can unilaterally deviate with respect to *any* set of edges adjacent to it.
 - During the second stage, the active node can bilaterally deviate with a node selected from *all* possible nodes, thus requiring all nodes to be aware of the global topology of the network. In our current dynamics, the active node can choose to unilaterally or bilaterally deviate, and can only do so with nodes in its local neighborhood.

Not only are the dynamics in [4] a very restricted version of those considered in this paper (as the type of deviations allowed during each stage are explicitly restricted), but also the nodes are required to have *global information* about the network (which is not realistic in most modern data-networks).

The main result of [4] is that the network converges to a stable and desirable equilibrium under certain conditions; we recover this result here, but for the more general model of dynamics.

- Further, in this paper we are able to prove a convergence result even in the setting where the link formation cost may be zero. This result is nontrivial because, as noted, the limiting topologies may not be trees (as is always the case in the analysis of [4]).

In contrast to our approach, which is inspired by the literature on network formation games, in recent years a large body of literature has considered efficiency loss due to self-interested behavior in routing-related applications; see, e.g., [8]. A related thread of literature has considered the efficiency of equilibria in *network design games*; see, e.g., [9, 10, 11]. In this line of work, a network is formed based on *unilateral* decisions of the players, and the cost of the resulting network is shared among the players. The goal of that body of research is to design mechanisms that incentivize agents to choose to use the network in an efficient manner.

Our work is also related to the literature on learning in games; see, e.g., [12, 13, 14] for surveys. In this literature the emphasis is on studying classes of dynamic methods that converge into the set of equilibria (e.g., correlated or Nash equilibria), without regard to efficiency. Our approach departs significantly from this literature, as we are interested in convergence to desirable equilibria.

Finally, there is an extensive body of research in the application of game theory to networks; see, e.g., [15] for a survey, and [16, 17] for a discussion of pricing in networks. In the application domain, our work is related to papers on topology formation in ad hoc networks; e.g., [18, 19, 20]. However, these works all consider Nash equilibria, whereas our focus is on *pairwise* interactions between nodes.

The rest of the paper is organized as follows. In Section 2 we present the notation used in the paper. In Section 3 we present the utility model considered. In Section 4 we define the static game, and recall the definition of pairwise Nash stability. In Section 6 we define the proposed two-stage dynamics. The main results are stated in Section 7 and are discussed in Section 8. The proofs of the main results are in online appendices in [21].

2 Notation

Let $G(V, E)$ be an undirected graph over V , with edge set E . We assume that V has n elements, and E has m elements. For a given pair of nodes u and v in V , we call $d(u, v; G)$ the distance in G (i.e., number of edges) from u to v . If u and v are in two distinct connected components, we set the distance $d(u, v; G)$ to be n . Finally, for a given node $u \in V$, we call $\delta_G(u)$ its degree in G .

3 Utility Model

We assume the nodes of the graph are strategic agents. We assume the cost to a node $v \in V$ is of the form:

$$C(v; G) = \pi\delta_G(v) + h(v; G), \quad (1)$$

for some positive function h . We assume that the cost is arbitrarily large (but finite) if G is disconnected; thus we will restrict our model and analysis to connected networks.

The function h models a situation where the cost to a node v is proportional to $f(v; G)$, the number of packets v forwards and receives in G .

We now define more precisely the function h .

Definition 1 (Routing-Related Cost). Let $v \in V$ be a node in a connected network G . Let $c_v > 0$ be given. We call c_v the per-unit routing cost associated to v . We define the routing-related cost to v in G as

$$h(v; G) = c_v f(v; G), \quad (2)$$

where $f(v, G)$ is determined by the routing policy and the traffic matrix. In our setting, we assume all-to-all traffic routed uniformly over all shortest paths.

Next, we assume that links in the network are the result of bilateral agreements between the participating nodes. We further assume that such an agreement induces a utility transfer, or payment, between the nodes participating in the agreement. Let \mathbf{P} be a n by n matrix where, for all $1 \leq i, j \leq n$, p_{ij} is the utility transfer from node i to node j . We call \mathbf{P} the *payment matrix*.

Let $G(V, E)$ be a given connected network topology, and let (P) be a given n by n matrix with real entries, and $P_{ii} = 0$ for all i . The total utility $U(i; G, \mathbf{P})$ to node i in G , given the payment matrix \mathbf{P} , is:

$$U(i; G, \mathbf{P}) = \sum_{j \neq i} (p_{ji} - p_{ij}) - C(i; G), \quad (3)$$

i.e., the total payments made to i , minus the total payments made by i , minus the cost of being in network topology G .

3.1 Contracting

In order to completely define the total utility to node i as in Equation 3, we need to establish how entries in the payment matrix \mathbf{P} are calculated. Recall that we assumed that a link $e = ij$ in G is the result of a bilateral agreement between i and j . We assume that the link arises due to a *contract* between i and j . Contracts are directional, i.e., the link ij may exist due to either the contract (i, j) or the contract (j, i) .

If $ij \notin G$, and the contract (i, j) is agreed upon by both i and j , then we assume that i pays j an amount $Q(i, j; G + ij)$; the function Q here is called the *contracting function*. This function gives the value of the contract formed between i and j , given the resulting network topology.

We believe two interpretations of the contracting function are reasonable. First, we might imagine that an external regulator has dictated that contracts between nodes must have pre-negotiated tariffs associated with them; these tariffs are encoded in the contracting function. Note that the regulator in this case dictates changes in the value of the contract as the surrounding network topology changes.

A second interpretation of the contracting function does not assume the existence of the regulator; instead, we presume that the value of the contracting function is the outcome of bilateral negotiation between the nodes in the contract. Note that the structure of our game assumes that this negotiation takes place *holding the network topology fixed*; i.e., the negotiation is used to determine the value of the contract, given the topology that is in place. One example is simply that $Q(i, j; G)$ is the result of a Rubinstein bargaining game of alternating offers between i and j , where i makes the first offer

[22]. If both players are infinitely patient, the resulting contracting function is identical to uniform cost sharing. More details can be found in [4].

We will be interested in contracting functions exhibiting two natural properties: *monotonicity* and *anti-symmetry*. Informally, monotonicity states that, given a network topology, the utility transfer associated with a contract is increasing in the burden associated to the contract proposed. Anti-symmetry asserts that $Q(i, j; G)$ must be equal to the negation of $Q(j, i; G)$. Anti-symmetric contracting functions ensure that directionality of the contract does not affect the value of the contract (i.e., (i, j) and (j, i) lead to the same utility transfer).

Assumption 1 (Anti-symmetry). *The contracting function is anti-symmetric, i.e., for all $u \neq v$ and for all G ,*

$$Q(u, v; G + uv) = -Q(v, u; G + uv)$$

Assumption 2 (Monotonicity). *We assume that the contracting function is monotone in the change of traffic cost to a node. In other words, let u, v and w be three distinct nodes such that $uv \notin G$ and $uw \notin G$. Let G_u, G_v and G_w be the connected components where u, v and w lie respectively. Then if $h(w; G_u \cup G_w + uv) - h(w; G_w) < h(v; G_v \cup G_u + uv) - h(v; G_v)$, then $Q(u, w; G + uv) < Q(u, v; G + uv)$; and if $h(w; G_u \cup G_w + uv) - h(w; G_w) = h(v; G_v \cup G_u + uv) - h(v; G_v)$, then $Q(u, w; G + uv) = Q(u, v; G + uv)$.*

3.2 State of the Game

From Equation 3 we see that the utility of all nodes is defined by G and \mathbf{P} . In order to keep track of the contracts in place at a given time in the network, we introduce the *contracting graph* Γ to be a directed graph over V such that $(i, j) \in \Gamma$ if nodes i and j have agreed to the corresponding contract.

Thus the state of the system is completely determined by the tuple (G, \mathbf{P}, Γ) .

Remark 1. Note that, from our contracting assumption, we have the following properties:

1. for all $i \neq j, ij \in G$ if and only if $(i, j) \in \Gamma$ or $(j, i) \in \Gamma$; and
2. for all $i \neq j, p_{ij} \neq 0$ implies $(i, j) \in \Gamma$.

4 Static Game

We can now formally define the game we consider. We use the same static game we defined in [4]. Note that variants of this game have been considered in previous literature (see [23, 24, 5]). We consider a network formation game in which each node selects nodes it wishes to connect to, as well as nodes it is willing to accept connections from. Formally, each node i simultaneously selects a subset $F_i \subseteq V$ of nodes i is willing to accept connections from, and a subset $T_i \subseteq V$ of nodes i wishes to connect to. We let $\mathbf{T} = (T_i, i \in V)$ and $\mathbf{F} = (F_i, i \in V)$ denote the composite strategy vectors.

An undirected link is formed between two nodes i and j if i wishes to connect to j (i.e., $j \in T_i$), and j is willing to accept a connection from i (i.e., $i \in F_j$). All edges that are formed in this way define the network topology $G(\mathbf{T}, \mathbf{F})$ realized by the strategy vectors \mathbf{T} and \mathbf{F} ; i.e., $j \in T_i, i \in F_j$ implies that $ij \in G(\mathbf{T}, \mathbf{F})$.

Further, if $i \in F_j$ and $j \in T_i$, then a *binding contract* is formed from i to j ; we denote this contract by (i, j) , and add it to the contracting graph $\Gamma(\mathbf{T}, \mathbf{F})$. The contracting graph captures the inherent directionality of link formation: in our model a link is only formed if one node asks for the link, and the target of the request accepts.

Finally, given a contracting function Q , as presented in Section 3.1 we define the payment matrix $\mathbf{P}(\mathbf{T}, \mathbf{F})$: if the formed network is G , then the contract (i, j) leads to a payment $Q(i, j; G)$ from i to j . This completely defines the state of the system. The utility of each node is defined as in Section 3 and thus the outcome of the game is well defined.

By an abuse of notation, and where clear from context, we will often use the shorthand $G = G(\mathbf{T}, \mathbf{F})$, $\Gamma = \Gamma(\mathbf{T}, \mathbf{F})$, and $\mathbf{P} = \mathbf{P}(\mathbf{T}, \mathbf{F})$ to represent specific instantiations of the network topology, contracting graph, and payment matrix, respectively, arising from strategy vectors \mathbf{T} and \mathbf{F} . We refer to a triple (G, Γ, \mathbf{P}) arising from strategic decisions of the nodes as a *feasible outcome* if there are strategy vectors \mathbf{T} and \mathbf{F} that give rise to (G, Γ, \mathbf{P}) .

5 Stability and Efficiency

In this section we define our solution concept, *pairwise Nash stability*, and formally define efficiency. While this is related to our earlier work in [4], the modification of our solution concept is an important change from our previous work: it allows us to present the dynamics in a succinct way (as will be clear in Section 6), and to better interpret the results (as will be clear in Section 8).

As is commonly observed in network formation games, Nash equilibrium lacks predictive power because link formation is inherently a *bilateral* process; thus we adopt the notion of *pairwise Nash stability* as our solution concept [2]. Informally, pairwise Nash stability requires that no unilateral deletion of contracts by a single node are profitable, and that no two nodes can simultaneously increase their utility by adding new contract(s) between them. In that sense, a network is pairwise Nash stable if it is a Nash network and pairwise stable (as originally defined by Jackson and Wolinsky in [3]).

Formally, suppose that the current strategy vectors are \mathbf{T} and \mathbf{F} , and the current network topology and contract graph are $G = G(\mathbf{T}, \mathbf{F})$ and $\Gamma = \Gamma(\mathbf{T}, \mathbf{F})$ respectively. First, suppose that node i attempts to unilaterally deviate. Then the strategy (T'_i, F'_i) if i after deviation is such that $(T'_i, F'_i) \subseteq (T_i, F_i)$.¹ Next, suppose that two nodes i and j attempt to bilaterally deviate; this involves changing the pair of strategies (T_i, F_i) and (T_j, F_j) together such that, after the deviation, $j \in T'_i$ & $i \in F'_j$ or $i \in T'_j$ & $j \in F'_i$. Any deviation will of course change both the network topology, as well as the contract graph.

¹ Although we restrict a unilateral deviation to only encompass the deletion of contracts, it can be shown that this is equivalent to allowing *any* unilateral deviation.

However, we assume that any contracts present both before and after the deviation *retain the same payment*. This is consistent with the notion of a contract: unless the deviation by i and j entails either breaking an existing contract or forming a new contract, there is no reason that the payment associated to a contract should change. With this caveat in mind, we formalize our definition of pairwise Nash stability as follows.

Definition 2. Assume Q is a contracting function. Given strategy vectors \mathbf{T} and \mathbf{F} , let $G = G(\mathbf{T}, \mathbf{F})$, $\Gamma = \Gamma(\mathbf{T}, \mathbf{F})$, and $\mathbf{P} = \mathbf{P}(\mathbf{T}, \mathbf{F})$. Given strategy vectors \mathbf{T}' and \mathbf{F}' , define $G' = G(\mathbf{T}', \mathbf{F}')$ and $\Gamma' = \Gamma(\mathbf{T}', \mathbf{F}')$. Define \mathbf{P}' according to:

$$P'_{k\ell} = \begin{cases} P_{k\ell}, & \text{if } (k, \ell) \in \Gamma' \text{ and } (k, \ell) \in \Gamma; \\ Q(k, \ell; G'), & \text{if } (k, \ell) \in \Gamma' \text{ and } (k, \ell) \notin \Gamma; \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

Then (\mathbf{T}, \mathbf{F}) is a pairwise Nash stable equilibrium if: (1) No unilateral deviation is profitable, i.e., for all i , and for all $\mathbf{T}' \subseteq \mathbf{T}$ and $\mathbf{F}' \subseteq \mathbf{F}$ that differ from \mathbf{T} and \mathbf{F} (respectively) only in the i 'th components,

$$U_i(\mathbf{P}, G) \geq U_i(\mathbf{P}', G');$$

and (2) no bilateral deviation is profitable, i.e., for all pairs i and j , and for all \mathbf{T}' and \mathbf{F}' that differ from $\mathbf{T} \subseteq \mathbf{T}'$ and $\mathbf{F} \subseteq \mathbf{F}'$ only in the i 'th and j 'th components,

$$U_i(\mathbf{P}, G) < U_i(\mathbf{P}', G') \implies U_j(\mathbf{P}, G) > U_j(\mathbf{P}', G').$$

Notice that (4) is a formalization of the discussion above. The first condition in the definition ensures no unilateral deviation is profitable, and the second condition ensures that if node i benefits from a bilateral deviation with j , then node j must be strictly worse off.

We will typically be interested in pairwise Nash stability of the network topology and contracting graph, rather than pairwise Nash stability of strategy vectors. We will thus say that a feasible outcome (G, Γ, \mathbf{P}) is a *pairwise Nash stable outcome* if there exists a pair of strategy vectors \mathbf{T} and \mathbf{F} such that (1) (\mathbf{T}, \mathbf{F}) is a pairwise Nash stable equilibrium; and (2) (\mathbf{T}, \mathbf{F}) give rise to (G, Γ, \mathbf{P}) . Note that by our definition of the game, for all i and j such that $(i, j) \in \Gamma$ we must have $P_{ij} = Q(i, j; G)$ in a pairwise Nash stable outcome.

We are also interested in system-wide performance from a global perspective, and for this purpose we must study the *efficiency* of pairwise Nash stable equilibria; we measure the efficiency of a network topology via the total value obtained by all nodes using that topology.

Given two feasible outcomes (G, Γ, \mathbf{P}) and $(G', \Gamma', \mathbf{P}')$, we say that (G, Γ, \mathbf{P}) Pareto dominates $(G', \Gamma', \mathbf{P}')$ if all players are better off in (G, Γ, \mathbf{P}) than in $(G', \Gamma', \mathbf{P}')$, and at least one is strictly better off. A feasible outcome is *Pareto efficient* if it is not Pareto dominated by any other feasible outcome. Since payoffs to nodes are *quasilinear* in our model, i.e., utility is measured in monetary units, it is not hard to show that a feasible outcome (G, Γ, \mathbf{P}) is Pareto efficient if $G \in \arg \min_{G'} S(G')$, where $S(G)$ is the *social cost function*:

$$S(G) = \sum_{i \in V} C_i(G).$$

We call such feasible outcomes *efficient*. (Note that, in particular, the preceding condition does not involve the contracting function; contracts induce zero-sum monetary transfers among nodes, and do not affect global efficiency.)

Let $V_0 = \{u \in V : \forall v \in V, c_u \leq c_v\}$, and for $u \in V_0$, let $c_{\min} = c_u$. In [25] we proved that, for $\pi > c_{\min}$, all efficient outcomes were stars centered around nodes in V_0 . Thus, in such settings, all efficient outcomes have the same number of edges.

An immediate consequence from the definition of pairwise Nash stability is that a node has to either delete edges or add an edge during a deviation, but not both. By assumption, all nodes experience an arbitrarily large cost when in a disconnected network. Thus we can prove the following important theorem.

Theorem 1 (Pairwise Nash Stable Networks). *Assume that G is connected, and assume $\pi > 0$. For any contracting function Q , there exists a pairwise Nash stable outcome (G, Γ, \mathbf{P}) if and only if G is a tree.*

Hence all pairwise Nash stable networks also have the same number of edges. Thus, whenever $\pi > c_{\min}$, we define the *efficiency ratio* of a given tree T as the ratio $S'(T)/S'(G_{\text{eff}})$ where S' is equal to $S - 2(n - 1)\pi$ (i.e. $S'(T)$ is the cost of routing traffic through T), and G_{eff} is the network topology in an efficient outcome.

6 Dynamics

Although the utility model discussed in Section 5 is essentially the same as that from [4] (with minor modifications), the dynamics considered in this paper constitute a major departure from our previous work in [4]. We describe our new dynamics in this section.

Before we begin, note a direct consequence of Theorem 1 is that the line network is pairwise Nash stable. It is easy to see that its efficiency ratio is linear in n , because its social cost is $O(n^3)$ while the social cost of an efficient network is $O(n^2)$. Thus the *price of anarchy* (as defined by Papadimitriou in [26]) is at least linear in n . Another consequence is that all efficient networks are pairwise Nash stable, whenever $\pi > c_{\min}$; thus it is important to try to select good equilibria (in terms of efficiency). The dynamics we consider are well matched to this purpose.

Let $\ell > 1$ be a given integer. We study discrete-time dynamics that proceeds in *rounds*. At each round k , an exogenous process (called an *activation process*) selects an active node $u_k \in V$. An important case that we study is the *uniform activation process*: at each round k , the active node is selected independently and uniformly at random from the set of nodes. Thus, under the uniform activation process, for all $k > 0$ and all $v \in V$, $\mathbb{P}[u_k = v] = 1/n$ independent of the past history of the activation process.

Let $(G^{(k)}, \mathbf{P}^{(k)}, \Gamma^{(k)})$ be the state of the network at the beginning of round k . The dynamics at round k proceeds in two *stages*.

During the first stage, the active node u_k selects a set of contracts (possibly empty) from $\Gamma^{(k)}$ it currently participates in, and removes them. All payments associated to those contracts are set to zero. If all contracts associated to an edge $u_k x \in G^{(k)}$ are removed, then the edge $u_k x$ is removed. Let $(G_1^{(k)}, \Gamma_1^{(k)}, \mathbf{P}_1^{(k)})$ be the resulting state of the network following stage 1 of round k .

During the second stage, the active node u_k either selects a new set of contracts (possibly empty) from $\Gamma_1^{(k)}$ it participates in and remove them, or it selects a node w in its ℓ -neighborhood from $G^{(k)}$, i.e., from among those nodes such that $d(u_k, w; G^{(k)}) \leq \ell$. In this case u_k proposes the contract (u_k, w) to w . If w accepts, the contract (u_k, w) is added to $\Gamma_1^{(k)}$, the edge $u_k w$ is added to $G_1^{(k)}$, and we set

$$p_{u_k w} = Q(u_k, w; G_1^{(k)} + u_k w).$$

Note that the active node only contemplates deviating with nodes in its ℓ -neighborhood, thus making our equilibrium selection process both *decentralized* and *local*.

We assume the dynamics are *myopic* in the following sense:

- the active node u_k selects its actions in order to maximize its utility at the end of the round; and
- the node w selected during the second stage accepts or rejects the contract in order to maximize its utility given the state at the end of the first stage, i.e., given $(G_1^{(k)}, \Gamma_1^{(k)}, \mathbf{P}_1^{(k)})$.

As a tie-breaking rule, we assume the following notion of “inertia.”

Assumption 3 (Inertia). *Let u_k be the active node at round k , and let $(G_1^{(k)}, \mathbf{P}_1^{(k)}, \Gamma_1^{(k)})$ be the state of the network after the first stage of round k . Let $W \subseteq V$ be the subset of nodes to whom u_k considers offering a contract at the second stage, i.e., such that the utility of u_k is maximized after the second stage. If $|W| > 1$, then u_k selects the node in W it was most recently connected to. If no such node exists, u_k picks one uniformly at random.*

Assumption 3 states that, if the active node has more than one optimal choice after the first stage, it will choose to deviate with the node it was most recently connected to.

A tie-breaking rule is necessary as the active node at a round, say u_k , may not have a unique utility-maximizing choice of a “partner” node at stage 2. Thus, in order to avoid oscillations induced by the possibility of multiple optimal choices, a tie-breaking rule must be assumed. While we have chosen a specific notion of inertia, we emphasize that many other assumptions can also lead to convergent dynamics. For instance, among utility-maximizing choices of w , if node u_k always chooses the node w with the highest degree, our convergence results remain valid.

Convergence of our dynamics is defined as follows.

Definition 3 (Convergence). *Given an initial state $(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)})$, we say that the dynamics converge if, almost surely, there exists $K > 0$ such that, for all $k > K$,*

$$(G^{(k)}, \mathbf{P}^{(k)}, \Gamma^{(k)}) = (G^{(k+1)}, \mathbf{P}^{(k+1)}, \Gamma^{(k+1)}).$$

7 Results

In this section we state and prove our main results. We interpret these results by analyzing the efficiency of the limiting topologies in our dynamics.

Theorem 2 (Convergence Theorem for $\pi > 0$). *Let $\ell \geq 2$ be given. Suppose that Assumptions [1](#) to [3](#) hold. Further, assume that $\pi > 0$.*

Let $(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)})$ be an outcome of the static game such that $G^{(0)}$ is connected. Assume that the activation process is such that, almost surely, all nodes are activated infinitely often. Then the dynamics started at $(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)})$ converge. Further, if the activation process is uniform, the convergence time is polynomial.

For a given realization of the activation process, let (G, \mathbf{P}, Γ) be the limiting state. Then:

1. *G is a tree where all internal (i.e., non-leaf) nodes are of minimum routing cost; and*
2. *the limiting state (G, \mathbf{P}, Γ) is pairwise Nash stable.*

In order to state our result when $\pi = 0$, we need two extra tie-breaking assumptions. The idea is that $\pi = 0$ can induce bilateral deviations where the increase in cost for adding a new edge is nil. For such deviations, we want the value of the corresponding contracts to be zero. This is consistent with our interpretation of contracts as a way incentivize nodes to accept connections even when their cost in the network increases. Thus, if the cost of adding a connection is zero, there should be no need for an incentive. That is the content of the following assumption.

Assumption 4 (Zero Value Contracts). *We assume that the contracting function yields a zero utility transfer if and only if there is no extra cost associated to adding the proposed contract. More formally, for all distinct nodes u and v , for all network topologies G , $Q(u, v; G) = 0$ if and only if $C(v; G + uv) = C(v, G)$.*

We now need to decide whether the active node should contemplate adding zero-value contracts. This is a tie-breaking rule as contracts can only be added during the second stage of the dynamics, and thus the utility of the active node would stay constant should it propose a zero-value contract. As a tie breaking rule, we assume that the active node would not propose a zero value contract [4](#).

Assumption 5 (Dynamics of Zero Value Contracts). *Let u_k be the active node, and let (G, \mathbf{P}, Γ) be the state of the network prior to the bilateral deviation considered. If the utility of u_k after successfully adding the link $u_k w$ (for any such w) is identical to that when in state (G, \mathbf{P}, Γ) , then u_k does not propose any contract to w .*

The motivation for Assumption [5](#) is that even if a contract has zero value, in reality there is some implicit “burden” to setting up a contract, so that establishing a contract

² Note that if a node is *proposed* a zero value contract we assume that it would accept it. The rationale behind Assumption [5](#) is that the active node can choose what node to bilaterally deviate with, which is not the case for a node being proposed a zero-value contract.

with no value is not desirable. Assumption 5 together with Assumption 4 implies that if the best action for the active node is to add a zero value contract, it prefers to pass.

In order to prove our second theorem, we also need to assume that $\ell \geq 3$, i.e. we allow nodes to be aware of their 3-neighborhood when considering their second stage deviation (recall that Theorem 2 requires only $\ell \geq 2$).

Theorem 3 (Convergence Theorem for $\pi = 0$). *Let $\ell \geq 3$ be given. Assume that Assumptions 1 to 5 hold. Further, assume that $\pi = 0$.*

Let $(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)})$ be an outcome of the static game such that $G^{(0)}$ is connected. Assume that the activation process is such that, almost surely, all nodes are activated infinitely often. Then the dynamics started at $(G^{(0)}, \mathbf{P}^{(0)}, \Gamma^{(0)})$ converge. Further, if the activation process is uniform, the convergence time is polynomial.

For a given realization of the activation process, let (G, \mathbf{P}, Γ) be the limiting state. Then the following hold:

1. *Let G' be the graph obtained by contracting all cliques in G and replacing them with a node whose per-unit routing cost is set to be that of the smallest routing cost in that clique. Then,*
 - (a) *G' is a tree;*
 - (b) *all internal nodes have minimum per-unit routing cost.*
2. *For any clique in G , only nodes with the smallest per-unit routing cost (among nodes of the same clique) can have edges to nodes outside of the clique.*
3. *The limiting state (G, \mathbf{P}, Γ) is pairwise Nash stable.*

Upon inspection of the dynamics considered in this paper, a generalization one might consider is to allow the active node to unilaterally or bilaterally deviate during *both* stages (recall that our dynamics allow only unilateral deviations during the first stage). However, as shown in [21], such dynamics may fail to converge under our assumptions.

8 Discussion of Results

In this section we provide a brief analysis of the results of Theorem 2, and a formal analysis of the results of Theorem 3.

In our previous work [4], under the same assumptions, we proved that a significantly restricted version of the dynamics considered in this paper converge to the same set of outcomes. In particular, in this paper activated nodes choose which links to break in the first stage of each round, whereas our model in [4] imposed random breaking of edges. Further, this paper analyzes a setting where nodes may be limited in their visibility of the network, and thus only contract with nodes in their ℓ -neighborhood; by contrast, the dynamics in [4] requires nodes be able to “see” the entire network (i.e., that $\ell = n$). In addition, we provide a novel proof technique via a Lyapunov argument that succinctly addresses these generalizations.

As argued in [4], our dynamics select good equilibria in that only nodes of minimum routing cost forward packets. An important special case is where there is a unique node i of minimum per-unit routing cost, i.e., such that $c_i < c_j$ for all $j \neq i$. In that case, our

dynamics converge to a star centered around node i , which is the most efficient pairwise stable outcome. In [27, 10], the term *price of stability* was coined to refer to the ratio of the efficiency of the best equilibrium to the optimum efficiency; thus our dynamics select an equilibrium that achieves the price of stability.

We now consider the case where $\pi = 0$ (cf. Theorem 3), with a view towards demonstrating that significantly different results are obtained when compared with the case where $\pi > 0$. First, we note that when $\pi = 0$, any outcome of the static game such that the network topology is the complete network is pairwise Nash stable. To see this, it suffices to note that when the network topology is the complete network, no node forwards any traffic, and thus the cost to all nodes is minimized. Next, by Assumption 4, it follows that all contracts in such network topology have a zero payment associated with them. Thus no unilateral deviation is profitable, and no bilateral deviation is possible as all edges are already part of the network. Recall that in [25] it is proved that, when $\pi < c_{\min}$, the only efficient outcomes of the static game are those where the network topology is the complete network. Thus, when $\pi = 0$, we conclude that *any efficient outcome is pairwise Nash stable*.

Surprisingly, allowing $\pi = 0$ can lead both to situations where *more* efficient outcomes are chosen than when $\pi > 0$, as well as situations where *less* efficient outcomes are chosen than when $\pi > 0$. First, from Theorem 2, note that when $0 < \pi < c_{\min}$, our dynamics *cannot* select an efficient outcome. By contrast, in [21] we select a set of parameters of the model, and construct a contracting function such that the efficient outcome is a fixed point of our dynamics when $\pi = 0$. Thus, by allowing $\pi = 0$, we can select more efficient outcomes than when $\pi > 0$.

However, allowing $\pi = 0$ can also make the *most inefficient* outcome selected worse. This is also shown in [21], where we select a set of parameters, and construct a contracting function such that the social cost of a fixed point when $\pi = 0$ is strictly larger than the social cost of *any* fixed point using the same parameters and contracting function when $0 < \pi < c_{\min}$. Thus, in this setting, by allowing $\pi = 0$ we can select *less* efficient outcomes than when $\pi > 0$.

Acknowledgments

The authors benefited from helpful conversations with John N. Tsitsiklis and Matthew O. Jackson. This work was supported by DARPA under the ITMANET Program, by the Canada Research Chairs Program, by the NSERC, and by the NSF under grant CMMI-0620811.

References

- [1] Jackson, M.O.: A survey of models of network formation: Stability and efficiency. Working Paper 1161, California Institute of Technology, Division of the Humanities and Social Sciences (2003)
- [2] Jackson, M.O.: The stability and efficiency of economic and social networks. Microeconomics 0211011, EconWPA (November, 2002)
- [3] Jackson, M.O., Wolinsky, A.: A strategic model of social and economic networks. Journal of Economic Theory 71(1), 44–74 (1996)

- [4] Arcaute, E., Johari, R., Mannor, S.: Network formation: Bilateral contracting and myopic dynamics. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 191–207. Springer, Heidelberg (2007)
- [5] Johari, R., Mannor, S., Tsitsiklis, J.N.: A contract-based model for directed network formation. *Games and Economic Behavior* 56(2), 201–224 (2006)
- [6] Jackson, M.O., Watts, A.: The existence of pairwise stable networks. *Seoul Journal of Economics* 14(3), 299–321 (2001)
- [7] Jackson, M.O., Watts, A.: The evolution of social and economic networks. *Journal of Economic Theory* 106(2), 265–295 (2002)
- [8] Roughgarden, T.: *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge (2005)
- [9] Lin Chen, H., Roughgarden, T., Valiant, G.: Designing networks with good equilibria. In: *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)* (2008)
- [10] Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-optimal network design with selfish agents. In: *Proceedings of ACM Symposium on the Theory of Computing*, pp. 511–520 (2003)
- [11] Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., Shabo, R.: On the price of stability for designing undirected networks with fair cost allocations. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) *ICALP 2006*. LNCS, vol. 4051, pp. 608–618. Springer, Heidelberg (2006)
- [12] Fudenberg, D., Levine, D.K.: *The Theory of Learning in Games*. MIT Press, Cambridge (1998)
- [13] Cesa-Bianchi, N., Lugosi, G.: *Prediction, Learning, and Games*. Cambridge University Press, Cambridge (2004)
- [14] Mannor, S., Shamma, J.: Multi-agent learning for engineers. *Artificial Intelligence*, 417–422 (2007); Special Issue on Foundations of Multi-Agent Learning
- [15] Altman, E., Boulogne, T., El-Azouzi, R., Jiménez, T., Wynter, L.: A survey on networking games in telecommunications. *Computers & Operations Research* 33(2), 286–311 (2006)
- [16] Falkner, M., Devetsikiotis, M., Lambadaris, I.: An overview of pricing concepts for broadband IP networks. *IEEE Communications Surveys* 3(2) (2000)
- [17] Briscoe, B., Darlagiannis, V., Heckman, O., Oliver, H., Siris, V., Songhurst, D., Stiller, B.: A market managed multiservice Internet (M3I). *Computer Communications* 26(4), 404–414 (2003)
- [18] Komali, R., MacKenzie, A.: Distributed topology control in ad hoc networks: A game theoretic perspective. In: *Proceedings of IEEE CCNC*, pp. 563–568 (2006)
- [19] Eidenbenz, S., Kumar, V., Zust, S.: Equilibria in topology control games for ad hoc networks. In: *Proceedings of International Conference on Mobile Computing and Networking*, pp. 2–11 (2003)
- [20] Fabrikant, A., Luthra, A., Maneva, E.N., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: *Proceedings of the Twenty-Second ACM Symposium on Principles of Distributed Computing (PODC 2003)*, Boston, Massachusetts, USA, July 13–16, 2003, pp. 347–351 (2003)
- [21] Arcaute, E., Johari, R., Mannor, S.: Local two-stage myopic dynamics for network formation games. Technical report, Stanford University Management Science and Engineering (2008)
- [22] Osborne, M.J., Rubinstein, A.: *Bargaining and Markets*. Academic Press, San Diego (1990)
- [23] Bloch, F., Jackson, M.O.: The formation of networks with transfers among players. *Journal of Economic Theory* 133(1), 83–110 (2007)

- [24] Corbo, J., Parkes, D.: The price of selfish behavior in bilateral network formation. In: Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing, PODC 2005, Las Vegas, NV, USA, July 17-20, 2005, pp. 99–107 (2005)
- [25] Arcaute, E., Johari, R., Mannor, S.: Network formation: Bilateral contracting and myopic dynamics. Technical report, Stanford University Management Science and Engineering (2007)
- [26] Papadimitriou, C.H.: Algorithms, games, and the internet. In: Proceedings on 33rd Annual ACM Symposium on Theory of Computing, Heraklion, Crete, Greece, July 6-8, 2001, pp. 749–753 (2001)
- [27] Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: 45th Symposium on Foundations of Computer Science (FOCS 2004), Rome, Italy, 17-19 October 2004, pp. 295–304 (2004)

Interference Games in Wireless Networks^{*}

Vincenzo Auletta, Luca Moscardelli, Paolo Penna, and Giuseppe Persiano

Dipartimento di Informatica ed Applicazioni “R. M. Capocelli”,
Via Ponte don Melillo - I-84084 Fisciano (SA), University of Salerno, - Italy

Abstract. We present a game-theoretic approach to the study of scheduling communications in wireless networks and introduce and study a class of games that we call *Interference Games*. In our setting, a player can successfully transmit if it “*shouts strongly enough*”; that is, if her transmission power is sufficiently higher than all other (simultaneous) transmissions plus the environmental noise. This physical phenomenon is commonly known as the Signal-to-Interference-plus-Noise-Ratio (SINR).

1 Introduction

We study *Interference Games* which arise in the context of wireless communications where multiple transmissions create interference and thus unnecessary energy loss for the nodes. Each node can be regarded as a player who has her own “profit” from successfully transmitting data, and a cost proportional to the energy spent for transmitting.

The scenario in which each player of the network acts independently so to optimize her own payoff (the “net profit” given by the energy loss and the success/unsuccess in transmitting) gives rise to an interesting class of games which we call Interference Games. Unlike the classical congestion games [14], in Interference Games there is a single resource (the physical media) but each player has a number of strategies available (the transmitting power). Players essentially compete for the media and, in a single slot, at most one player can transmit successfully. Indeed, a player transmits successfully if her signal strength at the receiver is larger than the sum of the signals of all other players plus the environmental noise (see Section 2). Though transmitting with higher power is more expensive, players may strategically decide to do so because they care more about successfully transmitting. This creates a mutual interference which may result in suboptimal performance like unnecessary energy consumption and/or transmissions failures (it may be well be the case that all players transmit with high power and thus they all fail).

1.1 Our Contribution

It is natural to ask how well does the system work if players optimize their own payoff, that is, if they only care about the success of their own transmission and the energy they

^{*} Research funded by the European Union through IST FET Integrated Project AEOLUS (IST-015964), and Information Communication Technologies Project FRONTS (project number 215270, Seventh Framework).

spent for it. To this aim, we study several notions of equilibria and their global efficiency. In particular, we consider *Nash equilibria* [12] where no player has an incentive to unilaterally change her strategy, *correlated equilibria* [23] where players have no incentive to deviate from a “suggested” strategy, and *sink equilibria* where players periodically perform best response and the whole system cycles through a “sink” consisting of a set of states of the game [6]. For all of these notions, it is possible to quantify the “system” performance in terms of *social welfare* and *fairness*. The former measures the overall “happyness” of the players, while the second one concerns how “equally” players have been treated. Finally, we consider *repeated games* [13], in which the “basic” Interference Game is played (possibly infinitely) many times. In such a context, we consider *subgame perfect equilibria* which provide a stronger solution concept compared to Nash equilibria (intuitively, the underlying “protocol” is also robust to deviations that occur for a finite amount of iterations).

We prove the following results on the existence and performance of the considered notions of equilibria for the case of a “perfectly symmetric” game in which all players value a successful transmission the same amount v , and they have the same set of strategies (see Section 2). We show that pure Nash equilibria do not exist if there are at least two transmission powers. Since mixed Nash equilibria and sink equilibria always exist [12,6], we consider these two notions. For two players, there exist sink equilibria with social welfare $v - k$, with k being the number of possible transmission powers, and these equilibria are also fair for odd k .

We show that every mixed Nash equilibrium has either bad social welfare or bad fairness (i.e., one of the two is equal zero). In contrast, we prove that correlated equilibria can be fair and attain a positive social welfare greater than $v - 2k$ (this improves to $v - k$ in the case of odd k). We also show optimal fair correlated equilibria for some specific games (namely, for $k \leq 3$).

Finally, we consider the case of infinitely *repeated* games with discount factor [13]. We prove that for the two-player scenario it is possible to obtain fair subgame perfect equilibria with optimal social welfare (i.e., $v - 1$). The result holds for the case in which every player knows only if her previous transmissions were successful or not.

1.2 Related Work

Fiat et al. [4] study contention resolution protocols for selfish agents aiming at accessing a broadcast channel. They focus on the scenario in which each player has one packet to transmit and she can choose either to transmit or not to transmit at each time slot (that is, each player has two possible strategies). They analyze the well known *Aloha* protocol and provide a new protocol being a Nash equilibrium for the game and having better performances (in terms to transmission delays) with respect to *Aloha*.

Adlakha et al. [1] study Bayesian Interference Games in a wireless scenario in which players select a power profile over the available bandwidth to maximize their own data rate (measured via Shannon capacity). They analyze Nash equilibria of the incomplete information game in which players are unaware of the interference they cause to the other ones.

In Timing Games [8,9,10] two players must decide when to make a single move at some time between 0 and T . The payoffs of the players usually depend on which

player moves first and/or the time that she moves. Though Timing Games can be seen as special Interference Games, in Timing Games the payoffs are always positive which is not true for Interference Games. This determines a different structure of equilibria.

Scheduling wireless communications under the Signal-to-Interference-plus-Noise-Ratio (SINR) model have been studied in [711].

2 Model and Definitions

In the Signal-to-Interference-plus-Noise-Ratio (SINR) model (see for instance [711]), a node α is successful in transmitting if and only if

$$\frac{p_\alpha/d_\alpha^a}{\text{Noise} + \sum_{\beta \neq \alpha} (p_\beta/d_\beta^a)} \geq B, \quad (1)$$

where d_α is the distance of node α from the receiver and p_α is the power of node α 's transmission.

We study the SINR model from a game theoretical point of view and introduce a class of games which we call *Interference Games*. There are n players corresponding to the nodes aiming to communicate. A strategy of a player α is an integral power transmission level in $\{0, 1, \dots, k\}$ and all players have the same set of strategies. Moreover, we denote by v how much a successful transmission is worth to a player (we assume this value to be the same for all players).

Given a strategy profile or *state* $s = (s_1, \dots, s_\alpha, \dots, s_n)$ of the game, player α is *successful* if s_α is larger than the sum of all other s_i 's. Notice that s_α is the power of α and the condition for being successful corresponds to the case in which all nodes are at the same distance from the receiver, $\text{Noise} > 0$, and $B = 1$ (see Equation 1). The *utility* or *payoff* $u_\alpha(s)$ of player α depends on her power consumption and the fact that her transmission is successful or not. Namely, if in s player α is successful and has used power p_α , then her payoff is $u_\alpha(s) = v - p_\alpha$. Otherwise, if in s player α is not successful and has used power p_α , then her payoff is $u_\alpha(s) = -p_\alpha$. If we deal with probabilistic choices, we are interested in the expected utility. Each player aims to maximize her own (expected) utility.

We consider the *social welfare* function $SW(s) = \sum_{\alpha \in N} u_\alpha(s)$ that is the sum of the payoff's of all players. The *fairness* of a state s is defined as the ratio between the minimum and the maximum (expected) utility of players; i.e., $\frac{\min_{\alpha \in N} u_\alpha(s)}{\max_{\alpha \in N} u_\alpha(s)}$; if the utilities of all the players are 0, the fairness is defined equal to 1. Moreover, we call *fair* a state with fairness equal to 1, and *unfair* a state with fairness equal to 0.

We now review the equilibrium notions that we use in this paper. A *pure Nash equilibrium* is a state in which no player can obtain a higher utility by changing her strategy, given the strategies of the other players. In a *mixed Nash equilibrium* we consider players picking a strategy independently according to some probability distribution (each player decides her own distribution). In a mixed Nash equilibrium no player can improve her expected payoff by changing her probability distribution, given the probability distributions of the other players. In *correlated equilibria* a "mediator" picks a state s according to some probability distribution and "suggests" strategy s_α to each player α .

Each player α is only aware of her suggested strategy and of the probability distribution used to pick the state. A probability distribution over the set S of all states is a correlated equilibrium if no player can improve her expected payoff by replacing her suggested strategy with a different one, given that the other players follow the suggested strategy (note that the expected payoff of α is *conditional* to the fact that player α has been suggested some strategy s_α). In *sink equilibria* we consider a so called state graph in which every node corresponds to a state of the game and there is a directed edge from s to s' if there is a player α such that $u_\alpha(s) < u_\alpha(s')$ and state s' is obtained from s by changing strategy s_α with some other strategy s'_α . Intuitively, edges corresponds to best response of some player and, in a sink equilibrium, players moves will “cycle” through some connected component (when the component has only one node we have a pure Nash equilibrium). More formally, sink equilibria are the strongly connected components of the state graph. Let Q be a sink equilibria and let $\pi : Q \rightarrow \mathbb{R}^+ \cup \{0\}$ be the steady state distribution of the random walk over states $q \in Q$ of the sink equilibrium. The (expected) social welfare of Q is the expected social value of states given by the steady distribution of the random walk over its states; i.e. $SW(Q) = \sum_{q \in Q} \pi(q)SW(q)$.

3 A Simple Interference Game

In this section we analyze a simple Interference Game characterized by $n = 2$ players and $k = 2$. The game is perfectly symmetric and a player α is successful if and only if $p_\alpha > p_\beta$, where β is the other player. Despite its simplicity, we can already derive some indications from this simple game. The utility matrix is

	0	1	2
0	0, 0	0, $v - 1$	0, $v - 2$
1	$v - 1, 0$	-1, -1	-1, $v - 2$
2	$v - 2, 0$	$v - 2, -1$	-2, -2

We start by proving that this simple game has no pure Nash equilibrium.

Theorem 1. *For any $k \geq 2$, the Interference Game has no pure Nash equilibrium, even for two players.*

Proof. Observe that the best response for a player to strategy $x < k$ of the other player is strategy $x + 1$ and the best response to strategy k is strategy 0. Therefore, the only possible pure equilibria are $(0, k)$ or $(k, 0)$. Since the best response to strategy 0 is strategy 1, such states are Nash equilibria only if $k = 1$. Similar arguments apply for the case of $n > 2$ players.

We now turn our attention to sink equilibria.

Theorem 2. *The Interference Game with $n = 2$ and $k = 2$ has a unique sink equilibrium with social welfare $v - 2$ and fairness 1.*

Proof. By recalling the above considerations on the best response moves, there exists a unique sink equilibrium given by the cycle:

$$(0, 1), (2, 1), (2, 0), (1, 0), (1, 2), (0, 2), (0, 1).$$

Since the equilibrium is a cycle, the steady distribution of the random walk is the uniform one, and it is easy to check that its social value is $v - 2$. Moreover, since for each state in the sink also its symmetric state is present, the fairness of the equilibrium is 1.

We continue our study by analyzing mixed Nash equilibria.

Theorem 3. *The Interference Game with $n = 2$ and $k = 2$ has a mixed Nash equilibrium with social welfare 0 and fairness 1.*

Proof. The equilibrium corresponds to the probability distribution $q = (q_0, q_1, q_2)$ with $q_0 = q_1 = 1/v$. To see that this is a Nash equilibrium, let $u_\alpha^{(q)}(i)$ be the payoff of player α when it plays strategy i , given that the other one plays according to probability distribution q . Clearly $u_\alpha^{(q)}(0) = 0$, while

$$u_\alpha^{(q)}(1) = q_0(v - 1) - q_1 - q_2 = q_0v - 1 = 0.$$

Similarly

$$u_\alpha^{(q)}(2) = q_0(v - 2) - q_1(v - 2) - 2q_2 = (q_0 + q_1)v - 2 = 0.$$

Since the payoff is constant for all three strategies, when both players play according to the probability distribution q , none has an incentive in unilaterally deviating. That is, q is a Nash equilibrium and the payoff of each node is 0; thus, also the social welfare is 0 and the equilibrium is fair.

We conclude the study of the case $n = 2$ and $k = 2$ by showing the best possible correlated equilibrium X , and proving that its social welfare is very close to the optimum. Each player receives a suggestion on the power to use for the transmission. We denote by $x(i, j)$ the probability that the first and the second players are suggested to transmit at power i and j , respectively. We will consider only symmetric distributions, that is, distributions for which $x(i, j) = x(j, i)$ that thus give fair correlated equilibria. We denote by q_{ij} the probability that player 2 receives suggestion j given that player 1 has received suggestion i , that is $q_{ij} = \frac{x(i,j)}{\sum_h x(i,h)}$. Since transmission of player 1 at power i is successful if and only if player 2 transmits at a power $j < i$, we have that the probability $Pr[j|i]$ that player 1 is successful at power j given that he was suggested to transmit at power i is equal to

$$Pr[j|i] = \sum_{h < j} q_{ih}$$

and the expected payoff $u_\alpha[j|i]$ of player α when transmitting at power j , given that he was suggested to transmit at power i , is equal to $vP_i[j|i] - j$. The definition of correlated equilibrium is that $u_\alpha[i|i] \geq u_\alpha[j|i]$.

Theorem 4. *For the Interference Game with $n = 2$ and $k = 2$, the optimal symmetric correlated equilibria has social welfare $v - 2 + \frac{v}{v^2 - 2}$.*

Sketch of proof. The following matrix turns out to be an optimal correlated equilibrium:

$$X := \begin{bmatrix} 0 & \frac{v-1}{v^2-2} & \frac{v^2-3v+2}{2v^2-4} \\ \frac{v-1}{v^2-2} & 0 & \frac{v-2}{2v^2-4} \\ \frac{v^2-3v+2}{2v^2-4} & \frac{v-2}{2v^2-4} & 0 \end{bmatrix}$$

thus proving the theorem.

4 Two Players and Arbitrarily Many Strategies

While for $k = 2$ the only Nash equilibrium has social welfare 0, it turns out that when k is odd there are Nash equilibria whose social welfare is strictly positive.

Theorem 5. *For every odd k , there exists a mixed Nash equilibrium with social welfare $v - k$.*

Sketch of proof. The following matrix is a mixed Nash equilibrium:

$$CF := \begin{bmatrix} 0 & \frac{2}{v} & 0 & \frac{2}{v} & 0 & \dots & \frac{2}{v} & 0 & 1 - \frac{k-1}{v} \\ 1 - \frac{k-1}{v} & 0 & \frac{2}{v} & 0 & \frac{2}{v} & \dots & 0 & \frac{2}{v} & 0 \end{bmatrix}.$$

Since CF is an unfair Nash equilibrium, we next investigate the existence of fair equilibria.

Theorem 6. *For any Interference Game, there exists a unique (fair) fully mixed Nash equilibrium, that is, a Nash equilibrium in which every player assigns nonzero probability to every strategy. Moreover, every fair Nash equilibrium has social welfare 0.*

Sketch of proof. In a fully mixed equilibrium, strategy 0 is in the support of every player which implies that the expected payoff of every player must be 0. Calculations show that the condition for having a Nash equilibrium impose that the probability distribution of each player is $q = (\frac{1}{v}, \frac{1}{v}, \dots, \frac{1}{v}, 1 - \frac{k}{v})$.

At Nash equilibrium, at least one player must have 0 in her support. Thus, in every fair Nash equilibrium 0 is in the support of all players and therefore the social welfare must be 0.

Correlated equilibria can be both fair and achieve good social welfare:

Theorem 7. *For any Interference Game there exists a fair correlated equilibrium with social welfare greater than $\max(0, v - 2k + 1)$.*

Sketch of proof. We modify the joint probability distribution of the Nash equilibrium given in the proof of Theorem 6 and obtain a correlated equilibrium given by the following matrix:

$$C = \begin{bmatrix} 0 & c2\lambda^2 & \dots & c\lambda^2 & c\lambda(1 - k\lambda) \\ 2c\lambda^2 & 0 & \dots & c\lambda^2 & c\lambda(1 - k\lambda) \\ \dots & \dots & \dots & \dots & \dots \\ c\lambda^2 & c\lambda^2 & \dots & 0 & c(\lambda(1 - k\lambda) + \lambda^2) \\ c\lambda(1 - k\lambda) & c\lambda(1 - k\lambda) & \dots & c(\lambda(1 - k\lambda) + \lambda^2) & 0 \end{bmatrix}$$

where $\lambda = 1/v$ and c is a suitable constant such that C is a probability distribution.

The social welfare can be further improved for even k :

Theorem 8. *For any Interference Game with k even there exists a fair correlated equilibrium with social welfare at least $v - k$. Moreover, for $k = 3$ there exists an optimal symmetric Correlated Equilibrium with social welfare $v - 3 + \frac{4v^2 - 11v + 7}{v^3 - v^2 - 4v - 5}$.*

We next generalize the sink equilibria described in Section 3. The main difference is that for odd k there exist two sink equilibria both with fairness less than 1.

Theorem 9. *The Interference Games with $n = 2$ and k even have a unique sink equilibrium with social welfare $v - k$ and fairness 1. The Interference Games with $n = 2$ and k odd have two sink equilibria with social welfare $v - k$ and fairness $\frac{(2k-2)v - 2k^2 + 2}{(2k+6)v - k^2 - 2k - 1}$.*

Sketch of proof. For $k = 3$, there are two sink equilibria:

$$(0, 1), (2, 1), (2, 3), (0, 3), (0, 1) \text{ and } (1, 0), (1, 2), (3, 2), (3, 0), (1, 0).$$

In each of them, one player has expected utility $\frac{v}{4} - 1$, and the other one $\frac{3}{4}v - 2$. Therefore, the fairness is $\frac{v-4}{3v-8}$. A similar argument generalizes to any even k .

5 Arbitrarily Many Players

The following theorem extends the results on Nash equilibria for two players given in Section 5.

Theorem 10. *There exists a fair Nash equilibrium with $n \geq 3$ players with social welfare equal to 0. Moreover, if k is odd, there exists an unfair Nash equilibrium with $n \geq 3$ players with social welfare equal to $v - k$.*

Sketch of proof. It is possible to show that, given a Nash equilibrium for the case of two players and in which at least one player having expected utility 0, it is possible to obtain a Nash equilibrium for $n \geq 3$ players by adding $n - 2$ players playing strategy 0 with probability 1. The theorem thus follows from the results on two players (Theorems 5,6).

Correlated equilibria achieve both fairness and good social welfare:

Theorem 11. *For any n and for odd k there exists a fair Correlated Equilibrium with social welfare $v - k$.*

6 Repeated Interference Games

In the repeated interference game, the same interference game is played (infinitely) many times and, at each repetition i , player α accumulates a new payoff $\delta^i \cdot u_\alpha(s^{(i)})$, where $s^{(i)}$ are the strategies played at repetition i and $\delta < 1$ is the discount factor. A simple protocol for two players consists in alternating transmissions, with the transmitting player using power 1; Every deviation from this results in a “punishment” phase in which both players transmit with maximal power for prescribed amount of time steps; Deviations from the punishment phase will “restart” of the punishment phase itself. This results in an optimal subgame perfect equilibrium:

Theorem 12. *For every v and k there exists $\bar{\delta} < 1$ such that the following holds. For any $\delta > \bar{\delta}$, the repeated Interference Game with v and k and discount factor δ has a fair subgame perfect equilibrium with expected payoff profile $((v - 1)/2, (v - 1)/2)$.¹*

The main idea is that if a player deviates from this punishment phase, this can be detected by the other player who sees that her transmission is successful (deviations from the “non-punishment” phase are detected because of transmission failure). This is sufficient for applying the result in [5].

References

1. Adlakha, S., Johari, R., Goldsmith, A.J.: Competition in wireless systems via bayesian interference games. CoRR, abs/0709.0516 (2007)
2. Aumann, R.J.: Subjectivity and correlation in randomized games. *Journal of Mathematical Economics* 1, 67–96 (1974)
3. Christodoulou, G., Koutsoupias, E.: On the price of anarchy and stability of correlated equilibria of linear congestion games. In: Brodal, G.S., Leonardi, S. (eds.) *ESA 2005*. LNCS, vol. 3669, pp. 59–70. Springer, Heidelberg (2005)
4. Fiat, A., Mansour, Y., Nadav, U.: Efficient contention resolution protocols for selfish agents. In: *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007*, New Orleans, Louisiana, USA, pp. 179–188 (2007)
5. Fudenberg, D., Maskin, E.: The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54(3), 533–554 (1986)
6. Goemans, M.X., Mirrokni, V.S., Vetta, A.: Sink equilibria and convergence. In: *FOCS*, pp. 142–154. IEEE Computer Society, Los Alamitos (2005)
7. Goussevskaia, O., Oswald, Y.A., Wattenhofer, R.: Complexity in Geometric SINR. In: *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, Montreal, Canada (September 2007)
8. Hendricks, K., Weiss, A., Wilson, C.: The war of attrition in continuous time with complete information. *International Economic Review* 29(4), 663–680 (1988)
9. Hendricks, K., Wilson, C.: Discrete versus continuous time in games of timing. Working Papers 85-41, C.V. Starr Center for Applied Economics, New York University (1985)
10. Lotker, Z., Patt-Shamir, B., Tuttle, M.R.: Timing games and shared memory. In: Fraigniaud, P. (ed.) *DISC 2005*. LNCS, vol. 3724, pp. 507–508. Springer, Heidelberg (2005)
11. Moscibroda, T., Wattenhofer, R., Zollinger, A.: Topology Control Meets SINR: The Scheduling Complexity of Arbitrary Topologies. In: *7th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*, Florence, Italy (May 2006)
12. Nash, J.F.: Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences* 36, 48–49 (1950)
13. Osborne, M.J., Rubinstein, A.: *A course in Game Theory*. MIT Press, Cambridge (1994)
14. Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)

¹ Every two rounds the “non-discounted” payoff of each player is $v - 1$.

Taxing Subnetworks^{*}

Martin Hoefer, Lars Olbrich, and Alexander Skopalik

Department of Computer Science, RWTH Aachen University, Germany

Abstract. We study taxes in the well-known game theoretic traffic model due to Wardrop. Given a network and a subset of edges, on which we can impose taxes, the problem is to find taxes inducing an equilibrium flow of minimal network-wide latency cost. If all edges are taxable, then marginal cost pricing is known to induce the socially optimal flow for arbitrary multi-commodity networks. In contrast, if only a strict subset of edges is taxable, we show NP-hardness of finding optimal taxes for general networks with linear latency functions and two commodities. On the positive side, for single-commodity networks with parallel links and linear latency function, we provide a polynomial time algorithm for finding optimal taxes.

1 Introduction

An important problem in traffic management is to set incentives for rational users to act in a favorable manner. An effective means to achieve this is to set appropriate taxes. In this paper, we study the problem of computing optimal taxes in the Wardrop model, a well-studied model for traffic routing with important applications in road networks and computer networks. In this model, we are given a network equipped with non-decreasing non-negative latency functions mapping flow on the edges to latency. For each of several commodities a fixed demand has to be routed between a source-sink pair. The cost of a flow assignment is the weighted sum of travel times between the source and target nodes. A flow that minimizes the total latency is called (*socially optimal*). A common interpretation of the Wardrop model is that flow is controlled by an infinite number of selfish users each of which carries an infinitesimal amount of flow. Each user aims at minimizing its path latency. An allocation, in which no user can improve its situation by unilaterally deviating from its current path is called *Wardrop equilibrium*. In general a Wardrop equilibrium is not socially optimal, i.e., it does not minimize the total latency. The inefficiency of selfish flows has been extensively studied in previous work [2, 18, 19, 21].

Taxing can be successful in improving total latency of equilibria. In this case users are assumed to minimize the sum of their latencies and taxes. A fundamental result is that using *marginal cost pricing* to tax every edge results in equilibrium flows that are optimal with respect to total latency [11]. A serious drawback of marginal cost pricing is that it requires *every* edge of the network to be taxable. In many situations there might be technical or legal restrictions that prevent an operator from imposing a tax on all

^{*} Supported by the DFG GK/1298 “AlgoSyn” and by the German Israeli Foundation (GIF) under contract 877/05.

edges. Therefore, we adjust the model to a more realistic case in which only a subset of edges can be taxed. The problem is to find a set of taxes for the subset of taxable edges that minimizes the total latency of the resulting Wardrop equilibrium. To the best of our knowledge, this generalization has not been considered before.

Taxing subnetworks can be more difficult and non-trivial. Consider a parallel link network of two links and linear latency functions. If one can tax only one edge, the latency cost is generally not monotone in the imposed tax. Using this insight, we carefully construct networks with one taxable edge and several distinct optimal taxes. A combination of these networks establishes NP-hardness of the problem for two commodities and linear latency functions in Sect. 3. On the other hand, for parallel link networks with linear latency functions, we derive a precise structural analysis of optimally taxed equilibrium flows in Sect. 4. This allows to construct a polynomial-time algorithm to find optimal taxes. Most proof details are omitted and will be given in the full version of the paper.

Related Work. There is a huge amount of work addressing the inefficiency of equilibria in the Wardrop model. Therefore, we only give a rough overview and concentrate on the classical results and recent developments. The game theoretic traffic model considered in this paper was introduced by Wardrop [24]. Beckmann et al. [1] observe that such an equilibrium flow is an optimal solution to a related convex program. They give existence and uniqueness results for traffic equilibria (see also [6] and [19]). Dafermos and Sparrow [6] show that the equilibrium state can be computed efficiently under some assumptions on the latency functions.

The inefficiency of Wardrop equilibria is a well-known phenomenon [16], which is exemplified by Braess paradox [2]. Bounding the inefficiency of equilibria, however, has only recently been considered, initiated by Koutsoupias and Papadimitriou [14], and for the Wardrop model by Roughgarden and Tardos [19]. Roughgarden [21] provides a cumulative overview of the most important results that have been obtained.

There are several approaches that have been proposed to address the inefficiency of equilibria. The effectiveness of taxes has been observed by Pigou [16] and generalized by Beckmann et al. [1]. They show that *marginal cost pricing* completely eliminates the inefficiency of selfish routing. Cole et al. [5] show existence of taxes inducing the optimal flow for single-commodity networks and heterogeneous users that value tax versus latency in an individual way. Fleischer [7] reduces the required taxes to linear functions. In the more general setting of multi-commodities, Fleischer et al. [8] and Karakostas and Kolliopoulos [10] independently prove the existence of optimal taxes.

Other approaches for coping with selfishness are, for example, proposed by Korilis et al. [13], who give methods for improving system performance by adding additional capacity to system resources. Cocchi et al. [3] study the role of various pricing policies in networks with selfish users. Roughgarden [20] studies designing networks that exhibit good performance when used selfishly and proves tight inapproximability results. Cole et al. [4] show hardness of computing taxes minimizing the total user disutility (latency plus tax) at equilibrium.

Korilis et al. [12] consider the problem of a Stackelberg leader, who in a first phase can fix the routes for a certain fraction of the demand. In a second phase, selfish users enter the system and route their own flow on top of the leader demand. The objective

of the leader is to minimize the resulting total cost of the total (both leader and selfish) flow. Roughgarden [17] shows that it is weakly NP-hard to compute the optimal leader strategy even for parallel links with linear latency functions. Kumar and Marathe [15] give a FPAS for this problem. Kaporis and Spirakis [9] show that for single-commodity networks the minimal fraction of flow needed by the leader to induce optimal cost can be computed in polynomial time. Subsequent papers [23, 22, 11] consider Stackelberg routing in different variants for more general networks.

2 Preliminaries

We consider Wardrop’s traffic model originally introduced in [24]. We are given a directed graph $G = (V, E)$ with vertex set V , edge set E , a set of commodities $[k] = \{1, \dots, k\}$ specified by source-sink pairs $(s_i, t_i) \in V \times V$, and flow demands $d_i > 0$. For single-commodity networks we normalize the demand to one. Considering only parallel edges, we speak of parallel link networks and denote the set of links by $[n] = \{1, \dots, n\}$. The edges are equipped with non-decreasing, continuous latency functions $\ell_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. We allow a set of non-negative taxes $\{\tau_e\}_{e \in T}$ to be imposed on a subset of edges $T \subset E$. We call edges in T *taxable* and edges in $N = E \setminus T$ *non-taxable*.

Let \mathcal{P}_i denote the admissible paths of commodity i , i. e., all paths connecting s_i and t_i , and let $\mathcal{P} = \bigcup_{i \in [k]} \mathcal{P}_i$. A non-negative path flow vector $(f_P)_{P \in \mathcal{P}}$ is *feasible* if it satisfies the flow demands $\sum_{P \in \mathcal{P}_i} f_P = d_i$ for all $i \in [k]$. Throughout this paper, we will consider only feasible path flow vectors. A path flow vector $(f_P)_{P \in \mathcal{P}}$ induces an edge flow vector $f = (f_e)_{e \in E}$ with $f_e = \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i: e \in P} f_P$. For single-commodity networks, we drop the index i . The latency of an edge $e \in E$ is given by $\ell_e(f_e)$ and the latency of a path P is given by the sum of the edge latencies $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$. The latency cost of a flow is defined as $C(f) = \sum_{P \in \mathcal{P}} \ell_P(f) f_P = \sum_{e \in E} \ell_e(f_e) f_e$. A flow f of minimal latency cost is called (*socially*) *optimal*. The cost of a path is defined as latency plus tax, i. e., $\ell_P(f) + \sum_{e \in P} \tau_e$. Finally, we call the quadruple $(V, T, N, (d_i))$ an *instance*.

A flow vector is considered stable when no fraction of the flow can improve its sustained cost by moving unilaterally to another path. Such a stable state is generally known as *Nash equilibrium*. In our model a flow is stable if and only if all used paths within a commodity have the same minimal cost, whereas unused paths may have larger cost. We call such a flow *Wardrop equilibrium*.

Definition 1. A feasible flow vector f is at Wardrop equilibrium if for every commodity $i \in [k]$ and paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ it holds that $\ell_{P_1}(f) + \sum_{e \in P_1} \tau_e \leq \ell_{P_2}(f) + \sum_{e \in P_2} \tau_e$.

In particular, without taxes, if f is at Wardrop equilibrium then all used paths in commodity i have equal latency $L_i(f)$ and the latency cost can be expressed as $C(f) = \sum_{i \in [k]} L_i(f) \cdot d_i$ (see [19, 24]). A classical result on taxing selfish flow, called *marginal cost pricing*, is that with taxes $\tau_e = x_e \cdot \ell'_e(x_e)$ for all $e \in E$ the resulting equilibrium flow minimizes the latency cost. With $\ell_e^*(x) = (x \cdot \ell_e(x))' = \ell_e(x) + x \cdot \ell'_e(x)$ denoting the marginal cost of increasing flow in edge e we have the following lemma.

Lemma 1 ([1, 6, 19]). *Let $(V, T, \emptyset, (d_i))$ denote an instance in which $x \cdot \ell_e(x)$ is a convex function for each edge e . Then a flow f minimizes the latency cost w.r.t. $(\ell_e)_{e \in T}$ if and only if it is at Wardrop equilibrium w.r.t. $(\ell_e^*)_{e \in T}$.*

In the restricted case with only a subset of edges being taxable such a result is obviously out of reach. This directly leads us to the following definition.

Definition 2. *Given an instance $(V, T, N, (d_i))$, a set of taxes $\{\tau_e\}_{e \in T}$ is called optimal, if there is an equilibrium flow f_τ w.r.t. $\ell + \tau$ with $C(f_\tau) \leq C(f_{\tau'})$ for all equilibrium flows $f_{\tau'}$ w.r.t. $\ell + \tau'$ for any $\{\tau'_e\}_{e \in T}$.*

3 NP-Hardness for Multi-commodity Networks

In this section we study the optimization problem of computing an optimal set of taxes. We show that this turns out to be NP-hard even for the two-commodity case with linear latency functions. We start with an observation which allows us to discretise the problem and enables us to prove the main result of this section.

Lemma 2. *There is a family of instances $(V, T, N_A, d_A)_{A \in \mathbb{N}}$ with parallel link networks allowing for two separated optimal tax values.*

Proof. Consider a parallel link network, in which two nodes s and t are connected via three links with $\ell_1(x) = x + A$ and $\ell_2(x) = \ell_3(x) = x$. Suppose we can only tax the third link. Set $d_A = A(1 + \frac{\sqrt{3}}{2})$. For tax $0 \leq \tau \leq A(1 - \frac{\sqrt{3}}{2})$, the total demand is split among links two and three at equilibrium. Since both used links are identical, $\tau = 0$ is optimal with an induced cost of $(\frac{7}{8} + \frac{\sqrt{3}}{2})A^2$. For $A(1 - \frac{\sqrt{3}}{2}) < \tau < A(1 + \frac{\sqrt{3}}{4})$ all links are used and the corresponding cost function $\frac{2}{3}\tau^2 - \frac{1}{3}A\tau + (\frac{11}{12} + \frac{\sqrt{3}}{2})A^2$ yields an optimal tax of $A/4$ with cost $(\frac{7}{8} + \frac{\sqrt{3}}{2})A^2$ as well. For $\tau \geq A(1 + \frac{\sqrt{3}}{4})$ the latency cost at equilibrium is $(\frac{11}{8} + \frac{3\sqrt{3}}{4})A^2$. Thus, both $\tau = 0$ and $\tau = A/4$ are optimal. \square

Theorem 3. *Given an instance $(V, T, N, (d_i))$, the problem of computing optimal taxes is NP-hard, even for only two commodities and linear latency functions.*

Proof. We reduce from the PARTITION problem: given n positive integers a_1, \dots, a_n , is there a subset $S \subseteq \{1, 2, \dots, n\}$ satisfying $\sum_{i \in S} a_i = \frac{1}{2} \sum_{i=1}^n a_i$? We will show that deciding the PARTITION problem reduces to deciding if a given 2-commodity instance $(V, T, N, (d_i))$ with latency functions admits taxes inducing a Wardrop equilibrium with a given cost. Given an arbitrary instance of PARTITION specified by positive integers a_1, \dots, a_n , we define an instance $(V_{\{a_i\}}, T_{\{a_i\}}, N_{\{a_i\}}, (d_{\{a_i\}}))$ as depicted in Fig. 1. Let the set of taxable edges T consist of the bold edges. Commodity one has a demand of $A = \prod_{i=1}^n a_i$ to route between $s_1 = v_1$ and $t_1 = v_{n+1}$, the second commodity has to route a demand of $\sum_i a_i$ between s_2 and t_2 . For $i \in [n]$ define the following constants: $A_{-i} = \prod_{j \neq i}^n a_j$, $D_i = \frac{2-4A_{-i}+A_{-i}^2}{4A_{-i}-2}$ and $E_i = 2a_i(D_i + 1) = \frac{AA_{-i}}{2A_{-i}-1}$. We show that $\{a_1, \dots, a_n\}$ is a YES instance if and only if there are taxes for instance $(V_{\{a_i\}}, T_{\{a_i\}}, N_{\{a_i\}}, (d_{\{a_i\}}))$ inducing a Wardrop equilibrium with cost of at most $C = \frac{n}{2}A^2 + \frac{7}{8}(\sum_i a_i)^2$. The idea is that the minimal latency cost is reached if and




Fig. 1. The network of an instance $(V_{\{a_i\}}, T_{\{a_i\}}, N_{\{a_i\}}, (d_{\{a_i\}}))$. The edges are labeled with the latency functions. Unlabeled edges have latency 0. Taxes can be imposed on the set of bold edges only.

only if the tax between v_i and v_{i+1} is 0 or a_i (inducing a latency cost of $A^2/2$ for this set of edges) and the sum of all taxes is exactly $\sum_i a_i/2$. \square

4 Parallel Links with Linear Latency Functions

We have seen that the latency cost is generally not monotone in the imposed tax even in case of linear latency functions and one taxable link. Further, such instances do not necessarily admit a unique optimal tax. These observations indicate that studying optimal taxes in parallel link networks might be intriguing. Our main goal in this section is to provide an algorithm for finding optimal taxes in single-commodity parallel link networks $(V, T, N, 1)$ in which every link $i \in [n]$ has a linear latency function $\ell_i(x) = a_i x + b_i$. This setting has been of special interest in the related problem of computing a Stackelberg leader strategy [17] described in the introduction. While this problem is already NP-hard in this setting, it may be surprising that we will be able to formulate a polynomial time algorithm for computing optimal taxes. Suppose the links are numbered by $N = \{1, \dots, k\}$ and $T = \{k + 1, \dots, n\}$, such that $b_1 \leq \dots \leq b_k$ and $b_{k+1} \leq \dots \leq b_n$. We use this labelling for convenience, but note that the ordering conditions apply only *within* N and T . We do not require $b_i \leq b_j$ for all $i \in N$ and $j \in T$ or any other restriction or relation between the links of N and T . W.l.o.g. we assume at most one constant latency link in $N \cup T$.

4.1 Candidate Supports Sets

A flow f is at Wardrop equilibrium if and only if there is a constant $L > 0$, s.t. all used links $i \in [n]$ have the same latency $L = \ell_i(f_i)$, whereas $L \leq \ell_{i'}(0) = b_{i'}$ for unused links $i' \in [n]$. Lemma 1 shows that a flow f is socially optimal if and only if there is a constant $C > 0$, s.t. $C = \ell_j^*(f_j) = 2a_j f_j + b_j$ for all used links $j \in [n]$, whereas $C \leq \ell_{j'}^*(0) = b_{j'}$ for unused links $j' \in [n]$.

Now consider an instance and increase the demand. The characterization yields that in equilibrium and in optimum links j will be filled with flow in order of increasing

b_j . Regarding cost the set of taxes will induce an equilibrium assigning flow to some link set $S \subset N \cup T$. All used non-taxable links have the same latency L . Since we allow for non-negative taxes only, the used taxable links will not have higher latency. This property allows us to parametrize the problem by the set of taxable and non-taxable links filled with flow. These sets turn out to be *candidate support sets* defined as follows.

Definition 3. *Every set of the form $S = \{1, \dots, l_1\} \cup \{k + 1, \dots, l_2\}$ with $1 \leq l_1 \leq k$ and $k + 1 \leq l_2 \leq n$ is called a candidate support set.*

Note that there are at most $n^2/4$ candidate support sets for any instance.

Lemma 4. *Let f denote a socially optimal flow for a parallel link network in which every edge is taxable. Then $\ell_1(f_1) \leq \ell_2(f_2) \leq \dots \leq \ell_n(f_n)$.*

Proof. The set of used links is of the form $\{1, \dots, l\}$ for some $l \leq n$. Since f is a minimal latency flow, all links $j \in \{1, \dots, l\}$ have equal marginal cost, and there is a constant $C > 0$ with $2a_j f_j + b_j = C$. Thus, $\ell_j(f_j) = a_j f_j + b_j = C/2 + b_j/2$. \square

Let us first argue that the consideration of candidate support sets is indeed sufficient to find optimal taxes. Imagine two separate commodities, routing demands d_N and $1 - d_N$ exclusively over N and T , resp. In such an instance, it would be optimal to set marginal cost taxes on T , and the set of used links form a candidate support set.

The difference to our setting is that demand can change between N and T , and thus we also need to ensure that latency and taxes create an equilibrium. If the optimal flow in T yields latencies smaller than L , then we can satisfy the latency constraint by setting appropriate non-negative taxes. Otherwise, the latency restriction reduces the flow on some used links. However, if the flow on a link is smaller than in the optimum due to the latency constraint, the marginal cost on this link is also smaller. Therefore, it is still optimal to fill the link with flow to the maximal possible extent (see Lemma 5). For all links not affected by the latency restriction, however, it is optimal to equalize the marginal costs, and the allocation of flow follows the ordering of offsets. In conclusion, the set of links allocated with flow remains a candidate support set.

4.2 Problem Parametrization

Fixing numbers n_S and t_S yields a candidate support set $S = N_S \cup T_S$ with $N_S = \{1, \dots, n_S\}$ and $T_S = \{k + 1, \dots, t_S\}$. For S denote by d_{N_S} and $1 - d_{N_S}$ the demand routed over N_S and T_S , respectively. $C_{N_S}(d_{N_S})$ is the latency cost for an equilibrium flow $(f_i)_{i \in N_S}$ of demand d_{N_S} . Denote by $C_{T_S}(1 - d_{N_S})$ the latency cost for an optimal flow $(f_j)_{j \in T_S}$ of demand $1 - d_{N_S}$ additionally fulfilling the latency restriction $\ell_j(f_j) \leq L(d_{N_S})$, where $L(d_{N_S})$ denotes the unique latency of all used links in N_S for a demand of d_{N_S} . Let $C(d_{N_S}) = C_{N_S}(d_{N_S}) + C_{T_S}(1 - d_{N_S})$ denote the latency cost of the flow.

The problem of finding a set of optimal taxes for a fixed set S can be formulated as follows: Minimize the cost function C , s.t. the flow for N is at equilibrium and the remaining flow on T is optimal subject to the additional constraint $\ell_j(f_j) \leq L(d_{N_S})$.

We will show that, if this minimization problem has a solution, the cost function $C(d_{N_S})$ is piecewise quadratic with at most n breakpoints and the optimal demand

distribution $(d_{N_S}^*, 1 - d_{N_S}^*)$ for N_S and T_S is efficiently computable. Iterating this for all possible sets S enables us to find optimal taxes.

We call a link $j \in T$ *full* w.r.t. some $L > 0$ if $f_j > 0$ and its latency equals the constraint value, i.e., if $\ell_j(f_j) = L$ or if $f_j = 0$ and $\ell_j(0) = b_j \geq L$. We call a link *relaxed* if $f_j > 0$ and $\ell_j(f_j) < L$. When shifting demand from N to T , the common latency L of used links in N decreases, while the demand on T increases. In the corresponding optimal flow on T respecting the constraint value, however, a full link never becomes relaxed. More formally, consider an instance (V, T, \emptyset, d) and let f denote the optimal flow respecting $\ell_i(f_i) \leq L$ for all i . With Lemma 4 we can assume the full links to form a set $\{p, \dots, n\}$ for some $p \geq 1$. Furthermore, assume there are $L' \leq L$ and $d' \geq d$ such that there is a flow of demand d' to T such that all used links have latency at most L' . For all non-constant links, we define $\ell_i^{-1}(L)$ to be the flow f_i such that $a_i f_i + b_i = L$ if $b_i \leq L$, and 0 otherwise.

Lemma 5. *The optimal flow f' respecting $\ell_i(f'_i) \leq L'$ for all i assigns $\ell_i^{-1}(L')$ flow to all non-constant links $i \in \{p_1, \dots, n\}$ for some uniquely defined $p_1 \leq p$.*

4.3 A Polynomial-Time Algorithm for Computing Optimal Taxes

Considering an optimal flow for an increasing demand, the links become used in order of their offsets. Lemmata 4 and 5 show that the links become full w.r.t. some bound in reverse order. Thus, we can determine the lower and the upper bound $d_{N_S}^{\min}$ and $d_{N_S}^{\max}$ for d_{N_S} such that the following holds: There is an equilibrium flow of demand d_{N_S} on N using exactly the links N_S and there is an optimal flow of demand $1 - d_{N_S}$ on T respecting the bound $L(d_{N_S})$ using exactly the links T_S .

Given a candidate support set S , we compute the optimal demand distribution $(d_{N_S}, 1 - d_{N_S})$. If such a distribution exists, we call S *feasible*. The corresponding demand interval $[d_{N_S}^{\min}, d_{N_S}^{\max}]$ can be computed in polynomial time by solving systems of linear equations.

Algorithm 1. OPTTAX $(V, T, N, 1)$

- 1: **for** every candidate support set S **do**
 - 2: **if** S feasible **then**
 - 3: compute the breakpoints $d_{N_S}^{\min} = d_{N_S k+1}, \dots, d_{N_S 1}, d_{N_S 0} = d_{N_S}^{\max}$
 - 4: $d_{N_S}^* \leftarrow \operatorname{argmin}_{0 \leq j \leq k} \min_{d_{N_S} \in [d_{N_S j}, d_{N_S j+1}]} C(d_{N_S})$
 - 5: **end if**
 - 6: **end for**
 - 7: $S^* \leftarrow \operatorname{argmin}_S C(d_{N_S}^*)$
 - 8: compute optimal flow on T_{S^*} respecting $L(d_{N_S^*}^*)$ with $\sum_{T_{S^*}} f_j^* = 1 - d_{N_S^*}^*$ and set $f_j^* := 0$ for $j \in T \setminus T_{S^*}$.
 - 9: set taxes $\tau_j \leftarrow L(d_{N_S^*}^*) - \ell_j(f_j^*)$ for $j \in T$
-

Lemma 6. *The cost function $C(d_{N_S})$ is piecewise quadratic for $d_{N_S} \in [d_{N_S}^{\min}, d_{N_S}^{\max}]$ with at most n breakpoints for every feasible candidate support set S . The breakpoints can be computed in polynomial time.*

Proof. We show that while C_{N_S} is a quadratic function, C_{T_S} and therefore C is piecewise quadratic with at most n breakpoints.

Suppose f is an equilibrium flow for N_S of demand d_{N_S} . There is some $L(d_{N_S}) > 0$ with $L(d_{N_S}) = a_i f_i + b_i$ for every $i \in N_S$. With $\sum_{N_S} f_i = d_{N_S}$, we infer that $L(d_{N_S})$ is linear and $C_{N_S}(d_{N_S}) = L(d_{N_S}) \cdot d_{N_S}$ is quadratic. Considering C_{T_S} , we need to respect the latency constraint for increasing $1 - d_{N_S}$. The cost function C_{T_S} turns out to be quadratic with at most n breakpoints. These breakpoints, i.e., the demand values for which the number of full links increases, can be calculated by solving systems of linear equations. \square

Given that restricting to candidate support sets is sufficient for finding optimal taxes, the following result holds.

Theorem 7. *Given an instance $(V, T, N, 1)$ with parallel links and linear latency functions, Algorithm OptTax(V,T,N,1) computes a set of optimal taxes $(\tau_j)_{j \in T}$ in polynomial time.*

References

- [1] Beckmann, M., McGuire, C.B., Winston, C.B.: Studies in the Economics of Transportation. Yale University Press (1956)
- [2] Braess, D.: Über ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung 12, 258–268 (1968)
- [3] Cocchi, R., Shenker, S., Estrin, D., Zhang, L.: Pricing in computer networks: motivation, formulation, and example. IEEE/ACM Transactions on Networking 1(6), 614–627 (1993)
- [4] Cole, R., Dodis, Y., Roughgarden, T.: How much can taxes help selfish routing? In: Proc. 4th EC, pp. 98–107 (2003)
- [5] Cole, R., Dodis, Y., Roughgarden, T.: Pricing edges for heterogeneous selfish users. In: Proc. 35th STOC, pp. 521–530 (2003)
- [6] Dafermos, S., Sparrow, F.T.: The Traffic Assignment Problem for a General Network. Journal of Research of the National Bureau of Standards, Series B 73(2), 91–118 (1969)
- [7] Fleischer, L.: Linear tolls suffice: New bounds and algorithms for tolls in single source networks. Theoretical Computer Science 348, 217–225 (2005)
- [8] Fleischer, L., Jain, K., Mahdian, M.: Taxes for heterogeneous selfish users in a multicommodity network. In: Proc. 45th FOCS, pp. 277–285 (2004)
- [9] Kaporis, A.C., Spirakis, P.G.: The price of optimum in stackelberg games on arbitrary single commodity networks and latency function. In: Proc. 18th SPAA (2006)
- [10] Karakostas, G., Kolliopoulos, S.G.: Edge pricing of multicommodity networks for heterogeneous selfish users. In: Proc. 45th FOCS, pp. 268–276 (2004)
- [11] Karakostas, G., Kolliopoulos, S.G.: Stackelberg strategies for selfish routing in general multicommodity networks. Algorithmica (to appear, 2008)
- [12] Korilis, Y.A., Lazar, A.A., Orda, A.: Achieving network optima using Stackelberg routing strategies. IEEE/ACM Transactions on Networking 5(1), 161–173 (1997)
- [13] Korilis, Y.A., Lazar, A.A., Orda, A.: Capacity allocation under noncooperative routing. IEEE Transactions on Automatic Control 42, 309–325 (1997)
- [14] Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Proc. 16th STACS, pp. 404–413 (1999)
- [15] Kumar, V.S.A., Marathe, A.: Improved results for Stackelberg scheduling strategies. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) ICALP 2002. LNCS, vol. 2380. Springer, Heidelberg (2002)

- [16] Pigou, A.C.: *The Economics of Welfare*. Macmillan, Basingstoke (1920)
- [17] Roughgarden, T.: Stackelberg scheduling strategies. In: Proc. 33rd STOC, pp. 104–113 (2001)
- [18] Roughgarden, T.: The Price of Anarchy is Independent of the Network Topology. In: Proc. 34th STOC, pp. 428–437 (2002)
- [19] Roughgarden, T., Tardos, É.: How Bad is Selfish Routing? *Journal of the ACM* 49(2), 236–259 (2002)
- [20] Roughgarden, T.: On the Severity of Braess’s Paradox: Designing Networks for Selfish Users is Hard. *Journal of Computer and System Sciences* 72(5), 922–953 (2004)
- [21] Roughgarden, T.: *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge (2005)
- [22] Sharma, Y., Williamson, D.P.: Stackelberg thresholds in network routing games or the value of altruism. In: Proc. 8th EC, pp. 93–102 (2007)
- [23] Swamy, C.: The effectiveness of Stackelberg strategies and tolls for network congestion games. In: Proc. 18th SODA, pp. 1133–1142 (2007)
- [24] Wardrop, J.G.: Some Theoretical Aspects of Road Traffic Research. In: Proc. of the Institute of Civil Engineers, Pt. II, pp. 325–378 (1952)

Anonymity-Proof Voting Rules

Vincent Conitzer

Departments of Computer Science and Economics
Duke University
Durham, NC, USA
conitzer@cs.duke.edu

Abstract. A (*randomized, anonymous*) voting rule maps any multiset of total orders (aka. *votes*) over a fixed set of alternatives to a probability distribution over these alternatives. A voting rule f is *false-name-proof* if no voter ever benefits from casting more than one vote. It is *anonymity-proof* if it satisfies voluntary participation and it is false-name-proof. We show that the class of anonymity-proof neutral voting rules consists exactly of the rules of the following form. With some probability $k_f \in [0, 1]$, the rule chooses an alternative uniformly at random. With probability $1 - k_f$, the rule first draws a pair of alternatives uniformly at random. If every vote prefers the same alternative between the two (and there is at least one vote), then the rule chooses that alternative. Otherwise, the rule flips a fair coin to decide between the two alternatives. We also show how the characterization changes if group strategy-proofness is added as a requirement.

1 Introduction

In many settings, a decision must be made on the basis of the preferences of multiple agents. Common examples include auctions and exchanges (where we must decide on an allocation of resources, as well as payments to be made or received by the agents) and elections (where we must decide on, say, one or more political representatives), but there are many other applications. A (*direct-revelation*) *mechanism* takes each agent's reported preferences as input, and produces a decision as output. An important issue is that self-interested agents will lie about their preferences if they perceive it to be to their advantage to do so. *Mechanism design* studies how to design mechanisms that produce good outcomes in spite of this. A key concept in mechanism design is that of *strategy-proofness*: a mechanism is strategy-proof if no agent can ever benefit from lying about her preferences. Strategy-proofness is roughly synonymous with *truthfulness* and *incentive compatibility*.¹ In mechanism design, attention is usually restricted to incentive compatible direct-revelation mechanisms. This is justified by a result known as the *revelation principle* [Gibbard, 1973; Green and Laffont, 1977; Myerson, 1979, 1981], which states (roughly) that, given that agents will misreport their preferences if

¹ To be more precise, strategy-proofness as the term is used here corresponds to *dominant-strategies* incentive compatibility. There are weaker notions of incentive compatibility, such as *Bayes-Nash* incentive compatibility, where *in expectation* over the other agents' preferences an agent is best off reporting her true preferences (assuming the others do so as well).

they perceive this to be to their benefit, anything that can be achieved by some mechanism can also be achieved by an incentive compatible direct-revelation mechanism²

In mechanism design, the spaces of possible outcomes and preferences often display a great deal of structure, which facilitates the designer's job. For example, in auctions and exchanges, it is often assumed that agents can make and receive payments, that their utility is linear in this payment, and that the effect of the payment on utility is independent of the rest of the outcome. This enables, for example, Vickrey-Clarke-Groves mechanisms [Vickrey, 1961; Clarke, 1971; Groves, 1973], which always choose the efficient allocation. However, such structure is not always available: for example, in an election, payments can typically not be made. If we do not assume any structure on the agents' preferences, then agents can rank the possible outcomes (aka. *alternatives*) in any possible way. These general settings, in which each agent ranks all the alternatives, and the mechanism chooses an alternative based on these rankings, are commonly referred to as *voting* settings. The rankings are the *votes*, and the mechanism is usually called a *voting rule*.

The revelation principle applies to voting settings just as it does to any other mechanism design setting, so we should ask which rules are strategy-proof. Gibbard [1977] provides a complete characterization of strategy-proof voting rules that are allowed to use randomization. (This characterization generalizes the better-known, earlier Gibbard-Satterthwaite theorem [Gibbard, 1973; Satterthwaite, 1975].) He shows that any strategy-proof rule is a randomization over *unilateral* rules, in which only one vote affects the outcome, and *duple* rules, in which only two alternatives have a chance of winning. (Because the overall rule is a randomization over such rules, it can still be the case that every voter affects the probability with which an alternative is chosen, and that every alternative has a positive probability of winning. Hence, Gibbard's characterization is not universally seen as a negative result [Barbera, 1979a].) He also provides some additional conditions on these rules to obtain an exact characterization of the strategy-proof voting rules.

However, strategy-proofness is often not sufficient. In open, anonymous environments such as the Internet, an agent can manipulate the mechanism in other ways. For one, if an agent does not participate in the mechanism, then the party running the mechanism (aka. the *center*) is not even aware of her existence. Perhaps more significantly, an agent can open multiple accounts and participate in the mechanism multiple times under different identifiers—and the center cannot know which identifiers correspond to the same agent. This led to the concept of *false-name-proofness* [Yokoo et al., 2004]. A mechanism is false-name-proof if an agent can never benefit from using multiple identifiers. Some positive and negative results on false-name-proofness have been obtained for combinatorial auctions and similar settings (e.g., [Yokoo et al., 2001; Yokoo, 2003; Yokoo et al., 2004, 2006]; [Rastegari et al., 2007]), but to our knowledge this concept has not yet been studied in voting settings.

In this paper, we define a (possibly randomized) voting rule to be *anonymity-proof* if it is false-name-proof, and it never hurts an agent to cast her (true) vote. Under the same

² To predict what will happen under a mechanism that is not incentive compatible, some solution concept from game theory must be used, and the version of incentive compatibility in the revelation principle depends on the choice of solution concept.

model as Gibbard [1977], we obtain a complete characterization of the anonymity-proof neutral voting rules. (A voting rule is *neutral* if it treats all alternatives symmetrically.) The proof is from first principles and (arguably) of reasonable length. The resulting class of voting rules is very limited (hence the result is mostly negative), but it does allow a modicum of responsiveness to the votes in cases where there is complete agreement among the voters on some pairs of alternatives. For example, in the special case where there are only two alternatives, the characterization tells us that if all votes prefer the same alternative, we can choose that alternative; but otherwise, we have to flip a fair coin to decide between them. This is in stark contrast to the case where we require only strategy-proofness, or even group strategy-proofness: for example, simply choosing the alternative that is preferred by more voters (the *majority* rule) is group strategy-proof.

1.1 Additional Motivation

Our primary reason for studying false-name-proofness in general social choice (voting) settings is that these settings lie at the heart of mechanism design, and hence provide the most natural starting point for a thorough study of the concept of false-name-proofness. Nevertheless, perhaps surprisingly, anonymous voting is in fact a very real and growing phenomenon on the Internet. It may seem that anonymous elections are unlikely to result in outcomes that reflect society's preferences well (and, in fact, this paper can be seen as a commentary on just how unlikely this is). However, it appears that in practice, often, the party organizing the election has more interest in publicity than in a properly chosen outcome; moreover, the convenience of anonymous Internet voting appeals to the voters as well.

A very recent example of this phenomenon is the "New 7 Wonders of the World" election, a global election that was organized by businessman Bernard Weber to elect contemporary alternatives to the ancient wonders. Anyone could vote, either by phone or over the Internet; for the latter, an e-mail address was required. One could also buy additional votes (of course, using another e-mail address was a much cheaper alternative). In spite of various irregularities (including unreasonably large numbers of votes in some cases) and UNESCO distancing itself from the election, the election seems to have attained some legitimacy in the public's mind.

For better or worse, mechanisms such as these are going to feature increasingly prominently in our economy and social infrastructure. Hence, the theory of mechanism design must be extended so that it can provide guiding principles to maximize the efficiency and trustworthiness of such mechanisms. The sooner this happens, the fewer bad mechanisms will take hold.

Our results also apply to Internet rating systems in which anonymous reviewers rate products, sellers, *etc.* Here, the set of alternatives is the set of possible (final, aggregate) ratings. It should be noted that in this context, it makes sense for agents' preferences to be restricted: for example, it makes little sense for an agent to prefer high > low > medium for a product's final rating. Specifically, *single-peaked* preferences [Black, 1948] are a natural restriction in this domain; we will discuss such preferences in the conclusion.

2 Definitions

Let X , $|X| = m$, be the set of *alternatives* over which the voters are voting. A voter's preferences are given by a total order \succ over the alternatives, together with a vector of utilities $\mathbf{u} = (u_1, \dots, u_m)$ where u_i is the voter's utility for the alternative that she ranks i th. (It is required that $u_i > u_{i+1}$ for all $1 \leq i \leq m$.) Each voter seeks to maximize her expected utility. As in Gibbard [1977], voters only report a total order (ranking) of the alternatives (not their utilities); a reported ranking is called a *vote*. Again as in Gibbard [1977], we do not allow for indifferences (real or reported) between alternatives. We will use the notation $v = a_1 \succ \dots \succ a_m$ for a vote. We will sometimes also use subsets in the order notation: for example, if $B = \{b_1, b_2, b_3\}$, then $a_1 \succ b_1 \succ b_2 \succ b_3 \succ a_2$ and $a_1 \succ b_3 \succ b_1 \succ b_2 \succ a_2$ are both of the form $a_1 \succ B \succ a_2$ (but, for instance, $a_1 \succ b_3 \succ b_2 \succ a_2 \succ b_1$ is not of this form). A *voting rule* f takes a multiset³ of votes V as input, and chooses the winning alternative based on these votes (possibly using randomization). Let $P_f(V, a)$ denote the probability with which f chooses a given votes V ; the function P_f defines the rule f . A voting rule is *neutral* if it treats all alternatives symmetrically—that is, if π is a permutation of the alternatives, then $P_f(\pi(V), \pi(a)) = P_f(V, a)$ (where $\pi(V)$ is the multiset that results from replacing each alternative a by $\pi(a)$ in each vote in V). In fact, the following weaker definition of neutrality will also suffice for our purposes: if a subset B of the alternatives is symmetric in V (that is, for any permutation π for which $\pi(a) = a$ for all $a \in X - B$, $\pi(V) = V$), then $P_f(V, b_1) = P_f(V, b_2)$ for all $b_1, b_2 \in B$. We are only interested in neutral voting rules⁴.

Definition 1. A voting rule f is *false-name-proof* if for any multiset of votes V , for any $v \in V, v = a_1 \succ \dots \succ a_m$, for any decreasing $\mathbf{u} = (u_1, \dots, u_m)$, and for any multiset of votes V' , we have $\sum_{j=1}^m P_f(V, a_j)u_j \geq \sum_{j=1}^m P_f(V \cup V', a_j)u_j$. That is, the voter corresponding to v cannot increase her expected utility by additionally casting votes V' .

It should be noted that under this definition, a voter who uses false names is assumed to cast at least one vote representing her true preferences. This only weakens the requirement. All of the rules in the characterization result of this paper are also false-name-proof in the stronger sense where none of the votes cast by the false-name voter are required to represent her true preferences. Hence, the characterization remains the same if this stronger requirement is used.

Definition 2. A voting rule f satisfies *participation* if for any multiset of votes V , for any $v \in V, v = a_1 \succ \dots \succ a_m$, for any decreasing $\mathbf{u} = (u_1, \dots, u_m)$, we have $\sum_{j=1}^m P_f(V, a_j)u_j \geq \sum_{j=1}^m P_f(V - \{v\}, a_j)u_j$. That is, the voter corresponding to v cannot increase her expected utility by not casting her vote.

³ This is implicitly assuming that every vote is treated equally; anything else would seem unreasonable in open, anonymous environments. Rules that treat every vote equally are commonly called *anonymous*; this is not to be confused with the definition of anonymity-proofness.

⁴ Sometimes rules that are not neutral are of interest, for example if one alternative is the incumbent and should be treated specially; but in most settings, only neutral rules are of interest.

Definition 3. A voting rule is anonymity-proof if it is false-name-proof and it satisfies participation.

Anonymity-proofness does not directly mention strategy-proofness. Thus, it may appear that even if a rule is anonymity-proof, it is possible that a voter can benefit from misreporting her preferences. However, all of the rules in the characterization result of this paper are also strategy-proof (this is implied by the fact that they satisfy the stronger version of false-name-proofness). Hence, the characterization remains the same if strategy-proofness is added as a requirement.

3 The Characterization of Anonymity-Proof Rules

In this section, we prove the main result. Showing that all the rules in the proposed class are anonymity-proof is not difficult; most of the proof consists of showing that all rules that are anonymity-proof are in the class. We prove the latter part using a sequence of six lemmas. Assuming the rule is anonymity-proof, these lemmas demonstrate how to transform any multiset of votes to a particular multiset of only two votes, without affecting one given alternative's probability of winning; and they demonstrate that this alternative's probability of winning in those two votes is as the theorem states.

The first lemma is a fundamental building block of the proof. It states that if we add a vote that agrees with an existing vote on the top k and bottom $l - k$ alternatives, then the probability of winning for each of those alternatives does not change.

Lemma 1. Consider a multiset of votes V , and suppose that for some $v \in V$, v is of the form $a_1 \succ \dots \succ a_k \succ B \succ a_{k+1} \succ \dots \succ a_l$. (Please note that l is equal to m only if B is empty.) Let v' (not necessarily in V) be another vote of the form $a_1 \succ \dots \succ a_k \succ B \succ a_{k+1} \succ \dots \succ a_l$ (that is, it is identical to v except for the internal ordering of B). Then, if f is anonymity-proof, for any $1 \leq i \leq l$, $P_f(V, a_i) = P_f(V \cup \{v'\}, a_i)$.

Proof. First, let us suppose that for some $1 \leq i \leq k$, $P_f(V, a_i) \neq P_f(V \cup \{v'\}, a_i)$. Without loss of generality, suppose that for any $1 \leq j < i$, $P_f(V, a_j) = P_f(V \cup \{v'\}, a_j)$. Consider the utility vector $\mathbf{u} = (1 - \epsilon, 1 - 2\epsilon, \dots, 1 - i\epsilon, (m - i)\epsilon, (m - i - 1)\epsilon, \dots, \epsilon)$. First, let us suppose that $P_f(V, a_i) < P_f(V \cup \{v'\}, a_i)$. Then, if the true preferences are given by V , the voter casting v has utility vector \mathbf{u} , and ϵ is sufficiently small, then the voter casting v has an incentive to cast v' as well. This is because (as $\epsilon \rightarrow 0$) she effectively seeks to maximize the probability of one of a_1, \dots, a_i winning, and casting v' as well does not affect the probabilities of a_1, \dots, a_{i-1} winning and increases that of a_i . On the other hand, suppose that $P_f(V, a_i) > P_f(V \cup \{v'\}, a_i)$. Then, if the true preferences are given by $V \cup \{v'\}$, the voter casting v' has utility vector \mathbf{u} , and ϵ is sufficiently small, then the voter casting v' has an incentive to not participate. This is because (as $\epsilon \rightarrow 0$) she effectively seeks to maximize the probability of one of a_1, \dots, a_i winning, and not participating does not affect the probabilities of a_1, \dots, a_{i-1} winning and increases that of a_i . Hence, for any $1 \leq i \leq k$, $P_f(V, a_i) = P_f(V \cup \{v'\}, a_i)$.

The case where $P_f(V, a_i) \neq P_f(V \cup \{v'\}, a_i)$ for some $k + 1 \leq i \leq l$ can be shown to contradict either false-name-proofness or participation by a symmetric argument (where, supposing without loss of generality that $P_f(V, a_j) = P_f(V \cup \{v'\}, a_j)$

for all $i < j \leq l$, the voter casting v or v' effectively tries to *minimize* the probability of one of the *last* $l - i + 1$ alternatives winning).

We obtain the following corollary, which states that it does not matter if the same vote is cast more than once. (This corollary is usually not powerful enough to use instead of the more general Lemma 1, but it provides some insight. In particular, for any fixed number of alternatives, this leaves only a finite number of multisets of votes to consider.)

Corollary 1. *For an anonymity-proof rule f , given that a vote is cast at least once, it does not matter how often it is cast.*

Proof. This follows from setting $B = \emptyset$ in Lemma 1.

The next few lemmas (2, 3, and 4) demonstrate how to transform any multiset of votes into a multiset of only two votes, without affecting one given alternative a 's probability of winning (assuming that the rule is anonymity-proof).

Lemma 1 allows us to prove the following lemma, which states that reordering the alternatives after a given alternative a in a vote, as well as reordering those before a , does not affect a 's probability of winning, unless we move alternatives past a .

Lemma 2. *Consider a multiset of votes V , and suppose that for some $v \in V$, v is of the form $B \succ a \succ C$. Let v' (not necessarily in V) be another vote of the form $B \succ a \succ C$ (that is, it is identical to v except for the internal ordering of B and C). Then, if f is anonymity-proof, $P_f(V, a) = P_f((V - \{v\}) \cup \{v'\}, a)$. That is, we can permute the alternatives on either side of a in a vote without affecting a 's probability of winning.*

Proof. Suppose first that we permute only C , that is, that B is ordered the same way in both v and v' . Then, we can apply Lemma 1 (letting a correspond to a_k in that lemma) to obtain $P_f(V, a) = P_f(V \cup \{v'\}, a)$, and similarly $P_f((V - \{v\}) \cup \{v'\}, a) = P_f(V \cup \{v'\}, a)$, hence $P_f(V, a) = P_f((V - \{v\}) \cup \{v'\}, a)$. The case where we permute only B can be proven symmetrically. But then, in the general case where both B and C are permuted, we can transform v into v' in two steps, as follows. Let v'' be the vote of the form $B \succ a \succ C$ that agrees with v on B but with v' on C . By the above, we have $P_f(V, a) = P_f((V - \{v\}) \cup \{v''\}, a) = P_f((V - \{v\}) \cup \{v'\}, a)$.

The next lemma shows that we *can* move an alternative b past a given alternative a in a vote, without affecting a 's probability of winning, *if* the other votes disagree on the relative ranking of a and b .

Lemma 3. *Consider a multiset of votes V , and suppose that for some $v \in V$, a is ranked before b . Additionally, suppose there is another vote $v' \in V$ that ranks a before b , and a third vote $v'' \in V$ that ranks b before a . Let v''' be a vote (not necessarily in V) that is obtained from v by improving b 's position, placing it somewhere ahead of a (while not changing the order in any other way). Then, if f is anonymity-proof, $P_f(V, a) = P_f((V - \{v\}) \cup \{v'''\}, a)$. That is, we can move b to the other side of a in a vote without affecting a 's probability of winning, if there are other votes that rank a before b and b before a .*

Proof. Let us first assume that a and b are adjacent in v and v''' . That is, a is ranked directly before b in v , and v''' is obtained from v simply by swapping a and b . By Lemma 1 (letting $\{a, b\}$ correspond to B in that lemma), for any alternative $c \notin \{a, b\}$, $P_f(V, c) = P_f(V \cup \{v'''\}, c)$, and also $P_f((V - \{v\}) \cup \{v'''\}, c) = P_f(V \cup \{v'''\}, c)$. Now, if we suppose that $P_f(V, a) < P_f(V \cup \{v'''\}, a)$, then, if the true preferences are given by V , the voter corresponding to v' would be better off casting v''' as well (since it will only affect the probabilities of a and b being elected, and v' prefers a). Conversely, if $P_f(V, a) > P_f(V \cup \{v'''\}, a)$, then the voter corresponding to v'' would be better off casting v''' as well. Hence, since f is false-name-proof, $P_f(V, a) = P_f(V \cup \{v'''\}, a)$. It similarly follows that $P_f((V - \{v\}) \cup \{v'''\}, a) = P_f(V \cup \{v'''\}, a)$ (since v' and v'' are still present in $(V - \{v\}) \cup \{v'''\}$). Hence, $P_f(V, a) = P_f((V - \{v\}) \cup \{v'''\}, a)$.

Now let us return to the general case where a and b are not necessarily adjacent in v and v''' . Let v'''' be the result of improving b 's position in v to just after a , and let v''''' be the result of swapping a and b in v'''' . Using Lemma 2, $P_f(V, a) = P_f((V - \{v\}) \cup \{v''''\}, a)$; using the above argument, $P_f((V - \{v\}) \cup \{v''''\}, a) = P_f((V - \{v\}) \cup \{v''''' \}, a)$; and using Lemma 2 again, $P_f((V - \{v\}) \cup \{v''''' \}, a) = P_f((V - \{v\}) \cup \{v'''\}, a)$.

In the next lemma, we use the previous lemmas to reduce a set of votes to a particular pair of votes, without affecting a 's probability of winning. (The proofs of the remaining lemmas and corollaries are omitted due to space constraint.)

Lemma 4. *Given a nonempty multiset of votes V and a distinguished alternative a , let B be the set of alternatives that are ranked before a by every vote in V , let C be the set of alternatives that are ranked before a by some votes in V and after a by others, and let D be the set of alternatives that are ranked after a by every vote in V . Let v (not necessarily in V) be a vote of the form $B \succ a \succ C \cup D$, and let v' (not necessarily in V) be a vote of the form $B \cup C \succ a \succ D$. Then, if f is anonymity-proof, $P_f(V, a) = P_f(\{v, v'\}, a)$.*

It should be noted that Lemma 4 does not cover the case where $V = \emptyset$; in this case, neutrality demands that an alternative be chosen uniformly at random. The next lemma characterizes the behavior of an anonymity-proof voting rule when only a single vote is cast.

Lemma 5. *Let $v = a_1 \succ \dots \succ a_m$. Let f be anonymity-proof and neutral, and let $p_f^i = P_f(\{v\}, a_i)$. Then, for some constant $0 \leq k_f \leq 1$, $p_f^i = k_f/m + (1 - k_f)(m - i) \cdot 2/(m(m - 1))$. That is, with probability k_f the rule chooses an alternative at random, and with probability $1 - k_f$ it draws a pair of alternatives at random and chooses the preferred one.*

The final lemma characterizes the probability of a winning in the special pair of votes from Lemma 4, using Lemma 5.

Lemma 6. *Let v be a vote of the form $B \succ a \succ C \cup D$, and let v' be a vote of the form $B \cup C \succ a \succ D$. Then, if f is anonymity-proof and neutral, $P_f(\{v, v'\}, a) = k_f/m + (1 - k_f)(2|D| + |C|)/(m(m - 1))$, where k_f is defined as in Lemma 5. That is, the probability that a wins is the same as under the following rule for selecting the*

winner: with probability k_f the rule chooses an alternative at random; with probability $1 - k_f$ it draws a pair of alternatives at random, and if every vote prefers the same alternative between the two, it chooses that alternative, otherwise it flips a fair coin to decide between the two alternatives.

Using the last three lemmas, the main result is now easy to prove. It states that any anonymity-proof neutral rule is either the rule that chooses an alternative at random, or the rule that draws two alternatives at random and runs the unanimity rule on these two alternatives, or a convex combination of these two rules.

Theorem 1. *The class of voting rules f that are anonymity-proof and neutral consists exactly of the following rules.*

- With some probability $k_f \in [0, 1]$, the rule chooses an alternative uniformly at random.
- With probability $1 - k_f$ it draws a pair of alternatives uniformly at random;
 - If every vote prefers the same alternative between the two (and there is at least one vote), then it chooses that alternative.
 - Otherwise, it flips a fair coin to decide between the two alternatives.

(All these rules are also false-name-proof in a stronger sense where the voter need not cast any vote with her true preferences, and this also implies that they are all strategy-proof.)

Proof. Let us first show that these rules indeed have the desired properties. They are clearly neutral. Conditional on a single random alternative being chosen, voters have no incentive to use false names or to not participate. Conditional on a random pair a, b of alternatives being drawn, there are four possibilities for a voter (who, without loss of generality, prefers a):

1. There are no other votes. In this case, the voter has a strict incentive to participate so that a is chosen, and no incentive to use false names.
2. All other votes prefer a . In this case, the voter has no incentive to use false names or not participate, since a will be chosen in any case.
3. All other votes prefer b . In this case, the voter has a strict incentive to participate so that at least a coin is flipped, and no incentive to use false names.
4. There are other votes that prefer a and other votes that prefer b . In this case, the voter has no incentive to use false names or not participate, since a coin will be flipped in any case.

We now show that there are no other rules with the desired properties. Let f be anonymity-proof and neutral. Lemma 5 defines k_f for this rule. Now, for an arbitrary multiset of votes V and an arbitrary alternative a , Lemma 4 shows how to convert V to a particular set of two votes $\{v, v'\}$, in a way that preserves a 's probability of winning, and also preserves a 's relationship to any other alternative b in the following sense:

- If all votes prefer a to b in V , the same is true in $\{v, v'\}$.
- If all votes prefer b to a in V , the same is true in $\{v, v'\}$.
- If some but not all votes prefer a to b in V , the same is true in $\{v, v'\}$.

Finally, Lemma 6 shows that for this set of two votes $\{v, v'\}$, alternative a 's probability of winning is as in the claim of this theorem. Because of the preservation properties of the conversion, this must also be true for the original set of votes V .

4 Discussion

In this section, we study some corollaries of the main result, and make some comparisons to rules that are only strategy-proof.

The characterization makes it clear that the optimal anonymity-proof rule (in any reasonable sense of the word “optimal”) is the one corresponding to $k_f = 0$, since this rule maximizes the probability that we can at least choose the better of two alternatives (if all votes agree). Even this rule is limited in the extent to which it can respond to the votes:

Corollary 2. *Under an anonymity-proof rule, the probability of any given alternative a winning is at most $2/m$ (for any multiset of votes). This probability is attained if and only if $k_f = 0$ and all votes rank a first.*

This is in sharp contrast to the class of strategy-proof rules. For example, it is strategy-proof to draw one of the votes at random and choose its most-preferred alternative (often referred to as the “random-dictator” rule). Under this rule, if an alternative ranks first in all votes, it will be chosen with probability 1. Also, within the class of strategy-proof rules, there is no rule that is clearly optimal. For example, it is also strategy-proof to draw a pair of alternatives at random, and choose the one that is preferred by more voters. Unlike the random-dictator rule, if there is an alternative that ranks first in all votes, this rule does not necessarily choose it; on the other hand, unlike the random-dictator rule, this rule does not run the risk of choosing an alternative that is ranked last by almost every vote (but first by a few).

Another sharp contrast between strategy-proof rules such as the above two and any anonymity-proof rule is the following. For the winning alternative not to be chosen uniformly at random, anonymity-proof rules require complete agreement on at least one pair of alternatives:

Corollary 3. *If V and a are such that for any $b \neq a$, there is a vote in V that prefers a to b , as well as one that prefers b to a , then for any anonymity-proof voting rule, $P_f(V, a) = 1/m$.*

5 Extension: Group Strategy-Proofness

A stronger notion than strategy-proofness is *group strategy-proofness*. A mechanism is group strategy-proof if there is never a coalition of agents that can jointly misreport their preferences so that they are all better off. An analogous result to Gibbard's characterization of strategy-proof voting rules has been given for group strategy-proofness [Barbera, 1979b].

Neither of group strategy-proofness and anonymity-proofness implies the other. For example, with two alternatives, the majority rule is group strategy-proof. On the other

hand, as it turns out, not all of the rules in Theorem 1 are group strategy-proof. The following theorem shows how the characterization in this paper changes if group strategy-proofness is added as a requirement.

Theorem 2. *The class of voting rules f that are anonymity-proof, group strategy-proof, and neutral consists exactly of the following rules.*

- For two alternatives, the rules that satisfy the conditions are the same as in Theorem 1.
- For three or more alternatives, only the rule that chooses an alternative uniformly at random satisfies the conditions.

Proof. For two alternatives, under any of the rules from Theorem 1 to increase the probability of one alternative winning, it is necessary to get some of the voters that prefer the other alternative to change their votes—but of course they have no incentive to do so. Hence, these rules are group strategy-proof.

For three or more alternatives, all we need to show is that if $k_f < 1$, then the rule is not group strategy-proof. (The $k_f = 1$ rule is group strategy-proof because it completely ignores the votes.) For three alternatives, consider the following profile of preferences: voter one prefers $a \succ b \succ c$, with utilities 3, 1, 0, respectively; voter two prefers $c \succ b \succ a$, also with utilities 3, 1, 0, respectively. If both voters vote truthfully, then there is no agreement on any pair of alternatives, so that the winner will be chosen uniformly at random, and each voter obtains an expected utility of $4/3$. However, if the voters cast the votes $a \succ c \succ b$ and $c \succ a \succ b$ instead, then the probability that b wins is $k_f/3$, whereas the probability for each of a and c is $k_f/3 + (1 - k_f)/2$. This results in an expected utility of $1(k_f/3) + 3(k_f/3 + (1 - k_f)/2) = 3/2 - k_f/6$ for each voter, which is strictly more than $4/3$ when $k_f < 1$. Hence the rule is not group strategy-proof. This example is easily extended to more than three alternatives (for example, by placing the additional alternatives at the bottom of each voter's preferences).

6 Future Research

Although Theorem 1 completely characterizes anonymity-proof neutral voting rules, much remains to be done in future research. The most natural next direction to take is to consider settings where the space of possible preferences is restricted. It is well-known that such restrictions can introduce very satisfactory *strategy-proof* rules. For example, in many settings there is a natural order on the alternatives (e.g., in political elections, we can order candidates by how far to the left of the political spectrum they are). In such a setting, a voter's preferences are said to be *single-peaked* if she always prefers alternatives that are closer to her most-preferred alternative to alternatives that are further away (when these alternatives are on the same side of the most-preferred alternative) [Black, 1948]. It is well-known that when preferences are single-peaked, choosing the most preferred alternative of the *median* voter (the voter that, if we sort the voters by their most preferred alternatives, ends up in the middle) is strategy-proof, and (if the number of voters is odd) this alternative will be preferred to any other alternative by more than half of the voters (i.e., it is the *Condorcet winner*). Single-peakedness

can only be of limited help for anonymity-proofness: for example, when there are only two alternatives, single-peakedness does not restrict preferences at all, so we cannot do anything more than in the general case. Specific application settings can also allow for more positive results, as has already been shown to be the case for combinatorial auctions. In a sense, such settings correspond to a very special way of restricting preferences. Other directions for future research include dropping the requirement of neutrality, and extending the result to allow voters to express indifferences.

Finally, if no good anonymity-proof mechanisms turn out to exist for a setting that we are interested in, then we need to consider other options. One natural solution is to verify agents' identities, that is, to check whether multiple preference reports came from the same agent. It is generally not necessary to verify the identities of all agents; rather, it suffices to verify those of a select few based on the submitted preference reports [Conitzer, 2007]. Another option is to suppose that each additional identifier used comes at a small cost to the manipulating agent. Much more positive results can be obtained in that setting [Wagman and Conitzer, 2008].⁵ In either case, the results in this paper provide a natural starting point for analysis. A final approach is to try to stop the problem at the source and make it impossible or impractical for an agent to sign up for more than one account. It seems difficult to do so without compromising the anonymity of the Internet, though it is not inconceivable: see Conitzer [2008] for one possible approach to achieving this using memory tests (which is, for now, far from practical).

Acknowledgements

I would like to thank Hervé Moulin for a helpful discussion; audiences at two Dagstuhl workshops, the INFORMS Annual Meeting, the World Congress of the Game Theory Society, Microsoft Research, Duke University, Hebrew University, Bar-Ilan University, and the California Institute of Technology for helpful feedback; and the Alfred P. Sloan Foundation and the National Science Foundation for support (through a Research Fellowship and award number IIS-0812113, respectively).

References

- Barbera, S.: Majority and positional voting in a probabilistic framework. *The Review of Economic Studies* 46(2), 379–389 (1979)
- Barbera, S.: A note on group strategy-proof decision schemes. *Econometrica* 47, 637–640 (1979)
- Black, D.: On the rationale of group decision-making. *Journal of Political Economy* 56(1), 23–34 (1948)
- Clarke, E.H.: Multipart pricing of public goods. *Public Choice* 11, 17–33 (1971)

⁵ Alternatively, it has been suggested to consider the setting where each voter can cast at most some constant k votes (at no cost). This setting seems much less problematic because we can simply *expect* every voter to cast k votes, so that all the false-name manipulations cancel out. (Admittedly, things are not quite this simple, because it may be optimal for a voter to make some of her k votes different from her other votes—but it is easy to see that there is no reason to do so in, say, a majority election between two alternatives).

- Conitzer, V.: Limited verification of identities to induce false-name-proofness. *TARK*, 102–111 (2007)
- Conitzer, V.: Using a memory test to limit a user to one account. *AMEC* (2008)
- Gibbard, A.: Manipulation of voting schemes: a general result. *Econometrica* 41, 587–602 (1973)
- Gibbard, A.: Manipulation of schemes that mix voting with chance. *Econometrica* 45, 665–681 (1977)
- Green, J., Laffont, J.-J.: Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* 45, 427–438 (1977)
- Groves, T.: Incentives in teams. *Econometrica* 41, 617–631 (1973)
- Myerson, R.: Incentive compatibility and the bargaining problem. *Econometrica* 41(1) (1979)
- Myerson, R.: Optimal auction design. *Mathematics of Operations Research* 6, 58–73 (1981)
- Rastegari, B., Condon, A., Leyton-Brown, K.: Revenue monotonicity in combinatorial auctions. *AAAI*, 122–127 (2007)
- Satterthwaite, M.: Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10, 187–217 (1975)
- Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16, 8–37 (1961)
- Wagman, L., Conitzer, V.: Optimal false-name-proof voting rules with costly voting. *AAAI*, 190–195 (2008)
- Yokoo, M., Sakurai, Y., Matsubara, S.: Robust combinatorial auction protocol against false-name bids. *Artificial Intelligence* 130(2), 167–181 (2001)
- Yokoo, M., Sakurai, Y., Matsubara, S.: The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. *Games and Economic Behavior* 46(1), 174–188 (2004)
- Yokoo, M., Matsutani, T., Iwasaki, A.: False-name-proof combinatorial auction protocol: Groves mechanism with submodular approximation. In: *AAMAS*, pp. 1135–1142 (2006)
- Yokoo, M.: The characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. *IJCAI*, 733–742 (2003)

Overlapping Coalition Formation

Georgios Chalkiadakis¹, Edith Elkind¹, Evangelos Markakis²,
and Nicholas R. Jennings¹

¹ School of Electronics and Computer Science, University of Southampton, United Kingdom
{gc2, ee, nrj}@ecs.soton.ac.uk

² Centre for Math and Computer Science (CWI) Amsterdam, The Netherlands
v.markakis@cwi.nl

Abstract. In multiagent domains, agents form coalitions to perform tasks. The usual models of cooperative game theory assume that the desired outcome is either the grand coalition or a coalition structure that consists of disjoint coalitions (i.e., a partition of the set of agents). However, in practice an agent may be involved in executing more than one task, and distributing his resources between several (not necessarily disjoint) coalitions. To tackle such scenarios, we introduce a model for cooperative games with overlapping coalitions. We then focus on concepts of stability in this setting. In particular, we define and study a notion of the core, which is a generalization of the corresponding notion in the traditional models of cooperative game theory. Under some quite general conditions, we characterize the elements of core. As a corollary, we also show that any element of the core maximizes the social welfare. We then introduce a concept of balancedness for overlapping coalitional games, and use it to characterize coalition structures that can be extended to elements of the core. Furthermore, we generalize the notion of convexity to our setting, and show that under some natural assumptions convex games have a non-empty core. To the best of our knowledge, this is the first paper to provide a generic model for overlapping coalition formation, along with a theoretical treatment of stability in this setting.

1 Introduction

In many settings, groups of agents have to form teams to perform certain tasks. In the game theory literature, this process is known as *coalition formation*. Traditionally, it is assumed that the desired outcome of this process is either the *grand coalition*, i.e., the set of all agents, or a *coalition structure* that consists of disjoint coalitions (i.e., a *partition* of the set of agents). Furthermore, most of this research focuses on *transferable utility* (TU games) in which there is no restriction on how agents can split the total payoff among themselves. In particular, agents from one coalition can make a payment to agents not in that coalition.

While the above assumptions are natural for some settings, there are many scenarios where they are not applicable. This is mainly for two reasons. First, in several scenarios it may only be possible to achieve the best outcome if agents can simultaneously belong to more than one coalition. In such circumstances, agents almost invariably need to distribute their resources between the coalitions in which they participate. Such “overlaps” are natural in a plethora of interesting settings: As a straightforward e-commerce example, consider online trading agents representing individuals or virtual enterprises,

and facing the challenge of forming coalitions and allocating their owners' capital to a variety of projects simultaneously. Such requirements are also common in systems requiring multirobot coordination, computational grid networks, and sensor networks (see, e.g., [22][15]). To date, however, there has been essentially no work on overlapping coalition formation, with just a few exceptions which we discuss in Section 3. In particular, we are not aware of any attempt to study coalitional stability in such settings. Second, it may not always be possible to split the value of a coalition with agents that do not even belong to that coalition, i.e., to allow cross-coalition transfers. Indeed, the inability of some of the agents to work together and share payoffs may be one of the primary reasons why the grand coalition does not form, and a particular coalition structure arises (for a detailed discussion, see the work of Aumann and Dreze [4]).

To address the above concerns, we introduce and study a model that explicitly takes *overlapping coalition formation (OCF)* into account. The model is applicable in situations where agents need to allocate different parts of their resources to simultaneously serve different tasks as members of different coalitions. Further, our work departs from the conventional transferable utility framework in several ways. First, there are no inherent superadditivity assumptions in our work, and hence the grand coalition does not always emerge. (Thus, our subsequent definition of the core incorporates coalition structures, unlike most traditional work in economics.) Second, we do not allow cross-coalitional transfers (this is realistic, since an agent not contributing to a coalition should not expect to receive payoff from it). Thus, though we do allow arbitrary transfers *within* coalitions, our games are not, technically speaking, games with fully transferable utility; rather, they can be considered as games with side-payments [4][19]. Finally, our model can take task (coalitional action) execution explicitly into account; this facilitates possible extensions to tackle coalition formation under uncertainty.¹

We then explore the stability concept of the *core* for such settings, and provide conditions for its existence. In particular, under some general assumptions, we first provide a characterization for outcomes, i.e., pairs of the form (overlapping coalition structure, imputation), to be in the core. Our proof is based on a graph-theoretic argument, which may be of independent interest. As a corollary of this theorem, we show that any outcome in the core maximizes the social welfare. Second, we characterize coalition structures that admit payoff allocations such that the resulting pair is in the core. This is done by generalizing the Bondareva-Shapley theorem to our setting (note that this theorem does not hold for arbitrary non-transferable utility games). Following that, we extend the notion of convexity in coalitional games to overlapping coalitions, and show that any convex OCF game has a non-empty core. Finally, we provide some natural extensions of our model, and suggest directions for future work.

2 Background

In this section, we provide a brief overview of the basic concepts in cooperative game theory regarding non-overlapping coalition structures. Let $N = \{1, \dots, n\}$ be a set of players (or “agents”). A subset $C \subseteq N$ is called a *coalition*. A *coalition structure (CS)* (in non-overlapping environments) is a partition of the set of agents.

¹ To ease notation, we only show how to incorporate coalitional actions in the model in Sec. 7.

Under the assumption of *transferable utility*, coalition formation can be abstracted into a fairly simple model. This assumption postulates the existence of a (divisible) commodity (e.g., “money”) that can be freely transferred among players. The role of the *characteristic function* of a *coalitional game with transferable utility (TU-game)* is to specify a single number denoting the worth of a coalition. Formally, a characteristic function $v : 2^N \mapsto \mathbb{R}$ defines the *value* $v(C)$ of each coalition C [32]. A transferable utility game is completely specified by the set of players N and the characteristic function v ; we write $G = (N, v)$.

While the characteristic function describes the payoffs available to coalitions, it does not prescribe a way of distributing these payoffs. This is captured by the notion of an *imputation*, defined as follows. We say that an *allocation* is a vector of payoffs $\mathbf{x} = (x_1, \dots, x_n)$ assigning some payoff to each $j \in N$. An allocation is *efficient* with respect to a coalition structure CS if $\sum_{j \in S} x_j = v(S)$ for all $S \in CS$. An allocation (x_1, \dots, x_n) is called an *imputation* if it is efficient, and satisfies *individual rationality*, i.e., $x_j \geq v(\{j\})$ for $j = 1, \dots, n$. $I(CS)$ denotes the set of all imputations for CS .

The non-overlapping core. When rational agents seek to maximize their individual payoffs, the *stability* of the underlying coalition structure becomes critical, as agents might be tempted to abandon agreements in pursuit of further gains for themselves. A structure is stable only if the outcomes attained by the coalitions and the payoff combinations agreed to by the agents satisfy both individual and group rationality. Given this, research in coalition formation has developed several notions of stability, among the strongest and the most well-studied ones being the *core* [18]. Taking coalition structures into account, the core of a TU game is a set of outcomes (CS, \mathbf{x}) , $\mathbf{x} \in I(CS)$, such that no subgroup of agents is motivated to depart from their coalitions in CS .

Definition 1. Let CS be a coalition structure, and let $\mathbf{x} \in \mathbb{R}^n$ be an allocation of payoffs to the agents. The core of a game (N, v) is the set of all pairs (CS, \mathbf{x}) such that $\mathbf{x} \in I(CS)$ and $\forall C \subseteq N$ it holds that $\sum_{j \in C} x_j \geq v(C)$.

Hence, no coalition would ever “block” the proposal for a core allocation. Unfortunately, the core is a strong notion, and there exist many games where it is empty.

The core definition above is essentially the definition provided in [24] (and is also very similar to the one coined in [17]). If we assume superadditivity of the characteristic function (i.e., $v(U \cup T) \geq v(U) + v(T)$ for any disjoint coalitions U and T) then in the definition above we may only consider outcomes where CS is simply the grand coalition and $\sum_{j \in N} x_j = v(N)$. The core definition then becomes the traditional definition that has been used in the vast majority of the economics literature (see [20]).

The environments of interest in our work however are primarily non-superadditive and we will not make any such assumption on the characteristic function. Indeed, there is a plethora of realistic application scenarios where the emergence of the grand coalition is either not guaranteed, might be perceivably harmful, or is plainly impossible (see, e.g., [24,23]). In addition to such motivations, Aumann and Dreze [4] also provide a thorough and insightful discussion on why coalition structures arise: they put forward a series of arguments on how this might happen, and explain that coalition structures may emerge naturally even in superadditive environments for a variety of reasons.

3 Related Work

As mentioned in the introduction, very little work exists on overlapping coalition formation settings. Here we discuss some notable exceptions, as well as some related work on the core in the context of non-overlapping coalition structures.

To begin, Shehory and Kraus present a setting for overlapping coalition formation in [28]. In their model, the agents have goals and capabilities, i.e., abilities to execute certain actions. To serve their goals, the agents have to participate in coalitions, to each of which they contribute some of their capabilities, which can thus be thought of as resources. The authors then propose heuristic algorithms that lead to the creation of overlapping coalition structures. However, the authors stop short of addressing the question of the stability of overlapping coalitions. Dang *et al.* [15] also examine heuristic algorithms for overlapping coalition formation to be used in surveillance multi-sensor networks. However, their work does not deal with payoff allocation issues, and does not attack the overlapping coalition formation problem from a game-theoretic perspective.

Conconi and Perroni [13] present a model of international multidimensional policy coordination in a *non-cooperative* setting: Agreement structures between countries can be overlapping, namely a country may participate in multiple agreements, by contributing any number of proposed “elementary strategies” (which can be regarded as being chosen from *discrete* sets of resources) to an agreement. They then introduce an equilibrium concept to describe stability in this setting. In contrast to our work, their setting is non-cooperative, and there is no attempt to globally characterize the set of stable agreements (as we do). While they introduce “negotiation tie-in” restrictions (i.e., requirements that the players must form agreements on multiple issues), they only use these as a tool to describe conditions for the stability of “joint global agreements” (where *all* players participate on agreements over *all* issues). Furthermore, their model can be seen as dealing with discrete rather than continuous resources (in comparison, we deal with continuous resources and our results also hold in the discrete case).

More recently, Albizuri *et al.* [1] presented an extension of Owen’s value [21] (which, in turn, can be thought of as a generalization of the *Shapley value* [25]) to an overlapping coalition formation setting. Specifically, they present an axiomatic characterization of their *configuration value*. Though they show through an example that the stability notion of the “large consistent set” [12] (which is a non-overlapping concept) can be applied to their configuration value, they do not further discuss other solution concepts. Moreover, in contrast to our approach, there is no restriction on the number of coalitions an agent might belong to. In particular, in the model of [1] there is no notion of resources that an agent needs to distribute across the coalitions he belongs to.

Regarding non-overlapping coalition structures as presented in Section 2, Sandholm and Lesser [24] examine the problem of allocating *computational resources* to coalitions. They do not restrict themselves to superadditive settings, but discuss the stability of coalition structures instead. In particular, they introduce a notion of bounded rational core that explicitly takes into account coalition structures. Apt *et al.* [23] also do not restrain themselves to problems where the outcome is the grand coalition only. Instead, they introduce various stability notions for abstract games and discuss simple transformations (e.g., split and merge rules) by which stable partitions of the set of players may emerge. However, they do not consider any extensions to overlapping coalitions.

4 Our Model

In this section we extend the traditional model of Section 2 to cooperative games with overlapping coalitions. In most scenarios of interest, even if overlapping coalitions are allowed, an agent would not be able to participate in all possible coalitions due to lack of time, cash flow, or energy. To model this, we assume that each agent possesses a certain amount of resources which he can distribute among the coalitions he joins. Without loss of generality, we can make a normalization and assume that each agent has one unit of resource: an agent's contribution to a coalition is thus given by the fraction of his resources that he allocates to it. We can also think of this as the agent's "participation level", or the fraction of time he devotes to a coalition. Of course, an agent may own several types of resources (e.g., time *and* money), and his contribution to a coalition would then be described by a vector rather than a scalar. Our model, and all of our results, extend to this more general setting in a straightforward manner. Nevertheless, for conciseness, we restrict our presentation to the single-resource setting.

As discussed above, in the non-overlapping model a coalition is a subset of agents, and a game is defined by its characteristic function $v : 2^N \mapsto \mathbb{R}$, representing the maximum total payoff that a coalition can get. In our setting, an *overlapping (or partial) coalition* is given by a vector $\mathbf{r} = (r_1, \dots, r_n)$, where r_j is the fraction of agent j 's resources contributed to this coalition ($r_j = 0$ means that j is not a member of the coalition). The *support* of a partial coalition \mathbf{r} is denoted by $\text{supp}(\mathbf{r})$ and is defined as $\text{supp}(\mathbf{r}) = \{j \in N \mid r_j \neq 0\}$. We can now define the games we will be considering in the rest of this work.

Definition 2. An OCF-game G with player set $N = \{1, \dots, n\}$ is given by a function $v : [0, 1]^n \rightarrow \mathbb{R}$, where $v(\mathbf{0}^n) = 0$.

Function v maps each partial coalition \mathbf{r} to the corresponding payoff. We denote this game by $G = (N, v)$, or, if N is clear from the context, simply by v . Clearly, a "classic" coalition S can now be represented as the vector e^S , where $(e^S)_j = 1$ for $j \in S$ and 0 otherwise. In the economics literature, these are sometimes called *crisp* coalitions, whereas coalitions of the form (r_1, \dots, r_n) with at least one r_j in $(0, 1)$ are referred to as *fuzzy* coalitions [7]. We will avoid this term in this work as in computer science the term "fuzzy" refers to other concepts. We instead refer to coalitions of this kind as *partial* coalitions, or simply coalitions.

In most scenarios of interest, v is *monotone*, i.e., satisfies $v(\mathbf{r}) \geq v(\mathbf{r}')$ for any \mathbf{r}, \mathbf{r}' such that $r_j \geq r'_j$ for all $j = 1, \dots, n$. Note that if v is monotone, we have $v(\mathbf{r}) \geq 0$ for any $\mathbf{r} \in [0, 1]^n$, since we set $v(\mathbf{0}, \dots, \mathbf{0}) = 0$. In what follows, we will explicitly indicate which of our results rely on the monotonicity of v .

We now need to specify what the outcomes of an OCF-game are. In the non-overlapping setting, an outcome is a pair (CS, \mathbf{x}) , where CS is a partition on N and \mathbf{x} is an imputation for CS . To extend this definition to our scenario, we start by introducing the notion of a coalition structure with overlapping coalitions:

Definition 3. For a set of agents $T \subseteq N$, a coalition structure on T is a finite list of vectors (*partial coalitions*) $CS_T = (\mathbf{r}^1, \dots, \mathbf{r}^k)$ that satisfies (i) $\mathbf{r}^i \in [0, 1]^n$; (ii)

² At the end of this section we discuss how our work differs from existing work on fuzzy games.

$\text{supp}(\mathbf{r}^i) \subseteq T$ for all $i = 1, \dots, k$; and (iii) $\sum_{i=1}^k r_j^i \leq 1$ for all $j \in T$. We will refer to k as the size of the coalition structure CS_T and write $|CS_T| = k$. Also, CS_T denotes the set of all coalition structures on T .

In the definition above, each $\mathbf{r}^i = (r_1^i, r_2^i, \dots, r_n^i)$ corresponds to some partial coalition (r_j^i being the fraction of the resources that agent j contributes to \mathbf{r}^i). The constraints state that every agent from T distributes at most one unit of his resources between the various coalitions he participates in (those may include the singleton coalition). This allows coalitions to be overlapping. Note that the coalition structure is a list rather than a set, i.e., it can contain two or more identical partial coalitions. Observe also that an agent is not required to allocate all of his resources, i.e., it can be the case that $\sum_{i=1, \dots, k} r_j^i < 1$. However, under monotonicity, we can assume that for each agent j we have $\sum_{i=1}^k r_j^i = 1$ (i.e., a coalition structure is a fractional partition of the agents).

We should note here that the introduction of overlapping coalition structures imposes some new technical challenges. While in the non-overlapping setting the number of different coalition structures is finite, in our setting there can be infinitely many different partial coalitions, and hence infinitely many coalition structures. In particular, this implies that it is impossible to find the social welfare-maximizing coalition structure by enumerating all candidate solutions (in fact the maximum may not even be attained). In contrast, in a non-OCF setting this approach is possible—though, in general, infeasible.

We now extend the definition of v to coalition structures by setting $v(CS) = \sum_{\mathbf{r} \in CS} v(\mathbf{r})$. Furthermore, for any $S \subseteq N$ we define $v^*(S) = \sup_{CS \in CS_S} v(CS)$. Intuitively, $v^*(S)$ is the least upper bound on the value that the members of S can achieve by forming a coalition structure. We say that v is *bounded* if $v^*(N) < \infty$; for most games of interest, v is likely to be bounded.

As in our setting the agents will not necessarily form the grand coalition, we will be interested in arguing about coalition structures from CS_N . The coalition structure will impose restrictions on admissible ways of distributing the gains: a payoff vector corresponds to an imputation *iff* it is obtained by distributing the value of each coalition:

Definition 4. Given a coalition structure $CS \in CS_N$, $|CS| = k$, an imputation for CS is a k -tuple $\mathbf{x} = (x^1, \dots, x^k)$, where $x^i \in \mathbb{R}^n$ for $i = 1, \dots, k$, such that

- (Payoff Distribution) for every partial coalition $\mathbf{r}^i \in CS$ we have $\sum_{j=1}^n x_j^i = v(\mathbf{r}^i)$ and $r_j^i = 0$ implies $x_j^i = 0$;
- (Individual Rationality) the total payoff of agent j is at least as big as what he can achieve on his own: $\sum_{i=1}^k x_j^i \geq v^*({j})$.

The set of all imputations for CS is denoted by $I(CS)$. Notice that in Definition 4 the profit from a task assigned to a partial coalition is only distributed among agents involved in executing it. Thus, no transfers of that payoff are allowed to outsiders. Now, the set of outcomes that is of interest to us is the set of *feasible agreements*:

Definition 5. A *feasible agreement* (or an outcome) for a set of agents $J \subseteq N$ is a tuple (CS, \mathbf{x}) where $CS \in CS_J$, $|CS| = k$ for some $k \in \mathbb{N}$, and $\mathbf{x} = (x^1, \dots, x^k) \in I(CS)$. We denote the set of all feasible agreements for J by $\mathcal{F}(J)$.

The payoff p_j of an agent j under a feasible agreement (CS, \mathbf{x}) is $p_j(CS, \mathbf{x}) = \sum_{i=1}^k x_j^i$. We write $\mathbf{p}(CS, \mathbf{x})$ to denote the vector $(p_1(CS, \mathbf{x}), \dots, p_n(CS, \mathbf{x}))$.

Definition of the core. As explained in Section 2, no group of agents should be able to profitably deviate from a configuration in the core. Hence, any definition of the core has to depend on the notion of permissible deviations used. In this section we consider a fairly straightforward such notion.

In the non-overlapping setting with coalition structures, a pair (coalition structure, imputation) is in the core if no set of agents can jointly profit by deviating and forming a coalition on their own (Def. 1). Similarly, in the overlapping coalitions scenario, a deviating group of agents S may also want to break its obligations to other agents, i.e., withdraw or reduce its contribution to coalitions that contain agents in $N \setminus S$. As the deviators do not take into account the interests of other agents, they cannot expect to receive payoffs from the coalitions that contain non-deviating agents and therefore might have been hurt by the deviation. We formalize this approach as follows.

Definition 6 (The overlapping core). A tuple (CS, \mathbf{x}) is in the core of an OCF-game $G = (N, v)$ (we write $(CS, \mathbf{x}) \in \text{core}(G)$), if for any set of agents $J \subseteq N$, any coalition structure CS_J on J , and any imputation $\mathbf{y} \in I(CS_J)$, we have $p_j(CS_J, \mathbf{y}) \leq p_j(CS, \mathbf{x})$ for some agent $j \in J$.

It is easy to see that when restricted to non-overlapping coalitions, this definition collapses to the traditional definition of the core.

Finally, we point out here the differences between our concept of the core and the Aubin core [7], which is the main solution concept in fuzzy games. An outcome in the core of an OCF game is not necessarily the grand coalition, but it can be an (overlapping) coalition structure. In contrast, in the Aubin core, the imputations and core elements are always the grand coalition along with a split of the payoff $v(N)$. Furthermore, in our definition a core outcome needs to be stable against any deviation of a set S to a (possibly overlapping) coalition structure. In the Aubin core, an outcome need only be stable against a deviation to a partial (“fuzzy”) coalition, but not necessarily against deviations to a coalition structure. In short, the formation of coalition structures (overlapping or not) is not addressed in the fuzzy games literature.

5 Core Characterization

In the previous section, we introduced a definition of the core for overlapping coalition formation games. We now proceed to provide a characterization of the set of outcomes in the core: essentially, an outcome is in the core if and only if under this outcome the total payments to each subset of agents match the maximum value that can be achieved by this subset. Our proof relies on several technical restrictions on the function v that defines the game. In particular, we require v to be continuous, monotone and bounded (observe that if a game is monotone and bounded, then $v^*(S) < \infty$ for any $S \subseteq N$), as well as to satisfy another natural restriction defined later. These assumptions allow us to avoid some pathological situations that may arise in our model at its generality, such as the supremum $v^*(N)$ being unachievable (e.g., if v is strictly concave in one of its arguments, it can be the case that no finite coalition structure can achieve $v^*(N)$).

Specifically, we say that a game (N, v) is *U-finite* if for any (CS, \mathbf{x}) such that $|CS| > U$ and $\mathbf{x} \in I(CS)$, there exists a (CS', \mathbf{y}) such that $|CS'| \leq U$, $\mathbf{y} \in I(CS')$,

and $p_j(CS, \mathbf{x}) \leq p_j(CS', \mathbf{y})$ for all $j = 1, \dots, n$ (i.e., for any outcome (CS, \mathbf{x}) with more than U coalitions there exists another outcome (CS', \mathbf{y}) with at most U coalitions that is weakly preferred to (CS, \mathbf{x}) by all agents). When this condition holds, we can assume that all coalition structures that arise in a game consist of at most U partial coalitions. This is a natural restriction in many practical scenarios, as it might be difficult for agents to maintain a very complicated collaboration pattern. It holds when, e.g., there is a bound on the number of partial coalitions each agent can be involved in. Another natural example is provided by a class of games where for any two partial coalitions \mathbf{r}, \mathbf{r}' such that $\text{supp}(\mathbf{r}) = \text{supp}(\mathbf{r}')$ and $r_j + r'_j \leq 1$ for any $j = 1, \dots, n$, we have $v(\mathbf{r} + \mathbf{r}') \geq v(\mathbf{r}) + v(\mathbf{r}')$. Note that in such games we can assume that no coalition structure contains two partial coalitions with the same support S , as it is at least as profitable for the players in S to merge these partial coalitions. (However, notice that this does not imply superadditivity, nor does it mean that the grand coalition necessarily emerges, as the criterion above refers only to coalitions with identical support.) Hence, any such game is 2^n -finite.

Remark 1. Note that in all of our results U can also be a function of n (as long as $U(n) < \infty$). Alternatively, instead of imposing the condition of U -finiteness on $v(\cdot)$, we could restrict the set of allowed outcomes (or potential deviations) to coalition structures with at most U partial coalitions. All of our results hold under this model as well.

We now state and prove the first of our main results.

Theorem 1. *Given a game (N, v) , where v is monotone, continuous, bounded, and U -finite for some $U \in \mathbb{N}$, an outcome (CS, \mathbf{x}) is in the core of (N, v) iff for all $S \subseteq N$*

$$\sum_{j \in S} p_j(CS, \mathbf{x}) \geq v^*(S). \tag{1}$$

Proof: For the “if” direction, suppose that (CS, \mathbf{x}) satisfies $\sum_{j \in S} p_j(CS, \mathbf{x}) \geq v^*(S)$ for all $S \subseteq N$. Assume for the sake of contradiction that (CS, \mathbf{x}) is not in the core, i.e., there exists a set S , a coalition structure $CS_S \in \mathcal{CS}_S$ and an imputation $\mathbf{y} \in I(CS_S)$ such that $p_j(CS_S, \mathbf{y}) > p_j(CS, \mathbf{x})$ for all $j \in S$. Then we have $v(CS_S) = \sum_{j \in S} p_j(CS_S, \mathbf{y}) > \sum_{j \in S} p_j(CS, \mathbf{x}) = v^*(S)$, a contradiction with the way $v^*(S)$ was defined.

For the “only if” direction, consider an outcome (CS, \mathbf{x}) that does not satisfy $\textcircled{1}$; we will show that (CS, \mathbf{x}) is not in the core. To begin, set $\mathbf{p} = \mathbf{p}(CS, \mathbf{x})$, and assume $\sum_{j \in S} p_j < v^*(S)$ for some $S \subseteq N$. To show that (CS, \mathbf{x}) is not in the core, we will construct a set S' , a coalition structure $CS_{S'} \in \mathcal{CS}_{S'}$ and an imputation $\mathbf{y} \in I(CS_{S'})$ such that $p_j(CS_{S'}, \mathbf{y}) > p_j$ for all $j \in S'$. Fix a set S that satisfies $\sum_{j \in S} p_j < v^*(S)$. Choose ε small enough so that $\sum_{j \in S} p_j < v^*(S) - \varepsilon$, and let $\mathcal{CS}_S^\varepsilon = \{CS_S \in \mathcal{CS}_S \mid v(CS_S) \geq v^*(S) - \varepsilon\}$. By definition of $v^*(S)$, there is an infinite sequence of coalition structures $CS^{(t)}$ that satisfies $\lim_{t \rightarrow \infty} v(CS^{(t)}) = v^*(S)$, so the set $\mathcal{CS}_S^\varepsilon$ is non-empty. Given a coalition structure $CS_S \in \mathcal{CS}_S$, an imputation $\mathbf{y} \in I(CS_S)$ and a respective payoff vector $\mathbf{q} = \mathbf{p}(CS_S, \mathbf{y})$, define the *total loss* $TL(CS_S, \mathbf{q})$ of (CS_S, \mathbf{q}) as $\sum_{j: p_j > q_j} (p_j - q_j)$. Set $TL_{\min} = \inf\{TL(CS_S, \mathbf{q}) \mid CS_S \in \mathcal{CS}_S^\varepsilon, \mathbf{y} \in I(CS_S), \mathbf{q} = \mathbf{p}(CS_S, \mathbf{y})\}$. First, we prove that there exists a coalition structure $CS \in \mathcal{CS}_S^\varepsilon$ and an imputation $\mathbf{y} \in I(CS_S)$ that achieve the total loss of TL_{\min} (the proof is omitted).

Lemma 1. *Under the theorem’s conditions, there exists a $CS_S \in CS_S^\varepsilon$, an imputation $\mathbf{y} \in I(CS_S)$ and a payoff vector $\mathbf{q} = \mathbf{p}(CS_S, \mathbf{y})$ s.t. $TL(CS_S, \mathbf{q}) = TL_{\min}$.*

Let (CS_S, \mathbf{y}) be an outcome that satisfies $v(CS_S) \geq v^*(S) - \varepsilon$, $TL(CS_S, \mathbf{p}(CS_S, \mathbf{y})) = TL_{\min}$, whose existence is guaranteed by Lemma 1. Set $\mathbf{q} = \mathbf{p}(CS_S, \mathbf{y})$. Let us now construct a directed graph Γ whose vertices are the agents and there is an edge from j to i if there exists a coalition in CS_S containing both j and i such that under \mathbf{y} , agent j gets a non-zero payoff from that coalition, i.e., for some $\mathbf{r}^k \in CS_S$ we have $r_j^k, r_i^k > 0$ and $y_j^k > 0$. Observe that if there is an edge (j, i) in Γ , we can change \mathbf{y}^k by increasing the payoff to i by a small enough δ and decreasing the payoff to j by the same value of δ without violating the constraints, i.e., we have $\mathbf{z} = (\mathbf{z}^1, \dots, \mathbf{z}^\ell) \in I(CS_S)$, where $\mathbf{z}^l = \mathbf{y}^l$ for $l \neq k$ and $\mathbf{z}^k = (y_1^k, \dots, y_j^k - \delta, \dots, y_i^k + \delta, \dots, y_n^k)$. Now, color all vertices of Γ as follows: a vertex j is red if the agent j is underpaid under \mathbf{y} , i.e., $q_j < p_j$, white if j is indifferent, i.e., $q_j = p_j$, and green if he is overpaid, i.e., $q_j > p_j$. As $\sum_{j \in S} p_j < v^*(S) - \varepsilon$ and $\sum_{j \in S} q_j = v(CS_S) \geq v^*(S) - \varepsilon$, the graph contains at least one green vertex. As argued above, if there is a path from a green vertex j to a red vertex i , we can transfer a small amount of payoff from j to i and hence decrease the total loss, which is a contradiction with our choice of (CS_S, \mathbf{y}) . Hence, given an arbitrary green vertex j , the set of all vertices reachable from j in the graph, which we denote by $R(j)$, can only contain green or white vertices.

We would now like to argue that the agents in $R(j)$ can successfully deviate from (CS, \mathbf{x}) . Indeed, let CS' be the coalition structure that consists of the coalitions that the agents in $R(j)$ form among themselves in CS_S . Clearly, the value of CS' is equal to the total value of the coalitions formed by these agents in CS_S . Note also that under (CS_S, \mathbf{y}) , the agents in $R(j)$ do not get any payoffs from coalitions that involve agents not in $R(j)$. Indeed, suppose that an $i \in R(j)$ gets a non-zero payoff from a coalition that involves an agent $k \notin R(j)$. Then in Γ there is an edge from i to k , a contradiction with how $R(j)$ was constructed. In other words, in CS_S , the payoffs that the agents in $R(j)$ get come only from the coalitions that they form among themselves, and yet these agents are all green or white, i.e., each of them is doing no worse than what he was doing under CS , and some of them (in particular, agent j) are doing strictly better. To finish the proof, let the agents in $R(j)$ distribute the payoffs in the same way as in (CS_S, \mathbf{y}) , except that player j transfers a small fraction of his payoffs to each of the white players in $R(j)$ (this is possible by construction). The last step ensures that each agent in $R(j)$ is strictly better off than in (CS, \mathbf{x}) . This demonstrates that (CS, \mathbf{x}) is not in the core, as required. \square

Remark 2. Note that we did not have to make use of the additional restrictions we imposed on v to prove the “if” direction of the theorem (these are used in the proof of Lemma 1). Hence, this implication holds for an arbitrary G .

It is easily verifiable that Theorem 1 holds in the non-overlapping case with coalition structures as well. The result is trivial to prove in that setting, as each agent’s payoffs come from just one coalition; in contrast, we had to use more involved combinatorial arguments for transferring payoffs among agents. We also get the following interesting result as a corollary:

Corollary 1. *By setting $S = N$ in the statement of Theorem 7 we conclude that any outcome in the core maximizes the social welfare.*

Characterizing the Core Coalition Structures. In Theorem 6 we saw a necessary and sufficient condition for a tuple (CS, \mathbf{x}) to belong to the core. Now, suppose that we are only given a structure $CS = (\mathbf{r}^1, \dots, \mathbf{r}^k)$ and we want to check whether there exists *some* payoff allocation \mathbf{x} such that (CS, \mathbf{x}) belongs to the core. Below we characterize the set of coalition structures CS that admit payoff allocations \mathbf{x} such that the corresponding tuple (CS, \mathbf{x}) belongs to the core. Our characterization can be seen as a generalization of the notion of *balancedness* in the context of overlapping coalition formation. In the classic setting, the analogous question is “when does the grand coalition admit a payoff allocation in the core”, answered by Bondareva and Shapley [626]. Before we proceed to our result, we define balancedness w.r.t. to a coalition structure:

Definition 7. *Fix a coalition structure $CS = (\mathbf{r}^1, \dots, \mathbf{r}^k)$, $k \in \mathbb{N}$, and let $K = \{1, \dots, k\}$. A collection of numbers $\{\lambda_S\}_{S \subseteq N}$, $\{\mu_i\}_{i \in K}$ is called balanced w.r.t. the given coalition structure CS , if and only if $\lambda_S \geq 0$ for all S , and $\sum_{S: j \in S} \lambda_S + \mu_i = 1$ for all $i \in K, j \in \text{supp}(\mathbf{r}^i)$.*

Definition 8. *A game is called balanced w.r.t. a coalition structure $CS = (\mathbf{r}^1, \dots, \mathbf{r}^k)$ if and only if for every collection $\{\lambda_S\}_{S \subseteq N}$, $\{\mu_i\}_{i \in K}$, which is balanced w.r.t. CS , it holds $\sum_S \lambda_S v^*(S) + \sum_{i=1}^k \mu_i v(\mathbf{r}^i) \leq v^*(N)$.*

The proof of the following theorem (omitted due to space constraints) is based on LP-Duality, and relies on the characterization result of Theorem 6 furthermore, the proof illustrates that the condition of balancedness introduced above arises rather naturally.

Theorem 2. *Let (N, v) be an OCF-game, where v is monotone, continuous, bounded, and U -finite for some $U \in \mathbb{N}$ and consider a coalition structure $CS = (\mathbf{r}^1, \dots, \mathbf{r}^k)$, for some $k \in \mathbb{N}$. There exists an imputation \mathbf{x} such that (CS, \mathbf{x}) belongs to the core if and only if the game is balanced w.r.t. CS .*

Remark 3. In the traditional superadditive setting, the condition of balancedness is a bit simpler and more intuitive. In our setting, the characterization leads to a slightly more complicated expression, essentially due to the fact that the linear program that describes core allocations for each coalition structure requires a larger set of constraints.

6 Convex OCF Games Have a Non-empty Core

In this section, we first generalize the notion of convexity to OCF-games and then proceed to show that it provides a sufficient condition for the non-emptiness of the core.

Recall that for classical TU-games convexity means that for $R \subseteq N$ and $S \subset T \subseteq N \setminus R$ it holds that $v(S \cup R) - v(S) \leq v(T \cup R) - v(T)$. Thus, convexity in the classic TU-games setting means that it is more useful for a coalition R to join a larger coalition than a smaller one. We now apply this intuition to our setting:

Definition 9. An OCF-game $G = (N, v)$ is convex if for each $R \subseteq N$ and $S \subset T \subseteq N \setminus R$ the following condition holds: for any $(CS^S, \mathbf{x}^S) \in \mathcal{F}(S)$, any $(CS^T, \mathbf{x}^T) \in \mathcal{F}(T)$, and any $(CS^{S \cup R}, \mathbf{x}^{S \cup R}) \in \mathcal{F}(S \cup R)$ that satisfies $p_j(CS^{S \cup R}, \mathbf{x}^{S \cup R}) \geq p_j(CS^S, \mathbf{x}^S) \forall j \in S$, there exists an outcome $(CS^{T \cup R}, \mathbf{x}^{T \cup R}) \in \mathcal{F}(T \cup R)$ s.t.

$$\begin{aligned} p_j(CS^{T \cup R}, \mathbf{x}^{T \cup R}) &\geq p_j(CS^T, \mathbf{x}^T) \quad \forall j \in T, \text{ and} \\ p_j(CS^{T \cup R}, \mathbf{x}^{T \cup R}) &\geq p_j(CS^{S \cup R}, \mathbf{x}^{S \cup R}) \quad \forall j \in R. \end{aligned}$$

This definition is similar in flavour to that in [29], where a generalization of convexity is defined in the context of stochastic cooperative games. The intuition behind this definition is as follows: Consider two fixed agreements, one on S and one on T respectively. Any time that there is a feasible agreement on $S \cup R$ that the members of S do not object to compared to their own agreement (i.e., all members of S are weakly better off than in their previous agreement), then there is a feasible agreement on $T \cup R$ such that (i) the members of T do not object to this agreement, compared to the previous agreement on T and (ii) the members of R weakly prefer this agreement to the agreement on $S \cup R$.

We now show that convexity is a sufficient condition for the non-emptiness of the core in analogy to the classic result on convex TU-games [27]. Here we only give an outline of the proof.

Theorem 3. *If an OCF-game $G = (N, v)$ is convex and v is continuous, bounded, monotone and U -finite for some $U \in \mathbb{N}$, then the core of this game is not empty.*

Proof Sketch: To prove the theorem, we explicitly construct an outcome (CS, \mathbf{x}) , and show that it belongs to the core. The construction proceeds in rounds. First let p_1 be the maximum payoff that agent 1 can achieve on his own. In round 2, we choose an agreement (CS_2, x_2) on $\{1, 2\}$ that maximizes agent 2's payoff subject to the constraint that agent 1 receives at least p_1 . Then in round 3 we pick an agreement (CS_3, x_3) that maximizes the payoff of agent 3 subject to the constraint that the other 2 agents are not worse off compared to the previous round. We continue in this manner till we reach an agreement (CS_n, x_n) in round n . We then prove that (CS_n, x_n) is in the core by using induction and exploiting the game's convexity property. \square

In the traditional setting, if a game is represented using oracle access for $v(S)$, there is a trivial algorithm for computing an element of the core in convex games. Indeed, one can set the payoff vector to be the vector of the marginal contributions of the agents for an arbitrary permutation of the set of agents. In our setting, our proof does yield a procedure for constructing an element of the core, however not a polynomial time one. Our procedure requires solving a series of optimization questions, which for arbitrary convex games are NP-hard. In the future, we would like to find classes of convex games where our proof yields a polynomial time algorithm. In particular, looking at our proof, this would be true for games in which we can solve in polynomial time the following problem: Given a set of agents $S \subseteq N$, a feasible agreement on S , (CS, \mathbf{x}) , and an agent $k \notin S$, find a feasible agreement (CS', \mathbf{y}) on $S \cup \{k\}$ that maximizes $p_k(CS, \mathbf{y})$ subject to the constraints: $p_j(CS', \mathbf{y}) \geq p_j(CS, \mathbf{x})$.

7 Conclusions, Extensions, and Future Work

In this paper we introduced a model of cooperative games that allows for overlapping coalitions and takes into account the need for resource allocation. In doing so, we generalize the usual models where either the grand coalition is the only desirable outcome or the outcomes are required to be partitions of the set of agents. Given our model, we defined and studied a notion of the core which is a generalization of the core in the traditional models of cooperative game theory. Under some quite general conditions, we provided a characterization for an outcome—that is, a (coalition structure, imputation) pair—to belong to the core. We also showed that any outcome in the core maximizes the social welfare. Further, we introduced balancedness for OCF games, defined balanced OCF games, and showed that a coalition structure CS admits an imputation x so that (CS, x) is in the core if and only if the game is balanced. Finally, we extended the notion of convexity to our setting and showed that convex games have a non-empty core. This is one of the very first attempts to provide a theoretical treatment of overlapping coalition formation, and, to the best of our knowledge, the first to present a generic model for overlapping coalition formation and study stability in a thorough manner.

Extensions. In order to not overload notation, we avoided modeling *coalitional actions* in our presentation so far. However, in realistic environments coalitions are formed to execute tasks, which can be represented as coalitional actions. This is easily incorporated in our model, as follows: A coalition is allowed to select an action from a (usually finite) action space \mathcal{A} . Without loss of generality, we assume that each coalition can undertake any action in \mathcal{A} .³ The value of a coalition is then determined by the resource contribution levels of its members *and* the action selected. Therefore, the characteristic function in our setting is then defined on (r, a) pairs, where $r = (r_1, \dots, r_n)$ is a vector of resources, and $a \in \mathcal{A}$ is an action. All of our definitions and results generalize readily to the situation where each coalition has a choice of actions (simply put, our presentation so far corresponds to a situation where each coalition had exactly one action available to it).

Another extension we have examined has to do with modeling the available resources. For ease of presentation it was assumed throughout the paper that there exists only one type of (continuous) resource. Nevertheless, all of our results still hold if we assume multiple types of resources. Moreover, we have also studied a “discrete” OCF setting, with agent resources taking values in a finite set (i.e., an agent cannot contribute an arbitrary percentage of his resources to a coalition). With discrete resources, the number of possible coalition structures is now finite (as a coalition in our setting is a collection of resources—see Section 4). All of our definitions and theorems carry through in this setting with minor differences in the arguments used in the proofs.

Finally, we have also investigated alternative notions of deviations and concepts of the core. Specifically, we have so far assumed that deviators cannot rely on getting any payoffs from their coalitions with non-deviators: in this sense, the deviators can be seen as “self-reliant”, and their form of deviation as “malicious” (since when deviating they pay no consideration to the harm inflicted on their partners). However, one can

³ The situation where this is not the case can be modeled by setting the value of the respective (coalition, action) pair to 0.

consider a more “benign” form of deviation, where deviating teams do not break their obligations to non-deviators (possibly due to restrictions imposed by the multiagent system architecture), but instead have to rearrange resources already used in coalitions with other deviating agents. In this case, it is natural to assume that the deviating agents get to keep their share of the profit from the coalitions with non-deviators. For lack of space, we refer further discussion of these issues to an extended version of this paper.

Future work. There exist many exciting open questions for future work. In particular, it would be interesting to investigate the alternative notion of stability proposed above. We also plan to study the computational complexity of the core. Even in superadditive settings (where the coalition structure is simply the grand coalition N), computing an allocation in the core or checking if the core is non-empty are NP-hard problems [11,31,16,14]. In the absence of superadditivity, there are even stronger lower bounds on the complexity of the problem [23]. Hence we can only hope to identify special classes of games where we can have efficient algorithms for computing core allocations. As noted earlier, an element of the core in convex games can be computed in the traditional setting simply by taking the vector of the marginal contributions of the agents for an arbitrary permutation of the set of agents. In our setting, even though our proof yields a procedure for constructing an element of the core, it requires solving a series of optimization questions, which for arbitrary convex games are NP-hard. It would be desirable to find classes of convex games where our proof yields a polynomial time algorithm.

We are also interested in finding processes that lead to the core in not necessarily convex games; though *randomized algorithms* such as the ones presented in [17] and [9] could trivially extend to the overlapping setting, they would be of little practical value here due to the huge space of potential overlapping configurations. Therefore, we are interested in finding ways to exploit known game structure to prune the search space for potential stable configurations. Another subject of future research is extending our model to allow for infinite coalition structures. Furthermore, it would be definitely interesting to establish links between outcomes in the core and outcomes of bargaining equilibria in overlapping coalitional bargaining games.

Finally, the incorporation of actions in our model allows for the investigation of action stochasticity and, more generally, uncertainty in an OCF setting. For instance, a coalitional action can be associated with a distribution over possible payoff outcomes resulting from its execution. This poses challenges to study such models from both a theoretical and a practical standpoint, since the introduction of uncertainty leads to several intricacies not readily resolved by the use of “deterministic” concepts and models, as the work of Suijs *et al.* [29,30], Blankenburg *et al.* [5], and Chalkiadakis *et al.* [8,9,10] demonstrates. On a related note, enriching our model description so as to capture type uncertainty [8,9,10] would allow for the ready translation of uncertainty regarding the types (capabilities) of players to coalitional value uncertainty, while also capturing the potential stochasticity of actions’ outcomes.

Acknowledgements. We would like to thank the participants of the 13th Coalition Theory Network workshop, and especially Tomasz Michalak, for valuable feedback regarding these ideas. This research was undertaken as part of the ALADDIN (Autonomous

Learning Agents for Decentralised Data and Information Networks) project, which is jointly funded by a BAE Systems and EPSRC strategic partnership (EP/C548051/1). Edith Elkind was also supported by ESRC under grant ES/F035845/1.

References

1. Albizuri, M., Aurrecochea, J., Zarzuelo, J.: Configuration values: Extensions of the coalitional Owen value. *Games and Economic Behavior* 57, 1–17 (2006)
2. Apt, K., Radzik, T.: Stable Partitions in coalitional games, Work.Paper (2006), <http://arxiv.org/abs/cs.GT/0605132>
3. Apt, K., Witzel, A.: A Generic Approach to Coalition Formation, Work.Paper (2007), <http://arxiv.org/abs/0709.0435>
4. Aumann, R., Dreze, J.: Cooperative Games with Coalition Structures. *International Journal of Game Theory* 3(4), 217–237 (1974)
5. Blankenburg, B., Klusch, M., Shehory, O.: Fuzzy Kernel-Stable Coalitions Between Rational Agents. In: *Proceedings of the Second International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2003)* (2003)
6. Bondareva, O.: Some Applications of Linear Programming Methods to the Theory of Cooperative Games. *Problemy Kibernetiki* 10, 119–139 (1963) (in Russian)
7. Branzei, R., Dimitrov, D., Tijs, S.: *Models in cooperative game theory*. Springer, Heidelberg (2005)
8. Chalkiadakis, G.: A Bayesian Approach to Multiagent Reinforcement Learning and Coalition Formation under Uncertainty. PhD thesis, Department of Computer Science, University of Toronto, Toronto (2007)
9. Chalkiadakis, G., Boutilier, C.: Bayesian Reinforcement Learning for Coalition Formation Under Uncertainty. In: *Proceedings of the 3rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)* (2004)
10. Chalkiadakis, G., Markakis, E., Boutilier, C.: Coalition Formation under Uncertainty: Bargaining Equilibria and the Bayesian Core Stability Concept. In: *Proceedings of AAMAS 2007* (2007)
11. Chvatal, V.: Rational behavior and computational complexity, Tech. Rep. SOCS-78.9, School of Computer Science, McGill University, Montreal (1978)
12. Chwe, M.: Farsighted Coalitional Stability. *Journal of Economic Theory* 63, 299–325 (1994)
13. Conconi, P., Perroni, C.: Issue Linkage and Issue Tie-In in Multilateral Negotiations, CESifo Work.Paper 601 (2001)
14. Conitzer, V., Sandholm, T.: Complexity of Determining Non-Emptiness of the Core. In: *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI 2003)* (2003)
15. Dang, V.D., Dash, R.K., Rogers, A., Jennings, N.R.: Overlapping coalition formation for efficient data fusion in multi-sensor networks. In: *Proceedings of the 21st National Conference on AI (AAAI 2006)*, pp. 635–640 (2006)
16. Deng, X., Papadimitriou, C.: On the complexity of cooperative solution concepts. *Mathematics of Operation Research* 19, 257–266 (1994)
17. Dieckmann, T., Schwalbe, U.: Dynamic Coalition Formation and the Core, Economics Department Working Paper Series, Department of Economics, National University of Ireland - Maynooth (1998)
18. Gillies, D.: Some Theorems on n-Person Games. PhD thesis, Department of Mathematics, Princeton University, Princeton (1953)

19. Greenberg, J., Kats, A.: Unilateral Transfers and Pareto Optimality. *Econometrica* 48(3), 777–779 (1980)
20. Osborne, M., Rubinstein, A.: *A course in game theory*. MIT Press, Cambridge (1994)
21. Owen, G.: Values of games with a priori unions. In: HERNN, R., MOSCHLIN, O. (eds.) *Essays in Honor of Oskar Morgenstern*. *Lecture Notes in Economics and Mathematical Systems*. Springer, Heidelberg (1977)
22. Patel, J., Teacy, W., Jennings, N.R., Luck, M., Chalmers, S., Oren, N., Norman, T., Preece, A., Gray, P., Shercliff, G., Stockreisser, P., Shao, J., Gray, W., Fiddian, N., Thompson, S.: Agent-based virtual organisations for the Grid. *Multiagent and Grid Systems* 1(4), 237–249 (2005)
23. Sandholm, T., Larson, K., Andersson, M., Shehory, O., Tohme, F.: Coalition Structure Generation with Worst Case Guarantees. *Artificial Intelligence* 111(1–2), 209–238 (1999)
24. Sandholm, T., Lesser, V.: Coalitions Among Computationally Bounded Agents. *Artif. Intelligence* 94(1) (1997)
25. Shapley, L.: A Value for n-Person Games. In: Kuhn, H., Tucker, A. (eds.) *Contributions to the Theory of Games II*, pp. 307–317. Princeton University Press, Princeton (1953)
26. Shapley, L.: On Balanced Sets and Cores. *Naval Research Logistics Quarterly* 14, 453–460 (1967)
27. Shapley, L.: Cores of Convex Games. *International Journal of Game Theory* 1, 11–26 (1971)
28. Shehory, O., Kraus, S.: Formation of overlapping coalitions for precedence-ordered task-execution among autonomous agents. In: *Proc. of the 2nd Intern. Conference on Multi-Agent Systems (ICMAS 1996)*, pp. 330–337 (1996)
29. Suijs, J., Borm, P.: Stochastic cooperative games: superadditivity, convexity and certainty equivalents. *Journal of Games and Economic Behavior* 27, 331–345 (1999)
30. Suijs, J., Borm, P., Wagenaere, A.D., Tijs, S.: Cooperative games with stochastic payoffs. *European Journal of Operational Research* 113, 193–205 (1999)
31. Tanin, A.: On the core of network synthesis games. *Mathematical Programming* 50, 123–135 (1991)
32. von Neumann, J., Morgenstern, O.: *Theory of Games and Economic Behavior*. Princeton University Press, Princeton (1944)

A Network-Based Asymmetric Nash Bargaining Solution

Edoardo Gallo

University of Oxford, Nuffield College, New Road, Oxford OX1 1NF, UK
edoardo.gallo@economics.ox.ac.uk

Abstract. This paper presents an evolutionary bargaining model between two groups of buyers and sellers. One buyer and one seller are randomly matched to play the Nash demand game: they choose a best reply based on information about past bargains coming from other members of their group. Information arrival is modeled as a Poisson process, and the rates of these processes form a weighted communication network. Over the long run, the stochastically stable division is the asymmetric Nash bargaining solution (ANB) with weights determined by the structure of the communication network in each group. The optimal networks for a group are (quasi)-regular networks without weak links.

Keywords: Network, noncooperative bargaining.

1 Introduction

Anybody who has ever traveled to a developing country has witnessed that bargaining among members of different groups is very much the norm in everyday life. In a variety of contexts, ranging from the marketplace to renting a form of transportation, there are members of a group of sellers constantly negotiating with members of several groups of buyers, e.g. the tourist and local populations. While less predominant, bargaining in developed countries is also common in, for instance, transactions involving illegal activities (e.g. prostitution, drugs) and wholesale markets.¹

Classical bargaining theory provides only a partial model for bargaining in this context. The seminal contributions by Nash ([2]) and Rubinstein ([3]) consider a game with two players, and the outcome is fully determined by the characteristics of these two agents. Specifically, in these and most subsequent models the outcome is determined by the functional form of each player's utility. If these models were naively applied to, say, predict the outcomes of the thousands of negotiations happening on a typical day in the marketplace then the models would predict a staggering multiplicity of outcomes, corresponding to the thousands of possible combinations of utility pairs.

¹ See the extended version of this paper ([1]) for references to empirical work on pricing in illegal and wholesale markets.

On the other hand, empirical and anecdotal evidence points to a remarkable level of convergence. Sellers bargaining on the price of the same product in independent pairwise interactions with different buyers will end up charging the same price or a limited set of prices even though there is no centralized coordination device. Over time norms on the price to charge get established and everyone follows them. This does not mean that only one price will be observed: the norm might be to charge multiple prices depending on which customer segment, e.g. tourist or local, a buyer belongs to. What is important to note is that these norms are at the *group* level, and therefore their exact form will depend on some properties of these groups, not just of their individual members.

Moreover, in the contexts studied in this paper the individual agents have limited knowledge of the process they are embedded in. Classical bargaining theory assumes players have common knowledge of the utility functions, or of the distribution of utility functions, and they know the structure of the game they are playing. On the other hand, in the typical transaction in a marketplace the buyer will have little or no knowledge of the seller, let alone having any clue of the functional form of his utility. As Young ([4]) argues, in these situations bargainers' expectations are not shaped by some sophisticated reasoning on the nature of the game and the strategy of the other player, but, more simply, by previous market transactions they have heard about as they wander in the marketplace.

This paper presents an evolutionary bargaining model where agents base their decisions on information about previous plays which circulates within the group they belong to. The group structure determines the information each agent in the group has access to, this in turn determines the agent's play, and, in the long-term, the establishment of a norm of play that each member of the group follows.

The first result in theorem 1 shows that the process without "mistakes" always converges to a convention independently from the initial state, as long as the network is not complete or close to complete. A convention means that each buyer always makes the same demand x and each seller always makes the same demand $1 - x$. The condition on the network structure guarantees that the information available to each player on the history of demands is sufficiently incomplete to avoid the whole process getting stuck in a cycle.

The second result in theorem 2 shows that the process with "mistakes" converges to a unique stable division which maximizes the asymmetric Nash bargaining solution (ANB) with weights that depend on the network structure. Specifically, the weights are determined by the player in each group with the least number and/or weakest communication links. A consequence of this result is that the optimal architectures for a group are communication networks where all the players are connected by strong links and have very similar number of connections.

Due to space constraints all the proofs of the above results are omitted in this paper, the interested reader is invited to consult the extended version of this paper ([1]). The extended version contains several further results derived from

the model presented here. First, it shows how changes in the network structure affect the shares a group obtains in the stable division. Second, it discusses the implications of the results for the observed pricing patterns in specific wholesale and illegal markets. Third, it explores how the predictions of the model change if the two groups belong to the same communication network, allowing in this way communication between buyers and sellers.

The main contribution of this paper is to investigate how the internal communication structure of a group determines the outcome of a bargaining process. This has implications for the optimal communication structure for a group. Moreover, further analysis of this model carried out in [1] studies the effects of changes in the communication structure on the solution. The predictions in [1] open up the possibility of empirical testing to verify whether network structure plays a significant role in determining the outcome of bargaining processes.

Finally, an overview of the related literature. In the bargaining literature, the most closely related model is by Young ([4]). In the economics of network literature there are a few papers which investigate how the structural position of one agent in a network affects her bargaining power, and here I will discuss one of the main contributions by Calvó-Armengol ([5]).

Young ([4]) builds an evolutionary bargaining model where individuals from two populations of bargainers are randomly matched to play the Nash demand game: they make demands by choosing best replies based on an incomplete knowledge of precedents. The model in this paper adopts the same bargaining procedure and the same assumptions on the agents' knowledge and rationality. However, the main substantial difference is that [4] exogenously imposes agents' sample size of past information and does not model the process by which agents receive the information to play the game. Thus, the system in [4] is simpler to analyze, but the predictions of the model do not depend on the internal communication structure of a group and his model does not allow the comparative statics analysis carried out in [1].

Calvó-Armengol ([5]) studies the noncooperative bargaining game by Rubinstein ([3]) for the case of n players connected by a graph which constrains the feasible bargaining pairs. He shows that there is a unique stationary subgame perfect equilibrium. When the population is homogeneous in time preferences, ex-post payoffs do not depend on the network structure. However, the communication structure is important for expected payoffs, and, after removing cost to delay through disagreement, he derives the players' bargaining power as a function of their relative positions in the network. The focus of our paper is different. Calvó-Armengol ([5]) investigates how the position of *one* agent in a network affects her *individual* payoffs, while the aim of this paper is to understand how the *overall* structural properties of the network influence the establishment of a norm for the *whole group*.

The rest of the paper is organized as follows. Section [2] presents the model. Section [3] shows that the process converges to a convention. Section [4] derives the bargaining solution. Section [5] derives the optimal network structure for a group. Section [6] concludes and presents further results in [1].

2 The Model

This section presents the main elements of the model: the adaptive play bargaining process, the network concepts and terminology used, and the Markov process which describes the evolution of the system.

Adaptive play bargaining process. Consider two finite, non-empty groups of individuals $B = \{1, \dots, n_B\}$ and $S = \{1, \dots, n_S\}$: the buyers and sellers. The two groups are separate, i.e. $B \cap S = \emptyset$. In each period t one buyer and one seller drawn at random meet to divide a pie of size normalized to one. They play the Nash demand game: b demands a fraction x_t and s demands a fraction y_t , if $x_t + y_t \leq 1$ then b and s get their demands, otherwise they get nothing. For mathematical convenience, assume that the set of possible divisions is discrete and finite. Let $\delta = 10^{-p}$ ($p \in \mathbb{Z}_+$) be the precision of the demands, and assume $x_t, y_t \in D$, where D is the set of all p -place decimal fractions that are feasible demands. The sequence $h = \{(x_1, y_1), \dots, (x_t, y_t)\}$ is the complete *global history* up to and including period t . Each agent remembers the last m rounds of the bargaining game that she has played, where m stands for the memory of the player.

Agents receive information to play the game as follows. Suppose player $b \in B$ is picked to play the game at $t + 1$: in the $\Delta t = 1$ time period she receives information from some of the other buyers in B about past bargaining rounds. Information arrival is modeled as a Poisson process. Specifically, in the $\Delta t = 1$ time interval, the probability $P(s_{bj}(\Delta t = 1) = k)$ that b receives a sample $s_{bj}(\Delta t = 1)$ of k past bargains from player j is equal to:

$$P(s_{bj}(\Delta t = 1) = k) = \frac{e^{-g_{bj}} g_{bj}^k}{k!} \tag{1}$$

where the rate g_{bj} of arrival of information can be interpreted as how often b receives information from j . By standard properties of Poisson processes, the probability that b receives a sample $s_b(\Delta t = 1)$ of k past bargains is equal to:

$$P(s_b(\Delta t = 1) = k) = \sum_{j \in B} P(s_{bj}(\Delta t = 1) = k) = \frac{e^{-g_b} g_b^k}{k!} \tag{2}$$

where $g_b = \sum_{j \in B} g_{bj}$, and the size of the sample k is the amount of information that b has prior to playing the bargaining game. Clearly, the expected amount of information that b has is equal to g_b .

Agents are boundedly rational as they are not aware of the game they are embedded in and they base their decision exclusively on the information they receive. Specifically, agents do not have prior knowledge or beliefs about the utility function of the other side, and they do not know the distribution of utility functions in the general population. Agent b chooses an optimal reply to the cumulative probability distribution $G(y)$ of the demands y_j made by sellers in his sample, where $G(y) = \frac{h}{s_b(t)}$ if and only if there are exactly h demands y_l in the sample $s_b(t)$ such that $y_l \leq y, \forall y \in D$.

Agent b has a concave and strictly increasing von Neumann-Morgenstern utility function $u(x)$. Assume that $u(x)$ is defined for all $x \in [0, 1]$ and that it is normalized so that $u(0) = 0$. Buyer b 's expected payoff is then equal to $Eu(x) \equiv u(x)G(1 - x)$. Thus, b chooses x_{t+1} so as to maximize $Eu(x)$, and if there are several values of x to choose from then each one of them is chosen with positive probability.

The set-up for seller s is analogous, and the utility function of the sellers will be denoted by $v(y)$.

Networks. A *weighted, undirected network* is represented by a symmetric matrix $[g_{ij}]^{n \times n}$, where $g_{ij} \in \mathbb{R}_+$. The entry g_{ij} is the rate at which i receives information about past bargains from j , it indicates the *strength* of the *communication link* between i and j . If $g_{ij} > 0$ then agents i and j are connected and they communicate directly with each other. If $g_{ij} = 0$ then i and j are not connected in the communication network. Throughout this paper let $g_{ii} \equiv \bar{g}$, i.e. an agent can also receive information from her own past experience and this channel is constant for all players.

The *neighborhood* of i in g is $L_i(g) = \{j \in N | g_{ij} > 0\}$, $d_i \equiv |L_i(g)|$ denotes the size of i 's neighborhood, or the *degree* of i . $g_i \equiv \sum_{j \in L_i(g)} g_{ij}$ is the *weighted degree* of i . The *complete network* is the network $g^C = \{g | g_{ij} > 0, \forall i, j \in N\}$. A *regular network* $\bar{g}_{d,a}$ is a network $\bar{g}_{d,a} = \{g | g_{ij} = 0, a; d_i(g) \equiv d; \forall i, j \in N; a \in \mathbb{R}_+\}$. Let $Z(g) = \min_Z \{|\bigcup_{i=1}^Z L_i| = n \mid \forall i, \exists j \neq i \text{ such that } L_i \cap L_j \neq \emptyset\}$. In words, $Z(g)$ is the minimum number of nodes in g such that the union of their (partly overlapping) neighborhoods covers the whole network. Clearly, $Z(g^C) = 1$.

Markov process. Let S be the state space, whose elements are sequences $\mathbf{s} = \{v_1, \dots, v_n\}$, where v_i is a vector of size m for each agent i , and $n \equiv n_B + n_S$. If $i \in B$ then $v_i = \{y_{k-m}^i, \dots, y_k^i\}$, i.e. the entries of v_i are the m last demands made by sellers in bargaining rounds involving i . Similarly, if $i \in S$ then $v_i = \{x_{k-m}^i, \dots, x_k^i\}$. Let $p_b(x | \mathbf{s})$ be agent b 's best-reply distribution, i.e. $p_b(x | \mathbf{s}) > 0$ if and only if demanding x is b 's best-reply to a sample received when the system is in state \mathbf{s} . Analogously, $p_s(y | \mathbf{s})$ is seller s 's best-reply distribution.

Assume that the process starts at an arbitrary time $t_0 > n \cdot m$, and denote the initial state by \mathbf{s}^0 . At each $t > t^0$ one pair of agents $(b, s) \in B \times S$ is drawn at random with probability $\pi(b, s)$, where $\pi(b, s) > 0, \forall (b, s) \in B \times S$. At time t , consider a state $\mathbf{s} = \{v_1, \dots, v_i, \dots, v_j, \dots, v_n\}$, where $v_i = \{y_{k-m}^i, \dots, y_k^i\}$, $v_j = \{x_{k-m}^j, \dots, x_k^j\}$. Define \mathbf{s}' to be a successor of \mathbf{s} if it has the form $\mathbf{s}' = \{v_1, \dots, v'_i, \dots, v'_j, \dots, v_n\}$. where $v'_i = \{y_{k-m+1}^i, \dots, y_{k+1}^i\}$ and $v'_j = \{x_{k-m+1}^j, \dots, x_{k+1}^j\}$. The transition probability $P_{ss'}$ of moving from state \mathbf{s} to state \mathbf{s}' is then equal to:

$$P_{ss'} = \sum_{b \in B} \sum_{s \in S} \pi(b, s) p_b(x_{t+1} | \mathbf{s}) p_s(y_{t+1} | \mathbf{s}) \tag{3}$$

This Markov process will be called an *evolutionary bargaining process with local information sharing*.

Mistakes. In the process described so far agents always give a best reply to the sample they happen to pick. In reality, people make mistakes for a variety of reasons: human beings are bad at computing probabilities and they might miscalculate the expected utility from an offer, they are prone to get “distracted,” they experiment, or sometimes they are outright irrational. The following is a more formal definition of a mistake.

Definition 1. Let $\mathbf{s} = \{v_1, \dots, v_i, \dots, v_j, \dots, v_n\}$ and let $\mathbf{s}' = \{v_1, \dots, v'_i, \dots, v'_j, \dots, v_n\}$ be a successor of \mathbf{s} , where $v_i = \{y_{k-m}^i, \dots, y_k^i\}$, $v_j = \{x_{k-m}^j, \dots, x_k^j\}$, $v'_i = \{y_{k-m+1}^i, \dots, y_{k+1}^i\}$ and $v'_j = \{x_{k-m+1}^j, \dots, x_{k+1}^j\}$. The demand x_{k+1}^j is a mistake by i if it is not a best response to any sample i could have received given that the system is in state \mathbf{s} . A mistake y_{k+1}^i by j is defined similarly.

Another concept that will be useful in the analysis of the perturbed process is the *resistance* in moving from one state \mathbf{s} to another state \mathbf{s}' .

Definition 2. Let \mathbf{s} and \mathbf{s}' be two states of the system. The resistance $r(\mathbf{s}, \mathbf{s}')$ is the least number of mistakes required for the system to go from state \mathbf{s} to \mathbf{s}' .

Note that if \mathbf{s}' is a successor of \mathbf{s} then $r(\mathbf{s}, \mathbf{s}') = 0, 1, 2$ as the maximum number of mistakes in any one-time transition is two, i.e. both the buyer and seller involved in that bargaining round make a mistake.

Now let ϵ be the absolute probability that agents in the model make mistakes, and let λ_b, λ_s be the relative probability that buyers and sellers do it respectively. Thus, $\epsilon\lambda_b$ and $\epsilon\lambda_s$ are the probabilities that buyers and sellers make a mistake. Denote by $q_b(x | \mathbf{s})$ the buyer’s conditional probability of choosing x given that the current state is \mathbf{s} and that she is not giving a best-response offer to the sample picked. Assume $\lambda_b, \lambda_s, \epsilon > 0$ and that $q_b(x | \mathbf{s}), q_s(y | \mathbf{s})$ have full support.

This process also yields a stationary Markov chain on S that can be described by the probability of moving from a state \mathbf{s} to a successor state \mathbf{s}' , similarly to equation (3) above. Assume that the process starts at an arbitrary time $t_0 > n \cdot m$, and denote the initial state by \mathbf{s}^0 . At each $t > t^0$ one pair of agents $(b, s) \in B \times S$ is drawn at random with probability $\pi(b, s)$, where $\pi(b, s) > 0, \forall (b, s) \in B \times S$. Let \mathbf{s} be the state at time t , and let \mathbf{s}' be a successor of \mathbf{s} , where \mathbf{s} and \mathbf{s}' are defined above. The transition probability $P_{\mathbf{s}\mathbf{s}'}^\epsilon$ of moving from state \mathbf{s} to state \mathbf{s}' is then equal to:

$$\begin{aligned}
 P_{\mathbf{s}\mathbf{s}'}^\epsilon = & \sum_{b \in B} \sum_{s \in S} \pi(b, s) [(1 - \epsilon\lambda_b)(1 - \epsilon\lambda_s)p_b(x_{t+1} | \mathbf{s})p_s(y_{t+1} | \mathbf{s}) + \\
 & + \epsilon\lambda_b(1 - \epsilon\lambda_s)q_b(x_{t+1} | \mathbf{s})p_s(y_{t+1} | \mathbf{s}) + \epsilon\lambda_s(1 - \epsilon\lambda_b)q_s(x_{t+1} | \mathbf{s})p_b(y_{t+1} | \mathbf{s}) + \\
 & + \epsilon^2\lambda_b\lambda_sq_b(x_{t+1} | \mathbf{s})q_s(y_{t+1} | \mathbf{s})] \tag{4}
 \end{aligned}$$

This Markov process will be called a *perturbed evolutionary bargaining process with local information sharing*. Note that $\lim_{\epsilon \rightarrow 0} P_{\mathbf{s}\mathbf{s}'}^\epsilon = P_{\mathbf{s}\mathbf{s}'}$.

3 Convergence

First, consider the unperturbed process P . The first step in the analysis is to define an appropriate concept of stability for this system, and to show that in the long-term the process will reach it. Intuitively, the system will be in a stable state if, after a certain time t , any buyer (seller) will always make the same demand x because in any sample she receives of previous sellers' (buyers') demands, the sellers (buyers) have always demanded $1 - x$. The following definition states this more formally.

Definition 3. *A state \mathbf{s} is a convention if any $v_i \in \mathbf{s}$ with $i \in B$ is such that $v_i = (1 - x, \dots, 1 - x)$, and any $v_j \in \mathbf{s}$ with $j \in S$ is such that $v_j = (x, \dots, x)$, where $x \in D$, $0 < x < 1$. Hereafter, denote this convention by \mathbf{x} .*

The following lemma shows that convention \mathbf{x} is an appropriate definition to work with because any \mathbf{x} is an absorbing state of P .

Lemma 1. *Every convention \mathbf{x} is an absorbing state of the Markov process P in (3).*

Proof: See [1].

The following theorem shows that if information is sufficiently incomplete then the process P converges to a convention. The incompleteness of information is delivered by the network structure: if the network is not complete, or close to complete, then some players cannot sample the history of the game for rounds played by individuals that do not belong to their neighborhoods.

Theorem 1. *If $Z(g^B), Z(g^S) \geq 2$ then from any initial state the evolutionary bargaining process converges almost surely to a convention.*

Proof: See [1].

Theorem 1 in [6] is a more general version of this statement applicable to any weakly acyclic n -person game. The goal is to show that from any initial state \mathbf{s} there is a positive probability p independent of t of reaching a convention within a finite number of steps. The requirement $Z(g^B) \geq 2$ means that there is no agent whose neighborhood includes the whole network and there are at least two agents b' and b_0 such that $g_{b'b_0} = 0$ and their intersection includes at least one agent b . The same requirement applies to the sellers' network, where agents s and s_0 are the equivalent of agents b' and b_0 respectively.

Now, consider the following path which happens with positive probability from any state \mathbf{s} at time t . First, b and s are picked to play the game for m consecutive periods, they draw samples σ and σ' respectively, they demand best-replies x and y respectively, and therefore they obtain a run $\rho = \{(x, y), \dots, (x, y)\}$ such that $v_b = (y, \dots, y)$ and $v_s = (x, \dots, x)$. Second, b' and s' are picked to play the game for m consecutive periods, they draw samples from v_b and v_s each time, they demand best-replies $1 - y$ and $1 - x$ respectively, and therefore they obtain a run $\rho' = \{(1 - y, 1 - x), \dots, (1 - y, 1 - x)\}$ such that $v_{b'} = (1 - x, \dots, 1 - x)$ and $v_{s'} =$

$(1 - y, \dots, 1 - y)$. Third, b_0 and s_0 are picked to play the game for m consecutive periods, they draw sample v_b and $v_{s'}$ each time, they demand best-replies $1 - y$ and y respectively, and therefore they obtain a run $\rho'' = \{(1 - y, y), \dots, (1 - y, y)\}$ such that $v_{b_0} = (y, \dots, y)$ and $v_{s_0} = (1 - y, \dots, 1 - y)$. Hereafter it is clear that there is a positive probability of reaching a convention $\mathbf{x} = (1 - y, y)$.

Second, consider the perturbed process P^ϵ . Given that the distribution q_b and q_s have full support, P^ϵ is irreducible. Thus, P^ϵ has a unique stationary distribution. Moreover, P^ϵ is strongly ergodic, i.e. $\forall \mathbf{s} \in S$, μ_s^ϵ is with probability one the relative frequency with which state \mathbf{s} will be observed in the first t periods as $t \rightarrow \infty$. The stability concept for this kind of perturbed process is a *stochastically stable convention*, which was introduced by [7].

Definition 4. A convention \mathbf{s} is stochastically stable if $\lim_{\epsilon \rightarrow 0} \mu_s^\epsilon > 0$. A convention \mathbf{s} is strongly stable if $\lim_{\epsilon \rightarrow 0} \mu_s^\epsilon = 1$.

Intuitively, in the long-run stochastically stable conventions will be observed much more frequently than unstable conventions when the probability ϵ of mistakes is small. A strongly stable convention will be observed almost all the time. The technique to compute the stochastically stable conventions is standard and it will not be explained in detail below, see [8] for an excellent introduction.

Construct a weighted, directed network $[r_{\mathbf{s}^i \mathbf{s}^j}]^{z \times z}$, where the nodes are the states $\mathbf{s} \in S$, the links are the resistances $r_{\mathbf{s}^i \mathbf{s}^j}$ connecting \mathbf{s}^i to \mathbf{s}^j , and z is the total number of states in S . Define an \mathbf{x} -tree $t_{\mathbf{x}} \in T_{\mathbf{x}}$ to be a collection of links in $[r_{\mathbf{s}^i \mathbf{s}^j}]^{z \times z}$ such that, from every node $\mathbf{x}' \neq \mathbf{x}$, there is a unique directed path to \mathbf{x} and there are no cycles. This construction leads to the definition of the concept of *stochastic potential* of a convention \mathbf{x} .

Definition 5. The stochastic potential $\gamma(\mathbf{x})$ of a convention \mathbf{x} is the least resistance among all $t_{\mathbf{x}} \in T_{\mathbf{x}}$. Mathematically:

$$\gamma(\mathbf{x}) = \min_{t_{\mathbf{x}} \in T_{\mathbf{x}}} \sum_{(\mathbf{x}', \mathbf{x}'') \in T_{\mathbf{x}}} r(\mathbf{x}', \mathbf{x}'') \tag{5}$$

Theorem 4 in [6] explains how to compute the stochastically stable states. The following is a special case of that result.

Theorem 2. Let μ^0 be a stationary distribution of the unperturbed process P . Then $\lim_{\epsilon \rightarrow 0} \mu_s^\epsilon = \mu^0$. Moreover, $\mu^0 > 0$, i.e. \mathbf{s} is stochastically stable, if and only if $\mathbf{s} = \mathbf{x}$ is a convention and $\gamma(\mathbf{x})$ has minimum stochastic potential among all conventions.

Proof. See [4]. □

4 Asymmetric Nash Bargaining Solution

Let us apply the methodology outlined above to find the division which the process will converge to. However, before proceeding with the analysis, a *mean-field assumption* is necessary to make the model more tractable. Given that

information arrival to a buyer/seller about past bargains is a Poisson process, there is a small probability that there are significant fluctuations in the total size of the sample received by the same buyer/seller. The variability of an agents' information sample over time poses significant challenges to an analytical investigation of the model, but, luckily, these fluctuations are small and therefore they can be ignored. Technically, assume that the size of the information sample of the buyer b is constant and equal to the amount of information b receives in expectation given the Poisson processes involving b , i.e. $s_b(t) \equiv \sum_{j \in B} g_{bj}$. The same assumption holds for the seller s . Hereafter, also assume that the individual memory $m \geq \max\{g_b, g_s\}$, where $b \in B$ and $s \in S$.

The first step is to compute the minimum resistance to moving from the convention \mathbf{x} to the basin of a different convention \mathbf{x}' . This is done in the following lemma.

Lemma 2. *For every $x \in D^0$ the minimum resistance to moving from x to a state in some other basin is $\lceil R(x) \rceil$, where:*

$$R(x) = \min_{b \in B} g_b \left(1 - \frac{u(x - \delta)}{u(x)} \right) \wedge \min_{s \in S} g_s \frac{v(1 - x)}{v(1 - \delta)} \wedge \min_{s \in S} g_s \left(1 - \frac{v(1 - x - \delta)}{v(1 - x)} \right) \tag{6}$$

Proof: See [\[1\]](#).

The intuition is as follows. First of all, some agents have to make mistakes in order for the system to move from one convention to a state in the basin of another convention. The agents who will switch with the least number of mistakes in their sample are the ones who receive the smallest samples. This explains the factors $\min_{b \in B} g_b$ and $\min_{s \in S} g_s$ in equation (6). Now, consider the case when some sellers make a mistake. The smallest mistake they can make is to demand a quantity δ more than the conventional demand $1 - x$. If they do this, buyers will attempt to resist up to the point when the utility from getting the conventional amount x some of the time, i.e. when sellers do not make a mistake, is equal to the utility from getting the lower amount $x - \delta$ all the time. This gives the first term in equation (6). The third term is the equivalent of the first one, but this time the buyers make a mistake and demand δ more than the conventional amount x .

Another possibility is that some buyers make a mistake, but this time they demand less than the conventional amount x . The "worst" mistake, from the buyers' point of view, would be to demand the minimum amount δ . If they do this, sellers will only switch at the point when the utility from getting the higher amount $1 - \delta$ some of the time, i.e. when buyers make a mistake, is higher than the utility from getting the conventional amount x all the time. This gives the second term in equation (6). The careful reader will point out that there should also be a fourth term, i.e. the equivalent of the second one with the roles of buyers and sellers reversed. This is true, but it is not included in equation (6) because this term is never strictly smaller than the last term.

The expression for $R(x)$ in (6) is the minimum of three monotone functions: the first two are strictly decreasing in x , while the last one is strictly increasing

in x . Thus, $R(x)$ is first strictly increasing and then strictly decreasing as x increases, so as $\delta \rightarrow 0$ it achieves its maximum at a unique value on the subset D .² Using this fact, the following theorem shows that there is a unique stable division, which is the division that maximizes the asymmetric Nash bargaining solution.

Theorem 3. *As $\delta \rightarrow 0$ there exists a unique stable division $(x, 1 - x)$, which maximizes the following asymmetric Nash bargaining solution:*

$$w^\beta(x)v^\sigma(1 - x) \tag{7}$$

where $\beta \equiv \min_{b \in B} g_b(g^B)$ and $\sigma \equiv \min_{s \in S} g_s(g^S)$.

Proof: See [1].

The proof of the theorem follows from two lemmas. The first lemma shows that a division $(x, 1 - x)$ is generically stable if and only if x maximizes the function $R(x)$ in equation (6). The second lemma shows that the maxima of $R(x)$ converge to the asymmetric Nash bargaining solution in (7) as $\delta \rightarrow 0$.

The intuition behind the solution is that if the precision δ is sufficiently small then over time the two groups will settle on a conventional division, which is the asymmetric Nash bargaining solution. This solution crucially depends on the communication networks that buyers and sellers use to learn about past bargaining rounds to determine what to demand once they are picked to play. More precisely, *ceteris paribus* (i.e. agents’ risk-aversion in the two groups is the same), the share a group gets hinges on the agent in the group with the least number and/or weakest communication links. The reason is that this agent will be the least informed when it comes to play the game, and therefore she will be the most susceptible to respond to mistakes from the other side. Over time, this susceptibility weakens the bargaining position of the whole group. As in standard bargaining models, the solution also depends on the utilities of the agents. *Ceteris paribus* (i.e. the least connected agents in each group have the same weighted degree), a group with less risk-averse agents will have a stronger bargaining position because agents who are less risk-averse are more likely to take chances, and therefore they are more demanding.

5 The Weakness of Weak Ties

Given the solution in theorem 3, a natural question is what is the optimal communication structure for a group of individuals that are engaging in this evolutionary bargaining process with another group. First, let us define a class of *quasi-regular networks*, which are "similar" to a given regular network.

Definition 6. *Consider the set G of undirected networks with n nodes and at most L links. Let $\bar{g}_{d,a}$ be the regular network with degree $d = \lfloor \frac{2L}{n} \rfloor$, i.e. the largest*

² Technically, $R(x)$ can achieve its maximum at one value x^* or at two values x^* and $x^* + \delta$. As $\delta \rightarrow 0$ these two values clearly converge to a unique maximum x^* .

regular network in G . The network $g \in G$ is a quasi-regular network generated by $\bar{g}_{d,a}$ if it can be obtained by randomly adding k links of any strength to $\bar{g}_{d,a}$, where $k \in [0, L - \frac{n}{2}]$.

A quasi-regular network is a network that is similar to a regular network in the sense that the links are distributed evenly among the nodes and there is minimal degree variation. Note that if $L/n \in \mathbb{N}$, i.e. the links can be exactly divided among the nodes, then the unique quasi-regular network is the generating regular network. If $L/n \notin \mathbb{N}$ then each node has at least as many links as in the largest feasible regular network, i.e. the generating regular network, and the remaining links are randomly assigned. The optimal communication structure for a group is a quasi-regular network, as the following corollary shows.

Corollary 1. Fix a communication network g^S for the sellers. Consider the set G of all possible communication structures g^B among the n_B buyers such that the total number of links is $L < \frac{n_B}{2}(n_B - 1)$ and the strength of each link is in the $[\underline{s}, \bar{s}]$ range, where $\underline{s}, \bar{s} \in \mathbb{R}$. The subset of networks $G_B \subset G$ that gives the highest share to buyers are the quasi-regular networks generated by the regular network $\bar{g}_{d,\bar{s}}$, where $d = \lfloor \frac{2L}{n_B} \rfloor$. The same statement holds reversing the roles of buyers and sellers.

Proof: See [□](#).

For illustrative purposes it is easier to give the intuition for the case of $L/n_B \in \mathbb{N}$. First, the optimal network must have communication links of maximum strength because they carry more information about past rounds, decreasing in this way buyers' susceptibility to sellers' mistakes. Second, the regular network is optimal because this is the unique network where the buyer with the lowest degree has the highest possible degree given the constraint L . If $L/n_B \notin \mathbb{N}$ then the regular network is still optimal but it is not unique anymore: a quasi-regular network generated by the regular network by randomly adding a few more links is also as good. Informally, the (quasi)-regular network is very "steady": it has no "weak points" that could be more susceptible to sellers' mistakes.

6 Conclusion and Extensions

This paper has presented an evolutionary bargaining model between two groups of buyers and sellers to explore the importance of social structure in determining bargaining outcomes. Here social structure means the communication network each agent relies on to gather information about past demands in order to play a best response. Over the long run, the stochastically stable division is the asymmetric Nash bargaining (ANB) solution with weights determined by the agent with the least weighted degree in each group. A consequence of the ANB solution is that the optimal networks for a group are (quasi)-regular networks: there is minimal variability in number of connections and agents are connected by strong bonds. Thus, close-knit networks with strong links are optimal in this

setting because they allow maximal sharing of information and they have no weak points that could increase the susceptibility to the other group's mistakes.

The extended version of this paper ([1]) contains other important results. Comparative statics analysis allows the exploration of how changes in the network structure affect the shares a group obtains in the stable division. The changes are modeled in terms of first and second order stochastic dominance shifts in the weighted degree distribution. The main result is that a group with a denser and more homogeneous communication structure will fare better.

Moreover, the extended version also solves the model for the case of buyers and sellers sharing the same communication network. The outcome of the bargaining process remains the same ANB solution. However, the consequences of the solution change. First, due to the modified set-up, the optimal networks for a group have to include all the agents and they are core-periphery networks. The members of the group that fares better form the core: a (quasi)-regular network with strong links. The members of the group that fares worse are at the periphery with few links per agent and no or very few connections amongst themselves. Second, a more homogeneous network, holding the mean degree constant, narrows down the difference between the two groups. In the limiting case of a regular network with homogeneous agents the stable division is 50-50.

Acknowledgments. I am very grateful to Peyton Young for his invaluable help and guidance throughout this project. Any remaining errors are the sole responsibility of the author.

References

1. Gallo, E.: Bargaining and Social Structure. Mimeo (2008)
2. Nash, J.F.: The Bargaining Problem. *Econometrica* 18(2), 155–162 (1950)
3. Rubinstein, A.: Perfect Equilibrium in a Bargaining Model. *Econometrica* 50(1), 97–109 (1982)
4. Young, P.H.: An Evolutionary Model of Bargaining. *Journal of Economic Theory* 59(1), 145–168 (1993)
5. Calvó-Armengol, A.: Bargaining Power in Communication Networks. *Mathematical Social Sciences* 41, 69 (2001)
6. Young, P.H.: The Evolution of Conventions. *Econometrica* 61(1), 57–84 (1993)
7. Foster, D., Young, P.H.: Stochastic Evolutionary Game Dynamics. *Theoretical Population Biology* 38, 219–232 (1990)
8. Young, P.H.: *Individual Strategy and Social Structure*. Princeton University Press, New York (2005)

How Public Opinion Forms

Fang Wu and Bernardo A. Huberman

Social Computing Lab, HP Labs, Palo Alto, CA 94304

Abstract. No aspect of the massive participation in content creation that the web enables is more evident than in the countless number of opinions, news and product reviews that are constantly posted on the Internet. Given their importance we have analyzed their temporal evolution in a number of scenarios. We have found that while ignorance of previous views leads to a uniform sampling of the range of opinions among a community, exposure of previous opinions to potential reviewers induces a trend following process which leads to the expression of increasingly extreme views. Moreover, when the expression of an opinion is costly and previous views are known, a selection bias softens the extreme views, as people exhibit a tendency to speak out differently from previous opinions. These findings are not only robust but also suggest simple procedures to extract given types of opinions from the population at large.

On reflection, it is rather surprising that people contribute opinions and reviews of topics which have already been extensively covered by others. While posting views is easy to understand when it involves no effort, like clicking on a button of a website, it is more puzzling in situations where it is costly, such as composing a review. If the opportunity to affect the overall opinion or rating diminishes with the number of published ones, why does anyone bother to incur the cost of contributing yet another review? From a rational choice theory point of view, if the utility to be gained does not outweigh the cost, people would refrain from expressing their views. And yet they do. This is reminiscent of the well analyzed voter's paradox [5,12,13], where a rational calculation of their success probability at determining the outcome of an election would make people stay home rather than vote, and yet they show up at the polls with high turnout rates (for a review, see [7]). In contrast to a political election, there is no concept of winning in online opinion systems. Rather, by contributing her own opinion to an existing opinion pool, a person affects the average or the distribution of opinions by a marginal amount that diminishes with the size of that pool.

Since user opinions play such an important role in trust building and the creation of consensus about many issues, there have been a number of recent of studies focused on the design, evaluation and utilization of online opinion systems [12,6,8] (for a survey, see [3]). It is surprising that with the exception of one study [10], little research has been done on the dynamic aspects of online opinion formation. It remains unclear, for example, whether reviews undergo any systematic changes as time goes on, or whether the opinions about given

political or societal views fluctuate a long time before reaching a final consensus. Thus the need to understand how online opinions are created and evolve in time in order to draw accurate conclusions from that data.

In this context we analyzed the dynamics of online opinion expression by analyzing the temporal evolution of a very large set of user views, ranging from 1.8 million online reviews of the 48,000 best selling books at [Amazon.com](#), to thousands of political resolutions voted on [Essembly.com](#) and the many arbitrary opinions offered for voting on [Jyte.com](#).

To start with, a forum where no historical data is available should exhibit no polarization of views as they are expressed over time. In order to test this as calibration for our study, we analyzed the votes of 16,660 resolves posted on the website [Essembly.com](#) from August 2, 2005 to December 12, 2006, among which 14,171 resolves received more than 21 votes. [Essembly.com](#) is a website that lets its users post and vote on political resolves by selecting one of the four choices: “agree”, “lean agree”, “lean against”, and “against”. A user does *not* see the voting results until she submits her own vote. When a user posts a new resolve, she is required to vote on it. The four voting options, from “agree” to “against”, are represented by $-1, -0.5, 0.5, 1$ in our database, respectively. When a user posts a new resolve, she tends to formulate it in a tone that sounds positive to her. As a consequence, 96.0% of all first votes are “agree”, while only 2.7% of all first votes are “against”. To remove this artificial bias, we discarded the first vote of each resolve and replaced all the remaining votes $\{X_n\}$ by $\{-X_n\}$ if $\sum X_n > 0$, where X_n denotes the quantified value of the n 'th vote. This way every vote is “agreed” by the majority.¹ This formed our final data set, which consists of 14,171 resolves, each having more than 20 votes.

We observed no clear trend in the expected value of EX_n . For each resolve we calculated a series of average votes: $\bar{X}_1, \dots, \bar{X}_{20}$, where \bar{X}_i is the average vote of the first i votes. We performed a linear regression of \bar{X}_n over n : $\bar{X}_n \sim kn + b$. The slope k reflects the overall trend of \bar{X}_i : if $k > 0$ the votes increase with time, if $k < 0$ they decrease. A histogram of the 14,171 slopes is shown in Fig. 1. A t -test of the null hypothesis “ $k = 0$ ” yields a p -value 0.064, which is not enough to reject the null. This confirms the absence of an overall dynamical trend.

To see the effects of other people’s opinions on the overall public opinion formation we studied [Jyte.com](#), a website that allows its users to make any claim they wish and let the community vote on it at no cost. The claims are wide ranging, from “Ocean exploration has more potential to benefit the human race than space exploration” to “Homeopathy shouldn’t be available on the NHS”. The web interface is simple and intuitive. Each claim is flanked by a positive button and a negative button and the numbers of total positive and negative votes are shown on the face of the buttons. Each user sees the numbers, makes up her mind, and submits her vote by clicking on one of the two buttons. Then the numbers get updated instantly.

¹ Another way to remove the bias is to randomly flip the sign of all votes for each resolve. The conclusion would be the same. Notice that the lack of normalization would lead to an artificially high initial value that would decay over time.




Fig. 1. The histogram of 14,171 slopes calculated from the **Essembly** data. It is not very clear whether most slopes are positive or negative.

We tracked the voting dynamics of all claims made in July 2007, among which 1,208 claims received no less than 10 votes. For each vote, we recorded $X_n = 1$ if the n 'th vote is positive, or $X_n = 0$ if it is negative. The quantity $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$ then represents the fraction of positive votes among the first n votes.

We constructed two data subsets from the 1,208 claims. The first set contains all the claims with less than say, 3 positive votes, among the first 10 (i.e. $\bar{X}_{10} < 0.3$). We call it the “negative” set. The second set contains all the claims with more than say, 7 positive votes, among the first 10 (i.e. $\bar{X}_{10} > 0.7$). We call it the “positive” set. These two sets contain the claims that are generally regarded as “very negative” and “very positive”, respectively. The negative set consists of 405 claims, and the positive set consists of 521 claims. The average \bar{X}_n for the two data sets are shown in Fig. 2(a) and (b). As we can see from the figure, the negative claims tend to become more negative as the voting goes on, and the positive claims tend to become more positive.² This shows that when people observe previous opinions before they express their own, they tend to follow the trend. As a result of this trend following, extreme views get reinforced and become increasingly more extreme.

Our last study focused on the situation where it is costly to express an opinion. We thus considered the online rating data for a large number of books collected from [Amazon.com](#). On [Amazon](#), a user observes the average rating of a book when she visits a book page (usually shown at the top, right under the title). If she decides to review a book, she is required to write a short paragraph of review in addition to a simple star rating. The average word count of [Amazon](#) reviews is 181.5 words [9]. Thus, the cost of opinion expression is high for [Amazon](#) compared with [Essembly](#) and [Jyte](#), and the person will only contribute if the gain from expressing an opinion is higher than the cost.

In the spirit of the voter's paradox, which assumes that the voter is more likely to vote when his vote is more probable to affect the outcome, we speculate

² This is independent of our choice of the values 3 and 7.

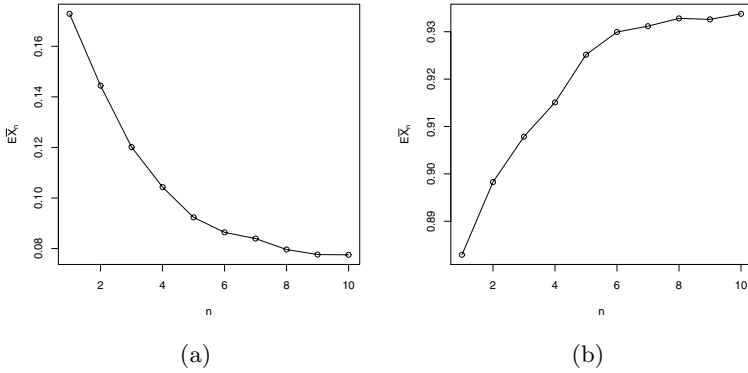


Fig. 2. The sample average fraction of positive votes $E\bar{X}_n$ as a function of the number of votes n , collected from **Jyte.com**. (a) Negative claims become more negative as time goes on. The average fraction decreases by an absolute amount 9% after 10 votes. (b) Positive claims become more positive as time goes on. The average fraction increases by an absolute amount 5% after 10 votes.

that in cases like **Amazon**, people will derive more utility the more they can influence the overall rating. To be precise, in cases where users’ opinions can be quantified and aggregated into an average value, the influence of an online opinion can be measured by how much its expression will change the average opinion [11]. Suppose that n users have expressed their opinions, X_1, \dots, X_n , on a given topic at a website, with X_i denoting the quantified value of the i ’th opinion. If the $(n + 1)$ ’th person expresses a new opinion X_{n+1} , it will move the average rating to

$$\bar{X}_{n+1} = \frac{n\bar{X}_n + X_{n+1}}{n + 1}, \tag{1}$$

and the absolute change in the average rating is given by

$$|\bar{X}_{n+1} - \bar{X}_n| = \frac{|X_{n+1} - \bar{X}_n|}{n + 1}. \tag{2}$$

We thus conjecture that a person is more likely to express her opinion whenever $|X_{n+1} - \bar{X}_n|$ is large — an opinion is likely to be expressed if it deviates by a significant amount from those already stated. Indeed, what is the point of leaving another 5-star review after one hundred people have already done so? This point has also been made within the “brag-and-moan” model [48] which assumes that consumers only choose to write reviews when they are very satisfied with the products they purchased (brag), or very disgruntled (moan). Note however, that the brag-and-moan model is static and thus predicts that \bar{X}_n is constant over time, in contradiction with the observed dynamical trends.

Our sample consisted of the top 4,000 best-selling titles in each of the following 12 categories, as of July 1, 2007: arts & photography, biographies & memoirs, history, literature & fiction, mystery & thrillers, reference, religion & spirituality,




Fig. 3. The average rating of 16,454 books on Amazon.com with more than 20 reviews. EX_n is the sample average rating of all the 16,454 n 'th ratings. As one can see from the figure, EX_n decreases by 0.4 stars in 20 steps. We did not obtain enough data from low selling books to show the opposite trend.

sports, travel, nonfiction, science, and entertainment. For each of the 48,000 books, a series of ratings was collected in time order, where each rating is an integer between 1 and 5 (number of stars). Among the 48,000 books, 16,454 books have no less than 20 ratings, and 11,920 have an average rating above 4.

We first checked the average rating of the 16,454 books as a function of n . As can be seen from Fig. 3, EX_n decreases almost linearly with n , so there is a clear dynamical trend in the ratings, which corroborates the observation reported in [10]. Later users indeed tend to write different reviews from those of earlier users. As opposed to what we observed in Jyote.com, the overall opinion tends to decrease away from the extreme ones [3].

Next we examined whether this dynamical trend is still prominent at the level of each individual book. Similar to our Essembly study, we performed a linear regression of \bar{X}_n over n : $\bar{X}_n \sim kn + b$. The histogram of 16,454 slopes (k) are shown in Fig. 4. As can be seen, most of the slopes are below zero. A t -test of the alternative hypothesis “ $k < 0$ ” yields a p -value less than 2.2×10^{-16} , which further confirms the declining trend.

Finally we measured directly how much one's rating deviates from the observed average rating. We plot the expected deviation

$$Ed_n = E|X_n - \bar{X}_{n-1}|. \tag{3}$$

as a function of n in Fig. 5. As can be seen, Ed_n increases with n . Since the expected deviation Ed_n of an i.i.d. sequence normally decreases with n , this increasing trend is indeed significant. This again supports our conjecture that those users who disagree from the public opinion will be more willing to express themselves and thus soften the overall opinion of a given book.

One point to be stressed is that the results do not imply that as time goes on the average perception of the book changes. Rather, from a large pool of readers

³ In a forthcoming paper we show that conditional expectations of Amazon ratings (given the average previous ones) follow the same trend.




Fig. 4. Histogram of the slopes of average book ratings on Amazon.com. Most of the slopes are negative, testifying a declining trend in the average ratings.




Fig. 5. The average deviation of Amazon ratings increases with the number of people

it is only those that want to make a difference with the prevailing opinion that choose to express themselves. This is seen when plotting the average “helpful ratio” as a function of star rating in Fig. 6 for users of Amazon. It can be seen that people find high ratings more helpful than low ratings, implying that the majority does not agree with this expression bias.

These results, made possible by the fact that the web presents a natural laboratory to study millions of opinions, show that in the process of expressing their views, people tend to follow different but regular patterns. When no information of previous views is available, the opinions expressed are drawn from a uniform distribution within the community. In cases where previous opinions are made known and it is painless to post a view, one observes either neutral opinions or a polarized consensus which reflect trend following by the group. In the latter case, opinions tend to reinforce previous ones and thus become more extreme. Finally there are many cases where expressing a view is costly, like when writing a book review. In this case people will tend to do so whenever they perceive they can offset the current view by presenting a differing one. Since the impact decreases with the number of posted opinions, the larger the pool, the more

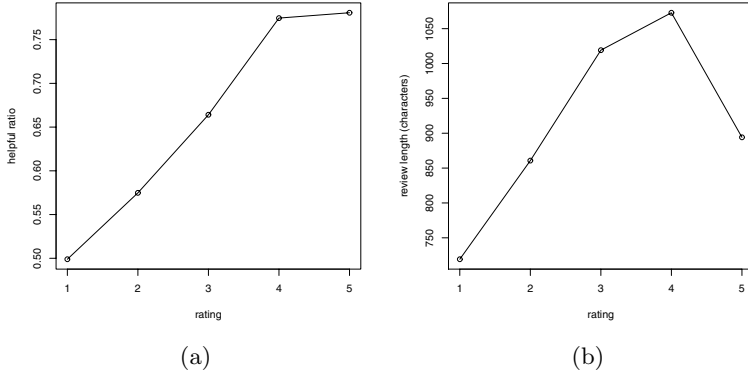


Fig. 6. (a) The average helpful ratio of five different star ratings. (b) The average review length of five different star ratings in the number of characters. The data is calculated for 4,000 bestselling mystery books. By comparing the two figures it is clear that people find high ratings more helpful not just because they are longer. For instance, 5-star reviews are on average shorter than 4-star and 3-star reviews but are nevertheless more helpful.

extreme the difference expressed. As a consequence one sees a softening of the prevailing view.

Besides explaining the observed data, these results show a cautionary note on the interpretation of public opinion. This is because a simple change in the order or frequency of given sets of views can change the ongoing expression in the community, and thus the perceived collective wisdom that new users will find when accessing that information.

References

1. Cosley, D., Lam, S.K., Albert, I., Konstan, J.A., Riedl, J.: Is seeing believing? How recommender interfaces affect users opinions. In: CHI 2003, Ft. Lauderdale, Florida, USA (April 2003)
2. Dellarocas, C.: Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In: ACM EC 2000, Minneapolis, Minnesota, USA (October 2000)
3. Dellarocas, C.: The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science* 49(10), 1407–1424 (2003)
4. Dellarocas, C., Narayan, R.: What motivates consumers to review a product online? A study of the product-specific antecedents of online movie reviews. In: Aberer, K., Peng, Z., Rundensteiner, E.A., Zhang, Y., Li, X. (eds.) WISE 2006. LNCS, vol. 4255. Springer, Heidelberg (2006)
5. Downs, A.: *An economic theory of democracy*. Harper & Row, New York (1957)
6. Gao, G., Gu, B., Lin, M.: The dynamics of online consumer reviews. In: Aberer, K., Peng, Z., Rundensteiner, E.A., Zhang, Y., Li, X. (eds.) WISE 2006. LNCS, vol. 4255. Springer, Heidelberg (2006)

7. Geys, B.: 'Rational' theories of voter turnout: A review. *Political Studies Review* 4, 16–35 (2006)
8. Hu, N., Pavlou, P.A., Zhang, J.: Can online reviews reveal a product's true quality? Empirical findings and analytical modeling of online word-of-mouth communication. In: *ACM EC 2006*, Ann Arbor, Michigan, USA (June 2006)
9. Ketzan, E.: Amazon.com and the new democracy of opinion (November 2002), http://www.themodernword.com/pynchon/papers_ketzan1.html
10. Li, X., Hitt, L.M.: Self selection and information role of online product reviews. In: Zhou, X., Su, S., Papazoglou, M.P., Orłowska, M.E., Jeffery, K. (eds.) *WISE 2004*. LNCS, vol. 3306. Springer, Heidelberg (2004)
11. Osborne, M.J., Rosenthal, J.S., Turner, M.A.: Meetings with costly participation. *American Economic Review* 90(4), 927–943 (2000)
12. Riker, W.H., Ordeshook, P.C.: A theory of the calculus of voting. *American Political Science Review* 62, 25–42 (1968)
13. Schuessler, A.A.: Expressive voting. *Rationality and Society* 12(1), 87–119 (2000)

A Game-Theoretic Analysis of Games with a Purpose

Shaili Jain and David C. Parkes

School of Engineering and Applied Science
Harvard University
{shailij,parkes}@eecs.harvard.edu

Abstract. We present a simple game-theoretic model for the ESP game, an interactive game devised to label images on the web, and characterize the equilibrium behavior of the model. We show that a simple change in the incentive structure can lead to different equilibrium structure and suggest the possibility of formal incentive design in achieving desirable system-wide outcomes, complementing existing considerations of robustness against cheating and human factors.

1 Introduction

Showcased by the early success of “Games with a Purpose” [3], human computation considers the possibility that networks of people can be leveraged to solve large-scale problems that are hard for computers. Work by von Ahn and others has shown the tremendous power that networks of humans possess to solve problems while playing computer games [4,7,5,6]. The ESP game is an example of such human computation; it is an interactive system that allows users to be paired to play games that label images on the web [4]. Users play the ESP game because it is an enjoyable game to play, with the added side-effect that they are doing useful work in the process. Subsequent work to the ESP game has included Peekaboom [7], Phetch [5], and Verbosity [6]. Hsu and colleagues [2,1] developed a simple game called PhotoSlap for determining content of images and provided a game-theoretic analysis for PhotoSlap.

While there has been incredible progress in the area of human computation, there is still much more potential. For “Games with a Purpose”, it seems especially appropriate to use game theory to better understand how to design incentives in order to achieve system-wide goals. For example, it appears anecdotally that during play of the ESP game that people coordinate on easy words and that the game is less effective in labeling less obvious, harder words. Through this line of work, we ultimately aim to show that proper incentive design along with appropriate system design is an important paradigm for human computation and peer production problems.

This paper aims to study behavior in the ESP Game through a game-theoretic light. We propose a simple model of the game and consider two different models of payoffs, namely *match-early* preferences (MEP) and *rare-words-first* preferences.

Match-early preferences model the setting in which players wish to complete as many rounds as possible and receive the same score irrespective of the words on which they match. The match-early preferences model is meant to reflect the current method of assigning scores to outcomes in the ESP game. Here we show that *low effort* is a Bayesian-Nash equilibrium for all distributions on word frequencies, with players focusing attention on high-frequency words. Rare-words-first preferences model the setting in which players wish to match on infrequent words before frequent words, we suppose because of appropriately designed incentives, and the speed with which a match is achieved is only a secondary consideration. We show that under this preference model, there is a significant difference in the equilibrium structure.

We briefly describe elements of the ESP game before introducing the model in the next section. In the ESP game, players are randomly paired with another player in the system for a set of 15 images. Players try to label as many images of the 15 as they can in the allotted 2.5 minutes. Players receive a fixed number of points after agreeing on a common word. In the set of 15 images, players get bonus points after agreeing on five images, ten images, and fifteen images in the same set. The only words that are used from the input streams are the first agreed word. An interesting feature of the ESP game, not modeled here, is the use of *Taboo words* [4]. Taboo words are words that are displayed next to the image that players cannot enter for the corresponding image. Taboo words are words that have been entered sufficiently many times in previous plays of the image and encourage players to enter different words so that the set of labels for an image can be extended [4]. Modeling the effect of the Taboo Words is an important direction for future work.

2 An ESP Model with Match-Early Preferences

We model the ESP game as a two-player, two-stage game of imperfect information. We focus on modeling one of the 15 rounds, and thus the game associated with a specific image. We model the ESP game with each player sampling words from a universe of possible words associated with the picture, to which we associate a frequency ordering. Players can vary the effort level that relates to how likely they are to sample frequent words as opposed to infrequent words. Then players decide which order to play their sampled words in the game. In the model of *match-early* preferences, we instead capture the strategic behavior of having 15 rounds under a time constraint by providing a preference for matching in an earlier location than a later location.

Let $d > 0$ denote the *dictionary size*, representing the number of words that each player will think of for the image at hand. We model a universe of words $U = \{w_1, w_2, \dots, w_n\}$, that represents all possible words to describe the image and the knowledge that the game designer is trying to learn. Each word has an

¹ In addition, von Ahn and Dabbish came up with a number of methods to circumvent cheating, where possible methods of cheating include players trying to be paired with themselves and global strategies such as entering “a” for every image [4].




Fig. 1. The game tree above represents the decision space of one player

associated frequency, where f_i denotes the frequency of word w_i . We assume that a player can rank the words sampled by frequency. The frequencies satisfy the property that $\sum_{i=1}^n f_i = 1$. We assume that the words in the universe are ordered according to decreasing frequency, that is $f_1 \geq f_2 \geq \dots \geq f_n$. The frequency of word i can be considered the frequency that the word would be mentioned if a population of humans were each asked to state d words related to the image. We assume that $1 < d < |U|$.

Though this game has no communication between players and thus is properly analyzed as a normal-form game, it is useful to talk about a first stage (choosing an effort level) and a second stage (choosing a permutation on the dictionary). In the first stage of the game, players privately choose an *effort level*: $E = \{L, M, H\}$ for *low*, *medium* or *high*. The choice of effort level determines the set of words in the universe from which a player samples her dictionary. If a player chooses L in the first stage of the game, the dictionary is sampled from the top $n_L > 0$ words (without replacement). We say that a player that chooses effort level L has universe U_L , where U_L is exactly the set of the highest n_L frequency words in U . If a player chooses effort M in the first stage of the game, the dictionary is sampled from the top $n_M > n_L > 0$ words. If a player chooses effort H , the dictionary is sampled from the top $n_H = n$ words (i.e., the entire universe). That is, word i in U is chosen with probability $f_{i,H} = f_i$.

Given word $x \in U$, $f_e(x)$ represents the frequency of word x given the player has chosen effort level e . In Figure 1, this sample is modeled as a move by nature and can be considered to be the point at which a player learns her “type”, namely her dictionary of words. Figure 1 represents the choices of a single player in the game, though both players are symmetric. Note that n_L, n_M, n_H , and d are parameters of the model and there is no cost associated with each of the first level actions. We establish that low effort is an equilibrium under match-early preferences even without introducing a cost, which would increase with effort and presumably increase the benefits of low effort. We leave introducing cost into this model for future work.

In the second stage, once each player privately learns her dictionary based on the effort level chosen, players choose a *permutation* on the words. This models the decision about the order in which a player should enter words. This order on a

player’s dictionary defines the second-stage action of each player and determines the *outcome* of the game. The outcome is defined by the first word that is in the ordered-list of both players and the location (where the location is defined to be the maximum value of the two positions where the word occurs in each ordered-list). In what follows, we refer to D_1 as the dictionary for player 1 and D_2 as the dictionary for player 2. The second stage strategy $s_1 \in S_1$ for player 1 defines a specific order $s_1(D_1)$ on the words in D_1 , for every possible dictionary. Likewise, player 2 has a second-stage strategy $s_2 \in S_2$ that defines an order $s_2(D_2)$ for every dictionary. We restrict our attention to strategies that involve playing all words in the dictionary since any strategy that does not involve playing all words is weakly dominated by one that involves playing all words. We also restrict our attention to *consistent* strategies, strategies for a player that do not change the relative ordering of elements depending on the player’s realized dictionary. In other words, a consistent strategy involves specifying a total ordering of elements on U and applying that total ordering to the realized dictionary.

A *complete strategy* for the ESP game is an ordered pair $\sigma_i = (e_i, s_i) \in E \times S_i = \Sigma_i$. This defines the play in both stages for all possible dictionaries. We focus here on pure strategies.

Definition 1. *Suppose player 1 outputs a list of words x_1, x_2, \dots, x_d and player 2 outputs a list of words y_1, y_2, \dots, y_d . If there exists $1 \leq i, j \leq d$ such that $x_i = y_j$, there was a match and this match occurred in location $\max(i, j)$. It is possible for two sequences to have more than one match, so we concern ourselves with the first match, that is the pair i, j that minimizes $\max(i, j)$ such that $x_i = y_j$.*

An *outcome* is an ordered pair $o = (w, l) \in (U \cup \phi) \times (\{1, \dots, d\} \cup \phi)$ where (ϕ, ϕ) indicates there was no match and the (w, l) pair otherwise indicates that the *first* match on word $w \in U$ in location $l \in \mathcal{L}$ where $\mathcal{L} = \{1, 2, \dots, d\} \cup \phi$. Let \mathcal{O} denote the set of possible outcomes. Since $s_1(D_1)$ and $s_2(D_2)$ specify orderings on given dictionaries, they induce an outcome: the location of the first match.

Let *outcome function* $g(s_1(D_1), s_2(D_2)) \in \mathcal{O}$ denote this outcome. The location (if any) of the first match is denoted by $g_l(s_1(D_1), s_2(D_2)) \in \mathcal{L}$.

Each player i has *valuation* $v_i(o)$ on outcome o , which induces a (weak) total preference ordering on outcomes. For match-early preferences, we require $(w_1, l_1) \equiv (w_2, l_1) \equiv \dots \equiv (w_n, l_1) \succ (w_1, l_2) \equiv (w_2, l_2) \equiv \dots \equiv (w_n, l_2) \succ \dots \succ (w_1, l_d) \equiv (w_2, l_d) \equiv \dots \equiv (w_n, l_d) \succ (\phi, \phi)$ for all players.

Let $\Pr(D_i|e_i)$ denote the probability of dictionary D_i given effort level e_i . Often times we write this as $\Pr(D_i)$ and leave the effort level implicit.

Definition 2. *The probability of first match in l_i given $s_1(D_1), s_2$, and distribution $\Pr(D_2)$, is $p(l_i, s_1(D_1), s_2) = \sum_{D_2} \Pr(D_2)I(g_l(s_1(D_1), s_2(D_2)) = l_i)$. Similarly, the probability of first match in l_i on w_j is $p(w_j, l_i, s_1(D_1), s_2) = \sum_{D_2} \Pr(D_2)I(g(s_1(D_1), s_2(D_2)) = (w_j, l_i))$. Often times we will abbreviate $p(l_i, s_1(D_1), s_2)$ as $p(l_i)$ and $p(w_j, l_i, s_1(D_1), s_2)$ as $p(w_j, l_i)$.*

Let $u_i(s_i(D_i), s_{2-i}(D_{2-i})) = v_i(g(s_1(D_1), s_2(D_2)))$ be the utility of player i given D_1, D_2 . Let $u_i(s_i(D_i), s_{2-i}) = \sum_{D_{2-i}} \Pr(D_{2-i})u_i(s_i(D_i), s_{2-i}(D_{2-i}))$ be the

expected (interim) utility of player i given D_i with respect to the distribution on all possible dictionaries of the other player, as induced by her effort level. Let $u_i(\sigma_i, \sigma_{2-i}) = \sum_{D_1} \sum_{D_2} \Pr(D_1|e_1) \Pr(D_2|e_2) u_i(s_i(D_i), s_{2-i}(D_{2-i}))$ be the expected (ex ante) utility of player i before dictionaries are sampled, given complete strategies $\sigma = (\sigma_1, \sigma_2)$.

We analyze the second stage of the game before analyzing the complete game. For this, consider the game induced by fixing top level effort levels (e_1, e_2) for the two players (the *second stage game conditioned on effort e_1 and e_2*). In the second stage, each player knows her own dictionary but not the dictionary of the other player. We can now define two useful equilibrium concepts:

Definition 3. *Strategy profile $s^* = (s_1^*, s_2^*)$ is an ex post Nash equilibrium of the second stage of the ESP game conditioned on effort levels e_1 and e_2 , if for every D_1 and every D_2 , we have:*

$$u_i(s_i^*(D_i), s_{2-i}^*(D_{2-i})) \geq u_i(s'_i(D_i), s_{2-i}^*(D_{2-i})), \quad \forall s'_i \neq s_i^*, \quad \forall i \in \{1, 2\} \quad (1)$$

Definition 4. *Strategy profile $s^* = (s_1^*, s_2^*)$ is a strict Bayesian-Nash equilibrium of the second-stage of the ESP game conditioned on effort levels e_1 and e_2 if for both players $i \in \{1, 2\}$, for every D_i ,*

$$u_i(s_i^*(D_i), s_{2-i}^*(D_{2-i})) > u_i(s'_i(D_i), s_{2-i}^*(D_{2-i})), \quad (2)$$

where the probability adopted in interim utility u_i for the distribution on the dictionary of player $2 - i$ is induced by the effort of that player in the first stage.

Definition 5. *Strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*) \in \Sigma_1 \times \Sigma_2$ is a strict Bayesian-Nash equilibrium of the ESP game if for players $i \in \{1, 2\}$, $u_i(\sigma_i^*, \sigma_{2-i}^*) > u_i(\sigma'_i, \sigma_{2-i}^*)$.*

Since the effort level chosen by each player is not visible to the other player, there is no need for a subgame perfect refinement.

3 Effort Level of Players under Match-Early Preferences

In this section, we analyze the equilibrium behavior under match-early preferences. We show that playing decreasing frequency in conjunction with low effort is a Bayesian-Nash equilibrium for the ESP game. First we see that playing playing words in order of decreasing frequency is not an ex-post Nash equilibrium for the second stage of the game.

Lemma 1. *Suppose that players are playing the same effort level and there are three words in the universe, w_1, w_2 , and w_3 with associated probabilities $f_{1,e}, f_{2,e}, f_{3,e}$, and with $f_{1,e} > f_{2,e} > f_{3,e}$. The second stage strategy profile $s = (s_1, s_2)$, where s_1 and s_2 are the strategies of playing words in order of decreasing frequency, is not an ex-post Nash equilibrium.*

Since playing words in order of decreasing frequency is not an ex-post Nash equilibrium, we focus instead on Bayesian-Nash equilibrium. The following

definition of *stochastic dominance* will enable equilibrium analysis for any valuation function that satisfies MEP. In Lemmas 2 and 3, we show that this notion of stochastic dominance is both sufficient and necessary for utility maximization.

Definition 6. Fixing effort levels e_1 and e_2 and fixing opponent's second-stage strategy s_2 , we say second-stage strategy s_1 with match vector $(p(l_1, s_1(D_1), s_2), p(l_2, s_1(D_1), s_2), \dots, p(l_d, s_1(D_1), s_2))$ stochastically dominates second-stage strategy s'_1 with match vector $(p(l_1, s'_1(D_1), s_2), p(l_2, s'_1(D_1), s_2), \dots, p(l_d, s'_1(D_1), s_2))$ with respect to dictionary D_1 if for every $1 \leq k \leq d$, $\sum_{a=1}^k p(l_a, s_1(D_1), s_2) \geq \sum_{a=1}^k p(l_a, s'_1(D_1), s_2)$. We say that the stochastic dominance property is strict if there exists a k such that $1 \leq k \leq d$ and $\sum_{a=1}^k p(l_a, s_1(D_1), s_2) > \sum_{a=1}^k p(l_a, s'_1(D_1), s_2)$.

Lemma 2. If strategy s_1 stochastically dominates strategy s'_1 with respect to dictionary D_1 , for fixed opponent strategy s_2 , then $u_1(s_1(D_1), s_2) \geq u_1(s'_1(D_1), s_2)$, for all valuations consistent with match-early preferences.

Lemma 3. If $u_1(s_1(D_1), s_2) \geq u_1(s'_1(D_1), s_2)$ for all valuations that are consistent with match-early preferences, then strategy s_1 must stochastically dominate strategy s'_1 with respect to D_1 , for fixed opponent strategy s_2 .

Lemmas 2 and 3 can be extended to show strict stochastic dominance implies strictly greater utility and vice versa, for all valuations consistent with MEP.

Lemma 4. Given player 2 plays her words in order of decreasing frequency, the probability of first match in location l_1 is strictly maximized when player 1 plays her most frequent word first, for all dictionaries D_1 and effort levels e_2 .

Lemma 5. For $1 \leq k < d$, given player 1 played her k highest frequency words first and player 2 plays her words in order of decreasing frequency, the probability of first match in locations l_1, \dots, l_k is strictly maximized when player 1 plays her $k + 1^{st}$ highest frequency word next, for all dictionaries D_1 and effort levels e_2 .

Lemmas 2, 4, and 5 establish that playing decreasing frequency is a strict best response to an opponent who plays decreasing frequency.

Theorem 1. The strategy profile consisting of players playing words in order of decreasing frequency (denoted (\downarrow, \downarrow)) is a strict Bayesian-Nash equilibrium of the second stage of the ESP game, conditioned on any choice of effort levels e_1 and e_2 , for any distribution over U and any valuation function that satisfies match-early preferences.

To show that playing L at the top-level along with playing decreasing frequency is a Bayesian-Nash equilibrium, we use the following definition of stochastic dominance for the top level of the game which fixes the equilibrium strategy for the bottom-level. The definition uses the following notation for a k -truncation of dictionary D : $D(k)$ is the set of k highest frequency words in D .

Definition 7. Fixing player 2’s complete strategy (e_2, s_2) , a complete strategy (e_1, s_1) stochastically dominates complete strategy (e'_1, s_1) for player 1 if:

$$\sum_{D_{1,e_1}} \Pr(D_{1,e_1} | e_1) \sum_{D_{2,e_2}} \Pr(D_{2,e_2} | e_2) \cdot I(g_l(s_1(D_{1,e_1}(k)), s_2(D_{2,e_2}(k))) = l_1, \dots, l_k) \geq \sum_{D_{1,e'_1}} \Pr(D_{1,e'_1} | e'_1) \sum_{D_{2,e_2}} \Pr(D_{2,e_2} | e_2) \cdot I(g_l(s_1(D_{1,e'_1}(k)), s_2(D_{2,e_2}(k))) = l_1, \dots, l_k) \quad \forall k$$

where $g_l(s_1(D_{1,e_1}(k)), s_2(D_{2,e_2}(k)))$ gives the outcome when second-stage strategies s_1 and s_2 act on $D_{1,e_1}(k)$ and $D_{2,e_2}(k)$ and $I(\cdot)$ is the indicator function. The dominance is strict if there exists a k such that the above inequality is strict.

Since Theorem 1 establishes that (\downarrow, \downarrow) is a strict Bayesian-Nash equilibrium of the second stage, for all effort levels, we set $(s_1, s_2) = (\downarrow, \downarrow)$ and we know that $I(g_l(s_1(D_{1,e_1}(k)), s_2(D_{2,e_2}(k))) = l_1, \dots, l_k) = I(D_{1,e_1}(k) \cap D_{2,e_2}(k) \neq \emptyset)$. Similar to Lemmas 2 and 3, we can show that stochastic dominance in Definition 7 is sufficient and necessary for utility maximization.

In order to establish stochastic dominance, we construct a randomized mapping for each dictionary that can be sampled when playing M to a number of dictionaries that can be sampled when playing L . Each dictionary in \mathcal{D}_M is mapped to a dictionary in \mathcal{D}_L that is at least as likely to match against the opponent’s dictionary, averaged over the distribution of all possible dictionaries for the opponent. This is shown in Lemma 6. In order to complete the proof, it is necessary to show that under the randomized mapping, no element in \mathcal{D}_L is mapped to with greater probability under the randomized mapping than under the original distribution over \mathcal{D}_L . This fact is shown in Lemma 7.

The randomized mapping h can be described as follows: Consider a dictionary $D \in \mathcal{D}_M$, $D = A \cup B$, where A is the set of “low words” and B is the set of “medium words” (in other words, $A = D \cap U_L$ and $B = D \cap (U_M - U_L)$). Under our randomized mapping, D is mapped to all dictionaries in $D_L \in \mathcal{D}_L$ such that $A \subset D_L$. In other words, D is mapped to dictionary $D_L \in \mathcal{D}_L$ with non-zero probability if and only if $A \subset D_L$. If $A \subset D_L$, then D is mapped to D_L with the same probability that you could would get D_L if you continued to sample individual words from U_M (without replacement) until you got d “low words”. Note that if D contains only medium words, D is mapped to all dictionaries in \mathcal{D}_L with non-zero probability. Likewise, if D contains only low words, D is mapped to only one dictionary in \mathcal{D}_L .

Lemma 6. For any $D_{1,M}$, where $D_{1,M}$ is a dictionary sampled with respect to the M effort level, for any h that satisfies the property that $D_{1,M}$ is mapped to a dictionary in \mathcal{D}_L that contains the set $D_{1,M} \cap U_L$, any effort level of player 2 and when players play decreasing frequency in the second stage, we have that:

$$\sum_{D_2} \Pr(D_2) \cdot I(h(D_{1,M})(k) \cap D_2(k) \neq \emptyset) \geq \sum_{D_2} \Pr(D_2) \cdot I(D_{1,M}(k) \cap D_2(k) \neq \emptyset) \quad \forall k \text{ and } D_{1,M} \tag{3}$$

In addition, the inequality is strict for all $k > k'$ when $h(D_{1,M}) \neq D_{1,M}$ and k' is the first coordinate where $h(D_{1,M})$ and $D_{1,M}$ differ.

Lemma 7 states the distribution obtained from sampling U_L directly is the same as the distribution obtained from sampling a medium dictionary, followed by the randomized mapping (sampling U_M until you get d low words).

Lemma 7. $\Pr(D_{1,L}|L) = \sum_{D_{1,M}} \Pr(D_{1,M}|M) \cdot \Pr(h(D_{1,M}) = D_{1,L})$

Lemma 8 uses Lemmas 6 and 7 to show that playing L stochastically dominates playing M , assuming players play decreasing frequency in the second stage. An identical argument can be used to show that playing L stochastically dominates playing H , assuming players play decreasing frequency in the second stage.

Lemma 8. For any effort level e_2 and when players play decreasing frequency in the second stage:

$$\sum_{D_{1,L}} \Pr(D_{1,L}|L) \sum_{D_2} \Pr(D_2) \cdot I(D_{1,L}(k) \cap D_2(k) \neq \emptyset) > \sum_{D_{1,M}} \Pr(D_{1,M}|M) \sum_{D_2} \Pr(D_2) \cdot I(D_{1,M}(k) \cap D_2(k) \neq \emptyset) \quad \forall k \tag{4}$$

Theorem 1 together with Lemma 8 gives us the following result.

Theorem 2. $((L, \downarrow), (L, \downarrow))$ is a strict Bayesian-Nash equilibrium for the complete game. Additionally, (L, \downarrow) is a strict best-response to both (M, \downarrow) and (H, \downarrow) .

4 The Effect of Rare-Words First Preferences

In this section, we consider the effect of modified preferences. We introduce a new model called *rare-words first* preferences and show some initial results regarding how equilibrium behavior is different under this new model.

Definition 8. Under rare-words first preferences, players prefer to match on rare words, with location as a secondary consideration. Any valuation function $v(o)$ that satisfies rare-words first preferences satisfies the following total ordering on outcomes: $(w_n, l_1) \succ (w_n, l_2) \succ \dots \succ (w_1, l_{d-1}) \succ (w_1, l_d) \succ (\phi, \phi)$.

This preference relation allows for a virtually identical definition of stochastic dominance as Definition 6 which in turn leads to results analogous to Lemmas 2 and 3, namely that stochastic dominance is both sufficient and necessary for utility maximization.

The following lemma is in stark contrast with the results in section 4, where we showed that (\downarrow, \downarrow) is a strict Bayesian-Nash equilibrium in the second stage, for all distributions over U , all valuation functions that satisfy MEP, and any pair of effort levels. Lemma 9 shows that we cannot say (\downarrow, \downarrow) is a Bayesian-Nash equilibrium for the second stage of the game for any distribution, without making more assumptions on the valuation function.

Lemma 9. *Consider any distribution over $U = \{w_1, w_2, \dots, w_n\}$ and suppose that player 2 is playing her words in order of decreasing frequency. For any dictionary of player 1, no consistent strategy of player 1 can stochastically dominate all other consistent strategies.*

Similarly, Lemma 10 shows that when a player is playing increasing frequency, we need to make more assumptions on the valuation function to discern the best-response in the space of consistent strategies.

Lemma 10. *Consider any distribution over $U = \{w_1, w_2, \dots, w_n\}$ and suppose that player 2 is playing her words in order of increasing frequency. For any dictionary of player 1, no consistent strategy of player 1 can stochastically dominate all other consistent strategies.*

We leave it as future work to characterize the set of sufficient and necessary conditions for which playing words in order of *increasing* frequency in conjunction with *high* effort for both players is a Bayesian-Nash equilibrium. Understanding the incentive structure that leads to high effort is important since it is one way the system designer can extend the set of labels for an image. In addition to this, future work should identify specific score functions that provide desirable equilibrium, and also move to analyze the larger context of a *system* of bilateral games with a view on understanding methods to induce large-scale desirable behavior, including such aspects as formalizing the role of Taboo Words and leveraging the entire sequence of words suggested by a player rather than just the particular match.

References

1. Chang, T.-H., Ho, C.-J., Jen Hsu, J.Y.: The photoslap game: Play to annotate. In: Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI), pp. 1966–1967 (2007)
2. Ho, C.-J., Chang, T.-H., Jen Hsu, J.Y.: Photoslap: A multi-player online game for semantic annotation. In: Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI), pp. 1359–1364 (2007)
3. von Ahn, L.: Games with a purpose. *IEEE Computer* 39(6), 92–94 (2006)
4. von Ahn, L., Dabbish, L.: Labeling images with a computer game. In: Proceedings of the 2004 Conference on Human Factors in Computing Systems (CHI), pp. 319–326 (2004)
5. von Ahn, L., Ginosar, S., Kedia, M., Liu, R., Blum, M.: Improving accessibility of the web with a computer game. In: Proceedings of the 2006 Conference on Human Factors in Computing Systems (CHI), pp. 79–82 (2006)
6. von Ahn, L., Kedia, M., Blum, M.: Verbosity: a game for collecting common-sense facts. In: Proceedings of the 2006 Conference on Human Factors in Computing Systems (CHI), pp. 75–78 (2006)
7. von Ahn, L., Liu, R., Blum, M.: Peekaboom: a game for locating objects in images. In: Proceedings of the 2006 Conference on Human Factors in Computing Systems (CHI), pp. 55–64 (2006)

Inapproximability of Combinatorial Public Projects

Michael Schapira¹ and Yaron Singer²

¹ The School of Computer Science and Engineering,
The Hebrew University of Jerusalem, Israel

`mikesch@cs.huji.ac.il`

² Computer Science Division, UC Berkeley, USA

`yaron@cs.berkeley.edu`

Abstract. We study the Combinatorial Public Project Problem (CPPP) in which n agents are assigned a subset of m resources of size k so as to maximize the social welfare. Combinatorial public projects are an abstraction of many resource-assignment problems (Internet-related network design, elections, etc.). It is known that if all agents have submodular valuations then a constant approximation is achievable in polynomial time. However, submodularity is a strong assumption that does not always hold in practice. We show that (unlike similar problems such as combinatorial auctions) even slight relaxations of the submodularity assumption result in non-constant lower bounds for approximation.

1 Introduction

There are various real-world settings in which a set of resources is chosen to collectively serve an entire community: In elections, for instance, candidates are chosen to serve a community of voters. States choose which roads to build for the benefit of their residents. Another interesting example is that of choosing overlay networks in the Internet [3,2] (for instance, in the context of inter-domain routing): A node in the network chooses the optimal subset of nodes to route traffic through, in a manner which is most beneficial to its clients [1]. Such an overlay - the subset of chosen nodes through which traffic is routed - is beneficial to different clients in different degrees. This creates the difficulty of maximizing the social welfare in this setting (as in the settings above).

The Combinatorial Public Project Problem (CPPP), recently presented and studied in [18], is an abstraction of such settings. In CPPP there are m resources and n agents, each with a valuation function defined over all subsets of resources, and the objective is to choose the k resources which maximize the social welfare of the agents. It is easy to see that CPPP is NP hard. However, in the case where each valuation function is submodular CPPP allows for good approximations [15,18] (within $1 - \frac{1}{e}$) [2]. A valuation function v is said to be *submodular* if for every two subsets of resources S, T we have $v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$.

¹ A formal description of the overlay networks setting is presented in Section 2.

² This bound is tight, see [8].

In this paper we study CPPP, focusing on computational and communication complexity lower bounds which result from relaxations of the above submodularity property. Indeed, out of the myriads of problems which can be depicted as combinatorial public projects, in many instances the submodularity constraint does not apply. We illustrate this point in Section 2, where we show how the difference between submodular and general valuations separates between two seemingly close Internet-related network design problems (overlay networks).

Our main result is showing that even small relaxations of the submodularity property in combinatorial public projects result in strong inapproximability bounds:

Theorem: Obtaining an approximation ratio of $m^{\frac{1}{4}-\epsilon}$ for CPPP with subadditive valuation functions requires exponential communication in m (for every constant $\epsilon > 0$).

A valuation function v is subadditive if for every two subsets of resources S, T we have that $v(S \cup T) \leq v(S) + v(T)$. In fact, we prove our lower bound for a more restricted class of valuations called fractionally-subadditive [8] (introduced in [17] and termed “XOS” there). This lower bound establishes that in terms of approximability there is a huge gap between the submodular and subadditive case. This is in stark contrast to what is known about the related combinatorial auctions problem, for which a 2-approximation is achievable for subadditive valuations [8]. We show that our lower bound is nearly tight by presenting a simple $O(m^{\frac{1}{2}})$ approximation algorithm for subadditive valuations. We leave the question of closing the gap between $m^{\frac{1}{4}}$ and $m^{\frac{1}{2}}$ open. Also, a big open question is showing that a similar lower bound is achievable in the computational-complexity model.

In [18] it is shown, for the case of CPPP with submodular valuations, that while a constant approximation ratio is possible, no such ratio is achievable via *truthful* algorithms. The $O(m^{\frac{1}{2}})$ approximation algorithm for subadditive valuations presented in this paper is truthful. Hence, our results show that the hardness of CPPP with subadditive valuations is due to computational reasons and not to the truthfulness requirement. That is, truthful approximation algorithms perform just, or nearly, as well as unrestricted algorithms.

Finally, we study the approximability of CPPP with general valuations. We prove two complementary lower bounds, one in the computational-complexity model and one in the communication-complexity model:

Theorem: Obtaining an approximation of $O(n^{\frac{1}{2}-\epsilon})$ in CPPP with general valuations, for any $\epsilon > 0$, is impossible unless $P = NP$. Obtaining an approximation of $O(n^{1-\epsilon})$ is impossible unless $P = ZPP$.

Theorem: Obtaining an approximation of $(1 - \epsilon)n$ for CPPP with general valuations requires exponential communication in m (for any $\epsilon > 0$ and for any $n \ll 2^m$).

2 Model and Motivation

In this section we formally present CPPP and our model. We motivate CPPP by presenting an Internet-related network design setting which is naturally formulated as a combinatorial public project. We use this example to illustrate the importance of considering relaxations of submodularity.

2.1 The Model

In CPPP there is a set of n agents $N = \{1, \dots, n\}$, a set of m resources $M = \{1, \dots, m\}$, and a parameter $1 \leq k \leq m$. Each agent has a valuation function (sometimes simply referred to as a *valuation*) $v_i : 2^{[m]} \rightarrow \mathcal{R}_{\geq 0}$. We make two standard assumptions regarding each valuation function v_i : That $v_i(\emptyset) = 0$ (normalized) and that for all $S \subseteq T \subseteq [m]$ it holds that $v_i(S) \leq v_i(T)$ (non-decreasing). The objective in CPPP is to find a subset of resources S^* of size k which maximizes the social-welfare. That is, we wish to find $S^* \in \operatorname{argmax}_{S \subseteq [m], |S|=k} \sum_i v_i(S)$.

We require algorithms for CPPP to run in time that is polynomial in the natural parameters of the problem – n and m . In some cases agents' valuations can be concisely represented (encoded in space polynomial in m and n). In these cases we wish to explore the *computational complexity* of the problem.

However, in general, the size of the “input” (the valuations) can be exponential in m , and so we must specify how it can be accessed. We take a “black box” (concrete complexity) approach (see [5]): Every valuation function is assumed to be represented by an oracle which can answer a certain type of queries, and we restrict our algorithms to ask polynomially (in both m and n) many such queries. We consider two standard types of queries:

- **Value queries:** A value query to a valuation v_i is in the form of a subset of resources $S \subseteq [m]$. The answer is simply $v_i(S)$. This is a natural type of query to use when designing algorithms for CPPP.
- **General queries:** A general query allows any type of questions (even computationally intractable ones), as long as each question is addressed to a single v_i . We only require that the “size” of the query not be too “large”. This is equivalent to Yao’s communication model [20] in which the different parties are computationally omnipotent and we measure the number of bits they must exchange to compute a given function (see an introduction to communication complexity in [12]). The immense strength of general queries is useful for showing impossibility results.

2.2 Overlay Networks

Consider the following Internet-inspired setting [3,2] we call the *overlay network problem*: We are given a network graph $G = (V, E)$, where $V = \text{SUM} \cup D$ (S, M, D are disjoint). We shall refer to the nodes in S as *source-nodes*, to the nodes in M as *potential overlay nodes*, and to the nodes in D as *destination-nodes*. Assume that

there is some metric function m that assigns a non-negative real value m_e to every $e \in E$ (assume that m 's value for a non-existent edge is ∞).

Each node $i \in S$ has a valuation function v_i that assigns a non-negative real value to every subset of potential overlays $M' \subseteq M$. Intuitively, for every $M' \subseteq M$, $v_i(M')$ specifies i 's desire to route through M' to the nodes in D . Each v_i is consistent with the metric m in the following sense: For every $M', M'' \subseteq M$, $v_i(M') \geq v_i(M'')$ iff the sum of the lengths of the shortest routes from i to all nodes in D that only go through nodes in M' is at most that obtained from routing only through nodes in M'' ³

The goal is to choose an overlay network from the potential overlay nodes. For a given parameter $1 \leq k \leq |M|$, we want to find a subset of M of size k that maximizes the social welfare of the source nodes. That is, we wish to find a set $M^* = \operatorname{argmax}_{M' \subseteq M, |M'|=k} \sum_{i \in S} v_i(M')$. It is easy to see that the overlay network problem is a special case of CPPP in which the agents are the source nodes, the resources are the potential overlays and the agents' valuations are induced by the network graph and the metric function ⁴. We distinguish between two versions of this problem, which illustrate the big differences between instances for which submodularity does, and does not, apply:

- **The submodular case:** Observe that if no two nodes in M are directly connected via an edge in E then the valuation functions of the source nodes in the overlay network problem are submodular. This is due to the fact that they exhibit the decreasing marginal values property that is known to be equivalent to submodularity: For every node $i \in S$, for every $M' \subseteq M'' \subseteq M$, and for every $a \in M$, it holds that $v_i(M' \cup \{a\}) - v_i(M') \geq v_i(M'' \cup \{a\}) - v_i(M'')$.
- **The non-submodular case:** In the more general case, in which nodes in M can be connected to one another, the valuation functions are no longer guaranteed to be submodular. For example, think of a network graph $G = (V, E)$ with four nodes: a single source node i , a single destination node d , and two potential overlay nodes a, b . Assume that $E = \{(i, a), (a, b), (b, d)\}$ and that m assigns a value of 0 to every edge in E . Let $v_i(\{a\}) = v_i(\{b\}) = 0$ and $v_i(\{a, b\}) = 1$. Observe that v_i is indeed consistent with m , but is neither submodular nor subadditive.

In this model the difference between general and submodular valuations originates from whether or not the nodes in M are connected to each other. These seemingly subtle and insignificant differences lead to violations of submodularity which motivate our study of CPPP with valuations which are not submodular.

³ Formally, for every node $i \in S$, and for every subset $M' \subseteq M$ we define $c_i(M')$ to be the sum of the lengths of the shortest paths from i to all nodes in D that only go through nodes in M' . That is, if we define $(\forall i \in S, d \in D) c_{id}(M')$ to be the length of the shortest route (given m) that has i as its first node, d as its last node, and all intermediate nodes in M' , then $c_i(M') = \sum_{d \in D} c_{id} M'$. $v_i(M') \geq v_i(M'')$ iff $c_i(M') \leq c_i(M'')$.

⁴ The artificial division of V into S, M , and D is only required for ease of exposition.

3 Subadditive Valuations

In this section we study CPPP with subadditive valuations. We prove our main result, which is showing a lower bound of $m^{\frac{1}{4}}$ for a class strictly contained in subadditive valuations. This result shows that relaxations of submodularity that may seem small at first glance, and are not too costly for other problems (e.g., combinatorial auctions [7,8]), lead to unreasonable approximation ratios for CPPP.

We show that this lower bound is nearly tight by presenting a simple \sqrt{m} approximation algorithm. This algorithm also has the advantage of being *truthful*. Hence, we show that for CPPP with subadditive valuations the gap between truthful and unrestricted algorithms in terms of approximation ratio is insignificant (this is in contrast to the submodular case in which the gap is huge [18]).

3.1 Lower Bound for Subadditive Valuations

We prove our result for fractionally-subadditive valuations [8] (defined in [17] and termed “XOS” there). This result is achieved in the communication complexity (general queries) model. Thus, we show that even if agents are computationally unlimited the number of bits they must exchange to obtain a good approximation is unreasonable. The proof uses a probabilistic construction of a collection of subsets of resources, which has useful combinatorial properties. These properties are then exploited in a reduction from a well known problem in communication complexity.

Definition 1. A valuation function v is said to be additive (linear) if for every $S \subseteq M$ $v(S) = \sum_{j \in S} v(\{j\})$.

Informally, a valuation function is fractionally-subadditive if it is the pointwise maximum over a set of additive (linear) valuations.

Definition 2. A valuation function v is said to be fractionally subadditive if there is a set of additive (linear) valuations $\{a_1, \dots, a_l\}$ such that for every $S \subseteq M$ $v(S) = \max_{r \in [l]} a_r(S)$.

The class of fractionally-subadditive valuation functions is known to be strictly contained in the class of subadditive valuations and to strictly contain all submodular valuations [17,13].

Theorem 3.1. Obtaining an approximation ratio of $m^{\frac{1}{4}-\epsilon}$ for fractionally-subadditive valuation functions requires exponential communication in m (for every $\epsilon > 0$).

Proof. Fix a small $\epsilon > 0$. We prove the theorem for the case $n = k = \sqrt{m}$. The proof is by reduction from the SET DISJOINTNESS problem. In the SET DISJOINTNESS problem, we have n parties. Each party $i = 1, 2, \dots, n$ holds a t -bit string which specifies a subset $A_i \subseteq \{1, \dots, t\}$. The parties are required to distinguish between the two following extreme cases:

1. $\bigcap_{i=1}^n A_i \neq \emptyset$
2. for every $i \neq j$ it holds that $A_i \cap A_j = \emptyset$

The SET DISJOINTNESS problem was studied in [1] where it was shown to require $\Omega(\frac{t}{n^4})$ communication complexity. In [16] Nisan shows a lower bound due to Radhakrishnan and Srinivasan of $\Omega(\frac{t}{n})$.

For our reduction, we wish to first show the existence of a exponentially large family of sets of items with a useful combinatorial property:

Definition 3. A collection F of subsets of $[m]$ is said to have the r -intersection property if, for every $S, T \in F$, $|S \cap T| \leq r$.

Lemma 3.2. There exists collection F of subsets of $[m]$ such that:

- F has the $2m^\epsilon$ -intersection property.
- For each $S \in F$ it holds that $|S| \geq \sqrt{m}$.
- $|F| \geq e^{\alpha m^\epsilon}$ for some constant $\alpha > 0$.

Proof. We prove the lemma via a probabilistic construction. Each set $S \in F$ is constructed by choosing uniformly at random each element $l \in [m]$ to be in S with probability $m^{\frac{\epsilon-1}{2}}$. Let S and T be two such sets in F . We wish to show that each such set is “large” and that the intersection between them is “small”, with very high probability. We will use the Chernoff bound.

Claim. (Chernoff Bound) Let X_1, \dots, X_m be a set of m independent random variables that take values in $\{0, 1\}$ such that for every l , $Pr[X_l = 1] = p$. Then, for any δ is in the range $[0, 2e - 1]$ we have that:

$$Pr\left[\sum_{l=1}^m X_l > (1 + \delta)pm\right] \leq e^{\frac{-\delta^2 pm}{3}} \tag{1}$$

$$Pr\left[\sum_{l=1}^m X_l < (1 - \delta)pm\right] \leq e^{\frac{-\delta^2 pm}{3}} \tag{2}$$

For every resource $l \in [m]$ we define a random variable X_l that is assigned a value of 1 if $l \in S \cap T$ and of 0 otherwise. Observe, that the probability that $X_l = 1$ is $m^{\epsilon-1}$, and by [1]:

$$Pr[|S \cap T| > 2m^\epsilon] = Pr[\sum_l X_l > 2m^\epsilon] < e^{\frac{-m^\epsilon}{3}}$$

Similarly, for every $S \in F$, define a random variable X_l that is assigned a value of 1 if $l \in S$ and 0 otherwise, by [2] we have that for any $\delta \in (0, 1)$:

$$Pr[|S| < (1 - \delta)m^{\frac{1+\epsilon}{2}}] = Pr[\sum_l X_l < (1 - \delta)m^{\frac{1+\epsilon}{2}}] < e^{\frac{-\delta^2 m^{\frac{1+\epsilon}{2}}}{3}}$$

We thus have that there is some constant $\alpha > 0$ such that:

$$Pr[|S \cap T| > 2m^\epsilon \text{ or } |S| < m^{\frac{1}{2}} \text{ or } |T| < m^{\frac{1}{2}}] < e^{-2\alpha m^\epsilon}$$

Since these bounds must hold for any S, T we get that as long as $|F|^2 \leq e^{2\alpha m^\epsilon}$ there is such a collection F . Thus, we know there exists a family of sets $F = \{S_1, \dots, S_t\}$ where $t = e^{\alpha m^\epsilon}$ for some constant $\alpha > 0$, with the property that every set is at least of size \sqrt{m} and $|S_i \cap S_j| \leq 2m^\epsilon \forall i \neq j \in t$.

Now, we show the reduction from the SET DISJOINTNESS problem. Let $1, 2, \dots, \sqrt{m}$ be the parties, and set $t = e^{\alpha m^\epsilon}$. Let A_i be the subset of $[t]$ held by party i . We construct an instance of CPPP with m resources, $n = \sqrt{m}$ agents (corresponding to the \sqrt{m} parties in the SET DISJOINTNESS problem), and set the number of resources to be chosen to $k = \sqrt{m}$. We identify each element $r \in [t]$ with a set S_r in the family F of subsets of $[m]$ described above. Each agent i constructs a valuation function v_i in the following manner: Let a_S denote the additive valuation that assigns a value of 1 to every resource in S and a value of 0 to every resource $j \notin S$. Let $v_i = \max\{a_{S_r} | r \in A_i\}$.

Observe that if $\bigcap_i A_i \neq \emptyset$ then there is a set S_r that has a corresponding additive valuation in all of the v_i 's. Hence, assigning a subset of S_r of size \sqrt{m} to the agents (simulated by the SET DISJOINTNESS parties) results in a social welfare value of m . What happens if for every two $i \neq j \in [t]$ $S_i \cap S_j = \emptyset$? We shall now show that in this case the optimal social welfare is $O(m^{\frac{3}{4}+\epsilon})$. This would mean that an approximation of $O(m^{\frac{1}{4}-\epsilon})$ to the CPPP with fractionally-subadditive valuations enables the distinction between the two extreme cases in the SET DISJOINTNESS problem. Therefore, we will then be able to conclude that $\Omega(\frac{t}{n})$ bits are required to do so (a number exponential in both n and m).

So, we are left with showing that if for every two $i \neq j \in [t]$ $S_i \cap S_j = \emptyset$ then the optimal social welfare is $O(m^{\frac{3}{4}+\epsilon})$. Assume, for the purpose of contradiction, that there is some set T of size \sqrt{m} such that the social welfare derived from T , $SW(T)$, is greater than $2m^{\frac{3}{4}+\epsilon}$. Let a_i be an additive valuation function of i for which v_i is maximized (for T). Observe that $SW(T) = \sum_{i \in [n]} a_i(T)$. Assume, w.l.o.g., that $T = \{1, \dots, \sqrt{m}\}$. For every resource $l \in T$, let x_l be the number of the a_i 's that assign a value of 1 to l . Observe that $SW(T) = \sum_{i \in [n]} a_i(T) = \sum_{l \in T} x_l$. Also observe that $\sum_{l \in T} x_l(x_l - 1) = \sum_{i \neq j} |S_i \cap S_j \cap T|$. Since we know that the cardinality of the intersection of every two sets cannot exceed $2m^\epsilon$ we now have that:

$$2m^{1+\epsilon} = 2n^2m^\epsilon \geq \sum_{i \neq j} |S_i \cap S_j \cap T| = \sum_{l \in T} x_l(x_l - 1)$$

Using elementary calculus, it is easy to show that $\sum_{l \in T} x_l^2 \geq m^{\frac{1}{4}} \sum_{l \in T} x_l$. (This is due to the fact that the worst case ratio is achieved when all x_l 's are equal.)

Combining the last two equations gives us that

$$SW(T) = \sum_{l \in T} x_l \leq 2m^{\frac{3}{4}+\epsilon}.$$

A contradiction.

3.2 A Truthful \sqrt{m} Approximation Algorithm

We show that the result stated in the above theorem is nearly tight by presenting a simple *truthful* algorithm which obtains a $\min\{k, \sqrt{m}\}$ approximation ratio (for any value of k and n) and requires at most $n \cdot m$ value queries.

The Algorithm:

- Arbitrarily partition $[m]$ into $r = \max\{\frac{m}{k}, \sqrt{m}\}$ disjoint subsets of equal size S_1, \dots, S_r .
- Ask each agent to specify her value for each of the different subsets S_t .
- Choose the subset S_t that maximizes the social welfare $\sum_i v_i(S_t)$.

This algorithm, similar to that presented in [5], is a simple maximal-in-range algorithm [5,6], and so it can be made truthful via VCG payments [19,41]. Observe that the algorithm indeed requires at most m value queries to be addressed to each of the n agents. Therefore, all that is left to show is that the algorithm provides the required approximation-ratio. We show this for subadditive valuations.

Proposition 3.3. *If v_1, \dots, v_n are subadditive then the algorithm provides an approximation ratio of $\min\{k, \sqrt{m}\}$.*

Proof. Let $k \leq \sqrt{m}$. Let O be a set of size k that maximizes the social welfare. Then, by (iterative use of) subadditivity, for every $i \in [n]$, $v_i(O) \leq \sum_{j \in O} v_i(\{j\})$. Hence, $\sum_{i \in [n]} v_i(O) \leq \sum_{i \in [n]} \sum_{j \in O} v_i(\{j\}) = \sum_{j \in [O]} \sum_{i \in [n]} v_i(\{j\})$. This implies that there is an element $j \in [m]$ such that the social welfare derived from j is at least $\frac{1}{|O|} = \frac{1}{k}$ of the optimal social welfare. This item j appears in one of the S_t 's, and so, because the valuations are non-decreasing, the social welfare derived from that S_t is also at least $\frac{1}{k}$ of the optimal social welfare. Since the algorithm optimizes over all the S_t 's it is bound to achieve the desired approximation ratio.

Let $k > \sqrt{m}$. Let O be the set of size k that maximizes the social welfare. Because the valuations are non-decreasing, $\sum_i v_i(O) \leq \sum_i v_i([m])$. Let $S_1, \dots, S_{\sqrt{m}}$ be some arbitrary partition of $[m]$ into \sqrt{m} disjoint subsets of size \sqrt{m} . Exploiting subadditivity in a way similar to that shown above implies that for one of these sets the social welfare is at least a $\frac{1}{\sqrt{m}}$ fraction of the social welfare for the entire set $[m]$. This concludes the proof of the proposition.

4 Inapproximability of CPPP with General Valuations

In this section we study CPPP with general valuations (but still normalized and non-decreasing). As the overlay networks example (Section 2), and others (elections etc.), illustrate, in many cases submodularity, and even subadditivity, do not apply. We prove strong inapproximability results for CPPP with general valuations in both the computational- and the communication-complexity models. In the communication-complexity model our lower bound is tight (a trivial matching upper bound exists).

Theorem 4.1. *Obtaining an approximation of $O(n^{\frac{1}{2}-\epsilon})$ to the social welfare in CPPP with general valuations, for any $\epsilon > 0$, is impossible unless $P = NP$. Obtaining an approximation of $O(n^{1-\epsilon})$ to the social welfare is impossible unless $P = ZPP$.*

Proof. We reduce from the Maximal Welfare Tree (MWT) problem studied in the context of distributed algorithmic mechanism design [10,9,14]. Our reduction preserves the hardness results for this problem as shown in [14]. In the MWT problem we are given a graph $G = (N, L)$ with a set of nodes N and links L . A unique destination node d is given and each node $a \in N \setminus \{d\}$ has a valuation function $v_a : P_a \rightarrow \mathcal{R}_{\geq 0}$, where P_a is used to denote the set of all simple paths from a to the destination d . The objective in MWT is to form a tree rooted in d which maximizes the social welfare, i.e., choose the tree T^* such that $T^* \in \operatorname{argmax}_{T \in T_L^d} \sum_{a \in N \setminus \{d\}} v_a(T)$, where T_L^d is the set of all possible trees in L rooted in d . We consider the special case of MWT in which for all $a \in N \setminus \{d\}$ we have $v_a : P_a \rightarrow \{0, 1\}$. It is known that for any $\epsilon > 0$ approximating MWT, even for this special case, within a factor of $O(n^{\frac{1}{2}-\epsilon})$ is impossible unless $P = NP$ and approximating within a factor of $O(n^{1-\epsilon})$ is impossible unless $P = ZPP$ [14].

The reduction from MWT is as follows: Given an instance of MWT such that the range of all valuation functions is $\{0, 1\}$, for each link $l \in L$ we associate a resource l' in CPPP and each node $a \in N \setminus \{d\}$ in MWT will correspond to an agent a' in CPPP. It remains to define the valuation function of a' . Note that since our interest is in showing a lower bound, we can adversarially set the number of chosen items to be $k = |N| - 1$. Now, let P_a^+ be the set of paths for which $v_a = 1$. For all $E \subseteq L$, the valuation function for the corresponding agent a' in CPPP is defined by:

$$v_{a'}(E) = \begin{cases} 1 & \exists P \in P_a^+ : P \subseteq E \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

Observe that choosing a tree T in MWT with social welfare value $SW(T) = c$ corresponds to choosing a set of resources that induces the same social welfare value in CPPP. Conversely, choosing a set of resources T' in CPPP s.t. $SW(T') = c'$ necessarily means that we can trim T' to a set of edges T which forms a routing tree with d as its source, and that we have exactly c' nodes which have routes to d in T , and hence $SW(T) = c'$ in MWT.

Theorem 4.2. *Obtaining an approximation ratio of $(1 - \epsilon)n$ for general valuations requires exponential communication in m (for any $\epsilon > 0$ and for any $n \ll 2^m$).*

Proof. For CPPP with general valuation functions, n agents, m items and a parameter $1 \leq k \leq m$ we show a lower bound of $\Omega\left(\binom{m}{k} \cdot n^{-1}\right)$ again by reducing from the SET DISJOINTNESS problem.

We construct an instance of CPPP with n agents in which no restrictions (except for being normalized and non-decreasing) apply to the agents' valuation functions. Let S_1, \dots, S_t be the (ordered) sets in the range of all possible allocations of size k . For each party i in SET DISJOINTNESS with the set $A_i \subseteq \{1, \dots, t\}$, we associate an agent i in CPPP with the following valuation function:

$$v_i(S_r) = \begin{cases} 1 & r \in A_i \\ 0 & \text{otherwise} \end{cases}$$

Observe that these valuation functions are indeed normalized and non-decreasing. Let S_l be the set which maximizes the social welfare, i.e., $l \in \operatorname{argmax}_{l \in [d]} |\{A_i \mid l \in A_i\}|$. To approximate the social welfare within a factor of $(1 - \epsilon)n$ for any $\epsilon > 0$, one must allocate some set S for which there are at least two agents i and j such that $v_i(S) = v_j(S) = 1$. Due to the above construction of the agents' valuation functions this necessarily implies deciding between the two extreme cases of the SET DISJOINTNESS problem. Thus, for $d = \binom{n}{k}$ we get a lower bound of $\Theta(\binom{n}{k} \cdot n^{-1})$ for CPPP with general valuation functions.

In the communication model a trivial matching upper bound of n exists: Query each agent i for her most valued set S_i of size k , and assign the agents a set $T \in \operatorname{argmax}_i v_i(S_i)$. It is easy to see that this indeed guarantees an n -approximation.

Acknowledgements

We warmly thank Christos Papadimitriou for many helpful and interesting discussions.

References

1. Alon, N., Matias, Y., Szegedy, M.: Journal of computer and system sciences. The space complexity of approximating the frequency moment 58(1), 137–137 (1999)
2. Andersen, D., Balakrishnan, H., Kaashoek, F., Morris, R.: Resilient overlay networks. In: SOSP 2001: Proceedings of the eighteenth ACM symposium on Operating systems principles, pp. 131–145. ACM, New York (2001)
3. Chun, B.-G., Fonseca, R., Stoica, I., Kubiatowicz, J.: Characterizing selfishly constructed overlay routing networks. In: INFOCOM (2004)
4. Clarke, E.H.: Multipart pricing of public goods. Public Choice 11, 17–33 (1971)
5. Dobzinski, S., Nisan, N., Schapira, M.: Approximation algorithms for combinatorial auctions with complement-free bidders. In: STOC (2005)
6. Dobzinski, S., Nisan, N.: Limitations of VCG-based mechanisms. In: STOC (2007)
7. Dobzinski, S., Schapira, M.: An improved approximation algorithm for combinatorial auctions with submodular bidders. In: SODA 2006 (2006)
8. Feige, U.: On maximizing welfare when the utility functions are subadditive. In: STOC 2006 (2006)
9. Feigenbaum, J., Papadimitriou, C.H., Sami, R., Shenker, S.: A BGP-based mechanism for lowest-cost routing. Distributed Computing 18(1), 61–72 (2005)
10. Feigenbaum, J., Sami, R., Shenker, S.: Mechanism design for policy routing. In: PODC, pp. 11–20 (2004)
11. Groves, T.: Incentives in teams. Econometrica, 617–631 (1973)
12. Kushilevitz, E., Nisan, N.: Communication Complexity. Cambridge University Press, Cambridge (1997)
13. Lehmann, B., Lehmann, D., Nisan, N.: Combinatorial auctions with decreasing marginal utilities. In: ACM conference on electronic commerce (2001)
14. Levin, H., Schapira, M., Zohar, A.: Interdomain routing and games. In: Proceedings of STOC 2008 (2008)
15. Nemhauser, G.L., Wolsey, L.A., Fisher, M.L.: An analysis of approximations for maximizing submodular set functions ii. Math. Programming Study 8, 73–87 (1978)

16. Nisan, N.: The communication complexity of approximate set packing and covering. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) ICALP 2002. LNCS, vol. 2380. Springer, Heidelberg (2002)
17. Nisan, N.: Bidding and allocation in combinatorial auctions. In: ACM Conference on Electronic Commerce (2000)
18. Papadimitriou, C.H., Schair, M., Singer, Y.: On the hardness of being truthful. In: FOCS (2008)
19. Vickrey, W.: Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 8–37 (1961)
20. Yao, A.C.-C.: Some complexity questions related to distributive computing. In: ACM Symposium on Theory of Computing (STOC), pp. 209–213 (1979)

Algorithms for Optimal Price Regulations

Alexander Grigoriev¹, Joyce van Loon^{1,*}, and Marc Uetz²

¹ Maastricht University, Quantitative Economics,
P.O. Box 616, NL-6200 MD Maastricht, The Netherlands

{a.grigoriev,j.vanloon}@ke.unimaas.nl

² University of Twente, Applied Mathematics,
P.O. Box 217, NL-7500 AE Enschede, The Netherlands

m.uetz@utwente.nl

Abstract. Since summer 2007, mobile phone users in the European Union (EU) are protected by a ceiling on the roaming tariff when calling or receiving a call abroad. We analyze the effects of this price regulative policy, and compare it to alternative implementations of price regulations. The problem is a three-level mathematical program: The EU determines the price regulative policy, the telephone operator sets profit-maximizing prices, and customers choose to accept or decline the operator's offer. The first part of this paper contains a polynomial time algorithm to solve such a three-level program. The crucial idea is to partition the polyhedron of feasible price regulative parameters into a polynomial number of smaller polyhedra such that a certain primitive decision problem can be written as an LP on each of those. Then the problem can be solved by a combination of enumeration and linear programming. In the second part, we analyze more specifically an instance of this problem, namely the price regulation problem that the EU encounters. Using customer-data from a large telephone operator, we compare different price regulative policies with respect to their social welfare. On the basis of the specific social welfare function, we observe that other price regulative policies or different ceilings can improve the total social welfare.

Keywords: Pricing problems, three-level program optimization, social welfare maximization, EU roaming regulation.

1 Introduction

It is of major importance to the European Union (EU) that European companies, governments and citizens play an important role at the realization of a world-economy based on knowledge. The EU tries to stimulate the development and use of new information and communication technology, and to enlarge the level of competition of the EU compared to other markets, e.g. United States and Japan. An important element of the European policy is to assure that ICT-services are

* Supported by METEOR, the Maastricht Research School of Economics of Technology and Organizations.

available and affordable for everyone. This contains for example telephony, fax, internet and free emergency numbers^[1]. However, especially the prices for making and receiving calls abroad, referred to as *roaming*, have been extremely high in the EU recently. A warning did not lead to a decrease in prices, and therefore the European Commission uses *price regulation* to force lower prices and more transparency in the market [2]. Currently, the EU considers the same instrument for *data roaming*, since the situation mirrors the one for voice roaming back in 2007. We analyze both algorithmically and practically the effects of the current price regulation on the *social welfare*, and compare it to alternative regulations.

We regard a general model in which a government tries to maximize *social welfare* through price regulation. This regulation should bound the producer (not necessarily a telephone operator) in setting the prices so as to protect the customers in the market. The producer determines the price of *items*. An item is not necessarily a physical product, but can for example also be a minute calling, internet connection, shipping, etc. We use a model in which there is only one producer that determines the pricing strategy, under given and known market competition. Obviously, there exist markets in which multiple producers operate and need to share a set of customers. However, for this research we make the simplifying assumption that, under given market regulations, producers reach an optimal market price, so we identify them with just one single producer.

1.1 Model

Let $K = \{1, \dots, m\}$ be the set of distinct item types a given producer wants to price. Let p_k be the price of items of type $k \in K$ that needs to be determined by the producer. Let $J = \{1, \dots, n\}$ be a set of potential customers. Every customer $j \in J$ has demand d_{jk} for item $k \in K$, which is the number of times customer j wants to purchase item k . For example, if item k represents calling abroad for one minute, then d_{jk} is the number of minutes j wants to call abroad. Or, if k is the start up of a process, then d_{jk} is 1 if j prefers to start the process and 0 otherwise. Every customer j requests a *contract* from the producer, which is specified by the total demand vector (d_{j1}, \dots, d_{jm}) . Once the item prices are determined, the price of the contract, $p(j)$, is defined by the following *affine function*

$$p(j) = d_{j0} + d_{j1}p_1 + \dots + d_{jm}p_m, \quad j \in J. \quad (1)$$

Note that the price of a contract is personal, due to a potential ‘entrance fee’ d_{j0} , and because it depends on the demand (d_{j1}, \dots, d_{jm}) of a single customer. We assume, e.g. through market research, that we know the customers’ demands d_{jk} for all item types $k \in K$, and valuation b_j , for all customers $j \in J$, which is the maximum amount customer j is willing to pay for her bundle. Therefore, we are faced with a purely algorithmic problem in contrast to mechanism design problems where the valuations are private information to the customers.

The pricing regime as defined in (1) is referred to as *affine pricing* in the economic literature; it was also discussed by Grigoriev et al. in [11]. It

¹ From: <http://www.europa-nu.nl>

is probably a more realistic model than the single item pricing model that was discussed in many papers on algorithmic profit maximizing pricing problems [2,3,7,8,9,10,12,13,14].

A solution to the problem is a price $p(j)$ for every customer $j \in J$, which is determined through a vector of item prices $p = (p_1, \dots, p_m)$ as given in (II). Every customer decides whether to accept this contract or not. Hereto, she sets binary variable w_j to 1 if she accepts, and 0 otherwise, in order to maximize her personal objective, denoted by $f_j^C(p)$. In this paper, we assume a linear objective function. For example, think of $f_j^C(p) = b_j - p(j)$. Let $w = (w_1, \dots, w_n)$ denote the strategies of all customers. Customers that accept the contract are referred to as *winners*, and the set of winners is defined by $W = \{j \in J : w_j = 1\}$. We assume that all items are available in unlimited supply, which is true for digital items for example.

The government protects the customers by means of regulative constraints. Let R denote the set of constraints imposed by the government. Throughout this paper, we assume that the number of regulations $|R|$ is constant. Every constraint $r \in R$ is defined by $g_r(p, \alpha_r) \leq 0$, where $\alpha = (\alpha_1, \dots, \alpha_{|R|})$ is a vector of price regulative parameters determined by the government. For example, a ceiling on item price p_k is implemented by letting $g_k(p, \alpha_k) = p_k - \alpha_k$. In this paper, function $g_r(p, \alpha_r)$ is restricted to be a *linear* function in α . We introduce a bilevel program in which the producer maximizes his objective $f^P(p, w)$ (e.g. revenue minus production costs) such that price vector $p = (p_1, \dots, p_m)$ satisfies the price regulative constraints. Every customer maximizes her objective $f_j^C(p)$. For a reference on bilevel programs, see e.g. [15].

$$\begin{aligned}
 \text{2LP : } & \max_{p \in \mathbb{R}_+^m} && f^P(p, w) \\
 & \max_{w_j \in \{0,1\}} && f_j^C(p)w_j \quad \forall j \in J \\
 & \text{s.t.} && g_r(p, \alpha_r) \leq 0 \quad \forall r \in R
 \end{aligned}$$

In the above mathematical program, price regulative parameters α and constraints $g_r(p, \alpha_r)$ are assumed to be given. This bilevel problem can be solved in polynomial time for a constant number of distinct items m , by a simple enumerative algorithm [11]. In this paper, we propose a *three-level* program where, on top of the two levels given by producer and customers, the government maximizes social welfare by modifying the price regulative parameters α . Also, instead of simply forbidding the violation of price regulative constraints $g_r(p, \alpha_r) \leq 0$, we introduce taxes $\tau = (\tau_1, \dots, \tau_{|R|})$. That is, if the producer's prices violate regulation $g_r(p, \alpha_r) \leq 0$, then he pays a penalty over the additional profit he receives by this violation. The actual penalty for violating regulation r is denoted by $f_r^{\text{TAX}}(p, w, \alpha, \tau) = g_r(p, \alpha_r)^+ \cdot \tau_r \cdot \bar{g}_r(w)$, where $g_r(p, \alpha_r)^+ = \max\{0, g_r(p, \alpha_r)\}$ is the amount of violation, τ_r is the tax (or penalty), and finally function $\bar{g}_r(w)$ needs to be specified for each type of price regulative constraint r . For example, if the regulation is a ceiling $g_k(p, \alpha_k) = p_k - \alpha_k$, it sounds reasonable to ask tax τ_k for each euro earned by violating the ceiling. Then $\bar{g}_r(w)$ would have to be defined as the total demand of all winners, $\sum_{j \in J} d_{jk}w_j$. Let us denote the total tax payment for a producer by

$$f^{\text{TAX}}(p, w, \alpha, \tau) = \sum_{r \in R} f_r^{\text{TAX}}(p, w, \alpha, \tau) = \sum_{r \in R} g_r(p, \alpha_r)^+ \cdot \tau_r \cdot \bar{g}_r(w). \quad (2)$$

The mathematical program (3LP) below shows the general structure of the three-level program that we consider.

$$\begin{aligned} \text{3LP : } & \max_{\alpha, \tau \in \mathbb{R}_+^{|R|}} && f^G(p, w, \alpha, \tau) \\ & \max_{p \in \mathbb{R}_+^m} && f^P(p, w) - f^{\text{TAX}}(p, w, \alpha, \tau) \\ & \max_{w_j \in \{0,1\}} && f_j^C(p)w_j \quad \forall j \in J \end{aligned}$$

If the government’s objective $f^G(p, w, \alpha, \tau)$ is *monotone* in the prices then we show that we can solve this program in polynomial time, given that the number of items m and the number of regulative constraints $|R|$ are constant.

Note that the lower two levels of 3LP can be seen as the Lagrangian of the bilevel program 2LP, and a strict price regulation as in the bilevel model (that is, forbidding to violate the regulative constraints) can still be implemented by letting τ_r be arbitrarily large for all $r \in R$ (then τ is not strictly a tax, but just an arbitrary penalty).

It is known that problem 2LP with a non-constant number of items m is hard to approximate within a semi-logarithmic factor in the number of customers n [8]. This means, that solving the three-level program 3LP has the same complexity already if the government’s objective is equal to the producer’s objective. As in [11], in this paper we make the assumption that the number of distinct item types m is constant. These are reasonable assumptions particularly for the applications that we aim at, since there the number of item types is very small (for example, price per minute for a call received, a call placed, and price per SMS). For a small number of items, the number of regulations is also assumed to be small.

1.2 Our Results

For a constant number of items m and regulative constraints $|R|$, we present a polynomial time algorithm to solve three-level program 3LP, under the restrictions that the government’s objective function is monotone in the item prices, and the price regulative constraints $g_r(p, \alpha_r)$ are linear in α for all regulations $r \in R$.

We explicitly define all functions in three-level program 3LP to optimize the social welfare for the specific problem faced by the European Union in regulating the roaming charges. We present polynomial time algorithms to find the optimal social welfare for the current EU policy (Theorem 3), optimization of the tax level only (Theorem 4), and optimization of the regulative parameters only (Theorem 5). After the theoretical results and description of the algorithms we perform a practical study in Section 4 to verify practical feasibility of the approach and to evaluate the result of different scenarios for implementing price regulations. Here, we use the actual price regulations set by the European Commission and investigate the EU policy in terms of social welfare.

2 Parameter and Tax Level Optimization

Consider three-level program 3LP, as specified in the introduction. In this section, we propose an algorithm to solve this program for which the government’s objective function is monotone in the prices, and the regulations $g_r(p, \alpha_r)$, $r \in R$, are linear in α (note that the function does not have to be linear in p).

Definition 1. *The set V of vertices is defined as all price vectors $p = (p_1, \dots, p_m)$ defined by m linearly independent constraints out of the $n + m$ constraints $f_j^C(p) = 0$, $j \in J$, and $p_k = 0$, $k \in K$.*

The government’s objective function f^G is monotone in the prices. Therefore, we straightforwardly derive the following theorem.

Theorem 1. *For any given vectors w , α and τ , the optimal price vector $p = (p_1, \dots, p_m)$ for the producer can only be at a vertex as defined in Definition 1.*

We propose an algorithm that solves problem 3LP by optimizing price regulative parameters α and the taxes τ simultaneously. Thereto, we partition the polyhedron of the price regulative parameters α into a polynomial number of smaller polyhedra, and solve a linear programming problem in each of those. These linear programs are defined in such a way that we can verify if a given price vector constitutes the producer’s optimum prices. This decision problem is in general non-linear. The trick here is to define the partition such that this decision problem becomes linear inside each of the small polyhedra. The optimum solution of a three-level mathematical program is eventually obtained by enumeration over all polyhedra and vertices, and evaluating the social welfare in each of them.

Consider the following decision problem.

Problem 1. *Are there price regulative vectors α and τ , such that vertex $v \in V$ with price vector $p^{(v)}$ maximizes the objective function $f^P - f^{\text{TAX}}$ for the producer?*

The main idea for the solution to Problem 1: Consider some arbitrary vertex $v \in V$, we would like to write constraints expressing the fact that vertex v maximizes the objective, namely

$$f^P(p^{(v)}, w^{(v)}) - f^{\text{TAX}}(p^{(v)}, w^{(v)}, \alpha, \tau) \geq f^P(p^{(u)}, w^{(u)}) - f^{\text{TAX}}(p^{(u)}, w^{(u)}, \alpha, \tau)$$

for all vertices $u \in V$. By definition, $f^{\text{TAX}}(p^{(v)}, w^{(v)}, \alpha, \tau)$ is nonlinear in the price regulative parameters. To linearize f^{TAX} we introduce a subdivision of $\mathfrak{R}_+^{|R|}$ into polyhedra A_l , $l = 1, \dots, L$. For a given vertex $p^{(v)}$, $v \in V$ and by linearity of $g_r(p^{(v)}, \alpha_r)$, there is a unique value for α_r , say a_r^v , where the sign of $g_r(p^{(v)}, \alpha_r)$ changes from ≤ 0 to > 0 . Doing this for all vertices $v \in V$ and all regulative constraints $r \in R$, we define a rectangular subdivision in $\mathfrak{R}_+^{|R|}$ for possible α ’s by $\alpha_r = a_r^v$. On each such defined α -rectangle A_l , we may now compare the producer’s objective in vertex $v \in V$ to the objective in all other vertices $u \in V \setminus v$. To do this, we split the penalty function f_r^{TAX} , $r \in R$, into two parts for every vertex $v \in V$ and polyhedron A_l : $x_r^{(v,l)}$ incorporates all terms in

the regulative constraint not multiplied with α_r and $y_r^{(v,l)}$ incorporates all terms multiplied with α_r . Therefore,

$$f_r^{\text{TAX}}(p^{(v)}, w^{(v)}, \alpha, \tau) = g_r(p, \alpha_r)^+ \cdot \tau_r \cdot \bar{g}_r(w) = x_r^{(v,l)} \tau_r + y_r^{(v,l)} \alpha_r \tau_r, \quad \forall r \in R.$$

We can derive a solution to Problem \square on A_l by solving the following mathematical program for every vertex $v \in V$.

$$\begin{aligned} f^{\text{P}}(p^{(v)}, w^{(v)}) - \left(\sum_{r \in R} x_r^{(v,l)} \tau_r + y_r^{(v,l)} \alpha_r \tau_r \right) &\geq \\ f^{\text{P}}(p^{(u)}, w^{(u)}) - \left(\sum_{r \in R} x_r^{(u,l)} \tau_r + y_r^{(u,l)} \alpha_r \tau_r \right) &\quad \forall u \in V \\ \check{\alpha}_r^{(l)} \leq \alpha_r \leq \hat{\alpha}_r^{(l)} &\quad \forall r \in R \\ \tau_r \geq 0 &\quad \forall r \in R. \end{aligned}$$

This quadratic program can be linearized by simple variable substitution $\phi_r = \alpha_r \tau_r$ for all $r \in R$. Therefore, for every vertex $v \in V$ and polyhedron A_l , Problem \square becomes a linear program (LP1) with variables τ_r and ϕ_r for all $r \in R$.

$$\begin{aligned} \text{LP1 : } \sum_{r \in R} \left(x_r^{(v,l)} - x_r^{(u,l)} \right) \tau_r + \left(y_r^{(v,l)} - y_r^{(u,l)} \right) \phi_r &\leq \\ f^{\text{P}}(p^{(v)}, w^{(v)}) - f^{\text{P}}(p^{(u)}, w^{(u)}) &\quad \forall u \in V \\ \check{\alpha}_r^{(l)} \tau_r \leq \phi_r \leq \hat{\alpha}_r^{(l)} \tau_r &\quad \forall r \in R \\ \tau_r \geq 0 &\quad \forall r \in R. \end{aligned}$$

On any polyhedron A_l this linear program is either infeasible, suggesting that there are no price regulative parameters in A_l that makes v the solution that maximizes the producer’s objective, or otherwise we obtain corresponding price regulative parameters in A_l . Eventually, a straightforward algorithm enumerating all vertices $v \in V$, checking feasible solutions for α ’s in $A_l, l = 1, \dots, L$, and picking the one that achieves the maximal social welfare, provides an optimal solution to the three-level program. Since the number of items m is constant, we have a polynomial number of vertices. As the number of regulative constraints $|R|$ is constant, we have only a polynomial number of polyhedra in $\mathfrak{R}_+^{|R|}$. For every polyhedron and vertex, we solve linear program LP1, deriving the following theorem.

Theorem 2. *Three-level program 3LP admits a polynomial time algorithm if the number of items m and the number of regulative constraints $|R|$ are constant.*

3 Optimization of European Regulation

In this section, we explicitly define functions $f^{\text{G}}, f^{\text{P}}, f^{\text{C}}$ and f^{TAX} to solve the problem faced by the European Union regarding the regulation on roaming. First of all, the EU sets ceilings on the prices. That is, $R = K$, and $g_k(p, \alpha_k) =$

$p_k - \alpha_k$ for all items $k \in K$. The objective of the producer is to maximize profit, defined as revenue minus costs. The revenue is the total payment by all winning customers. As for the costs, let c_k be the cost of providing one unit of k to a customer. The producer’s cost to serve customer j is denoted by $c(j) = d_{j1}c_1 + \dots + d_{jm}c_m$. The customers accept a contract if its price does not exceed the valuation, that is, the objective of customer $j \in J$ is defined as $(b_j - p(j))w_j$. If $p(j) \leq b_j$ this function is maximized by setting $w_j = 1$ and thus accepting the contract, and 0 otherwise.

In the simplest setting, the government regulates the prices by forbidding to violate the constraint $p_k \leq \alpha_k$ for every item $k \in K$. We model this by the following bilevel program.

$$\begin{aligned} & \max_{p \in \mathbb{R}_+^m} && \sum_{j \in J} (p(j) - c(j))w_j \\ & \max_{w_j \in \{0,1\}} && (b_j - p(j))w_j \\ & \text{s.t.} && p_k \leq \alpha_k \quad \forall k \in K \end{aligned}$$

To find the prices that will lead to the optimal profit for the producer, we use the affine pricing algorithm introduced in [11], in which we incorporate the price regulative constraints.

Theorem 3 ([11]). *For given price regulative constraints $p_k \leq \alpha_k$ which must not be violated, profit maximizing prices can be computed in polynomial time, given that the number of distinct item types m is constant.*

As already discussed in the introduction, we study if there are other price regulative strategies that might lead to an increase in *social welfare*. Thus, let us first proceed with a definition of the social welfare function we believe to be appropriate for modeling the roaming regulation problem. According to utilitarians such as Jeremy Bentham and John Stuart Mill, society should aim to maximize the total *utility* of individuals, aiming for “the greatest happiness for the greatest number”. Thus, the government strives to set the price regulative parameters so as to maximize the social welfare, defined as the sum of utilities. The utility of the producer is the total revenue minus costs. We assume that the producer is risk-neutral, and thus the marginal utility is equal for every extra euro earned. On the contrary, consumers have a concave utility function in general, which means that they are assumed to be risk-averse. A concave utility function induces that a gain in wealth conveys a smaller increase to utility than the reduction in utility imparted by a loss in wealth of equal magnitude, that is, *diminishing marginal utility*. Another property of a concave utility function is that a customer with a low valuation may value one unit of money more than a customer with a high valuation. In other words, the marginal utility of a euro to a ‘poor’ customer is likely to exceed the marginal utility of a euro to a ‘rich’ customer [17, Chapter VII]. Daniel Bernoulli [6] first proposed a utility function that is equal to the natural logarithm of wealth. A logarithmic function is monotonically increasing and the marginal utility function is monotonically decreasing, which are the two basic mathematical properties that consumer utility functions have to satisfy [4]. In the words of Savage [16], “no other function has been suggested as a better prototype for Everyman’s utility function”. Based on this discussion, we model

the utility of customer $j \in J$ as $\ln(b_j - p(j) + 1)w_j$, where the addition of 1 is solely to have a positive function. Conclusively, the social welfare, and thus the government’s objective function, is defined as $\sum_{j \in J} (\ln(b_j - p(j) + 1) + p(j) - c(j))w_j$. In Section 3.1, we furthermore include a tax payment for violating the price regulative constraints, similar as in Section 2. However, this does not change the social welfare function, as the tax is paid by the producer to the government. Thus, the producer’s utility is decreased by the same amount as the government’s utility is increased, also known as *transferable utility*. This assumption is justified when the producer and the government have a common currency that is valued equally by both. Another reason not to include the tax payment in the social welfare is that it is a punishment to the producer, not to the society as a whole.

3.1 Price Regulation by Tax

In this section, the price regulative constraints $p_k \leq \alpha_k$ are not enforced by law, but their violation is penalized via tax (or penalty) τ_1 . We define $\bar{g}_k(w) = \sum_{j \in J} d_{jk}w_j$ for all items $k \in K$. The penalty function is defined as

$$f_k^{\text{TAX}}(p, w, \alpha, \tau) = g_k(p, \alpha_k)^+ \cdot \tau_1 \cdot \bar{g}_k(w) = \sum_{j \in J} d_{jk}(p_k - \alpha_k)^+ \tau_1 w_j, \quad \forall k \in K.$$

Thus, given vector α , the government determines tax level τ_1 to maximize social welfare.

$$\begin{aligned} \max_{\tau_1 \geq 0} f^G &= \sum_{j \in J} (\ln(b_j - p(j) + 1) + p(j) - c(j))w_j \\ \max_{p \in \mathbb{R}_+^m} f^P - f^{\text{TAX}} &= \sum_{j \in J} (p(j) - c(j) - \sum_{k \in K} d_{jk}(p_k - \alpha_k)^+ \tau_1) w_j \\ \max_{w_j \in \{0,1\}} f_j^C w_j &= (b_j - p(j))w_j \end{aligned}$$

For any given $\alpha_k, k \in K$, consider the arrangement of linear equalities defined in \mathbb{R}^m by the *valuation constraints* $p(j) = b_j$ for every customer $j \in J$ (that is, $f_j^C = 0$), nonnegativity constraints $p_k = 0$ and price regulative constraints $p_k = \alpha_k$ for every item $k \in K$.

Definition 2. A vertex $v \in V$ is defined as a price vector $p = (p_1, \dots, p_m)$ that satisfies m linearly independent constraints out of the $n + 2m$ constraints $p(j) = b_j, j \in J, p_k = \alpha_k$, and $p_k = 0, k = 1, \dots, m$.

As a direct consequence of Theorem 1, for any given vector α and the given social welfare function f^G , profit maximizing price vectors $p = (p_1, \dots, p_m)$ can only be vertices as defined in Definition 2. More specifically, one can easily check that the necessary Karush-Kuhn-Tucker conditions (see e.g. [5]) do not hold in any point except the vertices.

Notice that a vertex $v \in V$ is most preferable to the producer if the profit after tax at this vertex, $f^P(p^{(v)}, w^{(v)}) - f^{\text{TAX}}(p^{(v)}, w^{(v)}, \alpha, \tau)$, is at least as high as at any other vertex. Let $W^{(v)} = \{j \in J : w_j^{(v)} = 1\}$. More precisely, for every vertex $u \in V \setminus v$, the tax level τ_1 must be such that

$$\sum_{j \in W^{(v)}} p^{(v)}(j) - c(j) - \sum_{k \in K} d_{jk}(p_k^{(v)} - \alpha_k)^+ \tau_1 \geq \sum_{j \in W^{(u)}} p^{(u)}(j) - c(j) - \sum_{k \in K} d_{jk}(p_k^{(u)} - \alpha_k)^+ \tau_1.$$

Note that all terms except τ_1 in the above inequality are known, as α_k is given and $p_k^{(v)}$ is defined for all $k \in K$ and $v \in V$. Let us denote $T^{(v)} = \sum_{j \in W^{(v)}} \sum_{k \in K} d_{jk}(p_k^{(v)} - \alpha_k)^+$. We rewrite the latter inequality and solve the following feasibility linear program (LP2) below for each vertex $v \in V$.

$$\begin{aligned} (T^{(v)} - T^{(u)}) \tau_1 &\leq \sum_{j \in W^{(v)}} (p^{(v)}(j) - c(j)) - \sum_{j \in W^{(u)}} (p^{(u)}(j) - c(j)) \quad \forall u \in V \setminus v \\ \tau_1 &\geq 0. \end{aligned}$$

Let $V^* \subseteq V$ be the set of vertices for which the above linear program has a feasible solution. Then, among all vertices in V^* we select the vertex v^* with the highest social welfare $\sum_{j \in W^{(v^*)}} \ln(b_j - p^{(v^*)}(j) + 1) + p^{(v^*)}(j) - c(j)$. The tax level τ_1 is obtained as a solution to the linear program for this particular vertex v^* . So we have proved:

Theorem 4. *For given price regulative constraints $p_k \leq \alpha_k$, the tax level τ_1 that maximizes the total social welfare, and the corresponding profit maximizing prices can be computed in polynomial time, given that the number of distinct item types m is constant.*

3.2 Parameter Optimization

So far we assumed *given* values of the price regulative parameters $\alpha_1, \dots, \alpha_m$. In this section we optimize these parameters under the regulation that the producer sets the price $p_k \leq \alpha_k$ for all $k \in K$. Hereto, we use the following model.

$$\begin{aligned} \max_{\alpha \in \mathbb{R}_+^m} & \sum_{j \in J} (\ln(b_j - p(j) + 1) + p(j) - c(j)) w_j \\ \max_{p \in \mathbb{R}_+^m} & \sum_{j \in J} (p(j) - c(j) - \sum_{k \in K} d_{jk}(p_k - \alpha_k)^+ \tau_1) w_j \\ \max_{w_j \in \{0,1\}} & (b_j - p(j)) w_j \quad \forall j \in J \\ \text{s.t.} & p_k \leq \alpha_k \quad \forall k \in K \end{aligned}$$

Since parameters $\alpha_k, k \in K$, are not given, let V denote the set of vertices as defined in Definition [1](#). For every vertex $v \in V$, let $\alpha_k = p_k^{(v)}$ for all $k \in K$. Let $U = \{u \in V : p_k \leq \alpha_k, \forall k \in K\}$. Then, vertex $u \in U$ is most preferable to the producer if the profit $\sum_{j \in W^{(u)}} p^{(u)}(j) - c(j) \geq \sum_{j \in W^{(u')}} p^{(u')}(j) - c(j)$ for all $u' \in U \setminus u$. Among all vertices that are most preferable given set U , we select the one with the highest social welfare and set the α -parameters accordingly.

Theorem 5. *For the regulation that forbids the producer to violate the constraints, the parameter vector α that maximizes the total social welfare, and the corresponding profit maximizing prices can be computed in polynomial time, given that the number of distinct item types m is constant.*

4 Computational Results

In the summer of 2007, the European Commission decided to implement a EU-wide ceiling on the international roaming tariffs. The maximum price for calling from abroad is € 0.5831, for receiving a call abroad is € 0.2856, and the maximum price an operator may charge another operator for using the network is € 0.3570. This latter is the cost of the operator for providing roaming service to the customers. In summer 2008, these prices are lowered to € 0.5474, € 0.2618 and € 0.3332, respectively. In summer 2009, they will decrease even further to € 0.5117, € 0.2261 and € 0.3094. The goal of this practical study is to analyze the effect of properly chosen parameters on social welfare, and the advantage of using taxes instead of forbidding to violate the price regulation. We use data from a telephone operator containing the phone usage of customers with a pre-paid subscription during one month, March 2007, thus this data set comes from the period before the introduction of the price regulation. The data contains, for each customer, the number of minutes and times each customer uses the mobile phone for different actions (e.g. calling within the home country or abroad, sending a text message, etc.). We determine for every customer which operator in the telephone market offers the cheapest possible total price for her complete contract. This price determines her *valuation* b_j . For this study, even though the data contains more information, we focus on optimizing the prices for roaming only, namely calling and receiving a call abroad. This is to say, we consider a problem in dimension 2, with prices p_1 for calling abroad, and p_2 for receiving a call abroad. We also impose the constraint that the price for receiving a call should not exceed the price for placing a call ($p_1 \geq p_2$). If customer $j \in J$ requests to call d_{j1} minutes from abroad and receives calls abroad for d_{j2} minutes, within a month, the price customer j has to pay is $p(j) = d_{j1}p_1 + d_{j2}p_2$.

4.1 Experiments

We apply the model and techniques from Section 3 to a data set containing 1366 customers. Also, we create one random sample out of this data set containing 500 customers. In the application of the first algorithm, we use the current price regulations imposed by the European Commission as described in the first paragraph of this section. The costs for calling from abroad (c_1) is also retrieved from the table, and the cost for the operator for a customer to receive a call, is half of this. We forbid the operator to violate this price regulation by law (LAW); i.e., a penalty $\tau_1 = \infty$. Second, we keep the price regulative parameters α_k as they are, but now we find a tax level $\tau_1 \geq 0$ which maximizes social welfare (TAX). Note that there can be a range of feasible tax levels achieving the maximal social welfare. This effect was also observed in our results. Third, we compute the optimal social welfare by optimizing over the price regulative parameters α_k (OPT).

Table 1 shows the total social welfare for all instances. A complete overview of the results is deferred to the full version of this paper. We summarize our conclusions from these results as follows. Introducing a tax (TAX) instead of

Table 1. Social welfare obtained using different algorithms with price regulations on making and receiving calls abroad

	Year	LAW	TAX	OPT
Complete data set $n = 1366$	2007	2834.20	3315.77	3315.77
	2008	2797.51	3443.59	3443.59
	2009	2809.15	3571.40	3571.40
Sample $n = 500$	2007	878.69	1030.63	1038.66
	2008	887.15	1054.92	1059.21
	2009	895.02	1074.73	1079.76

enforcing the price regulation by law (LAW) leads to an increase in the social welfare. This suggests that a more liberal price regulative policy might have the potential to improve social welfare. The tax levels that are obtained using the algorithm are non trivial. In extreme cases (not observed here, though) a producer might be able to participate in the market where it would not be profitable to do so if violation of price regulations was forbidden. Not surprisingly, the social welfare is maximal when the α -parameters are optimized (OPT). In the sample, the social welfare is strictly larger in the latter case. For the complete data set, both algorithms yield the same social welfare. Concluding, it seems that the current EU practice does not yield the optimum, at least not with the data set and experimental setup that we use here.

5 Conclusion

First, we think it is an interesting result in its own that the given three-level program can indeed be solved in polynomial time by making use of linear programming techniques. Even though techniques are comparably simple and crucially use the fact that the dimension m is constant, we believe it is not straightforward to come up with a polynomial time algorithm.

Second, on the more economic side, our computational results suggest that a more liberal price regulative policy, namely taxation instead of regulation by law may lead to an increase in social welfare. But of course, this conclusion cannot be made hard as it depends very much on the choice of the social welfare function.

References

1. European commission, information society, media directorate general: EU roaming regulation: Using your mobile abroad in the EU is now much cheaper (factsheet 59) (June 2007)
2. Aggarwal, G., Feder, T., Motwani, R., Zhu, A.: Algorithms for multi-product pricing. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 72–83. Springer, Heidelberg (2004)
3. Balcan, M.F., Blum, A.: Approximation algorithms and online mechanisms for item pricing. In: Proceedings of the 7th ACM Conference on Electronic Commerce, pp. 29–35. ACM, New York (2006)

4. Baz, J., Chacko, G.: Financial derivatives: pricing, applications, and mathematics. Cambridge University Press, Cambridge (2004)
5. Bazaraa, M.S., Sherali, H.D., Shetty, C.M.: Nonlinear programming: Theory and algorithms, 2nd edn. John Wiley & Sons, Inc., Chichester (1993)
6. Bernoulli, D.: Specimen theoriae novae de mensura sortis (Translated from Latin into English by L. Sommer in 1954). *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, vol. 1738, pp. 175–192
7. Briest, P., Krysta, P.: Single-minded unlimited supply pricing on sparse instances. In: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1093–1102, ACM-SIAM (2006)
8. Demaine, E.D., Feige, U., Hajiaghayi, M.T., Salavatipour, M.R.: Combination can be hard: Approximability of the unique coverage problem. In: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 162–171, ACM-SIAM (2006)
9. Elbassioni, K., Sitters, R., Zhang, Y.: A quasi-ptas for profit-maximizing pricing on line graphs. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) *ESA 2007*. LNCS, vol. 4698, pp. 451–462. Springer, Heidelberg (2007)
10. Grigoriev, A., van Loon, J., Sitters, R., Uetz, M.: Optimal pricing of capacitated networks. *Networks* (to appear, 2008)
11. Grigoriev, A., van Loon, J., Sviridenko, M., Uetz, M., Vredeveld, T.: Bundle pricing with comparable items. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) *ESA 2007*. LNCS, vol. 4698, pp. 475–486. Springer, Heidelberg (2007)
12. Grigoriev, A., van Loon, J., Sviridenko, M., Uetz, M., Vredeveld, T.: Optimal bundle pricing with monotonicity constraint. *Operations Research Letters* 36, 609–614 (2008)
13. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1164–1173, ACM-SIAM (2005)
14. Hartline, J.D., Koltun, V.: Near-optimal pricing in near-linear time. In: Dehne, F., López-Ortiz, A., Sack, J.-R. (eds.) *WADS 2005*. LNCS, vol. 3608, pp. 422–431. Springer, Heidelberg (2005)
15. Marcotte, P., Savard, G.: ch. 7 - Bilevel Programming: A Combinatorial Perspective. In: *Graph theory and combinatorial optimization*, pp. 191–218. Springer, Heidelberg (2005)
16. Savage, L.J.: *The foundations of statistics*, 2nd edn. Courier Dover Publications (1972)
17. Shavel, S.: *Foundations of economic analysis of law*. Harvard University Press (2004)

Improving the Efficiency of Load Balancing Games through Taxes^{*}

Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos

Research Academic Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece

Abstract. In load balancing games, there is a set of available servers and a set of clients; each client wishes to run her job on some server. Clients are selfish and each of them selects a server that, given an assignment of the other clients to servers, minimizes the latency she experiences with no regard to the global optimum. In order to mitigate the effect of selfishness on the efficiency, we assign taxes to the servers. In this way, we obtain a new game where each client aims to minimize the sum of the latency she experiences and the tax she pays. Our objective is to find taxes so that the worst equilibrium of the new game is as efficient as possible. We present new results concerning the impact of taxes on the efficiency of equilibria, with respect to the total latency of all clients and the maximum latency (makespan).

1 Introduction

Load balancing games are special cases of the well-known *congestion games* introduced by Rosenthal [18]. A congestion game Π consists of a set E of resources, each resource e having a non-negative and non-decreasing latency function f_e defined over non-negative numbers, and a set of n players. Each player i has a weight (or demand) w_i and can select among a set of permissible strategies $S_i \subseteq 2^E$ (where each strategy of player i is a set of resources). In general, players may follow *mixed strategies*, i.e., use a probability distribution over their permissible strategies. An assignment $A = (A_1, \dots, A_n)$ is a vector of strategies, one (possibly mixed) strategy for each player. We mostly refer to pure assignments, i.e., assignments where each player selects a single strategy with probability 1. The cost of a player i at an assignment A is defined as $cost_i(A) = \sum_{e \in A_i} f_e(n_e(A))$, where $n_e(A)$ is the total weight of players using resource e in A . The *social cost* of an assignment can be either the weighted total cost over all players or the maximum latency (makespan) over all resources. A pure (resp., mixed) assignment is a *pure* (resp., *mixed*) *Nash equilibrium* if no player has any incentive to unilaterally deviate to another strategy,

^{*} This work was partially supported by the European Union under IST FET Integrated Project 015964 AEOLUS and by a “Caratheodory” grant from the University of Patras.

i.e., $cost_i(A) \leq cost_i(A_{-i}, s)$ for any player i and for any pure (resp., mixed) strategy s , where (A_{-i}, s) is the assignment obtained if just player i deviates from A_i to s . In *linear congestion games*, the latency function of resource e is of the form $f_e(x) = \alpha_e x + \beta_e$ with non-negative constants α_e and β_e . We use the terms *weighted* and *unweighted* to distinguish between the cases where the clients have different or identical weights.

Load balancing games are congestion games where the strategies of players are singleton sets. In load balancing terminology, we use the terms *server* and *client* instead of the terms resource and player. The set of strategies of a client contains the servers that are permissible for the client. A load balancing game is called *symmetric* when all servers are permissible for any client. Usually, the servers of load balancing games have linear latency functions; an important special case is that of *related servers* where the latency function of server j is of the form $f_j(x) = \alpha_j x$, with $\alpha_j > 0$. Motivated by [8], we use the term *graph balancing games* to denote asymmetric load balancing games where each client is unweighted and has at most two permissible servers, and all servers have identical linear latency functions.

Since players act selfishly, load balancing games may reach assignments that do not minimize the social cost. We use the notion of the *price of anarchy* introduced in [15,17] to quantify the degradation of the overall system performance. In particular, the price of anarchy of a game Π is the maximum over all pure (or mixed) Nash equilibria of the ratio of the social cost of a pure (or mixed) Nash equilibrium over the social cost of the optimal assignment.

A vast amount of the literature (see [20,25] and the references therein) studies the complexity of computing equilibria of best and worst social cost and provides bounds on the price of anarchy for various games that can be thought of as special cases of congestion games such as load balancing games, when the social cost is defined as the makespan or the weighed total latency. Awerbuch et al. [1] and, independently, Christodoulou and Koutsoupias [4] prove tight bounds on the price of anarchy of congestion games with respect to the weighted total latency. Among other results concerning polynomial latency functions, they show that the price of anarchy of pure Nash equilibria in unweighted linear congestion games is 5/2 while for mixed Nash equilibria or pure Nash equilibria of weighted players it is 2.618. These bounds carry over to load balancing games [2] and can be improved for interesting special cases [2,16,22]. The price of anarchy of weighted load balancing games on m related servers is $\Theta(\frac{\log m}{\log \log \log m})$ [7] over mixed Nash equilibria with respect to the makespan. A better tight bound of $\Theta(\frac{\log m}{\log \log m})$ is known for pure Nash equilibria as well as for mixed Nash equilibria at identical servers [7,14].

In order to downscale the effect of selfishness to performance, we assign taxes to the servers. Formally, a *tax function* $\delta : E \times Q^+ \rightarrow Q^+$ assigns a tax $\delta_j(w)$ to each client of weight w that wishes to use server $j \in E$. Furthermore, we assume that clients are not equally sensitive to taxes. In particular, client i has a *tax sensitivity* $\gamma_i > 0$. Assuming selfish behavior of the clients, we obtain a new *extended game* (Π, δ) where each client now aims to minimize the expected

latency she experiences plus her disutility due to the taxes she pays at the server she uses. This disutility equals $\gamma_i \delta_j(w_i)$ when client i selects server j . Again, an assignment y is a pure Nash equilibrium for the extended game if no player has an incentive to unilaterally change her strategy, i.e., $\text{cost}_i(y) + \gamma_i \delta_s(w_i) \leq \text{cost}_i(y_{-i}, s^*) + \gamma_i \delta_{s^*}(w_i)$ for any client i that is on server s under the assignment y , where y_{-i}, s^* is the assignment produced when client i moves from s to s^* .

Like in our previous work [3] on the topic and motivated by [6], we consider both *refundable* and *non-refundable* taxes. In the former case, we assume that the collected taxes can be feasibly returned (directly or indirectly) to the players (e.g., as a “lump-sum refund”) and therefore the overall system disutility depends only on the social cost. However, refunding the collected taxes could be logistically or economically infeasible; the latter case models this scenario. We will say that a function $\delta : E \times Q^+ \rightarrow Q^+$ is a ρ -pure-efficient refundable tax for the load balancing game Π if the social cost for any pure Nash equilibrium of the extended game is at most ρ times the social cost of the optimal assignment. Similarly, a function $\delta : E \times Q^+ \rightarrow Q^+$ is a ρ -pure-efficient non-refundable tax for the load balancing game Π if the social cost plus the total disutility due to taxes at any pure Nash equilibrium is at most ρ times the social cost of the optimal assignment. Similar definitions apply to the case of mixed Nash equilibria.

The problem of computing optimal taxes has received significant attention in the economics and transportation science literature; the main underlying model in these studies is that of *non-atomic* congestion games [21]. These games differ from the atomic games that we consider in that each player controls an infinitesimal amount of demand, and, therefore, the actions of a single player cannot affect the overall system performance. The results about taxes in non-atomic congestion games (see for example [5,6,10,13]) do not carry over to the atomic model. In our previous work [3], we presented (among several negative and positive results on the influence of taxes, under the assumption that all clients are equally sensitive to taxes) 2-mixed-efficient refundable taxes with respect to the weighted total latency for linear atomic congestion games and a pure-optimal tax function for symmetric load balancing games; this latter result was extended by Fotakis and Spirakis [12] to also hold for network congestion games on series-parallel graphs. Swamy [23] studied more general (e.g., polynomial) latency functions for the case of atomic congestion games with splittable demands and presented taxes that ensure that the optimal assignment is a pure Nash equilibrium.

In this paper we show the following results concerning non-refundable taxes. For the case of graph balancing games and unweighted clients with different sensitivities, we present an 1.618-pure-efficient tax function. This is the first class of asymmetric load balancing games for which an upper bound better than 2 is achieved, while we note that the lower bound of 11/10 presented in [3] also holds for graph balancing games. Recall that the price of anarchy of these games can be at least 2.012 [22]. Our tax function exploits the structure of graph balancing games and also uses the optimal assignment which can be computed in polynomial time. Then, we consider symmetric load balancing games with unweighted clients and servers with polynomial latency functions of degree p .

We prove a negative result that no non-refundable tax function can be better than $\frac{(p+1)^{1+1/p}}{(p+1)^{1+1/p}-p}$ -pure-efficient, i.e., $O\left(\frac{p}{\ln p}\right)$ -pure-efficient. Note that this lower bound matches the known upper bound on the price of anarchy of these games which is a corollary of the relation to symmetric non-atomic congestion games in [11] and the upper bounds of [19]. Next, we focus on the makespan as the social cost. For the case of pure Nash equilibria and weighted clients on m related servers, we present a 2-pure-efficient tax function, greatly improving upon the $\Theta\left(\frac{\log m}{\log \log m}\right)$ bound on the price of anarchy presented in [7]. The tax function is defined using a particular fractional schedule of clients to servers. We also present a lower bound that shows that this tax function is best possible. Finally, for mixed Nash equilibria we observe that the introduction of taxes does not mitigate significantly the impact of selfishness, since no better than $O\left(\frac{\log m}{\log \log m}\right)$ -mixed-efficient taxes exist, even for games with unweighted clients on identical servers.

The rest of the paper is structured as follows. We begin by presenting, in Section 2, our result concerning graph balancing games. We continue in Section 3 with the negative result about non-refundable taxes in symmetric load balancing games with polynomial latency functions. The results concerning the objective of minimizing the makespan are presented in Section 4.

2 Efficient Taxes for Graph Balancing Games

In this section, we present 1.618-pure-efficient tax functions for graph balancing games. This is the first subclass of asymmetric load balancing games which is proved to have better than 2-pure efficient taxes. The tax function is simple and exploits the structure of the game. We will assign very small taxes to the servers so that each server is assigned a different tax. So, although we prove the result ignoring the total taxes paid by the clients, this quantity can become arbitrarily small and our result carries over to non-refundable taxes by adding an ϵ factor to the efficiency.

Consider a graph balancing game with a set of clients \mathcal{U} (with $|\mathcal{U}| = n$) and let $\hat{\delta}$ be such that $0 < \hat{\delta} \leq 1/\max_i \gamma_i$, where γ_i is the tax sensitivity of client i . First, we compute an optimal assignment and denote by o_j the number of clients using server j in this assignment. This computation can be done in polynomial time by a natural reduction to a minimum cost flow problem on a single-source network, similar to the reduction presented in [9] for computing an equilibrium of symmetric congestion games that minimizes Rosenthal’s potential. Then, we consider the graph having a node for each server j and an edge between two different nodes j_1 and j_2 for each client that has servers j_1 and j_2 as permissible servers. For each such edge corresponding to a client i , we define the edge’s *optimal node* to be the endpoint corresponding to the server that client i uses in the optimal assignment. We compute an orientation of the edges so that the corresponding directed graph is acyclic. Then, either this directed graph or the one in which all edges have opposite directions have the following property: at most half of the edges point to their non-optimal node. We select the orientation

that has this property and assign different taxes from $\left\{ \frac{1}{n}\hat{\delta}, \frac{2}{n}\hat{\delta}, \dots, \frac{n-1}{n}\hat{\delta}, \hat{\delta} \right\}$ to the nodes/servers so that for any edge directed from j_1 to j_2 , it is $\delta_{j_1} > \delta_{j_2}$.

Now, consider any pure Nash equilibrium of the extended game and let n_j denote the number of clients using server j . For a client i , denote by j_1 the server that client i uses in the pure Nash equilibrium and let j_2 be the other permissible server of client i ($j_1 = j_2$ if the client has only one permissible server). Since no client has an incentive to change her strategy, it is $n_{j_1} + \gamma_i \delta_{j_1} \leq n_{j_2} + 1 + \gamma_i \delta_{j_2}$. This means that $n_{j_1} \leq n_{j_2} + 1 + \gamma_i (\delta_{j_2} - \delta_{j_1})$ and $n_{j_1} \leq n_{j_2} + 1$ since n_{j_1} and n_{j_2} are integers and $\gamma_i (\delta_{j_2} - \delta_{j_1}) < 1$ by the definition of the tax function. This inequality holds for any client and we conclude that any pure Nash equilibrium of the extended game is also a pure Nash equilibrium for the original game.

We now show that any pure Nash equilibrium for the extended game is a $\frac{1}{2}$ -PNE for the original game, i.e., a pure Nash equilibrium that satisfies the property

$$\sum_j n_j^2 \leq \sum_j \left(n_j o_j + \frac{o_j}{2} \right). \tag{1}$$

For each client i , we denote by j_i and j'_i the servers she uses in the pure Nash equilibrium and in the optimal assignment, respectively. Denote by S the set of clients i such that $j_i = j'_i$. Denote by F the set of clients i such that $j_i \neq j'_i$ and $\delta_{j_i} > \delta_{j'_i}$. Then, the condition $n_{j_i} + \gamma_i \delta_{j_i} \leq n_{j'_i} + 1 + \gamma_i \delta_{j'_i}$, implying that client i has no incentive to use the server she uses in the optimal assignment, implies that $n_{j_i} \leq n_{j'_i}$, since n_{j_i} and $n_{j'_i}$ are integers and $\delta_{j_i} > \delta_{j'_i}$. The condition $n_{j_i} \leq n_{j'_i} + 1$ holds for any client i not belonging in S and F since the pure Nash equilibrium for the extended game is also a pure Nash equilibrium for the original game.

By the definition of the tax function, we have that $|\mathcal{U} \setminus (F \cup S)| \leq |\mathcal{U}|/2 = \frac{1}{2} \sum_j o_j$. By considering the equilibrium conditions for all clients, we have

$$\begin{aligned} \sum_j n_j^2 &= \sum_{i=1}^n \sum_{j:j_i=j} n_j \\ &= \sum_{i \in S} \sum_{j:j_i=j} n_j + \sum_{i \in F} \sum_{j:j_i=j} n_j + \sum_{i \in \mathcal{U} \setminus (F \cup S)} \sum_{j:j_i=j} n_j \\ &\leq \sum_{i \in S} \sum_{j:j'_i=j} n_j + \sum_{i \in F} \sum_{j:j'_i=j} n_j + \sum_{i \in \mathcal{U} \setminus (F \cup S)} \sum_{j:j'_i=j} (n_j + 1) \\ &= \sum_{i=1}^n \sum_{j:j'_i=j} n_j + |\mathcal{U} \setminus (F \cup S)| \\ &\leq \sum_j n_j o_j + \frac{1}{2} \sum_j o_j \\ &= \sum_j \left(n_j o_j + \frac{o_j}{2} \right). \end{aligned}$$

This completes the proof of inequality \square .

In our analysis, we will consider all $\frac{1}{2}$ -PNE for the original graph balancing game, and we will show that their price of anarchy is at most $\frac{1+\sqrt{5}}{2}$. We will need the following technical claim.

Lemma 1. *For any non-negative integers x and y ,*

$$\frac{3 - \sqrt{5}}{4}x^2 + \frac{3 + \sqrt{5}}{4}y^2 \geq xy + \frac{3(\sqrt{5} - 1)}{4}y - \frac{3\sqrt{5} - 5}{4}x.$$

Theorem 1. *For any graph balancing game, the tax function described above is a $\frac{1+\sqrt{5}}{2} \approx 1.618$ -pure-efficient tax.*

Proof. We will show that the price of anarchy of any $\frac{1}{2}$ -PNE of a graph balancing game is at most $\frac{1+\sqrt{5}}{2}$. Again, we denote by n_j and o_j the number of clients in server j in the $\frac{1}{2}$ -PNE and in the optimal assignment, respectively. By inequality [\(II\)](#) and since $\sum_j n_j = \sum_j o_j$, we have that the social cost is

$$\begin{aligned} \sum_j n_j^2 &\leq \sum_j \left(n_j o_j + \frac{o_j}{2} \right) \\ &= \sum_j \left(n_j o_j + \frac{o_j}{2} \right) + \frac{3\sqrt{5} - 5}{4} \sum_j o_j - \frac{3\sqrt{5} - 5}{4} \sum_j n_j \\ &= \sum_j \left(n_j o_j + \frac{3(\sqrt{5} - 1)}{4} o_j - \frac{3\sqrt{5} - 5}{4} n_j \right) \\ &\leq \sum_j \left(\frac{3 - \sqrt{5}}{4} n_j^2 + \frac{3 + \sqrt{5}}{4} o_j^2 \right) \\ &= \frac{3 - \sqrt{5}}{4} \sum_j n_j^2 + \frac{3 + \sqrt{5}}{4} \sum_j o_j^2 \end{aligned}$$

where the first equality follows since $\sum_j n_j = \sum_j o_j$, the second and third equalities are obvious, and the second inequality follows by Lemma [II](#). We obtain that the price of anarchy is

$$\frac{\sum_j n_j^2}{\sum_j o_j^2} \leq \frac{\frac{3+\sqrt{5}}{4}}{1 - \frac{3-\sqrt{5}}{4}} = \frac{1 + \sqrt{5}}{2}. \quad \square$$

Broadening the class of load balancing games that admit better than 2-pure-efficient taxes (or even pure-optimal taxes) is an interesting open problem.

3 Non-refundable Taxes in Symmetric Load Balancing

We now proceed to answer in a negative way a question posed in [\[3\]](#) concerning non-refundable taxes in symmetric load balancing games, i.e., whether taxes can diminish the effect of selfishness. Our following theorem suggests that taxes do not help in the case of symmetric load balancing with polynomial latency

functions of degree p , since for any tax function, the price of anarchy of the extended game in these games is $\frac{(p+1)^{1+1/p}}{(p+1)^{1+1/p-p}} \in O\left(\frac{p}{\ln p}\right)$. Clearly, our result also demonstrates that the known upper bound on the price of anarchy of such games (without taxes) is tight.

Theorem 2. *For any $p \geq 1$ and any $\epsilon > 0$, there exists a symmetric load balancing game with polynomial latency functions of degree p that does not admit better than $(\rho - \epsilon)$ -pure-efficient non-refundable taxes where $\rho = \frac{(p+1)^{1+1/p}}{(p+1)^{1+1/p-p}} \in O\left(\frac{p}{\ln p}\right)$.*

Proof. Let $k \geq 2$ be an integer and define $\lambda = \frac{2k^{p+1} - k^p - (k-1)^{p+1}}{k}$. We have $\lambda = k^p + k^p \left(\left(1 - \frac{1}{k}\right) - \left(1 - \frac{1}{k}\right)^{p+1} \right)$ and, since $k \geq 2$ and $p \geq 1$, it is $k^p < \lambda < 2k^p$. Define $y^* = k - \left\lfloor \left(\frac{\lambda}{p+1}\right)^{1/p} \right\rfloor$. Since $p \geq 1$ and $\lambda < 2k^p$, it is $1 \leq y^* \leq k$.

Consider a game with k clients where each client j has $\gamma_j = 1$, and $k + 1$ servers $0, 1, \dots, k$. Server 0 has latency function x^p while each of the other k servers has latency function λx^p . The assignment in which server 0 has $k - y^*$ clients, y^* among the other servers have exactly one client and any other server is empty has cost

$$opt = (k - y^*)^{p+1} + y^* \lambda.$$

In the absence of taxes, the assignment where all clients select server 0 is a pure Nash equilibrium since each of them has a latency of k^p and, in case a client decides to choose another server, she would face latency $\lambda > k^p$. The cost of this equilibrium is $cost = k^{p+1}$ and the price of anarchy is

$$PoA \geq \frac{cost}{opt} = \frac{k^{p+1}}{(k - y^*)^{p+1} + y^* \lambda}.$$

Therefore, in order to avoid this assignment as an equilibrium of the extended game, we have to assign taxes in such a way that at least one client has an incentive to change her choice. So, without loss of generality, we assume that there is a tax function δ , for which it holds that $\delta_0(w) = \alpha$ and $\delta_j(w) = 0$, for any $1 \leq j \leq k$. Note that, for any $\alpha \leq \lambda - k^p$, the aforementioned assignment remains a pure Nash equilibrium of the extended game, since any client at server 0 would have a cost of $k^p + \alpha \leq \lambda$. Now, assume that $\alpha = \lambda - k^p + \epsilon$ for any $\epsilon > 0$. Then, any client would have a incentive to leave server 0 and move to another server. Then, assuming that one client moves, the total cost $cost'$ (latency plus taxes) of the resulting assignment would be

$$\begin{aligned} cost' &= (k - 1)^{p+1} + \alpha(k - 1) + \lambda \\ &> (k - 1)^{p+1} + (\lambda - k^p)(k - 1) + \lambda \\ &= (k - 1)^{p+1} + \lambda(k - 1) - k^p(k - 1) + \lambda \\ &= \lambda k + (k - 1)^{p+1} - k^{p+1} + k^p \\ &= k^{p+1} \\ &= cost. \end{aligned}$$

Applying similar reasoning, it is not hard to see that by increasing $\delta_0(w) = \alpha$ even more so that more clients have an incentive to leave server 0, the total cost similarly increases. Therefore, the total cost is minimized by setting $\alpha = 0$.

Observe that $\lim_{k \rightarrow \infty} \frac{\lambda}{k^p} = 1$ and $\lim_{k \rightarrow \infty} \frac{y^*}{k} = 1 - \left(\frac{1}{p+1}\right)^{1/p}$. Hence,

$$\begin{aligned} \lim_{k \rightarrow \infty} \frac{k^{p+1}}{(k - y^*)^{p+1} + y^* \lambda} &= \lim_{k \rightarrow \infty} \frac{1}{\left(1 - \frac{y^*}{k}\right)^{p+1} + \frac{y^* \lambda}{k^{p+1}}} \\ &= \frac{1}{\left(\frac{1}{p+1}\right)^{1+1/p} + 1 - \left(\frac{1}{p+1}\right)^{1/p}} \\ &= \frac{(p+1)^{1+1/p}}{(p+1)^{1+1/p} - p} \\ &= \rho. \end{aligned}$$

Hence, for any $\epsilon > 0$, by setting k to a sufficiently large value, we obtain that the price of anarchy becomes at least $\rho - \epsilon$. □

4 Minimizing the Makespan

In this section we focus on the makespan as the social cost. We consider the well-known case where servers are related, i.e., server j has latency function $\alpha_j x$. Our upper bound uses tax functions that assign to each server a tax of either 0 or ∞ . In this setting, there is no difference between refundable and non-refundable taxes since no client is assigned to a server where it has to pay an infinite tax. Furthermore, the tax sensitivity of each client does not affect her behavior.

Denote by n the number of clients and by m the number of servers. We assume that the servers are sorted in non-decreasing order of α_i (i.e., $\alpha_i \leq \alpha_{i+1}$) and clients are sorted in non-increasing order of their weight (i.e., $w_i \geq w_{i+1}$). We define the following procedure that produces fractional schedules of makespan $T \geq \frac{\sum_i w_i}{\sum_i 1/\alpha_i}$. Observe that the quantity $\frac{\sum_i w_i}{\sum_i 1/\alpha_i}$ is a lower bound on the makespan of any fractional schedule; the numerator is the total weight of the clients and the denominator is the ‘‘capacity’’ of all servers.

1. set $j = 1, i = 1$, and $t = 0$;
2. while $i \leq n$ do
3. if $T - t \geq \alpha_j w_i$ then
4. put the remaining weight of client i at server j ;
5. set $t = t + \alpha_j w_i$ and $i = i + 1$;
6. else
7. put weight $\frac{T-t}{\alpha_j}$ of client i at server j ;
8. set $w_i = w_i - \frac{T-t}{\alpha_j}, j = j + 1$, and $t = 0$;

What the above procedure is doing is to consider each client (according to their ordering) and put as much of her weight as possible to the server of smallest

index so that the latency does not exceed T . This will end up with a fractional schedule in which there exists a server j' such that the latency of all servers $j \leq j'$ is exactly T , (if $j' < m$) the latency of server $j' + 1$ is at most T , and the latency of all servers $j > j' + 1$ (if any) is 0. Each client occupies consecutive servers and, furthermore, at most one client may have non-zero weights in two specific consecutive servers.

Given a value of T , the schedule produced by the procedure above is called *2-feasible* if for any client i and any two consecutive servers j and $j + 1$, it holds that $\alpha_j w_i^j + \alpha_{j+1} w_i^{j+1} \leq T$, where w_i^j denotes the weight of client i assigned to server j . We start with the value $T = \frac{\sum_i w_i}{\sum_i 1/\alpha_i}$ and run the procedure. If the schedule produced is 2-feasible, we stop. Otherwise, we increase T until the schedule produced by the procedure is 2-feasible. Let T^* be the corresponding value of T (i.e., the minimum value for which the schedule produced is 2-feasible). Clearly, if $T^* > \frac{\sum_i w_i}{\sum_i 1/\alpha_i}$, there will be at least one client i and two consecutive servers j and $j + 1$ such that $\alpha_j w_i^j + \alpha_{j+1} w_i^{j+1} = T^*$.

We now describe the tax function. We partition the clients into groups according to their weight, so that two clients i_1 and i_2 belong to same group when $w_{i_1} = w_{i_2}$. We denote by w_g^* the weight corresponding to group g . Let S_g denote the set of servers that contain non-zero weights of clients belonging to group g in the fractional schedule of makespan T^* . If $|S_g| = 1$, then we set $\delta_j(w_g^*) = 0$ for the server $j \in S_g$ and $\delta_{j'}(w_g^*) = \infty$ for any other server $j' \notin S_g$. Otherwise, when $|S_g| > 1$, we distinguish between two cases depending on whether the last server of S_g (i.e., the one with the larger index) contains only clients of group g or also clients of different groups. In the first case, we set $\delta_j(w_g^*) = 0$ for any server $j \in S_g$ and $\delta_{j'}(w_g^*) = \infty$ for any other server $j' \notin S_g$, while in the second case, we set $\delta_j(w_g^*) = 0$ for the $|S_g| - 1$ servers of S_g with smallest index and $\delta_{j'}(w_g^*) = \infty$ for any other server j' . In any case, we denote with Δ_g the set of servers j for which $\delta_j(w_g^*) = 0$.

We show the following result.

Theorem 3. *For any symmetric load balancing game on related servers, the above tax function is 2-pure-efficient with respect to the makespan.*

Proof. Consider the 2-feasible fractional schedule of makespan T^* produced as above. We first show that the optimal assignment has makespan at least T^* . This clearly holds if $T^* = \frac{\sum_i w_i}{\sum_i 1/\alpha_i}$. Otherwise, there will be a client i and two consecutive servers j and $j + 1$ such that $\alpha_j w_i^j + \alpha_{j+1} w_i^{j+1} = T^*$. Then, all clients with smaller index than i are fractionally scheduled at servers $1, \dots, j$ which have latency exactly T^* . In any integral schedule, either all of the clients $1, \dots, i$ will be scheduled to servers $1, \dots, j$ or some of them will be scheduled at some server with larger index than j . In the first case, the makespan will be at least T^* since the total weight of clients assigned to servers $1, \dots, j$ does not decrease compared to the fractional schedule. In the second case, a client of weight at least w_i will be assigned to a server j' with $\alpha_{j'} \geq \alpha_{j+1} \geq \alpha_j$. This server will have latency at least $\alpha_{j'} w_i \geq \alpha_j w_i^j + \alpha_{j+1} w_i^{j+1} = T^*$.

Now, we show that there exists an integral schedule with makespan at most $2T^*$ in which each client in group g selects a server from the set Δ_g . The clients that have non-zero weight in the server of S_j with the smallest index are scheduled in this server. Each other client of group g for which the server with largest index containing a non-zero amount of her weight in the fractional schedule is j is scheduled at server $j - 1$. In this way, the total weight of any server j may increase by at most the weight of clients in server $j + 1$ in the fractional schedule. Since $\alpha_j \leq \alpha_{j+1}$, the latency at server j will not exceed $2T^*$.

Observe that the tax function essentially divides the original game into subgames in the following sense. In any pure Nash equilibrium, the clients of group g with $|\Delta_g| = 1$ are forced to use the server of Δ_g . The clients of group g with $|\Delta_g| > 1$ play a symmetric game with linear latency functions at server j of the form $\alpha_j x + \beta_j$. Here, β_j denotes the latency at server j due to clients not belonging to group g which are forced to use server j . Furthermore, by the definition of the tax function, the sets Δ_g with size more than 1 are disjoint and, hence, the corresponding sets of clients do not interfere. It is not hard to see that any equilibrium in each subgame of clients of group g has the minimum possible integral makespan, i.e., at most $2T^*$. This completes the proof of the theorem. \square

The next theorem states that this tax function is best possible. The proof is omitted; it will appear in the final version.

Theorem 4. *For any $\epsilon > 0$, there exists a load balancing game on m identical servers that does not admit better than $(2 - \epsilon)$ -pure-efficient taxes with respect to the makespan.*

Unfortunately, taxes cannot significantly improve the price of anarchy with respect to the makespan over mixed Nash equilibria. To show this, we use a construction that we have also used in [3] to lower-bound the efficiency of taxes at mixed Nash equilibria with respect to the total latency. The construction applies to symmetric load balancing games with identical clients and identical servers and the proof follows by a standard balls-to-bins argument.

Consider a tax function δ . Without loss of generality, we assume that $\delta_j \leq \delta_{j'}$ for $j < j'$. Let k be equal to m if $1 + \frac{\sum_{i=1}^{m-1} \delta_i}{m-1} > \delta_m$, otherwise k is equal to the largest integer such that $\frac{m-1 + \sum_{j=1}^k \delta_j}{k} \leq \delta_{k+1}$. Let $D = \sum_{j=1}^k \delta_j$. Consider the following assignment y for all clients. Client i uses server j with probability $y_{ij} = \frac{1}{k} + \frac{D}{k(m-1)} - \frac{\delta_j}{m-1}$ if $j \leq k$ and $y_{ij} = 0$ otherwise. Notice that all clients have the same probability distribution. It can be verified that y is a mixed Nash equilibrium of the extended game.

In order to compute the expected makespan, it suffices to observe that it is the expectation of the maximum number of balls at any bin when m balls are thrown independently at m bins according to the probability distribution y . It is well-known (e.g., see [24]) that this expectation is minimized to $\Theta\left(\frac{\log m}{\log \log m}\right)$ when y is the uniform distribution. Thus, we obtain the following statement.

Theorem 5. *There exists a symmetric load balancing game on m servers that does not admit better than $\Omega\left(\frac{\ln m}{\ln \ln m}\right)$ -mixed-efficient taxes with respect to the makespan.*

Note that this bound matches the price of anarchy of symmetric load balancing with identical servers [7,14]. The price of anarchy for related servers is slightly higher [7]. We leave as an open problem whether taxes can improve the price of anarchy with respect to the makespan in this particular case and, more importantly, in the more general case of congestion games.

References

1. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005), pp. 57–66 (2005)
2. Caragiannis, I., Flammini, M., Kaklamanis, C., Kanellopoulos, P., Moscardelli, L.: Tight bounds for selfish and greedy load balancing. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 311–322. Springer, Heidelberg (2006)
3. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Taxes for linear atomic congestion games. In: Azar, Y., Erlebach, T. (eds.) ESA 2006. LNCS, vol. 4168, pp. 184–195. Springer, Heidelberg (2006)
4. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005), pp. 67–73 (2005)
5. Cole, R., Dodis, Y., Roughgarden, T.: Pricing network edges for heterogeneous selfish users. In: Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC 2003), pp. 521–530 (2003)
6. Cole, R., Dodis, Y., Roughgarden, T.: How much can taxes help selfish routing? *Journal of Computer and System Sciences* 72(3), 444–467 (2006)
7. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms* 3(1) (2007)
8. Ebenlendr, T., Krčal, M., Sgall, J.: Graph balancing: a special case of scheduling unrelated parallel machines. In: Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008), pp. 483–490 (2008)
9. Fabrikant, A., Papadimitriou, C., Talwar, K.: On the complexity of pure equilibria. In: Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC 2004), pp. 604–612 (2004)
10. Fleischer, L., Jain, K., Mahdian, M.: Tolls for heterogeneous selfish users in multi-commodity networks and generalized congestion games. In: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2004), pp. 277–285 (2004)
11. Fotakis, D.: Stackelberg strategies for atomic congestion games. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) ESA 2007. LNCS, vol. 4698, pp. 299–310. Springer, Heidelberg (2007)
12. Fotakis, D., Spirakis, P.: Cost-balancing tolls for atomic network congestion games. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 179–190. Springer, Heidelberg (2007)

13. Karakostas, G., Kolliopoulos, S.: Edge pricing of multicommodity networks for heterogeneous selfish users. In: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2004), pp. 268–276 (2004)
14. Koutsoupias, E., Mavronicolas, M., Spirakis, P.: Approximate equilibria and ball fusion. *Theory of Computing Systems* 36(6), 683–693 (2003)
15. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) *STACS 1999*. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
16. Lücking, T., Mavronicolas, M., Monien, B., Rode, M.: A new model for selfish routing. In: Diekert, V., Habib, M. (eds.) *STACS 2004*. LNCS, vol. 2996, pp. 547–558. Springer, Heidelberg (2004)
17. Papadimitriou, C.: Algorithms, games and the internet. In: Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001), pp. 749–753 (2001)
18. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
19. Roughgarden, T.: The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences* 67(2), 341–364 (2003)
20. Roughgarden, T.: Routing games. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
21. Roughgarden, T., Tardos, E.: How bad is selfish routing? *Journal of the ACM* 49(2), 236–259 (2002)
22. Suri, S., Tóth, C., Zhou, Y.: Selfish load balancing and atomic congestion games. *Algorithmica* 47(1), 79–96 (2007)
23. Swamy, C.: The effectiveness of Stackelberg strategies and tolls for network congestion games. In: Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007), pp. 1133–1142 (2007)
24. Vöcking, B.: How asymmetry helps load balancing. *Journal of the ACM* 50(4), 568–589 (2003)
25. Vöcking, B.: Selfish load balancing. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)

Network Formation and Routing by Strategic Agents Using Local Contracts

Elliot Anshelevich^{1,*} and Gordon Wilfong^{2,**}

¹ Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY

² Bell Labs, Murray Hill, NJ

Abstract. In the Internet, Autonomous Systems (ASes) make contracts called Service Level Agreements (SLAs) between each other to transit one another's traffic. ASes also try to control the routing of traffic to and from their networks in order to achieve efficient use of their infrastructure and to attempt to meet some level of quality of service globally. We introduce a game theoretic model in order to gain understanding of this interplay between network formation and routing. Player strategies allow them to make contracts with one another to forward traffic, and to re-route traffic that is currently routed through them. This model extends earlier work of [3] that only considered the network formation aspect of the problem. We study the structure and quality of Nash equilibria and quantify the prices of anarchy and stability, that is, the relative quality of a centralized optimal solution versus that of the Nash equilibria.

1 Introduction

The current Internet consists of tens of thousands of sub-networks known as *Autonomous Systems* (ASes), with each AS under the control of a single administrative authority. The task of this authority is to balance ensuring connectivity with the rest of the network against its own particular economic goals for managing the traffic entering and leaving the AS. These sometimes conflicting goals of connectivity versus traffic engineering are typically reflected in the local contracts called *Service Level Agreements* (SLAs) formed with neighboring ASes, which are essentially agreements to forward traffic. Without an appropriate SLA, no traffic would get from one AS to another, and so this system of business relationships is at the heart of Internet connectivity. The economic goals of an AS are also encoded into its routing policies, which are expressed in its local configuration of the Border Gateway Protocol (BGP) where BGP is the interdomain routing protocol in common use today. The local configuration of BGP determines the traffic into and out of the AS. These resulting traffic patterns may cause an AS to modify existing SLAs or create new ones. Thus there is an interplay between the contracts formed and the traffic patterns resulting from the encoding of the implied policies into

* This research supported in part by the NSF Postdoctoral Fellowship in Mathematical Sciences and by NSERC Discovery grant 342457-07.

** This research supported in part by NSERC Discovery grant 342457-07.

BGP. The major goal of this paper is to develop an understanding of this interplay between network formation and traffic routing in a network where contracts are formed only between immediate neighbors.

This paper builds on the earlier work of [3] in which a network formation game called the *Local Contract Formation Game* (LCFG) was defined. LCFG ignored the routing aspect of the resulting network, assuming that each traffic demand could only be sent along a given fixed route. In our paper, we allow routing choices as part of a player's strategy in addition to the choice of which contracts to form. We also incorporate edge (or node) capacity constraints into our model. While these new features add complexity to the model, they are an important step towards modeling real world constraints and degrees of freedom faced by entities in today's Internet.

Some aspects of LCFG as defined in [3] are retained in our model. For instance, both models have players representing typical Internet entities such as ASes, enterprise networks, and residential customers. Also, relationships between neighboring players are designed to capture three major properties of real-world business relationships seen in the Internet. The first of these properties is that relationships are strictly locally negotiated [14,15]. That is, players negotiate contracts only with their neighbors. While these contracts take into account global issues such as the topology of the rest of the network and traffic demands throughout the network, they are formed strictly by negotiations between two neighboring players. Another property is that two possible types of contracts can be formed: customer-provider or peer-peer [7,10]. Lastly, both models incorporate penalties such as those found in typical real world SLAs for failure to meet the conditions of a contract [13]. In particular, these penalties are designed to penalize a provider that accepts payment from a customer, but fails to satisfy the transit service for which the customer formed the contract.

Our Results. We focus on understanding the structure and quality of Nash equilibria, and on quantifying the prices of anarchy and stability¹, which compare the quality of a centrally-defined optimal solution with the quality of the worst and best Nash equilibrium, respectively. The detailed list of our results can be found in Section 3.

Related Work. The tradeoff between the benefits to a player (such as an AS) from routing its demands versus the loss of utility from transiting other traffic, was explored in several papers (e.g. [9,18,21]). Typically, however, the concentration has been on short-term routing and pricing schemes, e.g., [12,17,20]. While short-term pricing and admission models are of significant interest, they rely on an underlying set of business (longer-term economic) relationships. Moreover, agreements between entities (ISPs, enterprise or residential customers) tend to be based on more rigid contracts such as fixed bandwidth, or peak bandwidth contracts [14,15]. This is largely due to the complexity (and expense) of monitoring IP traffic at a packet or flow level.

¹ See [19] for a definition of price of anarchy (known there as the coordination ratio), and [2] for a definition of price of stability.

Motivated by this, several game theoretical models have addressed the strongly related notion of network formation. Some of these do not look at contract formation, but instead assume that edges have intrinsic costs [1,2,5,8]. On the other hand, contract formation models of networks have been heavily addressed as well, mostly in the economics literature [16]. This body of work mostly addresses questions distinct from those studied here. In particular, none of these combine routing with contract formation, consider customer-provider and peering contracts, or measure the impact of provider penalties. In [6] a very general model of network formation is considered, and the model of [18] focuses on solutions where all demands are satisfied, and on “pairwise-stable” equilibria.

2 Local Contract and Routing Formation Game

The Local Contract and Routing Formation Game (LCRFG) is the game we study in this paper. An instance of LCRFG consists of a directed capacitated graph $G = (V, E)$ and a set of demands D , both of which we describe below.

We represent the network under consideration as a *capacitated* directed graph $G = (V, E)$ where each (directed) edge $e \in E$ has a capacity on the amount of traffic that can be transited along e . Nodes in V are the players of the game, and represent entities that might wish to establish contracts with one another to forward their traffic. An edge represents that there is a physical connection between the nodes (e.g., a fiber connection) and hence the capacity of the edge represents the actual capacity of this connection. We can also include node capacities without changing our results. The physical connection between nodes implies that there is the potential for the nodes on that edge to form a contract to forward their traffic directly rather than through intermediate nodes.

A *demand* d represents the desire of a source node $s(d)$ to be able to have a unit of traffic routed to a destination node $t(d)$. A demand d also includes *values* $\lambda_{s(d)}(d)$ and $\lambda_{t(d)}(d)$ that represent the value that the source and destination nodes respectively place on having the demand traffic successfully routed. We describe below what it means for a demand to be successfully routed.

For each edge $e = (v, u)$ and a destination node t , u can *request* some amount, say c , from v . If v “accepts” this request, then v and u form a contract having the meaning that v pays u the amount c , and u is obligated to forward a unit of traffic flow from v to t . The amount c is called the *cost* of the contract. If v accepts the request then the contract is said to be an *active contract*. Let $X_e(t)$ be the set of all such active contracts for edge e and destination t . The capacity of an edge e is *violated* if $\sum_t |X_e(t)|$ is greater than the capacity of e . We use the notation $\chi_e(t)$ to denote an element of the set $X_e(t)$ and $c(\chi_e(t))$ to indicate the cost of the contract $\chi_e(t)$. Then we define $c(X_e(t)) = \sum_{\chi_e(t) \in X_e(t)} c(\chi_e(t))$. In the context of such a contract $\chi_e(t)$, node u is said to be the *provider* and v is the *customer*. We show that our results hold with other types of customer-provider contracts in Section 3, as well as with the introduction of peering contracts.

One aspect of a node’s strategy in an instance of LCRFG is then to decide which requests to make for contract payments, and which requests to accept.

Another aspect of a node's strategy is to determine how to route demands over active contracts. We do not limit our model to a single routing protocol, but instead leave open what exact routing protocol is being used, in order to be most general. The only assumption we make about routing decisions is that they satisfy the following constraints.

- (R1)** *The routing does not violate the contracts.* That is, if a demand d is routed from s to t , then the edges $e = (u, v)$ that are included in its path must have a contract $\chi_e(t)$ between u and v . Moreover, if x demands with destination t are forwarded from u to v , then there must be at least x contracts in $X_e(t)$ between u and v (since each contract $\chi_e(t)$ only obligates v to get a single demand to its destination t).
- (R2)** *If a valid route exists between a source and destination of a demand d , then demand d is successfully routed.* By a valid route, we mean a path of contracts $\chi_e(t(d))$ between $s(d)$ and $t(d)$, with the capacity provided by these contracts still available (not taken up by other demands). We call a successfully routed demand *active*.

The basic goal of a player v is to establish contracts with its neighbors so as to manage the trade-off between getting its demands (i.e., the demands having v as an endpoint) successfully routed versus allowing the demands of others to pass through it. The strategy of a player v consists of making (and accepting) requests for contract payments, and of choosing the routes for demands passing through v , as long as the routing constraints **(R1)** and **(R2)** are satisfied. The details of player strategies are addressed in Section 2.1. We now describe the utilities of the players, and the possible equilibria.

Utilities. The utility for a player v will include the payments v makes (or receives) for contracts and the values of active demands for which v is an endpoint. The utility will also involve a *transit cost* for active demands that include v in their path as well as certain *penalties*.

Transit Costs: Define $t(v)$ (the *cost of transiting*) for a node v to be the number of active demands routed through v .

Penalties: A provider must pay penalties to its customers if it fails to meet its obligations. There are many penalty systems that we could consider in this context. Just as with the routing protocol, instead of defining exactly what the penalty system is, we just give required conditions for it. Our results work for any penalty system meeting these conditions. Specifically, a penalty system must be such that:

- If a demand d is active, then no one pays penalties because of it.
- If a demand d that was previously active and routed through node v becomes inactive, but there is still a valid route from $s(d)$ to v , then v must pay a penalty to its customer on this valid route.

Notice that the second condition above assumes some notion of time passing. This assumption is unnecessary, as we can define many static penalty systems

that obey the above conditions. For example, we can say that all providers with valid routes from $s(d)$ pay penalties for an inactive demand. Or we could make a more complicated system and say that inactive demands still have routes assigned to them, and the last provider in such a route that fails to forward this demand to the next hop must pay the penalty. In either case, the conditions above are satisfied, since if a provider decided to prevent a demand from getting to its destination (without re-routing or changing contracts), it would have to pay a penalty to its customer for failing in its provider duties. We set the size of the penalty for failing to form a route for demand d to be $\lambda_{s(d)}(d) - 1$. For justification of our penalty system, and the intuition behind our definition of penalties, see [3].

The Utility Function: Given the endpoint values for active demands, the transiting costs, payments, and the penalties, the utility of node v is:

$$utility(v) = \sum_{d \in D^{end}(v)} \lambda_v(d) - t(v) + \sum_{e=(w,v),t} c(X_e(t)) - \sum_{e=(v,w),t} c(X_e(t)) - \sum_d p_v(d)$$

In other words, a node v gains the value of $\lambda_v(d)$ for each active demand that it originates (this set of demands is denoted by $D^{end}(v)$), loses 1 for every demand it transits, gets payment $c(X_e(t))$ according to the contract it makes with its neighbor on e (either positive if v is paid or negative if it pays) and loses $p_v(d)$ for penalties (either positive or negative depending on whether it pays or receives the penalty).

The utility function defined above applies to the case when for each edge e incident to a node v , the number of active contracts $\sum_t |X_e(t)|$ does not violate the capacity constraint of e . The second case is when the number of active contracts $\sum_t |X_e(t)|$ for some e incident to v violates the capacity of e . In this case, v is assigned a negative utility. It is easy to see that in a Nash equilibrium no node is involved with contracts that violate capacities, since they could always do better by doing nothing and receiving zero utility.

2.1 Player Strategies and Nash Equilibria

As mentioned above, a strategy of a player v consists of making (and accepting) requests for contract payments, and of possibly affecting the routes for demands passing through v , as long as the routing constraints (R1) and (R2) are satisfied. Many concrete games can be defined that fit into this framework (see [4] for some examples). Instead of focusing on only one such game, however, we choose to leave our framework as general as possible. In fact, our results hold for *any* game where a player v only has power over its contracts and over the demand routes passing through v . As long as player v cannot affect anything else by changing its strategy, we will show the existence of good Nash equilibria. To carefully define an equilibrium, we do not need to define precisely the strategies of the game, but must define the set of possible *player deviations*, with a solution being an equilibrium exactly when no player would want to take one of their possible deviations.

Given a solution \mathcal{S} (i.e., active contracts, their costs, and a routing), a *deviation* for a node v is a solution where v changes its actions while all others remain as in \mathcal{S} . We will see later that the only deviations of interest will not result in the ability to activate currently inactive demands. Specifically, a node v can deviate by performing the following actions, possibly both at once.

- Changing the costs of its requests for customer contracts, and changing its decisions on the acceptance/ non-acceptance of provider contracts.
- Re-routing any of the active demands that pass through v .

This means that v controls all of the contracts it makes, and all of the routes that pass through it. Just as our other assumptions about the routing protocol, this is a somewhat general assumption, and many routing protocols obey this condition. In the most powerful deviation that we allow, a node would be able to re-route by canceling all routes that pass through it (and thereby freeing their capacity), and then forming new routes for these demands instead. Our results will hold for any weaker deviations as well, such as ones where v can only change some of the routes that pass through it, instead of all of them.

3 Good Equilibria and Price of Stability

Definite Nash Equilibria and Edge-Cutting Deviations. A solution where all players v set all $c(\chi_e(t))$ to a number M so large that accepting such a contract guarantees the accepting node would have negative utility, results in no active edges, and is a Nash equilibrium (NE) with all nodes having a utility of 0. Moreover, for every NE there is an equivalent one where the payments demanded on inactive edges are M . We can also assume that there are no active contracts without demands being routed on these contracts, as such contracts would have no reason to exist. Call such a NE *definite*. Since our primary interest is to determine how well the players (and society) can do via stable solutions, we may restrict attention without loss of generality to such definite NEs. In a definite NE, no new contracts can be formed since all requests for inactive contracts are too high to be accepted. Thus we can now simply think of deviations as a node “cutting” some of the active contracts incident to it, as well as re-routing the demands that pass through it. Moreover, in a definite NE, the capacity of each active contract is fully saturated, and so new demands cannot be routed without active ones being re-routed.

Nash equilibria as good as OPT. The *social welfare*² of a given set of active contracts equals the total value of the demands that are active minus the total transit cost incurred. We now compare the quality of OPT (the solution that maximizes social welfare) with the quality of Nash equilibria. First, we consider the case where only the source $s(d)$ of a demand d is rewarded for d becoming active. All proofs can be found in the full version of this paper [4].

² Recall that the social welfare of a solution is the sum of the players’ utilities.

Theorem 1. *If $\lambda_{t(d)}(d) = 0$ for all demands d , then there exists a Nash equilibrium that is as good as OPT, and so the price of stability is 1.*

We now consider the more general version of our model, where both the source and the destination of a demand may benefit from this demand being successfully routed. For this model, we consider the quality of Nash equilibria using two different objective functions: the social welfare, and the *cut-loss* objective function, which is the total transit cost incurred plus the total value of the demands that are not connected. This objective function is also studied in [3][18]. While in Theorem 1 we showed that there always exists an equilibrium as good as the optimal solution, this is no longer true if destinations also benefit from active demands. In fact, [3] gives examples where for the social welfare objective, all equilibria may be arbitrarily far away from OPT.

For the cut-loss objective, however, there always exist good Nash equilibria. In the proof of the following theorem we not only show the existence of such equilibria, but also how to find them efficiently starting with an approximation to OPT. For this result, we assume that $\lambda_{s(d)}(d) = \lambda_{t(d)}(d)$ for all d . If this were not the case, then instead of a factor of 2, we obtain a price of stability that depends on the ratio between $\lambda_{s(d)}(d)$ and $\lambda_{t(d)}(d)$.

Theorem 2. *With respect to the cut-loss objective function, the price of stability is at most 2 in the case where for each demand d , $\lambda_{s(d)}(d) = \lambda_{t(d)}(d)$.*

Creating good equilibria. While for the social welfare objective, all Nash equilibria may be bad compared to OPT, we can still *create* good equilibria using similar methods as in [3]. For example, if we give incentives to certain nodes, we can form an equilibrium that is as good as the centralized optimal solution (see [3] for discussion about different types of incentives). For example, we can do this by increasing the $\lambda_{s(d)}(d)$ for every demand by a factor of 2. As in Theorem 2, we assume $\lambda_{s(d)}(d) = \lambda_{t(d)}(d)$ for all d . If this were not the case, then instead of increasing $\lambda_{s(d)}(d)$ by a factor of 2, the results hold if we set $\lambda_{s(d)}(d)$ to $\lambda_{s(d)}(d) + \lambda_{t(d)}(d)$, which is still a factor of 2 increase in total.

Theorem 3. *Let E^* be the set of active contracts in OPT, and D_{OPT} be the active demands in OPT. If we increase $\lambda_{s(d)}(d)$ by a factor of 2 for every d , then E^* together with D_{OPT} induce a Nash equilibrium (that is, there is an equilibrium with active contracts E^* and active demands D_{OPT}).*

Model Variations. All of the above results hold for a variety of extensions to our model (see [4]). These extensions include more realistic customer-provider contracts (where the provider is obligated to forward traffic to its customer, not just from its customer), and the addition of peering contracts. We also show that, with a slightly more constrained routing protocol, our results apply to contract systems where the contracts do not specify the destination. In other words, in this type of contract the provider would be obligated to forward all traffic from its customer, no matter where this traffic was destined.

Acknowledgments

We deeply thank Bruce Shepherd for illuminating discussions and for his many contributions.

References

1. Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-optimal network design with selfish agents. In: STOC (2003)
2. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The Price of Stability for Network Design with Fair Cost Allocation. In: FOCS (2004)
3. Anshelevich, E., Shepherd, B., Wilfong, G.: Strategic Network Formation through Peering and Service Agreements. In: FOCS 2006 (2006)
4. Anshelevich, E., Wilfong, G.: Network Formation and Routing by Strategic Agents using Local Contracts, <http://www.cs.rpi.edu/~eanshel>
5. Bala, V., Goyal, S.: A noncooperative model of network formation. *Econometrica* 68, 1181–1229 (2000)
6. Bloch, F., Jackson, M.O.: The Formation of Networks with Transfers among Players. *J. Economic Theory* 133(1), 83–110 (2007)
7. Erlebach, T., Hall, A., Schank, T.: Classifying Customer-Provider Relationships in the Internet. In: Proceedings in Informatics, pp. 52–64 (September 2002)
8. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C.H., Shenkar, S.: On a network creation game. In: PODC (2003)
9. Feigenbaum, J., Papadimitriou, C., Sami, R., Shenkar, S.: A BGP-based Mechanism for Lowest-Cost Routing. *Distributed Computing* 18, 61–72 (2005)
10. Gao, L., Rexford, J.: Stable internet routing without global coordination. In: Proceedings of ACM SIGMETRICS 2000, Santa Clara, CA (June 2000)
11. Griffin, T., Shepherd, F.B., Wilfong, G.: The Stable Paths Problem and Interdomain Routing. *IEEE/ACM Transactions on Networking* 10(2), 232–243 (2002)
12. He, L., Walrand, J.: Dynamic provisioning of service level agreements between interconnected networks. In: Conference on Stochastic Networks, June 19–24 (2002)
13. http://hosting.bellsouth.net/bellsouthhosting/s/s.dll?spage=cg/news/service_level.htm
14. <http://www.verio.com/access/pricing.cfm>
15. <http://www.xo.com/products/smallgrowing/internet/dia/oc3.html>
16. Jackson, M.: A survey of models of network formation: stability and efficiency. In: Demange, G., Wooders, M. (eds.) *Group Formation in Economics: Networks, Clubs and Coalitions*. Cambridge Univ. Press, Cambridge
17. Johari, R., Tsitsiklis, J.: Routing and Peering in a Competitive Internet. In: Proceedings of the IEEE Conference on Decision and Control (2004)
18. Johari, R., Mannor, S., Tsitsiklis, J.: A contract-based model for directed network formation. *Games and Economic Behavior* 56(2), 201–224 (2006)
19. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, pp. 404–413 (1999)
20. Laffont, J.J., Marcus, J.S., Rey, P., Tirole, J.: Internet connection and the off-net pricing principle. *The RAND Journal of Economics* 34(2), 370–390
21. Shepherd, B., Wilfong, G.: Multilateral Transport Games. In: INOC 2005 (2005)
22. Wang, H., Xie, H., Yang, Y.R.: On the stability of interdomain inbound traffic engineering. In: SIGCOMM Poster Session (2005)

Network Creation Games with Disconnected Equilibria

Ulrik Brandes¹, Martin Hoefer^{2,*}, and Bobo Nick¹

¹ Department of Computer & Information Science, University of Konstanz
`{ulrik.brandes,bobo.nick}@uni-konstanz.de`

² Lehrstuhl Informatik I, RWTH Aachen University
`mhoefer@cs.rwth-aachen.de`

Abstract. In this paper we extend a popular non-cooperative network creation game (NCG) [11] to allow for disconnected equilibrium networks. There are n players, each is a vertex in a graph, and a strategy is a subset of players to build edges to. For each edge a player must pay a cost α , and the individual cost for a player represents a trade-off between edge costs and shortest path lengths to all other players. We extend the model to a *penalized game* (PCG), for which we reduce the penalty for a pair of disconnected players to a finite value β . We prove that the PCG is not a potential game, but pure Nash equilibria always exist, and pure strong equilibria exist in many cases. We provide tight conditions under which disconnected (strong) Nash equilibria can evolve. Components of these equilibria must be (strong) Nash equilibria of a smaller NCG. But in contrast to the NCG, for the vast majority of parameter values no tree is a stable component. Finally, we show that the price of anarchy is $\Theta(n)$, several orders of magnitude larger than in the NCG. Perhaps surprisingly, the price of anarchy for strong equilibria increases only to at most 4.

1 Introduction

The study of distributed network creation with selfish agents has attracted much research interest from various disciplines. A general framework for such an approach was proposed by Jackson and Wolinsky [14]. In their games there are n players and each player is a vertex in a graph. A strategy consists of choosing which incident edges to build. Depending on the network structure there is a payoff for each player, and players adjust their strategy to maximize their payoff. A general finding was that there are games, in which no efficient network is stable for a concept of pairwise stability, which requires bilateral consent to construct a connection. The extensions and adjustments to this model are numerous [13]. In particular, several works extended the model to unilateral link creation and the Nash equilibrium as stability concept [5,9]. A particularly interesting variant was proposed by Fabrikant et al. [11]. In their network creation game (NCG) the

* Supported by DFG-Graduiertenkolleg 1298 “AlgoSyn” at RWTH Aachen University.




Fig. 1. Price of anarchy in the NCG

cost of creating an edge is fixed to a parameter α . Edge creation is unilateral, and the cost for a player is a trade-off between edge costs and structural network position measured by shortest path distances to all other players. In [2,8,11] the inefficiency of Nash equilibria was quantified using the price of anarchy [15], the ratio of the cost of the worst Nash equilibrium over the cost of a social optimum state. The presently known results on the price of anarchy are summarized in Fig. 1. Other equilibrium concepts were also studied, e.g. pairwise stable equilibria [7], or strong equilibria [3], as well as extensions to more general edge costs or different player cost trade-offs [1,10,17]. In network analysis [6], the inverse of the sum of shortest path lengths is one of the most commonly used measures of centrality known as *closeness* [12]. A problem with closeness is that global connectivity is required for the scores to be comparable. This means that in the NCG for moderate to high edge costs the trade-off is distorted by the enforcement of connectivity. Thus, it was not surprising that trees proved to be a prominent equilibrium structure [11].

In this paper, we remedy this problem by replacing the infinite cost of not being connected by a finite penalty β . This corresponds directly to a variant of closeness centrality proposed by Botafogo et al. [18], and it was suggested as an open problem in [11]. For special values of β it is closely related to a measure called *radiality* [19]. Our *penalized network creation game* (PCG) is introduced in Sect. 2. Since the cost of connected equilibria is the same as in the NCG, we are interested in existence, structure, and cost of disconnected Nash equilibria. If β is large, Nash equilibria of the PCG are similar to those of the NCG, in particular, they are connected. For smaller β , disconnected Nash equilibria evolve, and an interesting insight gained from Sect. 3 is that the prevalent tree structures of the NCG are absent whenever $\beta > 2$ or $\alpha > 1$ (see Theorem 3). In addition, we consider the price of anarchy in Sect. 4. There are parameter values, for which disconnected Nash equilibria appear but the social optimum is connected, which could lead to an unbounded price of anarchy. However, we show that the price of anarchy in the PCG is always bounded by $O(n)$. In addition, Theorem 4 reveals cases with a *tight* matching lower bound of $\Omega(n)$. This bound is strictly larger than any bounds for the NCG. In Sect. 5 we consider players that can play joint coordinated deviations and strong equilibria. Unless α and β are within a small range, the social optimum is also a strong equilibrium (Theorem 5). In Theorem 6 we prove that the price of anarchy for strong equilibria is at most 4. This reveals that in the PCG regular Nash equilibria can be several orders of magnitude more costly than strong equilibria, a question which is still unsolved for the NCG. Due to spacial constraints proofs are sketched or omitted and will be given in the full version of the paper.

2 The Model and Initial Results

The network connection game (NCG) is a tuple (V, α) and can be described as follows. The set of players V is the set of vertices of a graph. Possible edges $\{i, j\} \in V \times V$ have cost α . A strategy s_i of a player i is a subset $s_i \subset V \setminus \{v\}$ and indicates, which edges player i chooses to build. In this way a strategy vector s induces a set of edges between the players. Given a strategy vector s the individual cost for a player i is $c_i(s) = \alpha|s_i| + \sum_{j \neq i} dist_s(i, j)$, where $\alpha > 0$ and $dist_s(i, j)$ is the length of a shortest-path in the undirected graph $G_s = (V, E_s)$ induced by the strategy vector s . Note that G_s is assumed to be undirected, i.e. each edge can be traversed in any direction, independent of which player pays for it. In the regular connection game $dist_s(i, j) = \infty$ if players i and j are in different components of G_s . In the *penalized network creation game* (PCG) we are given a penalty value $\beta > 1$, and $dist_s(i, j) = \beta$ for players i and j in different components. A pure *Nash equilibrium* (NE) is a state s , in which no player can unilaterally decrease her cost c_i by changing her strategy s_i . We will restrict our attention to pure equilibria throughout. The *social cost* $c(s)$ of a state s is simply $c(s) = \sum_{i \in V} c_i(s)$. The social optimum state s^* is a state with minimum social cost. Note that for the cost of a state it does not matter, which of the two players connected by an edge chose to pay for it, and hence we will sometimes use the graph G_s for s . States that play an important role in the analysis of the PCG are the empty state $s_\emptyset = (\emptyset, \dots, \emptyset)$, s_K corresponding to the complete graph, in which each edge $\{i, j\}$ with $i \neq j$ is paid by player $\min\{i, j\}$, and s_Z corresponding to a center-sponsored star, in which one player purchases edges to all other players. Fabrikant et al. [11] show that there is always a pure NE in the NCG and mention that it might be found by iterative improvement steps. Finding a best-response for a player in a NCG, however, was shown NP-hard [11], and this translates to the PCG for sufficiently large penalty cost. In addition, we show that better-response dynamics may cycle, hence the game is no potential game [16]. As the dynamics involve no disconnectivities, the result follows directly for the PCG. Nevertheless, in the PCG there is always a pure NE. This serves as a first insight to motivate the further study of the properties of pure NE in the PCG.




Fig. 2. NCG with $k = 4$ and $4 < \alpha < 6$ with cycling better response iteration. Black dots indicate the player who pays for the edge.

Theorem 1. *Every PCG has a pure Nash equilibrium, but neither NCG nor PCG are potential games.*

Proof. We first disprove the existence of a potential function. For any $\alpha > 3$ choose an integer k with $k < \alpha < \frac{3k}{2}$. Now construct a strategy combination for $n = 4k$ players as depicted in Fig. 2. The following steps each represent a strict improvement for the players: (1) player 4 removes edge e_1 , (2) player 2 removes

edge e_2 , (3) player 4 builds edges e_1 and e_2 . This leads into an isomorphic state, and allows to construct an infinite improvement path. For the proof of existence it can be shown that the following states are NE: for $\alpha \geq \beta - 1$ the empty network s_\emptyset , for $1 \leq \alpha < \beta - 1$ the center-sponsored star s_Z , in which all edges are bought by the center node, and for $\alpha < 1$ and $\alpha < \beta - 1$ the state s_K . \square

3 Disconnected Equilibria

In this section we consider existence and structural properties of disconnected NE in the PCG. First, we clarify the existence of disconnected equilibria.

Theorem 2. *For $\alpha \geq \beta - 1$ the empty graph is always a disconnected NE. For $0 < \alpha < \beta - 1$ there is no disconnected NE.*

Proof. The first part follows from Theorem [1](#). For the second part consider a player v in a disconnected NE s . Let n_v be the size of the component of the graph G_s , in which v is located. Now suppose v changes her strategy by connecting to all $n - n_v$ players in other components. Then the change is $\alpha(n - n_v) + (n - n_v) - \beta(n - n_v) = (n - n_v)(\alpha - (\beta - 1)) < 0$. Hence, under these conditions every player in a disconnected state can decrease her individual cost. \square

The theorem provides a tight characterization using the empty graph. An interesting issue, however, is to explore whether non-empty disconnected NE are possible, because in many cases the empty graph represents a rather unrealistic prediction for a stable network. Note that a component of k players in a disconnected NE of a PCG with given α and β must be a NE in the corresponding NCG with α and k players. A prominent structure that has been identified as NE in the NCG are trees.

Trees. Tree graphs are a structure whose appearance is wide-spread in the NCG [\[2,11\]](#). The following analysis shows that this property is only due to the requirement that a NE must be connected. The following discussion reveals that in the PCG these structures can appear only in very special cases.

Lemma 1. *For $\beta > 2$ every non-singleton player v in a disconnected NE has at least one incident edge that was created by a different player $w \neq v$.*

Proof. Consider a player v in a component C with k players, who pays for all her d_v incident edges. As we have a NE, it is not profitable for v to disconnect from C , i.e. $\alpha d_v + \sum_{w \in C} \text{dist}(v, w) \leq \beta(k - 1)$. Consider a different player $v' \notin C$ that chooses to connect to all neighbors of v . This must not be profitable, so $\alpha d_v + \sum_{w \in C} \text{dist}(v, w) + 2 \geq \beta k$. Adding the inequalities yields $\beta \leq 2$. \square

Lemma 2. *Suppose there is a disconnected NE with a component C of $k > 1$ vertices. If $\alpha > (k - 1)(\beta - 2) + 1$, then for every player v there is an incident edge paid by a different player $w \neq v$.*

Proof. Suppose there is a player v that pays for all her $d_v \geq 1$ incident edges. As v does not want to remove all edges, we have $\alpha d_v + \sum_{w \in C} \text{dist}(v, w) \leq \beta(k-1)$, and thus $\alpha \leq \frac{1}{d_v} (\beta(k-1) - \sum_{w \in C} \text{dist}(v, w))$. Every pair of non-neighbor vertices in C has a distance of at least 2, so $\sum_{w \in C} \text{dist}(v, w) \geq 2(k-1) - d_v$. Substitution yields $\alpha \leq (k-1)(\beta-2) + 1$ as desired. \square

Theorem 3. *For $\beta > 2$ or $\alpha > 1$ no component of a disconnected NE is a tree.*

Proof. The first bound is a direct consequence of Lemma 1 and the fact that for a tree $|E| = |V| - 1$. Thus, for disconnected NE with tree components $\beta \leq 2$, and the second bound follows with Lemma 2. \square

Non-empty Equilibria. It can be shown that the appearance of currently known NE topologies from the NCG as components in disconnected NE of the PCG is quite limited. The existence of disconnected NE, however, is guaranteed by the empty network. This raises the question under which conditions on α and β non-empty disconnected NE can evolve. We first present a positive result.

Lemma 3. *For $3 \leq \alpha \leq 4$ and $\beta \leq (\alpha + 11)/5$ a cycle C_5 of 5 vertices can be a component of a disconnected NE.*

In contrast to the restricted interval, for which we can show existence, there is an unbounded region of parameter values, for which the empty network is the only disconnected network - in particular if α or β are large compared to n .

Lemma 4. *In a non-empty disconnected NE let n_l be the minimum size and diam_l the minimum diameter of any non-singleton component. Then (1) $\alpha < 12n_l \log n_l$, (2) $\beta \leq 1 + 2 \cdot \text{diam}_l$, (3) $\beta < 1 + 14\sqrt{n_l \log n_l}$, and (4) if $n > 6$, then $\beta < n/2$.*

Proof. We only prove the first three bounds here. For the first bound consider $\alpha \geq 12n_l \log n_l$ and a component with n_l players. This component must represent a NE in a NCG with the same α and n_l players, and thus according to 2 must be a tree. This contradicts Theorem 3 and the bound follows. Now consider a non-empty disconnected NE s for $\beta > 2$, and let C be a non-singleton component. As C is no tree, it must contain at least one cycle. Let U be a smallest of all cycles in C , and let v be an arbitrary player that constructed some edge e of U . Denote by s' the state that evolves if player v removes edge e . Note that by this removal no additional pair of players gets disconnected. As s is a NE, we have $\alpha \leq \sum_{w \in C} (\text{dist}_{s'}(v, w) - \text{dist}_s(v, w))$. As we have chosen U to be of minimum size, all shortest distances between vertices of U are given by the paths along the cycle. Thus, there is always a vertex u , for which the distances in s and s' are the same. This yields $\text{dist}_{s'}(v, w) \leq \text{dist}_{s'}(v, u) + \text{dist}_{s'}(u, w) = \text{dist}_s(v, u) + \text{dist}_s(u, w)$ for all $w \in C$. With $n_C = |C|$ we can conclude $\alpha \leq 2n_C \cdot \text{diam}(C) - \sum_{w \in C} \text{dist}_s(v, w)$. On the other hand, no vertex outside C must be able to profit from a connection to v , hence $\alpha + n_C + \sum_{w \in C} \text{dist}(v, w) \geq n_C \beta$. The last two inequalities deliver the second bound. We know from 1 that $\text{diam}(C) \leq \sqrt{4\alpha + 1}$. Together with the bounds (1) and (2) shown above this implies the third bound. \square

In contrast to these bounds, we have not been able to derive any non-empty disconnected NE for values of $\beta > 3$. This led us to formulate the following conjecture. Note that our bounds imply that if the conjecture is false, then there must be non-tree NE in the NCG with a diameter of size $\omega(1)$. This seems quite unlikely, as all non-tree NE found so far have diameter at most 3.

Conjecture 1 (Constant Penalty Conjecture). There is a constant β' such that for $\beta > \beta'$ the only disconnected NE is s_\emptyset .

4 Price of Anarchy

In this section we consider the price of anarchy in the PCG. We first consider the social optima of the game. For $\alpha \leq \min\{2, 2\beta - 2\}$ the complete graph s_K is the optimum. For $\alpha \leq 2$ and $\alpha \geq 2\beta - 2$ the empty graph s_\emptyset is the optimum. s_\emptyset remains the optimum for $\alpha \geq 2$ and $\alpha \geq \beta n - 2(n - 1)$. For the remaining range the star s_Z is the optimum. For $\alpha < \beta - 1$ we have seen in Theorem 2 that no disconnected NE exists. In addition, it can be shown that in this case a finite penalty for disconnectivity cannot disrupt any NE of the NCG. Hence, for this parameter range the price of anarchy is identical to the NCG. In general, however, the price of anarchy in the PCG can be strictly larger than for the NCG. Fig. 3 provides an overview of the bounds we obtained. Note that all these bounds are in $O(n)$ for the respective parameter values. We concentrate on the case $\max\{2, \beta - 1\} < \alpha < \beta n - 2(n - 1)$, in which disconnected NE can appear and the star is the social optimum.

Theorem 4. For $2\beta - 2 \leq \alpha \leq n\beta - 2(n - 1)$ the price of anarchy is bounded by $\Theta\left(\frac{n\beta}{\alpha}\right)$ for $\alpha \geq 12n \log n$ and $O\left(5^{\sqrt{\log n}} \log n + \frac{n\beta}{\alpha+n}\right)$ for $\alpha < 12n \log n$. For $\beta - 1 \leq \alpha \leq 2\beta - 2$ the price of anarchy is $\Theta(\min\{\beta, n\})$.

Proof. For the proof of the first bound consider $\alpha \geq 12n \log n$. According to Lemma 4 in this case every NE is either connected or s_\emptyset . For $\alpha \geq 12n \log n$ all connected NE have a constant price of anarchy 2, while s_\emptyset leads to an increase and proves our first bound: $\frac{c(s_\emptyset)}{c(s_Z)} = \frac{\beta n}{\alpha + 2(n-1)} \in \Theta\left(\frac{n\beta}{\alpha}\right)$. This bound increases from $\Theta(1)$ to $\Theta(n)$ if α drops from $n\beta - 2(n - 1)$ to $2\beta - 2$. It also shows that the price of anarchy induced by s_\emptyset is never more than $O(n)$ for $s^* = s_Z$ and $\alpha \geq \beta - 1$. Another range, for which s_\emptyset is the most expensive NE, is $\beta - 1 \leq \alpha \leq 2\beta - 2$ with $\beta \geq 7$. Then any directly connected pair induces a cost of $\alpha + 2 \leq 2\beta$. Any indirectly connected pair in a NE induces a cost $2 \text{dist}_s(v, w) \leq 2\sqrt{4\alpha + 1} \leq 2\sqrt{8\beta - 7} \leq 2\beta$. Thus, the cost of 2β induced by s_\emptyset is maximal for every pair of players. $c(s_\emptyset)/c(s_Z)$ characterizes the price of anarchy and results in $\Theta(\min\{\beta, n\})$, which proves the third bound. For the remaining range with $\alpha < 12n \log n$ there might be worse disconnected NE than s_\emptyset . However, components of these NE must be connected NE of smaller NCGs. We bound the price of anarchy for these NE by the fraction for s_\emptyset plus the maximum factor of any component NE in the corresponding NCG. With the bound of $5^{\sqrt{\log n}} \log n \in o(n^\epsilon)$ on the price of anarchy for the NCG 8 this proves our second bound $O(\max\{5^{\sqrt{\log n}} \log n, \min\{n, \beta\}\})$, which is at most $O(n)$. \square




Fig. 3. Price of anarchy in the PCG

5 Strong Equilibria

In this section we assume agents are able to jointly deviate to different strategies. As stability concept we consider the strong equilibrium [4], in which no coalition C of players can decrease the cost for each of its members by taking a joint deviation. More formally, if a state s is a strong equilibrium (SE), then for each coalition of players C and each possible strategy profile s'_C for the players in C it holds that if there is a player $i \in C$ with $c_i(s'_C, s_{-C}) < c_i(s)$, then there is another player $j \in C$ with $c_j(s'_C, s_{-C}) \geq c_j(s)$. The price of anarchy for SE is a direct adaption of the price for NE and was studied in [3] for the NCG. The next theorem summarizes structural and existence properties of SE in the PCG. It shows, in particular, that with the exception of a small range of parameter values strong equilibria always exist in the PCG. Finally, the main result in this section is a general constant upper bound on the price of anarchy for SE in the PCG.

Theorem 5. *For $\alpha < \beta - 1$ the SE of the PCG are exactly the SE of the corresponding NCG. For $\alpha \geq \beta - 1$ the social optimum in the PCG is a SE for all parameter values except $\beta < 3$, and $\beta n - 2n + 2 - (\beta - 1) < \alpha < \beta n - 2n + 2$.*

Theorem 6. *The price of anarchy for SE in the PCG is at most 4.*

References

1. Ackermann, H., Briest, P., Fanghänel, A., Vöcking, B.: Who should pay for forwarding packets? In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 208–219. Springer, Heidelberg (2007)
2. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., Roditty, L.: On Nash equilibria for a network creation game. In: Proc. 17th SODA, pp. 89–98 (2006)

3. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: Proc. 18th SODA, pp. 189–198 (2007)
4. Aumann, R.: Acceptable points in general cooperative n-person games. In: Contributions to the Theory of Games IV. Annals of Mathematics Study, vol. 40, pp. 287–324 (1959)
5. Bala, V., Goyal, S.: A non-cooperative model of network formation. *Econometrica* 68, 1181–1229 (2000)
6. Brandes, U., Erlebach, T. (eds.): Network Analysis: Methodological Foundations. LNCS Tutorial. Springer, Heidelberg (2005)
7. Corbo, J., Parkes, D.: The price of selfish behavior in bilateral network formation. In: Proc. 24th PODC, pp. 99–107 (2005)
8. Demaine, E., Hajiaghayi, M.T., Mahini, H., Zadimoghaddam, M.: The price of anarchy in network creation games. In: Proc. 26th PODC, pp. 292–298 (2007)
9. Dutta, B., Jackson, M.: The stability and efficiency of directed communication networks. *Review of Economic Design* 5(3), 251–272 (2000)
10. Eidenbenz, S., Kumar, A., Züst, S.: Equilibria in topology control games for ad hoc networks. In: Proc. DIALM-POMC Workshop Foundations of Mobile Comp., pp. 2–11 (2003)
11. Fabrikant, A., Luthera, A., Maneva, E., Papadimitriou, C., Shenker, S.: On a network creation game. In: Proc. 22nd PODC, pp. 347–351 (2003)
12. Freeman, L.: Centrality in social networks: Conceptual clarification. *Social Networks* 1(3), 215–239 (1979)
13. Jackson, M.: A survey of models of network formation: Stability and efficiency. In: Demange, G., Wooders, M. (eds.) *Group Formation in Economics; Networks, Clubs and Coalitions*, ch. 1, pp. 11–57. Cambridge University Press, Cambridge (2004)
14. Jackson, M., Wolinsky, A.: A strategic model of social and economic networks. *Journal of Economic Theory* 71(1), 44–74 (1996)
15. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Proc. 16th STACS, pp. 404–413 (1999)
16. Monderer, D., Shapley, L.: Potential games. *Games and Economic Behavior* 14, 1124–1143 (1996)
17. Moscibroda, T., Schmid, S., Wattenhofer, R.: On the topologies formed by selfish peers. In: Proc. 25th PODC, pp. 133–142 (2006)
18. Rivlin, E., Botafogo, R., Shneiderman, B.: Navigating in hyperspace: designing a structure-based toolbox. *Comm. ACM* 37(2), 87–96 (1994)
19. Valente, T., Foreman, R.: Integration and radially: measuring the extent of an individual’s connectedness and reachability in a network. *Social Networks* 20(1), 89–105 (1998)

Randomized Truthful Mechanisms for Scheduling Unrelated Machines^{*}

Pinyan Lu and Changyuan Yu

Institute for Theoretical Computer Science, Tsinghua University
Beijing, 100084, P.R. China
{lpy,yucy05}@mails.tsinghua.edu.cn

Abstract. In this paper, we consider randomized truthful mechanisms for scheduling tasks to unrelated machines, where each machine is controlled by a selfish agent. Some previous work on this topic focused on a special case, scheduling two machines, for which the best approximation ratio is 1.6737 [5] and the best lower bound is 1.5 [6]. For this case, we give a unified framework for designing universally truthful mechanisms, which includes all the known mechanisms, and also a tight analysis method of their approximation ratios. Based on this, we give an improved randomized truthful mechanism, whose approximation ratio is 1.5963. For the general case, when there are m machines, the only known technique is to obtain a $\frac{2m}{2}$ -approximation truthful mechanism by generalizing a γ -approximation truthful mechanism for two machines [6]. There is a barrier of $0.75m$ for this technique due to the lower bound of 1.5 for two machines. We break this $0.75m$ barrier by a new designing technique, rounding a fractional solution. We propose a randomized truthful-in-expectation mechanism that achieves approximation of $\frac{m+5}{2}$, for m machines.

For the lower bound side, we focus on an interesting family of mechanisms, namely *task-independent* truthful mechanisms. We prove a lower bound of $11/7$ for two machines and a lower bound of $\frac{m+1}{2}$ for m machines with respect to this family. They almost match our upper bounds in both cases.

1 Introduction

Mechanism design, an important area both in Game Theory and Computer Science, has received extensive study in the past few years. It is usually used to design a protocol for achieving some global objective, however requiring the interaction of some selfish agents. To deal with this, the most common solution concept is “truthfulness”, where the mechanism is designed so that for any participant agent, reporting his/her private data truthfully to the mechanism will always maximize his/her own utility, no matter how other agents act. We also focus on truthful mechanisms in this paper.

^{*} Supported by the National Natural Science Foundation of China Grant 60553001 and the National Basic Research Program of China Grant 2007CB807900, 2007CB807901.

The study of the algorithmic aspect of mechanism design was initiated by Nisan and Ronen in their seminal paper “Algorithmic Mechanism Design” [8]. Some computational properties such as good approximation ratios and polynomial running time are studied in mechanism design setting. Nisan and Ronen’s work mainly focused on a fundamental problem in computer science, scheduling unrelated machines. In a scheduling problem, there are n tasks to be allocated to m machines, which are controlled by selfish agents. The objective is to allocate the tasks so that the maximum completion time of these machines (called *makespan*) is minimized. A mechanism for the scheduling problem consists of two algorithms, the allocation algorithm and the payment algorithm. Our main interest is on the approximation ratio of the allocation algorithm. Nisan and Ronen proposed a deterministic truthful mechanism with an approximation ratio of m . Moreover, they proved a lower bound of 2 for all the deterministic truthful mechanisms. Randomization is always more powerful, and this is also true for this scheduling problem. They provided a randomized truthful mechanism with approximation ratio of 1.75 for two machines. Recently Mu’alem and Schapira gave a lower bound of $2 - 1/m$ for randomized truthful mechanisms [6]. They also generalized the 1.75 approximation mechanism for two machines to a $0.875m$ -approximation mechanism for m machines. In a previous work [5], we improved Nisan and Ronen’s result by a 1.67-approximation randomized truthful mechanism for two machines, together with a $0.837m$ -approximation mechanism for m machines using Mu’alem and Schapira’s technique in [6].

A fractional variant of truthful scheduling unrelated machines was first considered by Christodoulou, Koutsoupias and Kovács in [2]. They gave a fractional truthful mechanism with approximation ratio of $(m+1)/2$, and a lower bound of $2 - 1/m$ for any fractional truthful mechanisms. They also defined a family of allocation algorithms named as *task-independent* algorithm, in which tasks are allocated independently. For the task-independent truthful fractional mechanisms, they proved a tight lower bound of $(m + 1)/2$.

1.1 Our Results

In this paper, we first propose a unified approach to design truthful mechanisms for two machines, which contains all the known truthful mechanisms. One main contribution is that we not only unify all the known mechanisms, but also give a unified and tight analysis method for their approximation ratios. Based on this, we are able to give a randomized mechanism for two machines, which is universally truthful and has an approximate ratio of 1.5963.

A natural question would be how far we can go with this unified approach. We answer this question by a lower bound of 1.5788 for this approach. Further more, we also prove a lower bound of $11/7$ for all the task independent randomized mechanisms which are truthful even in a weaker version, i.e., truthful in expectation. So substantial new techniques are required to significantly improve our results.

For the general case, when there are m machines, the only known technique is to obtain a $\frac{\gamma m}{2}$ -approximation truthful mechanism by generalizing a γ -approximation truthful mechanism for two machines [6]. However the lower bound of 1.5 for scheduling two machines gives a barrier of $0.75m$ for this technique. We break this $0.75m$ barrier by a new designing technique. First, we adopt a truthful fractional mechanism with ratio $(m + 1)/2$ by Christodoulou, Koutsoupias and Kovács [2]. We add into this mechanism an important threshold so that it satisfies certain “bid condition”, which is essential for us to bound the loss of approximation ratio during the rounding process. Then we use a rounding technique in [4] to get a randomized mechanism, which is still truthful in expectation, and only loses little in approximation ratio. We finally obtain a randomized mechanism which is truthful in expectation and achieves an approximation ratio of $(m + 5)/2$.

We also give a lower bound of $(m + 1)/2$ for all task independent randomized mechanisms. This result shows that we really need some new techniques to break this $0.5m$ barrier.

2 Preliminaries and Notations

In this section we review some definitions and results on mechanism design and scheduling problem. More details can be found in [8]. In the following, for a generic matrix $a = (a_{ij})$, we use a_i to denote the i -th row of the matrix, and a_{-i} to denote the matrix obtained from a deleting a_i . We also use (v, a_{-i}) to denote the matrix obtained from a by replacing a_i with vector v . We use R_+ to denote the set of nonnegative real numbers.

In a scheduling problem, there are n tasks and m machines, where each machine $i \in [m]$ needs t_{ij} units of time to perform task $j \in [n]$. We usually use the matrix $t = (t_{ij})$ to denote an instance of the scheduling problem. In this paper, we consider that each machine is controlled by a strategic player. We assume that player i privately knows t_i , and we call the vector t_i player i 's type. After each player i declares his/her type, an allocation algorithm x will decide an allocation of all the tasks. We assume that all the players are selfish and want to perform as less tasks as possible, so players may misreport their types. We use $b_i \in R_+^n$ to denote player i 's reported type, and call it player i 's bid. Obviously b_i may not equal to t_i if that helps in player i 's interest. To avoid this lying issue, we introduce the payment algorithm p into a mechanism. Formally, a mechanism $M = (x, p)$ consists of two parts:

- **An allocation algorithm:** The allocation algorithm x , given the input of players' bid matrix $b = (b_1, \dots, b_m)$, outputs an allocation denoted by a matrix $x = (x_{ij})$. x_{ij} is 1 if task j is assigned to machine i , and 0 otherwise. In the fractional scheduling case, x_{ij} satisfies $0 \leq x_{ij} \leq 1$ and denotes the fraction of task j assigned to machine i . Every task must be completely assigned, hence $\sum_{j \in [n]} x_{ij} = 1, \forall i \in [m]$. Notice that each x_{ij} can be viewed as a function of b .

- **A payment algorithm:** The payment algorithm p , given the input of players' bid matrix b , outputs a vector $p = (p_1, \dots, p_m)$, where p_i denotes the money that player i receives from the mechanism. Each p_i can also be viewed as a function of b .

Randomized mechanism is defined to be a distribution of several deterministic mechanisms. In randomized mechanism, x_{ij} is a random variable denoting whether task j is assigned to machine i . For simplicity, we also use x_{ij} to denote $Pr(x_{ij} = 1)$ when the context is clear.

Now we specify the utility of each player. We use the quasi linear utility, which means the utility u_i of player i with type t_i over an allocation x and money p_i is defined as:

$$u_i(x, p_i | t_i) = p_i - \sum_{j \in [n]} x_{ij} t_{ij}.$$

Since x and p_i are both functions of bid matrix b , we can also write the utility as

$$u_i(b | t_i) = p_i(b) - \sum_{j \in [n]} x_{ij}(b) t_{ij}.$$

Recall that we want to solve the issue of lying about types, we are interested in truthful mechanisms. A mechanism $M = (x, p)$ is truthful if for each player i , reporting his/her true type will maximize his/her own utility. Formally, for any i , any bids b_{-i} of all other players, we have

$$u_i((t_i, b_{-i}) | t_i) \geq u_i((b_i, b_{-i}) | t_i), \quad \forall b_i \in R_+^n$$

For randomized mechanism, there are two versions of truthfulness. The stronger version is universally truthful, which requires the mechanism to be truthful when fixing all the random bits. The weaker version is truthful in expectation, which only requires that for each player, reporting his/her true type will maximize his/her own expected utility.

For a truthful mechanism M , we may assume that all the players will report their true types, hence $b = t$. Now, how can we evaluate the performance of mechanism's allocation algorithm x ? We consider the makespan, which is the maximum load of all the machines. Given input t , the makespan of mechanism M is denoted by $l_M(t)$, and $l_M(t) = \max_{i \in [m]} \sum_{j \in [n]} x_{ij} t_{ij}$. We use $l_{opt}(t)$ to denote the optimum, and $l_{opt}(t) = \min_x \max_{i \in [m]} \sum_{j \in [n]} x_{ij} t_{ij}$. A mechanism M is called c -approximation mechanism if for any instance t , we have $l_M(t) \leq c \cdot l_{opt}(t)$. For randomized mechanism M , we require $E[l_M(t)] \leq c \cdot l_{opt}(t)$, where the expectation is over the random bits used in the mechanism.

To sum up, we aim at designing (randomized) truthful mechanism with small approximation ratio. By the way, we also require the algorithms of the mechanism to be polynomial computable. When designing a mechanism, there are already several results about the characterization of truthfulness, which may help us to get rid of the payment issue. We mainly use Archer and Tardos' monotone

theorem for one parameter mechanism in [1]. In the one parameter case, each player i only has a single value as his/her type (i.e. the speed of machine i). Similar result is obtained in [7] for the auction setting.

Theorem 1. ([7,1]) *In a one parameter scheduling mechanism, an allocation algorithm admits a payment scheme to make the mechanism truthful if and only if it is monotone decreasing. In this case, the mechanism is truthful if and only if the payments $p_i(b_i, b_{-i})$ are of the form*

$$h_i(b_{-i}) + b_i x_i(b_i, b_{-i}) - \int_0^{b_i} x_i(u, b_{-i}) du$$

where the h_i are arbitrary functions, and x^i are the allocation functions (algorithm).

In this paper, we also consider the lower bound of approximation ratio for a special family of mechanisms, i.e. task independent truthful mechanisms. We first define task independent mechanisms.

Definition 1. *A deterministic mechanism M is task independent, if for any bid matrices b, b' such that $b_{ij} = b'_{ij}$ for any $i \in [m]$, then the allocation of task j does not change, i.e. $x_{ij}(b) = x_{ij}(b'), \forall i \in [m]$.*

For randomized mechanisms, there are also two versions of task independence. One is weak task independent randomized mechanism, which is a distribution over several task independent deterministic mechanisms. The other is (strong) task independent randomized mechanism, which satisfies that not only the allocation of task j does not change when j 's column of b is not changed, but also all the random variables x_{ij} are independent between different tasks. In this paper, we consider the stronger version.

The following theorem is a main tool used in proving lower bound.

Theorem 2. (Monotone theorem[8]) *In any truthful mechanism, the allocation algorithm must satisfy the following monotone property: for any two bids b and b' which differ only on machine i , the corresponding allocation $x(b)$ and $x' = x(b')$ satisfy*

$$\sum_{j=1}^m (x_{ij} - x'_{ij})(b_{ij} - b'_{ij}) \leq 0.$$

We remark that for randomized mechanism, the monotone property of the allocation algorithm still holds, which is proved implicitly in [6]. In our paper, we only use the following corollary for task independent randomized truthful mechanisms.

Corollary 1. *For any task independent randomized truthful mechanism M , any two bid matrices b, b' where b' is obtained from b by only changing b_{ij} to b'_{ij} ,*

then we have $(x_{ij}(b) - x_{ij}(b'))(b_{ij} - b'_{ij}) \leq 0$, where x_{ij} denotes the probability of assigning task j to machine i .

3 Scheduling Two Machines

Most of the previous works on this topic are for scheduling two machines. In this section, we first propose a unified framework for all the known mechanisms. Based on this framework, we give an improved truthful mechanism. Then we also explore the limitation of this approach by showing an almost tight lower bound for all the task-independent truthful mechanisms.

3.1 Unified Randomized Truthful Mechanisms M_f

Let $f : R_+ \rightarrow [0, 1]$ be a non-decreasing monotone function, satisfying $f(0) = 0$ and $\lim_{x \rightarrow \infty} f(x) = 1$. Then we have a randomized mechanism M_f for scheduling two machines based on f . Noticing that this kind of function f can be viewed as a cumulative distribution function for a random variable in R^+ , we have the following formal description of the mechanism M_f :

Input: The reported bid matrix b .
Output: A randomized allocation x and a payment $p = (p_1, p_2)$.
Allocation and Payment Algorithm:
 $x_{1j} \leftarrow 0, x_{2j} \leftarrow 0, j = 1, 2 \dots, n; p_1 \leftarrow 0; p_2 \leftarrow 0$.
 For each task $j = 1, 2 \dots, n$ do
 Choose $s_j \in R_+$ randomly according to function f such that $Pr(s_j \leq u) = f(u)$.
 if $b_{1j} \leq s_j^{-1}b_{2j}$,
 $x_{1j} \leftarrow 1, p_1 \leftarrow p_1 + s_j^{-1}b_{2j}$;
 else
 $x_{2j} \leftarrow 1, p_2 \leftarrow p_2 + s_j b_{1j}$.

This unified mechanism M_f is actually a generalization of Nisan and Ronen’s Biased MinWork Mechanism in a continuous setting. For the truthfulness, we have the following theorem.

Theorem 3. *For any non-decreasing monotone function $f : R_+ \rightarrow [0, 1]$, where $f(0) = 0$ and $\lim_{x \rightarrow \infty} f(x) = 1$, mechanism M_f is universally truthful.*

Proof. To prove that the mechanism M_f is universally truthful, we only need to prove that it is truthful when the random sequence $\{s_j\}$ is fixed. Since the utility of an agent equals the sum of the utilities obtained from each task and our mechanism is task-independent, we only need to consider the case of one task. In this case, say s_j is fixed and there is only one task j , the allocation algorithm is

monotone decreasing and the payment makes the mechanism truthful, according to the theorem [11](#) (with function $h_1(b_{2j}) = s_j^{-1}b_{2j}$ and $h_2(b_{1j}) = s_j b_{1j}$).

Now we demonstrate the power of our unified designing approach by showing that every known mechanism can be viewed as a mechanism M_f with respect to some function f .

$$f_1(x) = \begin{cases} 1, & x \geq 1, \\ 0, & 0 \leq x < 1; \end{cases} \quad f_2(x) = \begin{cases} 1, & x \geq \frac{4}{3}, \\ \frac{1}{2}, & \frac{3}{4} \leq x < \frac{4}{3}, \\ 0, & 0 \leq x < \frac{3}{4}; \end{cases} \quad f_3(x) = \begin{cases} 1, & x \geq \alpha, \\ r, & \beta \leq x < \alpha, \\ \frac{1}{2}, & \frac{1}{\beta} \leq x < \beta, \\ 1-r, & \frac{1}{\alpha} \leq x < \frac{1}{\beta}, \\ 0, & 0 \leq x < \frac{1}{\alpha}. \end{cases}$$

M_{f_1} is exactly the *Min Work Mechanism* proposed by Nisan and Ronen [8](#). This is indeed a deterministic mechanism, whose approximation ratio is 2, and it is the best determinate mechanism. M_{f_2} is the *Biased Min Work Mechanism* also proposed by Nisan and Ronen [8](#), whose approximation ratio is 1.75. Then we improved their result by M_{f_3} in our previous work [5](#). By taking $\alpha = 1.4844, \beta = 1.1854, r = 0.7932$ in f_3 , we have a randomized truthful mechanism with approximation ratio of 1.6737.

We can see that all the previous functions f are distribution functions of some discrete random variables. One essential reason is that we can apply a “task reducing” technique [8,5](#), then analyze the performance using a case by case method. However the number of subcases increased dramatically if we consider a more complicated function f . One of our main contribution in this paper is that we not only propose the unified framework M_f , but also provide a performance analysis method.

Theorem 4. *For any non-decreasing monotone function $f : R_+ \rightarrow [0, 1]$, the approximation ratio of the mechanism M_f is exactly $\max_{\alpha_1, \alpha_2 \in R_+} F(\alpha_1, \alpha_2)$, where $F : R_+ \times R_+ \rightarrow R$ is defined as following (Here $r_1 = f(\alpha_1)$ and $r_2 = f(1/\alpha_2)$)*

$$F(\alpha_1, \alpha_2) = (1 + \alpha_2)r_1r_2 + r_1(1 - r_2) + (1 + \alpha_1)(1 - r_1)(1 - r_2) + \max\{\alpha_1, \alpha_2\}r_2(1 - r_1).$$

By this theorem, we can easily estimate the approximation ratio of a given mechanism M_f . In particular, by choosing $f(x) = 1 - \frac{1}{2x^{2.3}}$, we can compute that its approximation ratio is 1.5963. The function we used here is only an illustration of our mechanism M_f . We also believe that there exists a better f , though very hard to find. It is also an interesting problem to explore the property of function f with which M_f can have smaller approximation ratio.

Theorem 5. *For $f(x) = 1 - \frac{1}{2x^{2.3}}$, the mechanism M_f for two machines is universally truthful and can achieve an approximation ratio of 1.5963.*

Before we prove the theorem [4](#), we first give a lemma, which gives an alternative description of the allocation in our mechanism M_f . Its proof is direct from the definition of our mechanism. Since we already proved that our mechanism M_f is truthful, we can also denote the bid as t in the following.

Lemma 1. *For any type matrix t of the two machines, M_f allocates each task independently and for each task $j = 1, 2 \dots, n$, if $t_{1j} = 0$, always allocate it to machine 1, otherwise allocates it to machine 1 with probability $f(t_{2j}/t_{1j})$ and to machine 2 with probability $1 - f(t_{2j}/t_{1j})$.*

Proof of Theorem [4](#): Fix any instance $t = (t_{ij})$, let l_{opt} be the optimal makespan. Let O_1, O_2 be the sets of tasks assigned to machine 1 and machine 2 respectively in an optimal solution. Then we have

$$l_{opt} = \max\left\{ \sum_{j \in O_1} t_{1j}, \sum_{k \in O_2} t_{2k} \right\}.$$

Now we estimate the expected makespan of our mechanism M_f , denoted by l^f . We use $l_i^f, i = 1, 2$, to denote the completion time of machine i , then $l^f = \max\{l_1^f, l_2^f\}$. Let M be a random variable such that $M = 1$ if $l_1^f \geq l_2^f$, and $M = 2$ otherwise. We also denote $Pr(M = 1, x_{1j} = 1)$ as P_j^1 and $Pr(M = 2, x_{2j} = 1)$ as P_j^2 in the following calculation. Then we have:

$$\begin{aligned} l^f &= \sum_{j \in [m]} t_{1j}P_j^1 + t_{2j}P_j^2 \\ &= \sum_{j \in O_1} t_{1j} \left(P_j^1 + \frac{t_{2j}}{t_{1j}}P_j^2 \right) + \sum_{k \in O_2} t_{2k} \left(\frac{t_{1k}}{t_{2k}}P_k^1 + P_k^2 \right) \\ &\leq \max_{j \in O_1} \left(P_j^1 + \frac{t_{2j}}{t_{1j}}P_j^2 \right) \cdot \sum_{j \in O_1} t_{1j} + \max_{k \in O_2} \left(\frac{t_{1k}}{t_{2k}}P_k^1 + P_k^2 \right) \cdot \sum_{k \in O_2} t_{2k} \\ &\leq l_{opt} \left(\max_{j \in O_1} \left(P_j^1 + \frac{t_{2j}}{t_{1j}}P_j^2 \right) + \max_{k \in O_2} \left(\frac{t_{1k}}{t_{2k}}P_k^1 + P_k^2 \right) \right) \\ &\leq l_{opt} \left(\max_{j \neq k} \left(P_j^1 + \frac{t_{2j}}{t_{1j}}P_j^2 + \frac{t_{1k}}{t_{2k}}P_k^1 + P_k^2 \right) \right) \end{aligned}$$

So the approximate ratio is bounded by the term

$$\max_{j \neq k} \left(P_j^1 + \frac{t_{2j}}{t_{1j}}P_j^2 + \frac{t_{1k}}{t_{2k}}P_k^1 + P_k^2 \right).$$

Fix any j, k , let $\alpha_1 = \frac{t_{2j}}{t_{1j}}, \alpha_2 = \frac{t_{1k}}{t_{2k}}$ and $P_{abc} = Pr(M = a, x_{bj} = 1, x_{ck} = 1)$, $a, b, c \in \{0, 1\}$. Then we can expand P_j^1 as $P_{111} + P_{112}$, since

$$Pr(M = 1, x_{1j} = 1) = Pr(M = 1, x_{1j} = 1, x_{1k} = 1) + Pr(M = 1, x_{1j} = 1, x_{2k} = 1).$$

Let $r_1 = Pr(x_{1j} = 1)$, $r_2 = Pr(x_{1k} = 1)$, then we have:

$$\begin{aligned}
 & P_j^1 + \frac{t_{2j}}{t_{1j}}P_j^2 + \frac{t_{1k}}{t_{2k}}P_k^1 + P_k^2 \\
 &= (P_{111} + P_{112}) + \alpha_1(P_{221} + P_{222}) + \alpha_2(P_{111} + P_{121}) + (P_{212} + P_{222}) \\
 &= (P_{111} + P_{112} + P_{212}) + \alpha_2(P_{111} + P_{121} + P_{221}) + (\alpha_1 - \alpha_2)P_{221} + (1 + \alpha_1)P_{222} \\
 &\leq Pr(x_{1j} = 1) + \alpha_2 Pr(x_{1k} = 1) + (\alpha_1 - \alpha_2)Pr(M = 2, x_{2j=1}, x_{1k} = 1) \\
 &\quad + (1 + \alpha_1)Pr(M = 2, x_{2j} = 1, x_{2k} = 1) \\
 &\leq Pr(x_{1j} = 1) + \alpha_2 Pr(x_{1k} = 1) + \max\{\alpha_1 - \alpha_2, 0\}Pr(x_{2j} = 1, x_{1k} = 1) \\
 &\quad + (1 + \alpha_1)Pr(x_{2j} = 1, x_{2k} = 1) \\
 &= (1 + \alpha_2)r_1r_2 + r_1(1 - r_2) + (1 + \alpha_1)(1 - r_1)(1 - r_2) + \max\{\alpha_1, \alpha_2\}r_2(1 - r_1) \\
 &= F(\alpha_1, \alpha_2)
 \end{aligned}$$

The first inequality is because $Pr(x_{1j} = 1) = P_{111} + P_{112} + P_{211} + P_{212}$ and so on. The second inequality is because $Pr(M = 2, x_{2j=1}, x_{1k} = 1) \leq Pr(x_{2j} = 1, x_{2k} = 1)$. By lemma [□](#), $r_1 = f(\alpha_1)$, $r_2 = f(1/\alpha_2)$, hence the approximation ratio is bounded by $\max_{\alpha_1, \alpha_2 \in R^+} F(\alpha_1, \alpha_2)$.

On the other direction, we use the following instance to show that our analysis of the approximation ratio is tight. We will use the following tables to illustrated tasks and their allocation throughout this paper. There are two tasks A and B . The left table shows the instance t , where $t_{1A} = 1$, $t_{1B} = \alpha_2$, $t_{2A} = \alpha_1$, $t_{2B} = 1$. The right table shows the allocation of this instance using our mechanism M_f : task A is assigned to machine 1 with probability r_1 , to machine 2 with probability $1 - r_1$, etc. Here $r_1 = f(\alpha_1)$ and $r_2 = f(1/\alpha_2)$.

	machine 1	machine 2
task A	1	α_1
task B	α_2	1

→

	machine 1	machine 2
task A	r_1	$1 - r_1$
task B	r_2	$1 - r_2$

For this instance, we have $l_{opt} \leq 1$ and the expected makespan produced by M_f is exactly $F(\alpha_1, \alpha_2)$. So the approximation ratio is at least $F(\alpha_1, \alpha_2)$. □

3.2 Lower Bound for Task Independent Mechanisms

In this section, we show a lower bound for all task independent truthful mechanisms. This lower bound for task independent randomized truthful mechanisms is especially interesting, since a recent work in [\[3\]](#) shows that any truthful mechanism for two machines is task independent, however in the weaker version. So any lower bound better than 1.5 in the weaker version would imply an improvement of the lower bound 1.5 for randomized mechanisms for two machines case.

Theorem 6. *For any task independent truthful mechanism for two machines, its approximation ratio cannot be less than $11/7(\approx 1.5714)$.*

Proof. Given any task independent truthful mechanism M , consider the following four instances(a is a constant to be specified later, and $a > 1$). We can assume that $r_1 \geq 1/2$, otherwise we relabel the machines in instance 1, and modify the other three instances respectively.

Instance 1:

	machine 1	machine 2
task 1	1	1
task 2	1	2

 \rightarrow

	machine 1	machine 2
task 1	r_1	$1 - r_1$
task 2	r_2	$1 - r_2$

For this instance, we have $l_M/l_{opt} = 2r_1 + (1 - r_1)r_2 + 3(1 - r_1)(1 - r_2) \geq 1 + r_1 \triangleq A_1$.

Instance 2:

	machine 1	machine 2
task 1	1	1
task 2	1	a

 \rightarrow

	machine 1	machine 2
task 1	r_1	$1 - r_1$
task 2	r_3	$1 - r_3$

For this instance, we have $l_M/l_{opt} = 2r_1r_3 - r_1 - ar_3 + a + 1 \triangleq A_2$.

Instance 3:

	machine 1	machine 2
task 1	a	a^2
task 2	1	a

 \rightarrow

	machine 1	machine 2
task 1	r_4	$1 - r_4$
task 2	r_3	$1 - r_3$

For this instance, we have $l_M/l_{opt} = (1 + \frac{1}{a})r_3r_4 - r_3 - ar_4 + a + 1 \triangleq A_3$.

Instance 4:

	machine 1	machine 2
task 1	a	a
task 2	$2a$	a

 \rightarrow

	machine 1	machine 2
task 1	r_5	$1 - r_5$
task 2	r_6	$1 - r_6$

For this instance, we have $l_M/l_{opt} = 2 - r_5 + 2r_5r_6 \geq 2 - r_5$.

Consider instance 3 and 4, we can change task 2's values in instance 3 to $2a$, a without affecting the allocation of task 1 since M is task independent. Then we decrease machine 2's value on task 1 from a^2 to a . By corollary [II](#), we know the probability that machine 2 gets task 1 should increase. That is to say, $1 - r_5 \geq 1 - r_4$. Then we have $l_M/l_{opt} \geq 2 - r_4 \triangleq A_4$.

To sum up, mechanism M 's approximation ratio is at least $\max\{A_1, A_2, A_3, A_4\}$ with the condition $r_1 \geq 1/2, a > 1$, where $A_1 = 1+r_1, A_2 = 2r_1r_3-r_1-ar_3+a+1, A_3 = (1 + \frac{1}{a})r_3r_4 - r_3 - ar_4 + a + 1, A_4 = 2 - r_4$. Choosing $a = 3/2$ and using a case-by-case analysis, we can prove that $\max\{A_1, A_2, A_3, A_4\} \geq 11/7$ for any r_1, r_2, r_3, r_4 with the assumption $r_1 \geq 1/2$.

4 Scheduling m Machines

First we give the framework of our mechanism for scheduling m machines, BOUNDED-SQUARE mechanism. (Here we only give the allocation algorithm.)

Input: The reported bid matrix $b = (b_{ij})$.

Output: A randomized allocation $X = (X_{ij})$.

Allocation Algorithm:

(1) For each task $j = 1, 2 \dots, n$ do

let $I_j \leftarrow \{i \in [m] : b_{ij} \leq 2 \min_{i \in [m]} b_{ij}\}$.

if $\min_{i \in [m]} b_{ij} = 0$, we assign task j among the machines in I_j with equal probabilities;

Otherwise we use the SQUARE allocation algorithm [2] in I_j :

For each machine $i = 1, 2 \dots, m$ do:

if $i \in I_j$, $x_{ij} \leftarrow \frac{\frac{1}{(b_{ij})^2}}{\sum_{s \in I_j} \frac{1}{(b_{sj})^2}}$, otherwise $x_{ij} \leftarrow 0$.

(2) Round (x_{ij}) to a randomized integer solution (X_{ij}) such that $E[X_{ij}] = x_{ij}, \forall i, j$. We will specify the method of rounding later.

In our BOUNDED-SQUARE mechanism, $x = (x_{ij})$ can be viewed as a fractional solution of the scheduling problem. It is adapted from the fractional mechanism SQUARE in [2]. However, we need some “bid condition” in order to bound the loss of performance due to the rounding process. Here we give a threshold of $2 \min_{i \in [m]} b_{ij}$ in the allocation, so if $x_{ij} > 0$, then $b_{ij} \leq 2 \min_{i \in [m]} b_{ij} \leq 2l_{opt}(b)$. This idea plays an essential role in our mechanism.

Regarding the truthfulness of our mechanism, the proof is based on the fact that the modified fractional mechanism is still truthful. The proof is similar as in [2] and omitted here.

Lemma 2. *For any rounding method satisfying $E[X_{ij}] = x_{ij}, \forall i, j$, there is a payment algorithm to make BOUNDED-SQUARE mechanism truthful in expectation.*

Now we begin to analysis the approximation ratio of our mechanism. Since our mechanism is already proved truthful, we can assume that the players will report their types truthfully, and use t instead of b . Given an instance t , we first show that this fractional solution approximates $l_{opt}(t)$ within a factor of $\frac{m+1}{2}$. The proof is also omitted.

Lemma 3. *Let $x = (x_{ij})$ be the fractional solution obtained in the BOUNDED-SQUARE mechanism, we have $\max_{i \in [m]} \sum_{j \in [n]} x_{ij} t_{ij} \leq \frac{m+1}{2} l_{opt}(t)$.*

For the rounding method, we use the algorithm proposed by Kumar et al. [4].

Lemma 4. *(Kumar et al. [4]) Given a fractional assignment x and a processing time matrix t , there exists a randomized rounding procedure that yields a random integer assignment X such that,*

1. for any i, j , $E[X_{ij}] = x_{ij}$.
2. for any i , $\sum_j X_{ij} t_{ij} < \sum_j x_{ij} t_{ij} + \max_{j: x_{ij} \in (0,1)} t_{ij}$ with probability 1.

In our mechanism, we already know that $\max_{j: x_{ij} \in (0,1)} t_{ij} \leq 2l_{opt}(t)$ due to the bid condition. So putting everything together (lemma 2, lemma 3, lemma 4), we have the following theorem.

Theorem 7. *The BOUNDED-SQUARE mechanism is truthful in expectation and has an approximation ratio of $\frac{m+5}{2}$.*

References

1. Archer, A., Tardos, É.: Truthful mechanisms for one-parameter agents. In: FOCS 2001: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science, Washington, DC, USA, p. 482. IEEE Computer Society Press, Los Alamitos (2001)
2. Christodoulou, G., Koutsoupias, E., Kovács, A.: Mechanism design for fractional scheduling on unrelated machines. In: Arge, L., Cachin, C., Jurdziński, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 40–52. Springer, Heidelberg (2007)
3. Dobzinski, S., Sundararajan, M.: On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In: EC 2008: Proceedings of the 9th ACM conference on Electronic commerce, pp. 38–47. ACM, New York (2008)
4. Kumar, V., Marathe, M., Parthasarathy, S., Srinivasan, A.: Approximation algorithms for scheduling on multiple machines. In: Annual IEEE Symposium on Foundations of Computer Science, vol. 0, pp. 254–263 (2005)
5. Lu, P., Yu, C.: An improved randomized truthful mechanism for scheduling unrelated machines. In: Albers, P., Weil, P. (eds.) 25th International Symposium on Theoretical Aspects of Computer Science (STACS 2008), Dagstuhl, Germany, pp. 527–538. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany (2008)
6. Mu’alem, A., Schapira, M.: Setting lower bounds on truthfulness: extended abstract. In: SODA 2007: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms, pp. 1143–1152. Society for Industrial and Applied Mathematics, Philadelphia (2007)
7. Myerson, R.: Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73 (1981)
8. Nisan, N., Ronen, A.: Algorithmic mechanism design (extended abstract). In: STOC 1999: Proceedings of the thirty-first annual ACM symposium on Theory of computing, pp. 129–140. ACM, New York (1999)

Optimal Mechanisms for Single Machine Scheduling

Birgit Heydenreich¹, Debasis Mishra², Rudolf Müller¹, and Marc Uetz³

¹ Maastricht University, Quantitative Economics, P.O. Box 616, 6200 MD Maastricht, The Netherlands

{b.heydenreich,r.muller}@ke.unimaas.nl

² Indian Statistical Institute, Planning Unit, 7, S.J.S. Sansanwal Marg, New Delhi - 110 016, India

dmishra@isid.ac.in

³ University of Twente, Applied Mathematics, P.O. Box 217, 7500 AE Enschede, The Netherlands

m.uetz@utwente.nl

Abstract. We study the design of optimal mechanisms in a setting where job-agents compete for being processed by a service provider that can handle one job at a time. Each job has a processing time and incurs a waiting cost. Jobs need to be compensated for waiting. We consider two models, one where only the waiting costs of jobs are private information (1-d), and another where both waiting costs and processing times are private (2-d). An optimal mechanism minimizes the total expected expenses to compensate all jobs, while it has to be Bayes-Nash incentive compatible. We derive closed formulae for the optimal mechanism in the 1-d case and show that it is efficient for symmetric jobs. For non-symmetric jobs, we show that efficient mechanisms perform arbitrarily bad. For the 2-d case, we prove that the optimal mechanism in general does not even satisfy IIA, the ‘independent of irrelevant alternatives’ condition. We also show that the optimal mechanism is not even efficient for symmetric agents in the 2-d case.

1 Introduction

The design of optimal auctions is recognized as an intriguing issue in auction theory; first studied by Myerson (1981) for single item auctions. In that setting, the goal is to maximize the seller’s revenue. We study the design of optimal auctions (or more precisely, mechanisms) in a setting where job-agents compete for being processed by a service provider that can only handle one job at a time.

Our results. We consider two cases. In the *one-dimensional* (1-d) case, jobs’ processing times are public information and a job’s weight is only known to the job itself. Publicly known probability distributions over a finite set of possible weights represent common beliefs about the weights. In the *two-dimensional* (2-d) case, both weights and processing times are private information of the jobs. In both cases we aim at finding Bayes-Nash incentive compatible mechanisms that minimize the expected expenses of the service provider. Given jobs’

reports about their private information, a mechanism determines both an order in which jobs are served, and for each job a payment that the job receives. The payment can be seen as a compensation for waiting. By a graph theoretic interpretation of the incentive compatibility constraints - as used e.g. by Rochet [12] and Malakhov and Vohra [7] - we are able to derive optimal mechanisms. For the one-dimensional case, we obtain closed formulae for modified job weights, and show that serving the jobs in the order of non-increasing ratios of modified weights over service times is optimal for the service provider, as long as a certain regularity condition is fulfilled. It turns out that the optimal mechanism is not necessarily efficient, i.e., in general it does not maximize total utility. But it does so if e.g. all jobs are symmetric. For non-symmetric jobs, we show by example that the objective can be arbitrarily far from optimal if we insist on efficiency. We also compare our optimal mechanism to the generalized VCG mechanism and see that expected payments differ even for the case of symmetric jobs. For the two-dimensional case, our main result is that the optimal mechanism generally does not satisfy a property called IIA, ‘independent of irrelevant alternatives’. That implies that the optimal mechanism cannot be expressed in terms of modified weights along the lines of the 1-d case. In fact, any kind of priority based list scheduling algorithm where the priorities of a job depend only on the characteristics of that job itself cannot in general be an optimal mechanism. We conclude that optimal mechanism design for the two-dimensional case is substantially more involved than two-dimensional mechanism design for auction settings, as studied in [7]. We also show that even for symmetric jobs, in the 2-d case the optimal mechanism is not efficient.

Related Work. Myerson [11] studies optimal mechanisms for single item auctions and continuous 1-dimensional type spaces. Here, optimal auctions are modified Vickrey auctions, i.e. modified efficient auctions. When regarding the seller as additional agent who bids zero in the original auction, his modified bid might become non-zero in the optimal auctions yielding a reservation price. For a comparison between Myerson’s and our results, see Section 3. In [4], the authors give an introduction to optimal mechanism design with 1-dimensional continuous types under dominant strategy incentive compatibility. Both Myerson’s and our optimal allocation rules turn out to be dominant strategy implementable as well, while they yield optimal mechanisms in the larger class of Bayes-Nash incentive compatible mechanisms. Malakhov and Vohra [7] regard optimal mechanisms for an auction setting with discrete 2-dimensional type spaces. The derived optimal mechanisms again employ the efficient allocation rule with modified bids. We show that their approach must fail in our setting. For details, we refer to Section 4. Armstrong [1] studies a multi-object auction model where valuations are additive and drawn from a binary distribution (i.e. high or low). He gives optimal auctions under specific conditions that reduce the type graph. From this paper it becomes evident that optimal mechanism design with multi-dimensional discrete types is difficult. For our model, we formalize this difficulty by showing that traditional approaches inevitably yield IIA-mechanisms and therefore must

fail. Other scheduling models have been looked at from a different angle in the economic literature. See, e.g., [8] for efficient and budget-balanced mechanism design in a 1-dimensional model and [9] for mechanisms that prevent merging and splitting of jobs.

2 Optimal Mechanisms for the 1-Dimensional Setting

Setting & Preliminaries. Consider a single machine which can handle one job at a time. Let $J = \{1, \dots, n\}$ denote the set of non-preemptive jobs. We regard jobs as selfish agents that act strategically. Each job j has a processing time p_j and a weight w_j . While p_j is publicly known, the actual w_j is private information to job j . We refer to the private information of a job as its type. Jobs share common beliefs about other jobs' types in terms of probability distributions. We assume discrete distribution of weights, that is, agent j 's weight w_j follows a probability distribution over the discrete set $W_j = \{w_j^1, \dots, w_j^{m_j}\} \subset \mathbb{R}$, where $w_j^1 \leq \dots \leq w_j^{m_j}$. Let φ_j be the probability distribution of w_j , that is, $\varphi_j(w_j^i)$ denotes the probability associated with w_j^i for $i = 1, \dots, m_j$. Let $\Phi_j(w_j^i) = \sum_{k=1}^i \varphi_j(w_j^k)$ be the cumulative probability up to w_j^i . Both φ_j and Φ_j are public information. We assume that jobs' weights are independently distributed. Let us denote by $W = \prod_{j \in J} W_j$ the set of all type profiles. For any job j , let $W_{-j} = \prod_{k \neq j} W_k$. Let φ be the joint probability distribution of $w = (w_1, \dots, w_n)$. Then $\varphi(w) = \prod_{j=1}^n \varphi_j(w_j^i)$ for $w = (w_1^{i_1}, \dots, w_n^{i_n}) \in W$. Let w_{-j} and φ_{-j} be defined analogously. For $w_j^i \in W_j$ and $w_{-j} \in W_{-j}$, we denote by (w_j^i, w_{-j}) the type profile where job j has type w_j^i and the types of all other jobs are w_{-j} .

A direct revelation mechanisms consists of an allocation rule f and a payment scheme π . Jobs have to report their weights and they might report untruthfully if it suits them. Depending on those reports, the allocation rule selects a *schedule*, i.e. an order in which jobs are processed on the machine. The payment scheme assigns a payment that is made to jobs in order to reimburse them for their waiting cost. The payments can be seen as a reimbursement for waiting.

Let $\mathfrak{S} = \{\sigma \mid \sigma \text{ is a permutation of } (1, \dots, n)\}$ denote the set of all feasible schedules. Then the allocation rule is a mapping $f: W \rightarrow \mathfrak{S}$. For any schedule $\sigma \in \mathfrak{S}$, let σ_j be the position of job j in the ordering of jobs in σ . Then, by $S_j(\sigma) = \sum_{\sigma_k < \sigma_j} p_k$, we denote the start time or waiting time of job j in σ . If job j has waiting time S_j and actual weight w_j^i , it encounters a valuation of $-w_j^i S_j$. If j additionally receives payment π_j , his total utility is $\pi_j - w_j^i S_j$, i.e., we assume quasi-linear utilities. Let us denote by $ES_j(f, w_j^i) := \sum_{w_{-j} \in W_{-j}} S_j(f(w_j^i, w_{-j})) \varphi_{-j}(w_{-j})$ the expected waiting time of job j if it reports weight w_j^i and allocation rule f is applied. Denote by $E\pi_j(w_j^i) := \sum_{w_{-j} \in W_{-j}} \pi_j(w_j^i, w_{-j}) \varphi_{-j}(w_{-j})$ the expected payment to j . We assume that jobs aim at maximizing their expected utility.

Definition 1. A mechanism (f, π) is Bayes-Nash incentive compatible if for every agent j and every two types $w_j^i, w_j^k \in W_j$

$$E\pi_j(w_j^i) - w_j^i ES_j(f, w_j^i) \geq E\pi_j(w_j^k) - w_j^k ES_j(f, w_j^k) \tag{1}$$

under the assumption that all agents apart from j report truthfully. If for allocation rule f there exists a payment scheme π such that (f, π) is Bayes-Nash incentive compatible, then f is called Bayes-Nash implementable. The payment scheme π is referred to as an incentive compatible payment scheme.

In order to account for individual rationality, we need to guarantee non-negative utilities for all agents that report their true weight. To that end, we add a dummy weight $w_j^{m_j+1}$ to the type space W_j for every agent j . We assume $ES_j(f, w_j^{m_j+1}) = 0$ and $E\pi_j(w_j^{m_j+1}) = 0$ for all $j \in J$. Furthermore, we impose the incentive constraints $E\pi_j(w_j^i) - w_j^i ES_j(f, w_j^i) \geq E\pi_j(w_j^{m_j+1}) - w_j^i ES_j(f, w_j^{m_j+1})$ implying that $E\pi_j(w_j^i) - w_j^i ES_j(f, w_j^i) \geq 0$ for any Bayes-Nash incentive compatible mechanism (f, π) . Therefore, the dummy weights together with the mentioned assumptions guarantee that individual rationality is satisfied along with the incentive constraints. The dummy weight can be interpreted as an option for any job not to take part in the mechanism.

Definition 2. An allocation rule f satisfies monotonicity w.r.t. weights or short monotonicity if for every agent $j \in J$, $w_j^i < w_j^k$ implies that $ES_j(f, w_j^i) \geq ES_j(f, w_j^k)$.

Theorem 1. An allocation rule f is Bayes-Nash incentive compatible if and only if it satisfies monotonicity w.r.t. weights.

The proof is standard and therefore omitted. We refer, e.g., to [10] for details.

The Type Graph. A useful tool for deriving optimal mechanisms is the type graph. It has been used earlier, e.g. in [6,7,10]. The type graph² T_f is defined for a fixed agent j . T_f has node set W_j and contains an arc from any node w_j^i to any other node w_j^k of length

$$\ell_{ik} = w_j^i [ES_j(f, w_j^k) - ES_j(f, w_j^i)].$$

Here, ℓ_{ik} represents the gain in expected valuation for agent j by truthfully reporting type w_j^i instead of lying type w_j^k . The incentive constraints for a Bayes-Nash incentive compatible mechanism (f, π) and job j can be read as

$$E\pi_j(w_j^k) \leq E\pi_j(w_j^i) + w_j^i [ES_j(f, w_j^k) - ES_j(f, w_j^i)] = E\pi_j(w_j^i) + \ell_{ik}.$$

That is, the expected payments $E\pi_j(\cdot)$ constitute a node potential in T_f . A standard result in graph theory says that these node potentials exist if and only if there is no negative cycle in the graph. That is, Bayes-Nash implementability of an allocation rule f is equivalent to the fact that the type graph T_f for any agent

¹ The exact definitions of the type graph might differ in the papers depending on the underlying model.

² We suppress the dependence on agent j in the notation and simply write T_f .

j has no negative cycle. We then say that the T_f 's satisfy the *non-negative cycle property*. Monotonicity is equivalent to the fact that there is no negative cycle of length two in T_f . We call this property the *non-negative two-cycle property*. It follows from

$$\begin{aligned} \ell_{ik} + \ell_{ki} &= w_j^i[ES_j(f, w_j^k) - ES_j(f, w_j^i)] + w_j^k[ES_j(f, w_j^i) - ES_j(f, w_j^k)] \\ &= (w_j^i - w_j^k)[ES_j(f, w_j^k) - ES_j(f, w_j^i)]. \end{aligned}$$

The last term is non-negative for all jobs j and any two types w_j^i and w_j^k iff monotonicity holds.

Optimal Mechanisms. It is well known that scheduling in order of non-increasing weight over processing time ratios minimizes the sum of weighted start times $\sum_{j=1}^n w_j S_j(f(w))$ for any type profile $w \in W$, and therefore maximizes the total valuation of all agents. This allocation rule is known as Smith's rule [13]. The optimal mechanism that we derive deploys a slightly different allocation rule, namely Smith's rule with respect to certain modified weights.

Our goal is to set up a mechanism that is Bayes-Nash incentive compatible and among all such mechanisms minimizes the expected total payment that has to be made to the jobs. Given any Bayes-Nash incentive compatible mechanism (f, π) , one can obviously substitute the payment scheme by its expected payment scheme yielding $(f, E\pi(\cdot))$ without losing Bayes-Nash incentive compatibility. Moreover, the expected total payment to the agents remains unchanged under the substitution. Therefore, we restrict focus to mechanisms in which agents always receive a payment that is equal to the expected payment given the agent's report and which is independent of the specific report of the other agents and of the actual allocation.

Note that, unlike e.g. in [11], in the discrete setting considered here revenue equivalence does not hold. Therefore, there are possibly multiple payment schemes that make an allocation rule incentive compatible. Let f be an allocation rule and let $\pi^f(\cdot)$ be a payment scheme that minimizes expected expenses for the machine among all payment schemes that make f Bayes-Nash incentive compatible. More specifically, $\pi_j^f(w_j^i)$ denotes the payment to agent j declaring weight w_j^i under this optimal payment scheme. Let $P^{min}(f) = \sum_{j \in J} \sum_{w_j^i \in W_j} \varphi_j(w_j^i) \pi_j^f(w_j^i)$ be the minimum expected total expenses for allocation rule f . The following lemma specifies the optimal payment scheme for a given allocation rule.

Lemma 1. *For a Bayes-Nash implementable allocation rule f , the payment scheme defined by*

$$\pi_j^f(w_j^{m_j+1}) = 0, \quad \pi_j^f(w_j^i) = \sum_{k=i}^{m_j} w_j^k [ES_j(f, w_j^k) - ES_j(f, w_j^{k+1})] \text{ for } i = 1, \dots, m_j$$

is incentive compatible, individually rational and minimizes the expected total payment made to agents. The corresponding expected total payment is given by

$$P^{min}(f) = \sum_{j \in J} \sum_{i=1}^{m_j} \varphi_j(w_j^i) \bar{w}_j^i ES_j(f, w_j^i),$$

where the modified weights \bar{w}_j are defined as follows

$$\bar{w}_j^1 = w_j^1, \quad \bar{w}_j^i = w_j^i + (w_j^i - w_j^{i-1}) \frac{\Phi_j(w_j^{i-1})}{\varphi_j(w_j^i)} \text{ for } i = 2, \dots, m_j.$$

The proof relies on the observation that minimal expected payments can be expressed as shortest path lengths in the type graph; we refer to the full version [5] for more details.

Given the minimum payments per allocation rule, we want to specify the allocation rule f which minimizes $P^{min}(f)$ among all Bayes-Nash implementable allocation rules.

Definition 3. *If $f \in \arg \min\{P^{min}(f) \mid f: W \rightarrow \mathfrak{S}, f \text{ Bayes-Nash implementable}\}$, then we call the mechanism (f, π^f) an optimal mechanism.*

We will need the following regularity condition that ensures Bayes-Nash implementability of the allocation rule in our optimal mechanism.

Definition 4. *We say that regularity is satisfied if for every agent j and $i = 2, \dots, m_j - 1$*

$$w_j^i + (w_j^i - w_j^{i-1}) \frac{\Phi_j(w_j^{i-1})}{\varphi_j(w_j^i)} \leq w_j^{i+1} + (w_j^{i+1} - w_j^i) \frac{\Phi_j(w_j^i)}{\varphi_j(w_j^{i+1})}.$$

This implies that $\bar{w}_j^i < \bar{w}_j^k$ whenever $w_j^i < w_j^k$.

Note that regularity is satisfied e.g. if the differences $w_j^i - w_j^{i-1}$ are constant and the distribution has a non-increasing reverse hazard rate.

Theorem 2. *Let the modified weights be defined as in Lemma 1. Let f be the allocation rule that schedules jobs in order of non-increasing ratios \bar{w}_j/p_j . If regularity holds, then (f, π^f) is an optimal mechanism.*

Proof. We show that f is Bayes-Nash implementable and minimizes $P^{min}(f)$ among all Bayes-Nash implementable allocation rules. For any allocation rule f , it is not hard to see that we can rewrite $P^{min}(f)$ as follows, using independence of weight distributions. Let $W'_j = W_j \setminus \{w_j^{m_j+1}\}$ and $W' = \prod_{j \in J} W'_j$.

$$\begin{aligned} P^{min}(f) &= \sum_{j \in J} \sum_{w_j^i \in W'_j} \varphi_j(w_j^i) \bar{w}_j^i ES_j(f, w_j^i) \\ &= \sum_{w \in W'} \varphi(w) \sum_{j \in J} \bar{w}_j S_j(f(w)). \end{aligned}$$

Thus, $P^{min}(f)$ can be minimized by minimizing $\sum_{j \in J} \bar{w}_j S_j(f(w))$ for every reported type profile w . This is achieved by scheduling in order of non-increasing

ratios \bar{w}_j/p_j . Under Smith's rule, the expected start time $ES_j(w_j)$ is clearly non-increasing in the modified weight \bar{w}_j . The regularity condition ensures that it is non-increasing in the original weights w_j . Therefore, Smith's rule with respect to modified weights satisfies monotonicity and is hence Bayes-Nash implementable by Theorem 1. This completes the proof. \square

It is not hard to see that the optimal allocation rule – Smith's rule with respect to modified weights – is even dominant strategy implementable, with the same total expected payment for the mechanism.

3 Optimality Versus Efficiency

For symmetric agents the optimal and the efficient allocation coincide.

Corollary 1. *If agents are symmetric, i.e. $W_1 = \dots = W_n$, $\varphi_1 = \dots = \varphi_n$ and $p_1 = \dots = p_n$ and if distributions are such that regularity holds, then the optimal mechanism is efficient.*

If weight distributions differ among agents or if agents have different processing times, then the optimal mechanism is in general not efficient. In fact, when restricting to efficient mechanisms, the total expected payment can be arbitrarily bad in comparison to the optimal one. This is illustrated by the following two examples; proofs can be found in the full version of this paper [5].

Example 1. Let there be two jobs 1 and 2 with $W_1 = \{M + 1\}$ and $W_2 = \{1, M\}$ for some constant M . Let $\varphi_2(1) = 1 - 1/M$, $\varphi_2(M) = 1/M$ and $p_1 = p_2 = 1$. Let Eff be the efficient and Opt be the optimal allocation rule. Then the ratio $P^{min}(Eff)/P^{min}(Opt)$ goes to infinity as M goes to infinity.

Remark 1. In the above, the ratio of the expected payments of the efficient versus the optimal allocation rule is analyzed. It is also easy to see that the expected ratio of the payments tends to infinity as M approaches infinity.

Example 2. Let there be two jobs 1 and 2 with the same weight distribution $W_1 = W_2 = \{1, M\}$, $\varphi_j(1) = 1 - 1/M$, $\varphi_j(M) = 1/M$ for $j = 1, 2$. Let $p_1 = 1/2$ and $p_2 = M/2 + 1$. Let Eff be the efficient and Opt be the optimal allocation rule. Then the ratio $P^{min}(Eff)/P^{min}(Opt)$ goes to infinity as M goes to infinity.

Remark 2. As in the first example, it is easy to see that the expected ratio of the payments tends to infinity as M approaches infinity.

Comparison to Myerson's result. For the single item auction and continuous type spaces, Myerson [11] has made similar observations: in his setting, the efficient auction is the Vickrey auction. The optimal auction can be seen as a modified Vickrey Auction with the seller submitting a bit himself. In our setting also, the allocation in the optimal mechanism is equivalent to the efficient allocation rule with respect to modified data. Nevertheless, in [11] the optimal and the efficient mechanism may differ. For the single item auction this can

be due to the seller keeping the item (even in the symmetric case) or because a bidder that has not submitted the highest bid can get the item in the asymmetric case. In our setting, the optimal and the efficient mechanism can only differ if agents are asymmetric, see Corollary 1 and Examples 1 and 2.

On the generalized VCG Mechanism. The VCG mechanism is due to Vickrey [14], Clarke [2] and Groves [3]. The allocation rule is the efficient one. In our setting this means scheduling in order of non-increasing ratios w_j/p_j . The payment scheme can be shown to be

$$\pi_j^{VCG}(w) = p_j \sum_{\substack{k \in J \\ \sigma_k < \sigma_j}} w_k,$$

where w is the reported type profile and σ the efficient schedule. As illustrated by examples 1 and 2, the allocation of the VCG mechanism can differ from the allocation of the optimal mechanism if agents are not symmetric. Moreover, if jobs are symmetric, the VCG mechanism still can be non-optimal in terms of payments. This is illustrated by the following example.

Example 3. There are two symmetric agents with $W_1 = W_2 = \{w^1, w^2\}$, $w^1 < w^2$, and $\varphi_j(w^1) = \varphi_j(w^2) = 1/2$ for $j = 1, 2$. Processing times are equal (w.l.o.g., $p_1 = p_2 = 1$). Then the expected expenses of the VCG mechanism can be shown to be strictly higher than those of the optimal mechanism.

4 The 2-Dimensional Setting

Setting and Notation. In contrast to the 1-dimensional setting, both weight and processing time of a job are now private information of the job. Hence j 's type is the tuple (w_j, p_j) . We assume public probability distribution information, i.e. $(w_j, p_j) \in W_j \times P_j$, where $W_j = \{w_j^1, \dots, w_j^{m_j}\}$ with $w_j^1 \leq \dots \leq w_j^{m_j}$ and $P_j = \{p_j^1, \dots, p_j^{q_j}\}$ with $p_j^1 \leq \dots \leq p_j^{q_j}$. Let φ_j be the probability distribution of j 's type, that is, $\varphi_j(w_j^i, p_j^k)$ denotes the probability associated with the type (w_j^i, p_j^k) for $i = 1, \dots, m_j$ and $k = 1, \dots, q_j$. Both φ_j and Φ_j are public. Distributions are independent between agents. Denote by $T = \prod_{j \in J} (W_j \times P_j)$ the set of all type profiles. For any job j , let $T_{-j} = \prod_{r \neq j} (W_r \times P_r)$ be the set of type profiles of all jobs except j . Let φ be the joint probability distribution of $(w_1, p_1, \dots, w_n, p_n)$. Then for type profile $t = (w_1^i, p_1^k, \dots, w_n^i, p_n^k) \in T$, $\varphi(t) = \prod_{j=1}^n \varphi_j(w_j^i, p_j^k)$. Let t_{-j} and φ_{-j} be defined analogously. For $(w_j^i, p_j^k) \in W_j \times P_j$ and $t_{-j} \in T_{-j}$, we denote by $((w_j^i, p_j^k), t_{-j})$ the type profile where job j has type (w_j^i, p_j^k) and the types of the other jobs are represented by t_{-j} . Denote by $ES_j(f, w_j^i, p_j^k) := \sum_{t_{-j} \in T_{-j}} S_j(f((w_j^i, p_j^k), t_{-j})) \varphi_{-j}(t_{-j})$ the expected waiting time of job j if he reports type (w_j^i, p_j^k) and allocation rule f is applied. Denote by $E\pi_j(w_j^i, p_j^k) := \sum_{t_{-j} \in T_{-j}} \pi_j((w_j^i, p_j^k), t_{-j}) \varphi_{-j}(t_{-j})$ the expected payment to j .

We assume that an agent can only report a processing time that is not lower than his true processing time and that a job is processed for his reported processing time. This is a natural assumption, since a job can add unnecessary work to achieve a longer processing time, but reporting a shorter processing time can easily be punished by preempting the job after the declared processing time (before it is actually finished).

Note that by regarding the processing time as private information, we introduce informational externalities: job j has a different valuation for a schedule if the processing time (and hence the type) of a job scheduled before j changes. In this regard, our model differs from the auction models studied in [11] and [7].

4.1 Bayes-Nash Implementability and the Type Graph

Definition 5. A mechanism (f, π) is called Bayes-Nash incentive compatible if for every agent j and every two types $(w_j^{i_1}, p_j^{k_1})$ and $(w_j^{i_2}, p_j^{k_2})$ with $i_1, i_2 \in \{1, \dots, m_j\}$, $k_1, k_2 \in \{1, \dots, q_j\}$, $k_1 \leq k_2$,

$$E\pi_j(w_j^{i_1}, p_j^{k_1}) - w_j^{i_1} ES_j(f, w_j^{i_1}, p_j^{k_1}) \geq E\pi_j(w_j^{i_2}, p_j^{k_2}) - w_j^{i_1} ES_j(f, w_j^{i_2}, p_j^{k_2}) \quad (2)$$

under the assumption that all agents apart from j report truthfully.

Note that by defining the incentive constraints only for $k_1 \leq k_2$, we account for the fact that agents can only overstate their processing time, but cannot understate it.

In order to ensure individual rationality, again add a dummy type t_j^d to the type space for every agent j , and let $ES_j(f, t_j^d) = 0$ and $E\pi_j(t_j^d) = 0$ for all $j \in J$. As in the 1-dimensional case, the dummy types together with the mentioned extra incentive constraints guarantee that individual rationality is satisfied along with the incentive constraints. Sometimes, it will be convenient to write $(w_j^{m_j+1}, p_j^k)$ for some $k \in \{1, \dots, q_j\}$ instead of t_j^d .

In the 2-dimensional setting, the type graph T_f of agent j has node set $W_j \times P_j$ and contains an arc from any node $(w_j^{i_1}, p_j^{k_1})$ to every other node $(w_j^{i_2}, p_j^{k_2})$ with $i \in \{1, \dots, m_j\}$, $i_2 \in \{1, \dots, m_j + 1\}$, $k \in \{1, \dots, q_j\}$, $k_1 \leq k_2$ of length

$$\ell_{(i_1 k_1)(i_2 k_2)} = w_j^{i_1} [ES_j(f, w_j^{i_2}, p_j^{k_2}) - ES_j(f, w_j^{i_1}, p_j^{k_1})].$$

Note that we have arcs only in direction of increasing processing times, since agents can only overstate their processing time. Furthermore, every node has an arc to the dummy type, but there are no outgoing arcs from the dummy type.

Definition 6. An allocation rule f satisfies monotonicity w.r.t. weights if for every agent $j \in J$ and fixed $p_j^k \in P_j$, $w_j^{i_1} < w_j^{i_2}$ implies that $ES_j(f, w_j^{i_1}, p_j^k) \geq ES_j(f, w_j^{i_2}, p_j^k)$.

Theorem 3. An allocation rule f is Bayes-Nash incentive compatible in the 2-dimensional setting if and only if it satisfies monotonicity with respect to weights.

Proof. The claim reduces to showing that in the type graph of any agent j the non-negative cycle property is equivalent to the non-negative two-cycle property. Since there is an arc from a node representing type $(w_j^{i_1}, p_j^{k_1})$ to the node representing type $(w_j^{i_2}, p_j^{k_2})$ if and only if $p_j^{k_1} \leq p_j^{k_2}$, cycles can only occur between nodes representing types with equal processing times. Hence, the proof is analogous to the 1-dimensional case. \square

Similar as in [7], one can show that some arcs in the type graph are not necessary, since the corresponding incentive constraints are implied by others. The reduced type graph of agent j contains only arcs that are necessary in that sense. A sketch of the reduced type graph is given in Figure 1. Expected payments correspond to node potentials in the reduced type graph. The reduced type graph comes handy particularly when considering our (counter) examples in the next subsection.




Fig. 1. Reduced type graph 2-d case

4.2 On Optimal Mechanisms

We start by quickly reviewing an approach to two-dimensional optimal mechanism design studied in [7]. Here, the authors regard a limited-supply multi-item auction, where each agent's type (i, j) is given by a marginal valuation i per item and a capacity j . Above that capacity, the agent has zero valuation for each additional item. The goal is revenue maximization. Bayes-Nash implementability is equivalent to the expected amount of items allocated to an agent being monotone in his reported value for i . Malakhov and Vohra [7] use the type graph approach to derive optimal mechanisms in this 2-d setting. Note, however, that the approach of [7], and also our approach for the 1-dimensional setting focus on one agent and its type graph. Hence, in terms of the scheduling model considered here, any optimal allocation rule derived this way is necessarily a modified Smith's rule with modified weights that can be computed from the characteristics (type report and distribution) of the agent itself. Such an allocation rule necessarily satisfies the following IIA property.

Definition 7. We say that an allocation rule f is independent of irrelevant alternatives (IIA) if the relative order of any two jobs j_1 and j_2 is the same in the schedules $f(t_1)$ and $f(t_2)$ for any two type profiles $t_1, t_2 \in T$ that differ only in the types of agents from $J \setminus \{j_1, j_2\}$.

In other words, the relative order of two jobs is independent of all other jobs. For the 2-d setting, this is not necessarily the case for optimal mechanisms.

Theorem 4. *The optimal allocation rule for the 2-dimensional setting does in general not satisfy IIA.*

Proof. The proof uses the following instance with three jobs. Job 1 has type $(1, 1)$, job 2 has type $(2, 2)$ and job 3 has type space $\{1.9, 2\} \times \{1, 2\}$. The probabilities for job 3's types are $\varphi_3(1.9, 1) = 0.8$, $\varphi_3(2, 2) = 0.2$ and $\varphi_3(1.9, 2) = \varphi_3(2, 1) = 0$ respectively. We show that the best allocation rule that satisfies IIA achieves a minimum expected total payment of at least 5.6, whereas there exists an allocation rule – violating IIA – with an expected total payment of 4.88. The details are contained in the full version of this paper [5]. \square

Theorem 4 shows that any list scheduling algorithm where the priority of a job can be computed from the characteristics of the job itself cannot be optimal in general. Moreover, the type graph approach must fail, since it focusses on a single agent. Hence, optimal mechanism design for our 2-dimensional setting is considerably more complicated than for the 1-dimensional setting and for traditional auction settings as described in [11] and [7].

One explanation for this complication may lie in the fact that the 2-d setting considered here in fact entails informational externalities, as opposed to the auction settings in [11] and [7]. On the other hand, the informational externalities introduced by private processing times are not the only cause for complications in the 2-dimensional setting: Consider the 1-dimensional setting, where only the processing times are private, but the weights are public information. It turns out that all allocation rules are implementable, even when we allow that jobs understate their processing times. The optimal payment to a job j that reports processing time p_j^k is equal to $w_j ES_j(f, p_j^k)$, and therefore the total payment to jobs for allocation rule f is equal to $P^{min}(f) = \sum_{j \in J} \sum_{k=1}^{q_j} \varphi_j(p_j^k) w_j ES_j(f, p_j^k)$. This is minimized by Smiths rule.

When there are only two agents present, then IIA is trivially satisfied. Recall that in the 1-dimensional case the optimal mechanism is efficient for symmetric agents and regular distributions and that the uniform distribution is regular. This is contrasted by the following theorem.

Theorem 5. *Even for two symmetric agents, 2×2 -type spaces and uniform probability distributions, the optimal mechanism is not efficient.*

Proof. We show that the efficient allocation is for some instances dominated by the w -rule, which schedules the job with the higher weight first. For details we refer to the full version of this paper [5]. \square

5 Conclusion

We have seen that the graph theoretic approach is an intuitive tool for optimal mechanism design, and yields a closed formula for the optimal mechanism in the 1-d discrete case. The same approach works for the continuous case, too.

Moreover, we have seen that in the two-dimensional case the canonical approach does not work and that optimal mechanism design seems to be considerably more complicated than in the traditional auction models. We leave it as an open problem to identify (closed formulae for) optimal mechanisms for the 2-d case. It is conceivable, however, that closed formulae don't exist.

Acknowledgements

Birgit Heydenreich was supported by NWO grant 2004/ 03545/MaGW 'Local Decisions in Decentralised Planning Environments'. Debasis Mishra acknowledges support by NWO (visiting grants B 46-593 and 040.11.032) and by METEOR, the Maastricht Research School of Economics of Technology and Organizations. We thank the anonymous referees for some helpful remarks.

References

1. Armstrong, M.: Optimal multi-object auctions. *Review of Economic Studies* 67, 455–481 (2000)
2. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* 11(1), 17–33 (1971)
3. Groves, T.: Incentives in teams. *Econometrica* 41, 617–631 (1973)
4. Hartline, J., Karlin, A.: Profit maximization in mechanism design. In: Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
5. Heydenreich, B., Mishra, D., Müller, R., Uetz, M.: Optimal Mechanisms for Single Machine Scheduling. Research Memorandum RM/08/033, Maastricht University (2008)
6. Heydenreich, B., Müller, R., Uetz, M., Vohra, R.: Characterization of revenue equivalence. *Econometrica* (forthcoming, 2008)
7. Malakhov, A., Vohra, R.: An optimal auction for capacity constrained bidders: a network perspective. *Economic Theory* (forthcoming, 2007)
8. Mitra, M.: Mechanism design in queueing problems. *Economic Theory* 17(2), 277–305 (2001)
9. Moulin, H.: On scheduling fees to prevent merging, splitting, and transferring of jobs. *Mathematics of Operations Research* 32, 266–283 (2007)
10. Müller, R., Perea, A., Wolf, S.: Weak monotonicity and Bayes-Nash incentive compatibility. *Games and Economic Behavior* 61(2), 344–358 (2007); Extended version available as Research Memorandum RM/05/040, Maastricht University (2005)
11. Myerson, R.: Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73 (1981)
12. Rochet, J.-C.: A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics* 16(2), 191–200 (1987)
13. Smith, W.: Various optimizers for single stage production. *Naval Research Logistics Quarterly* 3, 59–66 (1956)
14. Vickrey, W.: Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* 16, 8–37 (1961)

Welfare Undominated Groves Mechanisms

Krzysztof Apt^{1,2}, Vincent Conitzer³, Mingyu Guo³, and Evangelos Markakis¹

¹ Centre for Math and Computer Science (CWI), Amsterdam, The Netherlands
{apt, vangelis}@cwi.nl

² University of Amsterdam, Institute of Language, Logic and Computation,
Amsterdam, The Netherlands

³ Duke University, Department of Computer Science, Durham, NC, USA
{conitzer, mingyu}@cs.duke.edu

Abstract. A common objective in mechanism design is to choose the outcome (for example, allocation of resources) that maximizes the sum of the agents' valuations, without introducing incentives for agents to misreport their preferences. The class of Groves mechanisms achieves this; however, these mechanisms require the agents to make payments, thereby reducing the agents' total welfare.

In this paper we introduce a measure for comparing two mechanisms with respect to the final welfare they generate. This measure induces a partial order on mechanisms and we study the question of finding minimal elements with respect to this partial order. In particular, we say a non-deficit Groves mechanism is *welfare undominated* if there exists no other non-deficit Groves mechanism that always has a smaller or equal sum of payments. We focus on two domains: (i) auctions with multiple identical units and unit-demand bidders, and (ii) mechanisms for public project problems. In the first domain we analytically characterize all welfare undominated Groves mechanisms that are anonymous and have linear payment functions, by showing that the family of optimal-in-expectation linear redistribution mechanisms, which were introduced in [6] and include the Bailey-Cavallo mechanism [12], coincides with the family of welfare undominated Groves mechanisms that are anonymous and linear in the setting we study. In the second domain we show that the classic VCG (Clarke) mechanism is welfare undominated for the class of public project problems with equal participation costs, but is not undominated for a more general class.

1 Introduction

Mechanism design is often employed for coordinating group decision making among agents. Often, such mechanisms impose taxes that agents have to pay to a central authority. Although maximizing tax revenue is a desirable objective in many settings (for example, if the mechanism is an auction designed by the seller), it is not desirable in situations where no entity is profiting from the taxes. Some examples include public project problems as well as certain resource allocation problems without a seller (e.g., the right to use a shared good on a given time slot, or the exchange of take-off slots among airline companies). In such cases, we would like to have mechanisms that minimize the sum of the taxes (or, even better, achieve budget balance, that is, the sum

of the taxes is zero), while maintaining other desirable properties, such as efficiency, strategy-proofness and non-deficit (i.e., the mechanism does not need to be funded by an external source).

The well-known VCG mechanism¹ is efficient, strategy-proof and incurs no deficit. More generally, the family of Groves mechanisms, which includes VCG, is a family of efficient and strategy-proof mechanisms. Unfortunately though, Groves mechanisms are not budget balanced. In fact, in sufficiently general settings, it is impossible to have a mechanism that satisfies efficiency, strategy-proofness, and budget balance [4].

We therefore consider the following problem: within the family of Groves mechanisms, we want to identify non-deficit mechanisms that are optimal with respect to the sum of the payments, i.e., we cannot lower the mechanism's payments without violating efficiency, strategy-proofness or the non-deficit property. Such a mechanism, in a sense, maximizes the agents' welfare (among efficient mechanisms²). To make this precise, we first introduce a measure for comparing two feasible mechanisms (mechanisms that are efficient, strategy-proof and satisfy the non-deficit property). We say that a feasible Groves mechanism M *welfare dominates* another feasible Groves mechanism M' if for every type vector of the agents, the sum of the payments under M is no more than the sum of the payments under M' , and this holds with strict inequality for at least one type vector. This definition induces a partial order on feasible Groves mechanisms and we wish to identify minimal elements in this partial order. We call such minimal elements *welfare undominated*. Other partial orders, as well as other notions of optimality, have recently been considered in other work on redistribution mechanisms (see Section [1.1]). The notion of optimality that we study here is different from the previously studied ones at both a conceptual and a technical level, as we illustrate below.

We study the question of finding welfare undominated mechanisms in two domains. The first is auctions of multiple identical units with unit-demand bidders. In this setting, it is easy to see that VCG is welfare dominated by other Groves mechanisms, such as the Bailey-Cavallo mechanism [11,2]. We obtain a complete characterization of linear and anonymous redistribution mechanisms that are minimal elements in this partial order: we show that a linear, anonymous Groves mechanism is welfare undominated if and only if it belongs to the class of *Optimal-in-Expectation Linear (OEL) redistribution mechanisms*, which include the Bailey-Cavallo mechanism and were introduced in [6]. The second domain is public project problems, where a set of agents must decide on financing a project (e.g., building a bridge). Here, we show that in the case where the agents have identical participation costs, no mechanism welfare dominates the VCG mechanism. On the other hand, when the participation costs can be different across agents, there exist mechanisms that welfare dominate VCG. In both domains, our proofs rely on some general properties we establish for anonymous mechanisms, which may be of independent interest (see Section [3]).

The omitted proofs appear in the full version of the paper.

¹ In this paper, "the VCG mechanism" refers to the Clarke mechanism (aka pivotal mechanism), not to any other Groves mechanism.

² By sacrificing efficiency, it is sometimes possible to drastically lower the payments, so that the net effect is an increase in the agents' welfare [5,3]. However, most of the prior work has focused on the case where efficiency is a hard constraint, and we will do so in this paper.

1.1 Related Work

Recently, there has been a series of works on redistribution mechanisms, which are Groves mechanisms that redistribute some of the VCG payment back to the bidders. Bailey and Cavallo [12] introduced a mechanism that welfare dominates VCG in some cases, such as single-item auctions, but coincides with VCG in some more general settings. We will refer to this mechanism as the BC mechanism from now on (in fact, Bailey’s mechanism is not always the same as Cavallo’s mechanism, but it is in the settings in which we study it). A special case of the BC mechanism was independently discovered by Porter *et al.* [14]. Cavallo also proved that the BC mechanism is optimal among the family of *surplus-anonymous* mechanisms; however, this is a quite restrictive class of mechanisms. Guo and Conitzer [8] solved for a worst-case optimal redistribution mechanism for multi-unit auctions with nonincreasing marginal values. Moulin [13] independently derived the same mechanism under a slightly different worst-case optimality notion (in the more restrictive setting of multi-unit auctions with unit demand only). These worst-case notions are different notions of optimality than the one we consider in this paper. Guo and Conitzer [6] also solve for mechanisms that maximize expected redistribution (in a certain class of mechanisms), when a prior is available. Another notion of optimality, which is closer to the one studied in this paper, was introduced in [7], namely the notion of *undominated* mechanisms. A mechanism is undominated if there is no other mechanism under which every *individual* agent pays weakly less for every type vector, and strictly less in at least one case. This is a weaker concept than ours, in the sense that for a mechanism that is undominated, there may still exist mechanisms that welfare dominate it (by increasing the payment from some agents to decrease the payments from other agents more). In the other direction, if a mechanism is welfare undominated, then it is also undominated. We believe that the notion we study in this paper is more appropriate when one is interested in the final welfare of the agents. Technically, welfare undominance appears much more challenging and seems to require different techniques.

2 Preliminaries

2.1 Tax-Based Mechanisms

We first briefly review tax-based mechanisms (see, e.g., [10]). Assume that there is a set of possible outcomes or *decisions* D , a set $\{1, \dots, n\}$ of players where $n \geq 2$, and for each player i a set of *types* Θ_i and an (*initial*) *utility function* $v_i : D \times \Theta_i \rightarrow \mathbb{R}$. Let $\Theta := \Theta_1 \times \dots \times \Theta_n$.

In a (direct revelation) mechanism, each player reports a type θ_i and based on this, the mechanism selects an outcome and a payment to be made by every agent. Hence a mechanism is given by a pair of functions (f, t) , where f is the decision function and $t = (t_1, \dots, t_n)$ is the tax function that determines the players’ payments, i.e., $f : \Theta \rightarrow D$, and $t : \Theta \rightarrow \mathbb{R}^n$.

We assume that the (*final*) *utility function* for player i is a function $u_i : D \times \mathbb{R}^n \times \Theta_i \rightarrow \mathbb{R}$ defined by $u_i(d, t_1, \dots, t_n, \theta_i) := v_i(d, \theta_i) + t_i$ (that is, utilities are *quasilinear*). For each vector θ of announced types, if $t_i(\theta) \geq 0$, player i *receives* $t_i(\theta)$, and if

$t_i(\theta) < 0$, he **pays** $|t_i(\theta)|$. Thus when the true type of player i is θ_i and his announced type is θ'_i , his final utility is

$$u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i) = v_i(f(\theta'_i, \theta_{-i}), \theta_i) + t_i(\theta'_i, \theta_{-i}),$$

where θ_{-i} are the types announced by the other players.

2.2 Properties of Tax-Based Mechanisms

We say that a tax-based mechanism (f, t) is

- **efficient** if for all $\theta \in \Theta$ and $d' \in D$, $\sum_{i=1}^n v_i(f(\theta), \theta_i) \geq \sum_{i=1}^n v_i(d', \theta_i)$,
- **budget-balanced** if $\sum_{i=1}^n t_i(\theta) = 0$ for all $\theta \in \Theta$,
- **feasible** if $\sum_{i=1}^n t_i(\theta) \leq 0$ for all θ , i.e., the mechanism does not need to be funded by an external source,
- **pay-only** if $t_i(\theta) \leq 0$ for all θ and all $i \in \{1, \dots, n\}$,
- **strategy-proof** if for all $\theta, i \in \{1, \dots, n\}$ and θ'_i ,

$$u_i((f, t)(\theta_i, \theta_{-i}), \theta_i) \geq u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i).$$

Tax-based mechanisms can be compared in terms of the final social welfare they generate ($\sum_{i=1}^n u_i((f, t)(\theta), \theta_i)$). More precisely, one can define the following two natural partial orders as a way to compare mechanisms. The first was introduced in [7]. The second is the concept that we introduce and study in this paper, which we believe is a more appropriate concept when one is interested in the final social welfare of the agents.

Definition 1. Given two tax-based mechanisms (f, t) and (f', t') we say that (f', t') **dominates** (f, t) (due to [7]) if

- for all $\theta \in \Theta$ and all $i \in \{1, \dots, n\}$, $u_i((f, t)(\theta), \theta_i) \leq u_i((f', t')(\theta), \theta_i)$,
- for some $\theta \in \Theta$ and some $i \in \{1, \dots, n\}$, $u_i((f, t)(\theta), \theta_i) < u_i((f', t')(\theta), \theta_i)$.

Definition 2. Given two tax-based mechanisms (f, t) and (f', t') we say that (f', t') **welfare dominates** (f, t) if

- for all $\theta \in \Theta$, $\sum_{i=1}^n u_i((f, t)(\theta), \theta_i) \leq \sum_{i=1}^n u_i((f', t')(\theta), \theta_i)$,
- for some $\theta \in \Theta$, $\sum_{i=1}^n u_i((f, t)(\theta), \theta_i) < \sum_{i=1}^n u_i((f', t')(\theta), \theta_i)$.

In this paper, we are interested only in Groves mechanisms, so that the decision function f is always efficient, and (welfare) dominance is strictly due to differences in the tax function t . Specifically, in this context we have that (f, t') dominates (f, t) (or simply t' dominates t) if and only if

- for all $\theta \in \Theta$ and all $i \in \{1, \dots, n\}$, $t_i(\theta) \leq t'_i(\theta)$, and
- for some $\theta \in \Theta$ and some $i \in \{1, \dots, n\}$, $t_i(\theta) < t'_i(\theta)$,

and t' welfare dominates t if

- for all $\theta \in \Theta$, $\sum_{i=1}^n t_i(\theta) \leq \sum_{i=1}^n t'_i(\theta)$, and
- for some $\theta \in \Theta$, $\sum_{i=1}^n t_i(\theta) < \sum_{i=1}^n t'_i(\theta)$.

For two tax-based mechanisms t, t' , it is clear that if t' dominates t , then it also welfare dominates t . The reverse implication, however, does not need to hold³

We now define a transformation on tax-based mechanisms originating from the same decision function. This transformation was originally defined in [1] and [2] for the specific case of the VCG mechanism and in [7] for feasible Groves mechanisms. We call it the **BCGC transformation** after the authors of these papers.

Consider a tax-based mechanism (f, t) . Given $\theta = (\theta_1, \dots, \theta_n)$, let $T(\theta)$ be the total amount of taxes, i.e., $T(\theta) := \sum_{i=1}^n t_i(\theta)$. For each $i \in \{1, \dots, n\}$ let⁴

$$S_i^{BCGC}(\theta_{-i}) := \max_{\theta'_i \in \Theta_i} T(\theta'_i, \theta_{-i}).$$

We then define the tax-based mechanism t^{BCGC} as follows:

$$t_i^{BCGC}(\theta) := t_i(\theta) - \frac{S_i^{BCGC}(\theta_{-i})}{n}.$$

The following observations generalize some of the results of [1][2][7].

Note 1.

- (i) Each tax-based mechanism of the form t^{BCGC} is feasible.
- (ii) If t is feasible, then either t and t^{BCGC} coincide or t^{BCGC} dominates t .

2.3 Groves Mechanisms

Each **Groves mechanism** is a tax-based mechanism (f, t) such that the following hold⁵:

- $f(\theta) \in \arg \max_d \sum_{i=1}^n v_i(d, \theta_i)$, i.e., the chosen outcome maximizes the initial social welfare.
- $t_i : \Theta \rightarrow \mathbb{R}$ is defined by $t_i(\theta) := g_i(\theta) + h_i(\theta_{-i})$,
- $g_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j)$,
- $h_i : \Theta_{-i} \rightarrow \mathbb{R}$ is an arbitrary function.

Intuitively, $g_i(\theta)$ represents the (initial) social welfare from the decision $f(\theta)$, when player i 's (initial) utility is not counted. We now recall the following result (e.g., [10]):

Groves Theorem Every Groves mechanism (f, t) , is efficient and strategy-proof.

For several decision problems the only efficient and strategy-proof tax-based mechanisms are Groves mechanisms. By a general result of [9] this is the case for both domains that we consider in this paper and explains our focus on Groves mechanisms.

A feasible Groves mechanism is **undominated** if there is no other feasible Groves mechanism that dominates it [7]. A feasible Groves mechanism is **welfare undominated** if there is no other feasible Groves mechanism that welfare dominates it. Welfare undominance is a strictly stronger concept than undominance, as is illustrated in Appendix A.

³ In Appendix A we provide an example of two tax-based mechanisms that illustrates this.

⁴ To ensure that the maximum actually exists we assume that each tax function t_i is continuous and each set of types θ_i is a compact subset of some \mathbb{R}^k .

⁵ Here and below $\sum_{j \neq i}$ is a shorthand for the summation over all $j \in \{1, \dots, n\}$, $j \neq i$.

A special Groves mechanism—the **VCG** or **Clarke** mechanism—is obtained using⁶

$$h_i(\theta_{-i}) := - \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j).$$

In this case,

$$t_i(\theta) := \sum_{j \neq i} v_j(f(\theta), \theta_j) - \max_{d \in D} \sum_{j \neq i} v_j(d, \theta_j),$$

which shows that the VCG mechanism is pay-only.

Following [2], let us now consider the mechanism that results from applying the BCGC transformation to the VCG mechanism. We refer to this as the Bailey-Cavallo mechanism or simply the BC mechanism. Let $\theta' := (\theta_1, \dots, \theta_{i-1}, \theta'_i, \theta_{i+1}, \dots, \theta_n)$, so $\theta'_j = \theta_j$ for $j \neq i$ and the i th player’s type in the type vector θ' is θ'_i . Then

$$S_i^{BCGC}(\theta_{-i}) = \max_{\theta'_i \in \Theta_i} \sum_{k=1}^n \left[\sum_{j \neq k} v_j(f(\theta'), \theta'_j) - \max_{d \in D} \sum_{j \neq k} v_j(d, \theta'_j) \right],$$

that is,

$$S_i^{BCGC}(\theta_{-i}) = \max_{\theta'_i \in \Theta_i} \left[(n - 1) \sum_{k=1}^n v_k(f(\theta'), \theta'_k) - \sum_{k=1}^n \max_{d \in D} \sum_{j \neq k} v_j(d, \theta'_j) \right]. \quad (1)$$

In many settings, we have that for all θ and for all i , $S_i^{BCGC}(\theta_{-i}) = 0$, and consequently the VCG and BC mechanisms coincide. Whenever they do not, by Note 1(ii) BC dominates VCG. This is the case for the single-item auction, as it can be seen that there $S_i^{BCGC}(\theta_{-i}) = -[\theta_{-i}]_2$, where $[\theta_{-i}]_2$ is the second-highest bid among bids other than player i ’s own bid.

3 Anonymous Groves Mechanisms

Throughout this paper, we will be interested in a special class of Groves mechanisms, namely, anonymous Groves mechanisms. We provide here some results about this class that we will utilize in later sections. We call a function $f : A^n \rightarrow B$ **permutation independent** if for all permutations π of $\{1, \dots, n\}$, $f = f \circ \pi$. Following [12] we call a Groves mechanism (determined by the vector of functions (h_1, \dots, h_n)) **anonymous** if

- all type sets Θ_i are equal,
- all functions h_i coincide and each of them is permutation independent.

Hence, an anonymous Groves mechanism is uniquely determined by a single function $h : \Theta^{n-1} \rightarrow \mathbb{R}$.

⁶ Here and below, to ensure that the considered maximum exist, we assume that f and each v_i are continuous functions and D and each θ_i are compact subsets of some \mathbb{R}^k .

In general, the VCG mechanism is not anonymous. But it is anonymous when all the type sets are equal and all the initial utility functions v_i coincide. This is the case in both of the domains that we consider in this paper.

For any $\theta \in \Theta$ and any permutation π of $\{1, \dots, n\}$ we define $\theta^\pi \in \Theta$ by letting

$$\theta_i^\pi := \theta_{\pi^{-1}(i)}.$$

Denote by $\Pi(k)$ the set of all permutations of the set $\{1, \dots, k\}$. Given a Groves mechanism $h := (h_1, \dots, h_n)$ for which the type set Θ_i is the same for every player (and equal to, say, Θ_0) we construct now a function $h' : \Theta_0^{n-1} \rightarrow \mathbb{R}$ by putting

$$h'(x) := \frac{\sum_{\pi \in \Pi(n-1)} \sum_{j=1}^n h_j(x^\pi)}{n!},$$

where x^π is defined analogously to θ^π .

Note that h' is permutation independent, so h' is an anonymous Groves mechanism.

The following lemma shows that some of the properties of h transfer to h' .

Lemma 1. *Consider a Groves mechanism h and the corresponding anonymous Groves mechanism h' . Let $G(\theta) := \sum_{j=1}^n v_j(f(\theta), \theta_j)$. Suppose that for all permutations π of $\{1, \dots, n\}$, $G(\theta) = G(\theta^\pi)$. Then:*

- (i) *If h is feasible, so is h' .*
- (ii) *If an anonymous Groves mechanism h^0 is welfare dominated by h , then it is welfare dominated by h' .*

The assumption in Lemma 1 of permutation independence of $G(\cdot)$ is satisfied in both of the domains that we consider in this paper. Basically, Lemma 1 says that if a Groves mechanism is not welfare undominated, then it must be welfare dominated by an anonymous Groves mechanism.

4 Multi-unit Auctions with Unit Demand

In this section, we consider auctions where there are multiple identical units of a single good and all players have unit demand, i.e., each player wants only one unit. (When there is only one unit, we have a standard single-item auction.) For this setting, we obtain an analytical characterization of all welfare undominated Groves mechanisms that are anonymous and have linear payment functions, by proving that the optimal-in-expectation linear redistribution mechanisms (OEL mechanisms) [6], which include the BC mechanism, are the only welfare undominated Groves mechanisms that are anonymous and linear. We also show that undominance and welfare undominance are equivalent if we restrict our consideration to Groves mechanisms that are anonymous and linear in the setting of multi-unit auctions with unit demand.

4.1 Optimal-in-Expectation Linear Redistribution Mechanisms

The optimal-in-expectation linear redistribution mechanisms are special cases of Groves mechanisms that are anonymous and linear. The OEL mechanisms are defined only for

multi-unit auctions with unit demand, in which there are m indistinguishable units for sale, and no bidder is interested in obtaining more than one unit. For player i , her type θ_i is her valuation for winning one unit. We assume all bids (announced types) are bounded below by L and above by U , i.e., $\Theta_i = [L, U]$. (L can be 0.)

The tax function t of an anonymous linear Groves mechanism is defined as $t_i(\theta) = t_i^{VCG}(\theta) + r(\theta_{-i})$ for all i and θ . Here t^{VCG} is (the tax function of) the VCG mechanism, and r is a linear function defined as $r(\theta_{-i}) = c_0 + \sum_{j=1}^{n-1} c_j [\theta_{-i}]_j$ (where $[\theta_{-i}]_j$ is the j th highest bid among θ_{-i}). For OEL, the c_j 's are chosen according to one of the following options (indexed by k , k is from 0 to n , and $k - m$ is odd):

$k = 0$:

$$c_i = (-1)^{m-i} \binom{n-i-1}{n-m-1} / \binom{m-1}{i-1} \text{ for } i = 1, \dots, m,$$

$$c_0 = Um/n - U \sum_{i=1}^m (-1)^{m-i} \binom{n-i-1}{n-m-1} / \binom{m-1}{i-1}, \text{ and } c_i = 0 \text{ for other } i.$$

$k = 1, 2, \dots, m$:

$$c_i = (-1)^{m-i} \binom{n-i-1}{n-m-1} / \binom{m-1}{i-1} \text{ for } i = k + 1, \dots, m,$$

$$c_k = m/n - \sum_{i=k+1}^m (-1)^{m-i} \binom{n-i-1}{n-m-1} / \binom{m-1}{i-1}, \text{ and } c_i = 0 \text{ for other } i.$$

$k = m + 1, m + 2, \dots, n - 1$:

$$c_i = (-1)^{m-i-1} \binom{i-1}{m-1} / \binom{n-m-1}{n-i-1} \text{ for } i = m + 1, \dots, k - 1,$$

$$c_k = m/n - \sum_{i=m+1}^{k-1} (-1)^{m-i-1} \binom{i-1}{m-1} / \binom{n-m-1}{n-i-1}, \text{ and } c_i = 0 \text{ for other } i.$$

$k = n$:

$$c_i = (-1)^{m-i-1} \binom{i-1}{m-1} / \binom{n-m-1}{n-i-1} \text{ for } i = m + 1, \dots, n - 1,$$

$$c_0 = Lm/n - L \sum_{i=m+1}^{n-1} (-1)^{m-i-1} \binom{i-1}{m-1} / \binom{n-m-1}{n-i-1}, \text{ and } c_i = 0 \text{ for other } i.$$

For example, when $k = m + 1$, we have $c_{m+1} = m/n$ and $c_i = 0$ for all other i . For this specific OEL mechanism, $t_i^{OEL}(\theta) = t_i^{VCG}(\theta) + \frac{m}{n} [\theta_{-i}]_{m+1}$. That is, besides paying the VCG payment, every player receives an amount that is equal to m/n times the $(m + 1)$ th highest bid from the other players. Actually, this is the BC mechanism for this setting.

One property of the OEL mechanisms is that the sum of the taxes $\sum_{i=1}^n t_i^{OEL}(\theta)$ is always less than or equal to 0 and it equals 0 whenever

- $[\theta]_1 = U$, if $k = 0$.
- $[\theta]_{k+1} = [\theta]_k$, if $k \in \{1, \dots, n - 1\}$.
- $[\theta]_n = L$, if $k = n$.

Using this property, we will prove that the OEL mechanisms are the only welfare undominated Groves mechanisms that are anonymous and linear.

4.2 Characterization of Welfare Undominated Groves Mechanisms That Are Anonymous and Linear

We first show that the OEL mechanisms are welfare undominated. (It has previously been shown that they are undominated [7], but as we pointed out, being welfare undominated is a stronger property.)

Theorem 1. *No feasible Groves mechanism welfare dominates an OEL mechanism.*

According to Lemma 11 we only need to prove this for the case of anonymous Groves mechanisms:

Lemma 2. *No feasible anonymous Groves mechanism welfare dominates an OEL mechanism.*

We now show that within the family of anonymous and linear Groves mechanisms, the OEL mechanisms are the only ones that are welfare undominated. Actually, they are also the only ones that are undominated, which is a stronger claim since being undominated is a weaker property.

Theorem 2. *If a feasible anonymous linear Groves mechanism is undominated, then it must be an OEL mechanism.*

Hence, we have the following complete characterization in this context:

Corollary 1. *A feasible anonymous linear Groves mechanism is (welfare) undominated if and only if it is an OEL mechanism.*

The above corollary also shows that if we consider only Groves mechanisms that are anonymous and linear in the setting of multi-unit auctions with unit demand, then undominance and welfare undominance are equivalent.⁷

5 Public Project Problem with Equal Participation Costs

We now study a well known class of decision problems, namely public project problems—see, e.g., [10,12,11].

Public project problem. Consider $(D, \Theta_1, \dots, \Theta_n, v_1, \dots, v_n)$, where

- $D = \{0, 1\}$ (reflecting whether a project is canceled or takes place),
- for all $i \in \{1, \dots, n\}$, $\Theta_i = [0, c]$, where $c > 0$,
- for all $i \in \{1, \dots, n\}$, $v_i(d, \theta_i) := d(\theta_i - \frac{c}{n})$,

In this setting a set of n agents needs to decide on financing a project of cost c . In the case that the project takes place, each agent contributes the same share, c/n , so as to cover the total cost. Hence the participation costs of all players are the same. When the players employ a tax-based mechanism to decide on the project, then in addition to c/n , each player also has to pay or receive the tax, $t_i(\theta)$, imposed by the mechanism.

By the result of Holmstrom [9], the only efficient and strategy-proof tax-based mechanisms in this domain are Groves mechanisms. To determine the efficient outcome for a given type vector θ , note that $\sum_{i=1}^n v_i(d, \theta_i) = d(\sum_{i=1}^n \theta_i - c)$. Hence efficiency here for a mechanism (f, t) means that $f(\theta) = 1$ if $\sum_{i=1}^n \theta_i \geq c$ and $f(\theta) = 0$ otherwise, i.e., the project takes place if and only if the declared total value that the agents have for the project exceeds its cost. We first observe the following result.

⁷ Thus, we have also characterized all undominated Groves mechanisms that are anonymous and linear. There is no corresponding result in [7].

Note 2. In the public project problem the BC mechanism coincides with VCG.

Proof. It suffices to check that in equation (1) it holds that $S_i^{BCGC}(\theta_{-i}) = 0$ for all i and all θ_{-i} . By the feasibility of VCG we have $S_i^{BCGC} \leq 0$, hence all we need is to show that there is a value for θ'_i that makes the expression in (1) equal to 0. Checking this is quite simple. If $\sum_{j \neq i} \theta_j < \frac{n-1}{n}c$, then we take $\theta'_i := 0$ and otherwise $\theta'_i := c$. \square

We now show that in fact VCG cannot be improved upon. Before stating our result, we would like to note that one ideally would like to have a mechanism that is budget-balanced, i.e., $\sum_i t_i(\theta) = 0$ for all θ , so that in total the agents only pay the cost of the project and no more. However this is not possible and as explained in [10, page 861-862], for the public project problem no mechanism exists that is efficient, strategy-proof and budget balanced. Our theorem below considerably strengthens this result, showing that VCG is optimal with respect to minimizing the total payment of the players.

Theorem 3. *In the public project problem there exists no feasible Groves mechanism that welfare dominates the VCG mechanism.*

As in Section 4, we first establish the desired conclusion for anonymous Groves mechanisms and then extend it to arbitrary ones by Lemma 1.

Lemma 3. *In the public project problem there exists no anonymous feasible Groves mechanism that welfare dominates the VCG mechanism.*

6 Public Project Problem: The General Case

The assumption that we have made so far in the public project problem that each player's cost share is the same may not always be realistic. Indeed, it may be argued that 'richer' players (read: larger enterprises) should contribute more. Does it matter if we modify the formulation of the problem appropriately? The answer is 'yes'. First, let us formalize this problem. We assume now that each (initial) utility function is of the form $v_i(d, \theta_i) := d(\theta_i - c_i)$, where for all $i \in \{1, \dots, n\}$, $c_i > 0$ and $\sum_{i=1}^n c_i = c$.

In this setting, c_i is the cost share of the project cost to be financed by player i . We call the resulting problem the **general public project problem**. It is taken from [11, page 518]. We first prove the following optimality result concerning the VCG mechanism.

Theorem 4. *In the general public project problem there is no pay-only Groves mechanism that dominates the VCG mechanism.*

It remains an open problem whether the above result can be extended to the welfare dominance relation. On the other hand, the above theorem cannot be extended to feasible Groves mechanisms, as the following result holds.

Theorem 5. *For any $n \geq 3$, an instance of the general public project problem with n players exists for which the BC mechanism dominates the VCG mechanism.*

By Theorem 4, the BC mechanism in the proof of the above theorem is not pay-only.

7 Summary

In this paper, we introduced and studied the following relation on feasible Groves mechanisms: a feasible Groves mechanism *welfare dominates* another feasible Groves mechanism if the total welfare (with taxes taken into account) under the former is at least as great as the total welfare under the latter, for any type vector—and the inequality is strict for at least one type vector. This dominance notion is different from the one proposed in [7]. We then studied welfare (un)dominance in two domains. The first domain we considered was that of auctions with multiple identical units and unit demand bidders. In this domain, we analytically characterized all welfare undominated Groves mechanisms that are anonymous and have linear payment functions. The second domain we considered is that of public project problems. In this domain, we showed that the VCG mechanism is welfare undominated if cost shares are equal, but also that this is not necessarily true if cost shares are not necessarily equal (though we showed that the VCG mechanism remains undominated in the weaker sense of [7] among pay-only mechanisms in this more general setting).

Acknowledgments

Guo and Conitzer thank the National Science Foundation and the Alfred P. Sloan Foundation for support (through award number IIS-0812113 and a Research Fellowship, respectively). The work of Vangelis Markakis was funded by the NWO project DIACo-DeM, No 642.066.604.

References

1. Bailey, M.J.: The demand revealing process: to distribute the surplus. *Public Choice* 91, 107–126 (1997)
2. Cavallo, R.: Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In: *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Hakodate, Japan, pp. 882–889 (2006)
3. Faltings, B.: A budget-balanced, incentive-compatible scheme for social choice. In: Faratin, P., Rodríguez-Aguilar, J.-A. (eds.) *AMEC 2004*. LNCS, vol. 3435, pp. 30–43. Springer, Heidelberg (2006)
4. Green, J., Laffont, J.-J.: Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* 45, 427–438 (1977)
5. Guo, M., Conitzer, V.: Better redistribution with inefficient allocation in multi-unit auctions with unit demand. In: *Proceedings of the ACM Conference on Electronic Commerce (EC)*, Chicago, IL, USA (2008)
6. Guo, M., Conitzer, V.: Optimal-in-expectation redistribution mechanisms. In: *AAMAS 2008: Proc. of 7th Int. Conf. on Autonomous Agents and Multi Agent Systems* (2008)
7. Guo, M., Conitzer, V.: Undominated VCG redistribution mechanisms. In: *AAMAS 2008: Proc. of 7th Int. Conf. on Autonomous Agents and Multi Agent Systems* (2008)
8. Guo, M., Conitzer, V.: Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, doi:10.1016/j.geb.2008.06.007
9. Holmstrom, B.: Groves' scheme on restricted domains. *Econometrica* 47(5), 1137–1144 (1979)

10. Mas-Collel, A., Whinston, M., Green, J.: Microeconomic Theory. Oxford University Press, Oxford (1995)
11. Moore, J.: General Equilibrium and Welfare Economics: An Introduction. Springer, Heidelberg (2006)
12. Moulin, H.: Axioms of Cooperative Decision Making. Cambridge University Press, Cambridge (1988)
13. Moulin, H.: Efficient, strategy-proof and almost budget-balanced assignment, Working Paper (March 2007)
14. Porter, R., Shoham, Y., Tennenholtz, M.: Fair imposition. Journal of Economic Theory 118, 209–228 (2004)

A Dominance Is Distinct from Welfare Dominance

In this appendix, we give two tax-based mechanisms t and t' (both feasible, anonymous Groves mechanisms) such that t' welfare dominates t , but t' does not dominate t . Consider a single-item auction with 4 players. We assume that for each player, the set of allowed bids is the same, namely, integers from 0 to 3. Let t^{VCG} be (the tax function of) the VCG mechanism. For all $\theta \in \{0, 1, 2, 3\}^4$, $\sum_{i=1}^4 t_i^{VCG}(\theta) = -[\theta]_2$. This is because for a single-item auction, the VCG mechanism is the second-price auction. We define t and t' as follows: **Function t** : For all θ , $t_i(\theta) := t_i^{VCG}(\theta) + h(\theta_{-i})$, where $h(\theta_{-i}) = r([\theta_{-i}]_1, [\theta_{-i}]_2, [\theta_{-i}]_3)$, and the function r is given in the table below. (We recall that $[\theta_{-i}]_j$ is the j th-highest bid among bids other than i 's own bid.) **Function t'** : For all θ , $t'_i(\theta) := t_i^{VCG}(\theta) + h'(\theta_{-i})$, where $h'(\theta_{-i}) = r'([\theta_{-i}]_1, [\theta_{-i}]_2, [\theta_{-i}]_3)$, and the function r' is given in the table below.

$r(0, 0, 0)$	0	$r'(0, 0, 0)$	0	$r(2, 2, 0)$	1/2	$r'(2, 2, 0)$	1/2	$r(3, 2, 1)$	1	$r'(3, 2, 1)$	19/24
$r(1, 0, 0)$	0	$r'(1, 0, 0)$	0	$r(2, 2, 1)$	0	$r'(2, 2, 1)$	1/4	$r(3, 2, 2)$	0	$r'(3, 2, 2)$	1/6
$r(1, 1, 0)$	1/4	$r'(1, 1, 0)$	1/4	$r(2, 2, 2)$	1/2	$r'(2, 2, 2)$	1/2	$r(3, 3, 0)$	2/3	$r'(3, 3, 0)$	5/6
$r(1, 1, 1)$	1/4	$r'(1, 1, 1)$	1/4	$r(3, 0, 0)$	0	$r'(3, 0, 0)$	0	$r(3, 3, 1)$	0	$r'(3, 3, 1)$	7/12
$r(2, 0, 0)$	0	$r'(2, 0, 0)$	0	$r(3, 1, 0)$	1/4	$r'(3, 1, 0)$	1/4	$r(3, 3, 2)$	1	$r'(3, 3, 2)$	5/6
$r(2, 1, 0)$	1/12	$r'(2, 1, 0)$	7/24	$r(3, 1, 1)$	0	$r'(3, 1, 1)$	1/4	$r(3, 3, 3)$	0	$r'(3, 3, 3)$	1/2
$r(2, 1, 1)$	0	$r'(2, 1, 1)$	1/6	$r(3, 2, 0)$	2/3	$r'(3, 2, 0)$	2/3				

With the above characterization, t' welfare dominates t (the total tax under t' is never lower, and in some cases it is strictly higher: for example, for the bid vector $(3, 2, 2, 2)$, the sum of the r_i is $1/2$, but the sum of the r'_i is 1). On the other hand, t' does not dominate t : for example, $r(3, 3, 2) = 1 > 5/6 = r'(3, 3, 2)$. In fact, no feasible Groves mechanism dominates t .

Redistribution of VCG Payments in Assignment of Heterogeneous Objects

Sujit Gujar and Narahari Yadati

Dept of Computer Science and Automation
Indian Institute of Science, Bangalore, 560012
{sujit,hari}@csa.iisc.ernet.in

Abstract. In this paper, we seek to design a Groves mechanism for assigning p heterogeneous objects among n competing agents ($n > p$) with unit demand, satisfying weak budget balance, individual rationality, and minimizing the budget imbalance. This calls for designing an appropriate rebate function. When the objects are identical, this problem has been solved by Moulin [1] and Guo and Conitzer [2]. However, it remains an open problem to design such a rebate function when the objects are heterogeneous. We propose a mechanism, HETERO and conjecture that HETERO is individually rational and weakly budget balanced. We provide empirical evidence for our conjecture through experimental simulations.

Keywords: Groves mechanism, Budget imbalance, Redistribution function, Moulin mechanism, Rebate function.

1 Introduction

Groves mechanisms [3,4,5] are widely used in practice, since they have attractive game theoretic properties such as dominant strategy incentive compatibility (DSIC) and allocative efficiency (AE). However, in general, a Groves mechanism need not be budget balanced. That is, the total transfer of money in the system may not be zero. So the system will be left with a surplus or deficit. Using Clarke's mechanism [4], we can ensure under fairly weak conditions that there is no deficit of money, that is the mechanism is weakly budget balanced. In such a case, the system or the auctioneer will be left with some money.

Often, the surplus money is not really needed in many social settings such as allocations by the Government among its departments, etc. Since strict budget balance cannot coexist with DSIC and AE (Green-Laffont theorem [6]), we would like to redistribute the surplus to the participants as far as possible, preserving DSIC and AE. This idea was originally proposed by Laffont [7]. The total payment made by the mechanism as a redistribution will be referred to the as the *rebate* to the agents.

In this paper, we consider the following problem. There are n agents and p distinct/ heterogeneous objects ($n \geq p > 1$). Each agent desires exactly one object out of these p objects. His valuation for any of the objects is independent of his valuations for the other objects. Valuations of the different agents are

independent. Our goal is to design a mechanism for assignment of the p objects among the n agents which is allocatively efficient, dominant strategy incentive compatible, and maximizes the rebate (which is equivalent to minimizing the budget imbalance). In addition, we would like the mechanism to satisfy feasibility and individual rationality. Thus, we seek to design a Groves mechanism for assigning p heterogeneous objects among n agents satisfying:

1. Feasibility (F) or weak budget balance. That is, the total payment to the agents should be less than or equal to the total received payment.
2. Individual Rationality (IR), which means that each agent's utility by participating in the mechanism should be non-negative.
3. Minimizes budget imbalance.

We call such a mechanism as *Groves redistribution mechanism* or simply *redistribution mechanism*.

Due to the Green-Laffont theorem [6], we cannot guarantee 100% redistribution at all type profiles. So a performance parameter to this function will be the worst case redistribution. That is, the fraction of the surplus which is guaranteed to be redistributed irrespective of the bid profiles. This fraction will be referred to as *efficiency* in the rest of the paper (*Note*: This efficiency is different from allocative efficiency). The advantage of worst case analysis is that, it does not require any distributional information on the type sets of the agents. Also, this function should be deterministic and anonymous. A redistribution function is said to be anonymous if two agents having the same bids, get the same rebate. So, the aim is to design an anonymous, deterministic Groves redistribution function which maximizes the efficiency and satisfies feasibility and individual rationality.

Our paper seeks to extend the results of Moulin [1] and Guo and Conitzer [2] who have independently designed a Groves mechanism in order to redistribute the surplus when objects are identical (homogeneous objects case). Their mechanism is deterministic, anonymous, and has maximum efficiency over all possible Groves redistribution mechanisms. We will refer to their mechanism as the *WCO* mechanism or *Moulin* mechanism.

1.1 Relevant Work

Due to the Green-Laffont impossibility theorem [6], it is impossible to achieve allocative efficiency, DSIC, and budget balance simultaneously. If we desire to have DSIC and budget balance properties, we have to compromise on allocative efficiency. If we are interested in preserving AE and DSIC, we have to settle down for non-zero a surplus or a non-zero deficit of the money in the system. So, the goal would be to minimize this budget imbalance in the system. To reduce budget imbalance, the surplus refund idea was originally proposed by Laffont [7] and developed by Bailey [8]. Cavallo [9], Moulin [1], Guo and Conitzer [2] designed a Groves redistribution mechanism for assignments of p homogeneous objects among $n > p$ agents with unit demand.

1.2 Contributions and Outline

Our objective in this paper is to design a Groves redistribution mechanism for assignment of heterogeneous objects with unit demand. To the best of our knowledge, this is the first attempt to design a redistribution mechanism for assignment of heterogeneous objects. In particular, our contributions is as follows: We propose a deterministic and anonymous rebate function, HETERO for the heterogeneous setting. The mechanism HETERO is conjectured to be individually rational, feasible, and worst case optimal. For some specific values of n and p , we show that the HETERO satisfies all these properties, when the valuations of the agents are binary.

The paper is organized as follows. In Section 2, we will review Moulin/WCO mechanism in brief. We then propose a redistribution mechanism, HETERO, in Section 3. We will describe some empirical results in Section 4. We will conclude with directions for future work in Section 5.

1.3 The Model and Notation

The notation used is summarized in Table 1. Note that, if bid profile b is implicit, we will use t, t_i, r_i, k , and v_i to indicate $t(b), t_i(b), r_i(b), k(b)$, and $v_i(k(b))$ respectively. In this paper, we assume that, the payment made by agent i is of the form $t_i(\cdot) - r_i(\cdot)$, where $t_i(\cdot)$ is agent i 's payment in the Clarke's pivotal mechanism. We refer to $\sum_i t_i$, as the total Clarke's payment or the surplus in the system.

Table 1. Notation

n	Number of agents
p	Number of objects
\mathbb{R}_+	Set of positive real numbers
b_i	Bid submitted by agent i , $= (b_{i1}, b_{i2}, \dots, b_{ip})$.
b	(b_1, b_2, \dots, b_n) , the bid vector
K	The set of all allocations of p objects to n agents, each getting at most one object
$k(b)$	An allocation, $k(\cdot) \in K$, where the bid profile is b
$v_i(k(b))$	Valuation of the allocation k to the agent i , when b is the bid profile
v	$v : K \rightarrow \mathbb{R}$, the valuation function, $v(k(b)) = \sum_{i=1}^n v_i(k(b))$
$t_i(b)$	Payment made by agent i in the Clarke's pivotal mechanism, when the bid profile is b
$t(b)$	The Clarke payment, that is, the total payment received from all the agents, $t(b) = \sum_{i=1}^n t_i$
t^{-i}	The Clarke payment received in the absence of the agent i
$r_i(b)$	Rebate to agent i when bid profile is b
e	The efficiency of the mechanism, $= \inf_{\theta: t \neq 0} \frac{\sum r_i(\theta)}{t(\theta)}$, where θ is vector of true valuations of the agents for the objects.

2 Optimal Worst Case Redistribution When Objects Are Identical

In this case, every agent i has the same value for each object, say v_i . Without loss of generality, we will assume, $v_1 \geq v_2 \geq \dots \geq v_n$. In Clarke’s pivotal mechanism, the first p agents will receive the objects and each of these p agents will pay v_{p+1} . So, the surplus in the system is pv_{p+1} . For this situation, Moulin [1] and Guo and Conitzer [2] have independently designed a redistribution mechanism.

Guo and Conitzer [2] maximize the worst case fraction of total surplus which gets redistributed. Their mechanism that is, WCO mechanism, coincides with Moulin’s feasible and individually rational mechanism. Their redistribution function is,

$$\begin{aligned} r_i^{WCO} &= c_{p+1}v_{p+2} + c_{p+2}v_{p+3} + \dots + c_{n-1}v_n & i = 1, \dots, p+1 \\ r_i^{WCO} &= c_{p+1}v_{p+1} + \dots + c_{i-1}v_{i-1} + c_i v_{i+1} + \dots + c_{n-1}v_n & i = p+2, \dots, n \end{aligned} \tag{1}$$

where,

$$c_i = \frac{(-1)^{i+p-1} (n-p) \binom{n-1}{p-1}}{i \binom{n-1}{i} \sum_{j=p}^{n-1} \binom{n-1}{j}} \left\{ \sum_{j=i}^{n-1} \binom{n-1}{j} \right\}; \quad i = p+1, \dots, n-1 \tag{2}$$

The efficiency of this mechanism is e^* , where e^* is given by,

$$e^* = 1 - \frac{\binom{n-1}{p}}{\sum_{j=p}^{n-1} \binom{n-1}{j}}$$

This is an optimal mechanism, since there is no other mechanism which can guarantee more than e^* fraction redistribution in the worst case.

The following theorem by Guo and Conitzer [2] will be used to design our mechanism.

Theorem 1. *Let, $x_1 \geq x_2 \geq \dots x_n \geq 0$. Then*

$$a_1x_1 + a_2x_2 + \dots a_nx_n \geq 0 \text{ iff } \sum_{i=1}^j a_i \geq 0 \quad \forall j = 1, 2, \dots, n$$

3 A Redistribution Mechanism for the Heterogeneous Setting

We should note that the homogeneous objects case is a special case of the heterogeneous objects case in which each bidder submits the same bid for all objects. Thus, we cannot expect any redistribution mechanism to perform better than

the homogeneous object case. For $n \leq p + 1$, the worst case redistribution is zero for the homogeneous case and so will be for the heterogeneous case. So, we assume $n > p + 1$. We propose a redistribution mechanism. We will be referring to it as *HETERO*.

When the objects are identical, the WCO mechanism is given by equation (II). We give a novel interpretation to it. Consider the scenario in which one agent is absent from the scene. Then the Clarke’s payment received is either pv_{p+1} or pv_{p+2} depending upon which agent is absent. If we remove two agents, the surplus is pv_{p+1} or pv_{p+2} or pv_{p+3} , depending upon which two agents are removed. Till $(n - p - 1)$ agents are removed, we get non-zero surplus. If we remove $(n - p)$ or more agents from the system, there is no need for any mechanism for assignment of the objects. So, we will consider the cases when we remove k agents, where, $1 \leq k < n - p$.

Now let $t^{-i,k}$ be the average payment received when agent i is removed along with k other agents that is, a total of $(k + 1)$ agents removed in which i is also removed. The average is taken over all possible k selections from the remaining $(n - 1)$ agents. We can rewrite the WCO mechanism in terms of $t^{-i}, t^{-i,k}$. Observe that, $t^{-i}, t^{-i,k}$ can be defined in heterogeneous settings as well. We propose to use a rebate function defined as,

$$r_i^H = \alpha_1 t^{-i} + \sum_{k=2}^{k=n-p-1} \alpha_k t^{-i,k-1} \tag{3}$$

where α_k are the weights given to the surplus generated when a total of k agents are removed from the system.

The Equivalence of HETERO and WCO when Objects Are Identical

It is desirable that HETERO should match with the WCO mechanism when the objects are homogeneous. So we need that at all type profiles, r_i^H in equation (3) is equal to r_i^{WCO} in equation (II). Since the rebate is a function of the remaining $(n - 1)$ bids, we can write it as, $r_i = f(x_1, x_2, \dots, x_{n-1})$ where x_1, x_2, \dots, x_{n-1} are bids without the agent i , in decreasing order. Note, in this case, that each bidder will be submitting a bid $b_i \in \mathbb{R}_+$.

Now, we can write, $t^{-i,k}, r_i^H$, and r_i in terms of x_1, x_2, \dots, x_{n-1} , as,

$$t^{-i,k-1} = \sum_{l=0}^{k-1} \frac{\binom{p+l}{p} \binom{n-p-2-l}{k-1-l}}{\binom{n-1}{k-1}} x_{p+1+l}$$

$$r_i^H = \sum_{k=1}^{k=n-p-1} \alpha_k t^{-i,k-1} \tag{4}$$

$$r_i^{WCO} = \sum_{l=0}^{n-p-1} c_{p+1+l} x_{p+1+l} \tag{5}$$

where, $c_i, i = p + 1, p + 2, \dots, n - 1$ are given by equation (2).

Consider the type profile $(x_1 = 1, x_2 = 1, \dots, x_{p+1} = 1, x_{p+2} = 0, \dots, x_{n-1} = 0)$. For HETERO to agree with WCO, the coefficients of x_{p+1} in equation (4) and equation (5) should be the same. Now consider the type profile $(x_1 = 1, x_2 = 1, \dots, x_{p+2} = 1, x_{p+3} = 0, \dots, x_{n-1} = 0)$. As the coefficients of x_{p+1} in equation (4) and equation (5) are the same, the coefficients of x_{p+2} should also be equal in equation (4) and equation (5).

Thus, the coefficients of $x_{p+1}, x_{p+2}, \dots, x_{n-1}$ in equation (4) and equation (5) should agree.

Let $L = n - p - 1$. Thus, for $i = p + 1, \dots, n - 1$,

$$c_i = \sum_{k=0}^{n-i-1} \alpha_{L-k} \frac{\binom{i-1}{p} \binom{n-i-1}{k}}{\binom{n-1}{p+1+k}} \tag{6}$$

The above system of equations yields, for $i = 1, 2, \dots, L$,

$$\alpha_i = \frac{(-1)^{(i+1)}(L-i)!p!}{(n-i)!} \chi \sum_{j=0}^{L-i} \left\{ \binom{i+j-1}{j} \sum_{l=p+i+j}^{n-1} \binom{n-1}{l} \right\}; i = 1, 2, \dots, L \tag{7}$$

where χ is given by,
$$\chi = \frac{\binom{n-p}{p-1}}{\sum_{j=p}^{n-1} \binom{n-1}{j}}$$

As the HETERO mechanism matches with the WCO when objects are identical, the HETERO mechanism satisfies individual rationality and feasibility in the homogeneous case. These two properties, however, remain to be shown in the heterogeneous case.

3.1 Individual Rationality of HETERO

Conjecture 1. *The HETERO mechanism is individually rational.*

Intuition Behind Individual Rationality of HETERO

We have to show that for each agent i , $r_i^H \geq 0$ at all type profiles. For convenience, we will assume i implicitly. So, say, $r_i^H = r$ and $\Gamma_1 = t^{-i}$, $\Gamma_j = t^{-i, j-1}$, $j = 2, \dots, L$. Now, the rebate is given by the equation, $r = \sum_j \alpha_j \Gamma_j$. We have to show that $r \geq 0$. Note that, $\Gamma_1 \geq \Gamma_2 \geq \dots \geq \Gamma_L \geq 0$. So, if $\sum_{i=1}^j \alpha_i \geq 0 \forall j = 1 \rightarrow L$, individual rationality would follow from Theorem 1. We observe that, in general, this is not true. The important observation is, though Γ_i 's are decreasing positive real numbers, they are related. For example, we can show that if $\Gamma_1 > 0$, then $\Gamma_2 > 0$. Thus, though α 's are alternately positive and negative, the relation among Γ 's would not make r to go negative and it will be within limits in such a fashion that total rebate to the agents will be less than or equal to total Clarke's payment. It remains to show individual

rationality analytically in the general case. However, we are only able to show in the following cases.

1. Consider the case when $p = 2$. (i.) If $n = 4$, $\alpha_1 = \frac{1}{4}$. (ii.) If $n = 5$, $\alpha_1 = 0.27273$, $\alpha_2 = -0.18182$. (iii.) If $n = 6$, $\alpha_1 = 0.29487$, $\alpha_2 = -0.25641$, $\alpha_3 = 0.12821$.
2. Consider the case when $p = 3$. (i.) If $n = 5$, $\alpha_1 = \frac{1}{5}$. (ii.) If $n = 6$, $\alpha_1 = 0.21875$, $\alpha_2 = -0.15625$. (iii.) If $n = 7$, $\alpha_1 = 0.23810$, $\alpha_2 = -0.21429$, $\alpha_3 = 0.11905$.

By Theorem [11](#), it follows that for the above cases, the proposed mechanism satisfies the individual rationality. We provide some empirical evidence for the conjecture in Section [4](#).

4 Experimental Analysis and Empirical Evidence

Solving equations [\(7\)](#) is a challenging task. Though the new mechanism is the extension of the Moulin or the WCO mechanism, yet, we are not able to prove individual rationality and feasibility of HETERO analytically. We therefore seek empirical evidence.

4.1 Simulation 1

We consider various combinations of n and p . For each agent, and for each object, the valuation is generated as a uniform random variable in $[0, 100]$. We run our simulation for the following combinations of n and p .

For $p = 2$, $n = 5, 6, \dots, 14$, for $p = 3$, $n = 7, 8, \dots, 14$ and for $p = 4$, $n = 9, 10, \dots, 14$. For each combination of n and $p = 2$, we generated randomly 100,000 bid profiles (for $p = 3, 4$, this number was 20,000) and evaluated our mechanism. We also kept track of the worst case performance of our mechanism over these 100,000 bid profiles. Our mechanism was feasible and individually rational in these 100,000 bid profiles. The efficiency of our mechanism is upper bounded by that of the WCO mechanism. We observed that the worst case performance over these 100,000 random bid profiles was the same as that of WCO. This is a strong indication that our mechanism will perform well in general.

4.2 Simulation 2: Bidders with Binary Valuation

Suppose each bidder has valuation for each object, either 0 or 1. Then there are 2^{np} possible bid profiles. We ran an experiment to evaluate our mechanism with all possible bid profiles of agents with binary valuations. We considered $p = 2$ and $n = 5, 6, \dots, 12$. We found that the mechanism is feasible, individually rational, and the worst case performance is the same as that of the WCO mechanism. Note, as indicated earlier, no mechanism can perform better than the WCO mechanism in the worst case. And our mechanism performs as well as the WCO. Thus, though there is no analytical proof with us, for binary valuation settings, for $p = 2$ and $n = 5, 6, \dots, 12$, our mechanism is worst case optimal.

5 Conclusion

We proposed a mechanism, namely HETERO, for general settings when the objects are heterogeneous and private values of an agent for these objects are independent of each other. The mechanism is deterministic, anonymous, and DSIC. The HETERO mechanism extends the Moulin/WCO mechanism. Though we have not analytically proved feasibility and individual rationality, we have sufficient empirical evidence to conjecture that our mechanism is feasible and individually rational.

We are currently working on the proving individual rationality and feasibility for the proposed HETERO mechanism. We strongly believe that the new mechanism is a worst case optimal. Our immediate goal is to prove this fact or design a mechanism which is worst case optimal.

References

1. Moulin, H.: Efficient, strategy-proof and almost budget-balanced assignment. Technical report, Northwestern University, Center for Mathematical Studies in Economics and Management Science (2007)
2. Guo, M., Conitzer, V.: Worst-case optimal redistribution of VCG payments. In: EC 2007: Proceedings of the 8th ACM conference on Electronic Commerce, pp. 30–39. ACM, New York (2007)
3. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16(1), 8–37 (1961)
4. Clarke, E.: Multi-part pricing of public goods. *Public Choice* 11, 17–23 (1971)
5. Groves, T.: Incentives in teams. *Econometrica* 41, 617–631 (1973)
6. Green, J.R., Laffont, J.J.: Incentives in Public Decision Making. North-Holland Publishing Company, Amsterdam (1979)
7. Laffont, J., Maskin, E.: A differential approach to expected utility maximizing mechanisms. In: Laffont, J.J. (ed.) *Aggregation and Revelation of Preferences* (1979)
8. Bailey, M.J.: The demand revealing process: To distribute the surplus. *Public Choice* 91(2), 107–126 (1997)
9. Cavallo, R.: Optimal decision-making with minimal waste: strategyproof redistribution of VCG payments. In: AAMAS 2006: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, pp. 882–889. ACM, New York (2006)

Bin Packing of Selfish Items

Guosong Yu^{1,2} and Guochuan Zhang^{1,3,*}

- ¹ Department of Mathematics, Zhejiang University
Hangzhou 310027, China
- ² Department of Mathematics, Nanchang University
Nanchang 330031, China
- ³ College of Computer Science, Zhejiang University
zgc@zju.edu.cn

Abstract. We study a bin packing game in which any item to be packed is handled by a selfish agent. Each agent aims at minimizing his sharing cost with the other items staying in the same bin, where the social cost is the number of bins used. We first show that computing a pure Nash equilibrium can be done in polynomial time. We then prove that the price of anarchy for the game is in between 1.6416 and 1.6575, improving the previous bounds.

Keywords: Bin packing, Nash equilibrium, price of anarchy.

1 Introduction

In the one-dimensional bin packing problem, given a list $L = (a_1, a_2, \dots, a_n)$ of items, each with a size $s(a_i) \in (0, 1]$, we are asked to pack them into a minimum number of unit-capacity bins such that the sum of items in each bin is at most one. This bin packing system has only one decision maker who assigns items into bins without considering the own interest of each item. However, in many systems users are likely to behave selfishly, namely each user aims to optimize his own performance without coordination with the other users. Basically, each user would like to either maximize the resources allocated to him or, alternatively, minimize his cost [7]. In this paper we consider such a bin packing system where every item is handled by a selfish player [1]. Equivalently we can regard each item as a user. Assume that all the bins have the same cost of one. The goal of each item is to minimize its own cost sharing with the other items in the same bin. Let $c(B)$ be the content of a bin B , i.e., the sum of the items packed in bin B . If an item a is in B , its cost is defined as $\frac{s(a)}{c(B)}$. Since item a wants to minimize its cost, it will detect another bin B' such that $s(a) + c(B') \leq 1$ and $\frac{s(a)}{s(a)+c(B')} < \frac{s(a)}{c(B)}$ (i.e. $c(B) < s(a)+c(B')$). The former equation ensures that the resulting packing is feasible after a migrates to the bin B' , while the latter shows that the cost of item a will decrease after the migration. Thus item a will change

* Research supported in part by NSFC (60573020) and NSF of Zhejiang Province (Y605353).

its bin. A stable state is a feasible packing in which no item is willing and is able to move alone. However, the social cost that we need to minimize is the number of bins used, which is corresponding to the standard bin packing problem.

A framework to analyze the class of problems with selfish players is that of non-cooperative games. In game theory, the Nash equilibrium [15] is a solution concept of a game involving two or more players, in which each player is assumed to know the equilibrium strategies of the other players, and no player has anything to gain by changing only his own strategy (i.e., by changing unilaterally). If each player has chosen a strategy and no player can benefit by changing his or her strategy while the other players keep theirs unchanged, then the current set of strategy choices and the corresponding payoffs constitute a Nash equilibrium. Nash equilibria for pure strategies, where each player chooses to play an action in a deterministic non-aleatory manner, are briefly referred to as pure Nash equilibria. Determining whether Nash equilibria exist, effectively computing them, and measuring the price of anarchy [13] (i.e., the ratio between the worst Nash equilibrium and the social optimum, that is the optimal solution that could be achieved if all players cooperated.) have attracted much research in computer science.

Previous results. Bin packing is one of the well-studied problems in combinatorial optimization. As for the best algorithmic results of theoretical interests de La Vega and Lueker [5] presented an asymptotic polynomial time approximation scheme, and Karmarkar and Karp [11] improved to an asymptotic fully polynomial time approximation scheme. A well-known efficient approximation algorithm is *FFD* (First-Fit Decreasing), which works as follows: Sort all items in non-increasing order of sizes; pick the items one by one and pack them into the first bin (the leftmost bin) into which they fit (an empty bin is needed in case that such a bin does not exist). Johnson [10] proved that the asymptotic performance ratio of *FFD* is $11/9$. The tight bound was recently achieved by Dósa [6]. However, none of the above algorithms always produces a packing which is a pure Nash equilibrium.

The bin packing system with selfish users was first studied by Bilò [1]. He proved that this bin packing game always converges to a pure Nash equilibrium starting from any feasible packing. More precisely if a packing does not achieve a Nash equilibrium, after a finite number of migrations a stable state could be reached. He also analyzed the effectiveness of the system by deriving an upper bound of $5/3$ and a lower bound of $8/5$ for the price of anarchy.

A relevant problem is the selfish routing game originally introduced in [13]. One model of this game evaluates the link congestion in a network consisting of two nodes connected by m parallel links. It assumed a collection of n users, each employing a strategy over m parallel links to control the shipping of its own assigned traffic. The price of anarchy for this model was analyzed in [14, 12, 4]. In [9], Fotakis et al. studied pure Nash equilibria, where each user chooses exactly one link, and mixed Nash equilibria, where the choices of each user are modeled by a probability distribution over links. They proved the existence of the pure

Nash equilibria and presented a polynomial-time algorithm determining a pure equilibrium. The algorithm has been improved in [8].

Our contribution. In this paper, we give a polynomial time algorithm to determine a pure Nash equilibrium, showing that computing a pure Nash equilibrium is an easy problem for the bin packing game. Moreover, we deal with the price of anarchy of the game. A new lower bound of $\sum_{l=1}^{\infty} \frac{1}{2^{l(l-1)/2}} \approx 1.6416$ and an upper bound of $(41 + \sqrt{145})/32 \approx 1.6575$ are provided, significantly improving both bounds in [1].

2 Computing a Nash Equilibrium

Bilò [1] showed that a pure Nash equilibrium can always be obtained after a finite number of migrations (bounded by an exponential number) starting from any feasible packing. He further noted that there is always an optimal packing which is a pure Nash equilibrium, and thus computing a best pure Nash equilibrium is NP-hard. However it was open if computing a pure Nash equilibrium is hard or not. In this section we will derive a polynomial time bin packing algorithm finding a Nash equilibrium, thus solving this problem.

As mentioned before bin packing has been extensively studied and there exist quite a lot of bin packing algorithms in the literature [3]. Since every bin packing algorithm aims to find a good global solution, almost none of them always produces a Nash equilibrium. However there is indeed an exception. It is the subset sum algorithm which works as follows: for a given list L of items, let $L_1 = L$ and $i = 1$. If L_i is empty, output $m = i - 1$ and stop; otherwise pack a bin as full as possible with the items in L_i , and let L_{i+1} be the unpacked items in L_i . Set $i = i + 1$ and repeat the above packing process. It is easy to verify that the packing is a pure Nash equilibrium. Unfortunately this approach employs an optimal algorithm for the subset sum problem. It is not polynomial unless $P = NP$. Caprara and Pferschy [2] showed that the worst-case ratio of the algorithm is in between $\sum_{k=1}^{\infty} \frac{1}{2^{k-1}} \approx 1.6067$ and $\frac{4}{3} + \ln \frac{4}{3} \approx 1.6210$.

In the following we will introduce a new approach to find a pure Nash equilibrium in polynomial time. Our main idea is to apply algorithm *FFD* (First-Fit-Decreasing) recursively. Basically the approach consists of three steps in each round. The first step applies the *FFD* algorithm to the unpacked items. The second step picks up a subset of smaller items from the bins and put them into a list for the next round packing. The third step migrates the packed items individually to decrease their own cost, until a pure Nash equilibrium is reached for the packed items. The key issue is to bound both the number of migrations and the number of rounds polynomially. At any time if an item a_i can detect a bin which can accommodate a_i and the cost of a_i will decrease after moving to this bin (while the other items do not change their bins), we say that a_i is *active* and the bin is *cheap*. A *migration* is a single movement of an item from its current

bin to another bin it fits so that its own cost is reduced, while the other items stay without any action.

Before we give the full picture of the proposed approach we want to deal with a crucial case by applying the *FFD* algorithm to a list of items $\{a_1, a_2, \dots, a_n\}$, where $s(a_1) \geq s(a_2) \geq \dots \geq s(a_n)$. Suppose that *FFD* takes m bins for packing the items. If $m = 1$ it is trivially a best Nash equilibrium. Now assume $m \geq 2$. Sort the bins in non-increasing order of their contents, i.e., $c(B_1) \geq c(B_2) \geq \dots \geq c(B_m)$. We first deal with a packing in which the following equation holds.

$$s(a_n) > 1 - c(B_{m-1}) \tag{1}$$

Recall that a Nash equilibrium is converged from the *FFD* packing after a finite number of migrations [1]. Let k be this number. We claim that the number k can be polynomially bounded. Let B_j^i be the j -th fullest bin after the i -th migration, for $i = 1, \dots, k$ and $j = 1, \dots, m$.

Lemma 1. *If Equation (1) holds, then we have*

$$c(B_j) \leq c(B_j^1) \leq \dots \leq c(B_j^k), \quad j = 1, \dots, m - 1, \tag{2}$$

and

$$c(B_m) > c(B_m^1) > \dots > c(B_m^k). \tag{3}$$

Proof. Let us consider the first migration. Equation (1) shows that there is no space in any of the first $m - 1$ bins for the smallest item a_n and thus neither for any items. We can claim that no item in the bin B_m is active and B_m is the only possible cheap bin. Suppose some item a_i migrates to B_m from his bin B_l . In this case $c(B_l) < s(a_i) + c(B_m) \leq 1$. Then

$$s(a_n) \leq s(a_i) \leq 1 - c(B_m) < 1 - (c(B_l) - s(a_i))$$

and

$$1 - (s(a_i) + c(B_m)) < 1 - c(B_l) < s(a_n).$$

Note that no bins other than B_l and B_m change. Therefore B_l becomes the smallest bin B_m^1 and B_m becomes larger and falls in the set of $\{B_1^1, \dots, B_{m-1}^1\}$ after the first migration (i.e., after a_i has migrated). Clearly

$$c(B_m) > c(B_m^1) = c(B_l) - s(a_i).$$

Now we compare the two number sets $\{c(B_1^1), \dots, c(B_{m-1}^1)\}$ and $\{c(B_1), \dots, c(B_{m-1})\}$. They have exactly $m - 2$ common numbers. The different numbers are $c(B_m) + s(a_i)$ (in the first set) and $c(B_l)$ (in the second set), where the former is larger. Thus we conclude that $c(B_j) \leq c(B_j^1)$, $j = 1, \dots, m - 1$.

We further note that $s(a_n) > 1 - c(B_j^1)$ for $j = 1, 2, \dots, m - 1$. Hence the same procedure applies for the next migrations that meet the lemma.

Lemma 2. *If Equation (1) holds, then the bin packing game will converge to a pure Nash equilibrium after at most $(m - 1)n$ migrations.*

Proof. Let t_i be the number of migrations for an item a_i when the bin packing game converges to a pure Nash equilibrium. To prove this lemma we only need to show that the number of migrations for each item is bounded by $m - 1$, i.e., $t_i \leq m - 1$ for $i = 1, 2, \dots, n$.

If $t_i = 1$ for all $i = 1, 2, \dots, n$, we are done since $m \geq 2$. Consider any item a_i which moves at least twice. By Lemma 1 when a_i migrates, the cheap bin is the least full bin, i.e., the m -th fullest bin. Assume that the cheap bin becomes the r_j -th fullest bin immediately after the j -th migration of item a_i . Clearly $r_j \leq m - 1$. In the following we want to prove that $r_{j+1} > r_j$ for $j = 1, 2, \dots, t_i - 1$.

For simplicity we denote B_c to be the cheap bin (the m -th bin) immediately before the j -th migration of item a_i , and B to be the bin after the migration, i.e., the r_j -th fullest bin. Similarly we let \hat{B}_c and \hat{B} be the cheap bin immediately before the $(j + 1)$ -st migration of item a_i and the bin afterward, respectively. We have

$$c(\hat{B}) = s(a_i) + c(\hat{B}_c) < s(a_i) + c(B_c) = c(B),$$

since $c(\hat{B}_c) < c(B_c)$ by Equation (3) in Lemma 1. Moreover by Equation (2) the content of the r_j -th fullest bin after the $(j + 1)$ -st migration of item a_i is not smaller than $c(\hat{B})$, that implies $r_{j+1} > r_j$.

Note that r_j can never be m . It shows that $t_i \leq m - 1$ and the lemma follows.

This lemma tells us that if Equation (1) holds one can find a pure Nash equilibrium in polynomial time. It is easy to verify that the following condition implies Equation (1): In the *FFD* packing the smallest item a_n is packed into the last bin. It implies that a_n does not fit into the first $m - 1$ bins. It further shows that a_n cannot be packed into the first $m - 1$ fullest bins. Thus Equation (1) holds.

However if a_n is not packed into the last bin, Lemma 2 does not work. Assume that a_j ($j < n$) is the smallest item packed into the last bin. Remove all items a_{j+1}, \dots, a_n , and put them back into a list L_1 . One can see that the list $L - L_1$ of items satisfies Equation (1) by regarding a_j as a_n . This observation motivates us to apply the algorithm *FFD* recursively. It gives the following algorithm.

Algorithm *RFFD* (Recursive First Fit Decreasing)

Step 0. Let $k_0 = 0$ and $i = 1$; Sort a_i ($i = 1, \dots, n$) in decreasing order of their sizes; Open n vacant bins.

Step 1. Re-index the open bins in decreasing order of their contents, i.e., $c(B_1) \geq c(B_2) \geq \dots \geq c(B_n)$; Pack $a_{k_{i-1}+1}, \dots, a_n$ with the First Fit Algorithm; Suppose B_{p_i} is the bin with largest index which accommodates some of the items in $\{a_{k_{i-1}+1}, \dots, a_n\}$. Let a_{k_i} be the smallest item packed in B_{p_i} . If $k_i < n$, then take a_{k_i+1}, \dots, a_n out of B_1, \dots, B_{p_i} .

Step 2. Determine and make all migrations until a pure Nash equilibrium is obtained.

Step 3. If $k_i = n$, then it is a pure Nash equilibrium for all items; Output the non-empty bins. Else set $i = i + 1$, go to Step 1.

Algorithm *RFFD* recursively uses the bin packing algorithm *FFD*. In each round at least one item is packed and thus the number of rounds is at most n .

Theorem 1. *Algorithm RFFD performs $O(n^3)$ migrations.*

Proof. To prove this theorem we only need to show that at each round i there are at most $(p_i - 1)k_i \leq n^2$ migrations. Let $q_i = \max\{p_j, j = 1, \dots, i\}$. Then q_i is the number of non-empty bins immediately after the i -th round. If $p_i \geq q_{i-1}$, then at most $(p_i - 1)k_i$ migrations is made at Step 2 according to Lemma 2.

Now we turn to the much more difficult case that $p_i < q_{i-1}$. In this case the item a_{k_i} is not packed into the last nonempty bin at Step 1. However we can show that the bins $B_{p_i+1}, \dots, B_{q_{i-1}}$ remain unchanged (no items in and no items out) at Step 2 in this round. Hence the migrations happen among the first p_i bins and thus Lemma 2 still applies.

We only deal with the case that $p_i < q_{i-1}$ with the smallest index i , i.e., $q_{i-1} = p_{i-1}$. The other cases can follow analogously.

Note that $a_{k_{i-1}}$ is the smallest item in $B_1, \dots, B_{p_{i-1}}$ before the i -th round. From Lemma 1 we get $1 - c(B_l) < s(a_{k_{i-1}})$, for $l = 1, \dots, p_{i-1} - 1$. Now consider the first migration in the i -th round (Step 2). We will show that this migration cannot be made by any item $a \in \{a_1, \dots, a_{k_{i-1}}\}$. Denote these items as *lazy* items.

Before Step 1 of Round i there is a pure Nash equilibrium that satisfies the following two properties: (1) There is no enough space at any bin $B_j \in \{B_1, \dots, B_{p_{i-1}-1}\}$ for any lazy item. (2) None of the lazy items packed in any bin $B_j \in \{B_1, \dots, B_{p_{i-1}-1}\}$ wants to migrate to the bin $B_{p_{i-1}}$. After Step 1, the content of $B_{p_{i-1}}$ remains unchanged while the content of any other bin does not decrease. It implies that the above two properties still hold for all lazy items. And thus none of them makes the first migration. An item $a \in \{a_{k_{i-1}}, \dots, a_{k_i}\}$, called a *new* item, migrates, but not to the bin $B_{p_{i-1}}$. It is because any new item is not larger than any lazy item. Even a lazy item has no interest to move to the bin $B_{p_{i-1}}$, neither does a new item. So the bin $B_{p_{i-1}}$ has not any change after the first migration. This good point ensures that no lazy items moves in the next migrations and again $B_{p_{i-1}}$ remains unchanged during the Step 2 of Round i .

We see that only new items may migrate. Then the number of migrations in Step 2 of the i -th round is at most $q_{i-1}(k_i - k_{i-1}) \leq n^2$ by Lemma 2. Recall that the number of the rounds is no more than n . There are $O(n^3)$ migrations.

Finally we estimate the time to determine a migration. In fact we can always maintain a cheap bin so that it takes $O(n)$ time to check which item is willing to move into this bin. Therefore the total running time of algorithm RFFD is $O(n^4)$.

Theorem 2. *RFFD is a polynomial time algorithm for computing a pure Nash equilibrium.*

3 Bounding the Price of Anarchy

In this section, we prove that the price of anarchy of this game is in between $\sum_{l=1}^{\infty} \frac{1}{2^{\frac{l(l-1)}{2}}} \approx 1.6416$ and $(41 + \sqrt{145})/32 \approx 1.6575$. We first deal with the lower

bound. Let ε and γ be sufficiently small positive numbers. We classify items into N types where N is sufficiently large. For $k = 1, 2, \dots, N$, an item a is of type- k if $\frac{1}{2^k} - \gamma < s(a) < \frac{1}{2^k} + \gamma$. In addition let $m_N > 0$ be a given integer, and for $k = 1, 2, \dots, N - 1$,

$$m_k = \begin{cases} 2m_2 + 1, & k = 1 \\ 2^k m_{k+1} - 1 & k \geq 2 \end{cases}$$

We are ready to present the list L of items. Let $L = L_1 \cup L_2 \cup \dots \cup L_N$, where L_1 consists of m_1 items of $1/2 + d_1\varepsilon$ and for $k = 2, \dots, N$, L_k is defined as

$$\underbrace{\left\{ \frac{1}{2^k} - e_k\varepsilon, \frac{1}{2^k} - b_{k,1}\varepsilon, \dots, \frac{1}{2^k} - b_{k,m_k}\varepsilon, \frac{1}{2^k} + c_{k,1}\varepsilon, \dots, \frac{1}{2^k} + c_{k,m_k}\varepsilon, \frac{1}{2^k} + d_k\varepsilon \right\}}_{m_k(2^k - 4)} \cup \underbrace{\left\{ \frac{1}{2^k} + d_k\varepsilon \right\}}_{m_k}$$

The parameters $e_k, b_{k,i}, c_{k,i}$ and d_k are iteratively specified below. Let $d_1 = 1, e_1 = b_{1,1} = \dots = b_{k,1} = c_{1,m_1} = b_{1,m_1} = 0$. For $k \geq 2$,

- $\delta_{k-1} = d_{k-1} + c_{k-1,m_{k-1}} - b_{k-1,m_{k-1}} - (2^{k-1} - 4)e_{k-1}$,
- if $k = 2$ then $e_k = 0$ else $e_k = (1/2) \sum_{l=1}^{k-1} d_l$,
- $c_{k,i} = b_{k,i} + (2^k - 4)e_k + \delta_{k-1} + 1, i = 1, \dots, m_k$
- $b_{k,i+1} = c_{k,i} + \sum_{l=1}^{k-1} d_l, i = 1, \dots, m_k - 1$,
- $d_k = c_{k,m_k}$.

One can have a packing using bins as many as the number of the items of $1/2 + d_1\varepsilon$, which is m_1 . On the other hand we can construct a packing of Nash equilibrium using $\sum_{k=1}^N m_k$ bins. As N goes to infinity, we conclude

Theorem 3. *The price of anarchy for the bin packing game is at least*

$$\sum_{l=1}^{\infty} \frac{1}{2^{\frac{l(l-1)}{2}}} \approx 1.6416.$$

Now we turn to the upper bound. For a list L of items, consider any pure Nash equilibrium for the bin packing game. Let the bins used be B_1, \dots, B_m , where $m \geq 4$ and $c(B_1) \geq \dots \geq c(B_{m-1}) \geq c(B_m)$. Let OPT be the number of bins used by an optimal packing. It is obvious that $c(B_{m-1}) > 1/2$. Let $x = (41 - \sqrt{145})/48 \approx 0.6033$ and $z = 0.6$. We aim at an (asymptotic) upper bound of $1/x = (41 + \sqrt{145})/32 \approx 1.6575$. By carefully investigating the packing structure of the Nash equilibrium we can prove the following two lemmas (due to the page limit we omit the proofs).

Lemma 3. *If $c(B_{m-1}) \geq z$, then $m < OPT/x + 3$.*

Lemma 4. *If $s(B_{m-1}) < z$, then $m < OPT/x + 3$.*

Theorem 4. *Any pure Nash equilibrium for the bin packing game uses at most $\lfloor OPT/x + 3 \rfloor$ bins.*

Proof. It follows from Lemmas [3](#) and [4](#).

References

1. Bilò, V.: On the packing of Selfish Items. In: Proceedings of the 20th International Parallel and Distributed Processing Symposium (IPDPS), pp. 25–29 (2006)
2. Caprara, A., Pferschy, U.: Worst-case analysis of the subset sum algorithm for bin packing. *Operations Research letters* 32, 159–166 (2004)
3. Coffman, E.G., Garey, M.R., Johnson, D.S.: Approximation algorithms for bin packing: A survey. In: Hochbaum, D.S. (ed.) *Approximation algorithms for NP-hard problems*, PWS (1996)
4. Czumaj, A., Vocking, B.: Tight bound for the worst case equilibria. In: Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Mathematics (SODA), pp. 413–420 (2002)
5. de la Vega, W.F., Lueker, G.S.: Bin packing can be solved within $1 + \varepsilon$ in linear time. *Combinatorica* 1, 349–355 (1981)
6. Dósa, G.: The tight bound of first fit decreasing bin-packing algorithm is $FFD(I) \leq (11/9)OPT(I) + 6/9$. In: Chen, B., Paterson, M., Zhang, G. (eds.) *ESCAPE 2007*. LNCS, vol. 4614, pp. 1–11. Springer, Heidelberg (2007)
7. Even-Dar, E., Kesselman, A., Mansour, Y.: Convergence time to Nash equilibria. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) *ICALP 2003*. LNCS, vol. 2719, pp. 502–513. Springer, Heidelberg (2003)
8. Feldmann, R., Gairing, M., Lucking, T., Monien, B., Rode, M.: Nashification and the coordination ratio for a selfish routing game. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) *ICALP 2003*. LNCS, vol. 2719, pp. 514–526. Springer, Heidelberg (2003)
9. Fotakis, D.A., Kontogiannis, S.C., Koutsoupias, E., Mavronicolas, M., Spirakis, P.G.: The structure and complexity of nash equilibria for a selfish routing game. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) *ICALP 2002*. LNCS, vol. 2380, pp. 123–134. Springer, Heidelberg (2002)
10. Johnson, D.S.: Near-optimal bin packing algorithms. PhD thesis. MIT, Cambridge, MA (1973)
11. Karmarkar, N., Karp, R.M.: An efficient approximation scheme for the one-dimensional bin-packing problem. In: Proceedings of the 23rd Annual IEEE Symposium on Foundation of Computer Science (FOCS), pp. 312–320 (1982)
12. Koutsoupias, E., Mavronicolas, M., Spirakis, P.: Approximate equilibria and ball fusion. In: Proceedings of the 9th International Colloquium on Structural Information and Communication Complexity (SIROCCO), pp. 223–235 (2002)
13. Koutsoupias, E., Papadimitriou, C.: Worst-case Equilibria. In: Meinel, C., Tison, S. (eds.) *STACS 1999*. LNCS, vol. 1563, pp. 387–396. Springer, Heidelberg (1999)
14. Mavronicolas, M., Sirakis, P.: The price of selfish routing. In: Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing (STOC), pp. 510–519 (2001)
15. Nash, J.F.: Non-cooperative games. *Annals of Mathematics* 54, 28–295 (1951)

Restricted Core Stability of Flow Games^{*}

Mao-cheng Cai¹ and Qizhi Fang²

¹ Academy of Mathematics and Systems Science, Chinese Academy of Science
Beijing 100080, China

² Department of Mathematics, Ocean University of China, Qingdao 266071, China
qfang@ouc.edu.cn

Abstract. In this paper, we introduce a kind of restricted core stability for flow games, which is a generalization of the core stability of simple flow games. We first give a characterization on the restricted core, and then propose a sufficient and necessary condition on the restricted core stability for flow games associated with general networks. This condition yields that testing the restricted core stability can be done in polynomial time.

Keywords: Flow game, core, stability, maximum flow, minimum cut.

1 Introduction

A cooperative (profit) game $\Gamma = (N, \gamma)$ consists of a player set $N = \{1, 2, \dots, n\}$ and a characteristic function $\gamma : 2^N \rightarrow R_+$ with $\gamma(\emptyset) = 0$, where $\gamma(S)$ ($S \subseteq N$) represents the profit achieved by the players in S without participation of other players. Different philosophies on the allocation of the total profit $\gamma(N)$ result in different solution concepts, *e.g.*, the core, the Shapley value, the nucleolus and the stable set.

A vector $x = (x_1, x_2, \dots, x_n)$ is called an *imputation* of Γ if $x(N) = \gamma(N)$ and $x_i \geq \gamma(\{i\})$ for each $i \in N$ (individual rationality); denote by $I(\Gamma)$ the set of imputations of Γ . The *core* $C(\Gamma)$ of game Γ is defined to be a set of imputations satisfying subgroup rationality, *i.e.*,

$$C(\Gamma) = \{x \in \mathbb{R}^n : x(N) = \gamma(N), \text{ and } x(S) \geq \gamma(S), \forall S \subseteq N\}.$$

In this paper we use the shorthand $x(S) = \sum_{i \in S} x_i$. The *stable set*, due to von Neumann and Morgenstern [13], is a set of imputations \mathcal{F} satisfying that any two imputations in \mathcal{F} do not dominate each other and any imputation not in \mathcal{F} can be dominated by some imputation in \mathcal{F} . Although the concept of stability is very useful in the analysis of bargaining situations, it seems difficult to investigate its properties and computational issues because of the complexity of its definition [5].

In general, the core and the stable set are different, however, Shapley [10] proved that for convex games, the core is the unique stable set. This result

^{*} This work is supported by NSFC (No.10571117, 10671108 and 10771200) and NCET.

motivated researchers to study the problem: when do the core and the stable set coincide, that is, when is the core stable? As far as the core stability for concrete cooperative game models is concerned, only a few results have been obtained, such as for assignment games [11] and minimum coloring games [2].

The network flow is one of the most widely studied optimization problems and has numerous applications, see [1]. *Flow games* were first introduced by Kalai and Zemel [7,8], which arose from the profit allocation problem related to the maximum flow in a network. Kalai and Zemel [7,8], and Deng *et al.* [4] showed that the cores of flow games are always nonempty, and an element in the core can be found in polynomial time. On the other hand, it was proved that checking whether a given imputation belongs to the core is *co-NP*-complete [6]. Recently, both Deng, Fang and Sun [3], and Potters, Reijnierse and Biswas [9] showed that the nucleolus of a simple flow game (where all arc capacities in the network are equal to one) can be computed in polynomial time, while for general flow games, computing the nucleolus is *NP*-hard [3].

The motivation of this work is to extend the core stability to general flow games. For this purpose, the concepts of the *restricted core* and *restricted core stability* of flow games are introduced. We first give a characterization on the restricted core, and then propose a sufficient and necessary condition on the restricted core stability for flow games associated with general networks. Furthermore, we show that testing the restricted core stability can be done in polynomial time. The known results on simple flow games, such as the core characterization and core stability condition, can be viewed as direct corollaries of our results.

The organization of the paper is as follows. In Section 2 we introduce the definitions of restricted core and restricted core stability of flow games. Section 3 is dedicated to a structural description of the restricted core. In Section 4 we focus on the characterization of the network for which the associated flow game possesses the restricted stable core.

2 Definitions

2.1 Flow Game

A flow network is denoted by $D = (V, E; c; s, t)$, where V is the vertex set, E is the arc set and $c : E \rightarrow \mathbb{R}_+$ is the arc capacity function, s (the source) and t (the sink) are two distinct vertices of D . An arc $e \in E$ from u to v is denoted by (u, v) , u and v are called the tail and the head of e , respectively. For two disjoint subsets $X, Y \subseteq V$, let $E(X, Y)$ denote the set of arcs of E with tails in X and heads in Y , $\delta_E^+(X) = E(X, \bar{X})$ and $\delta_E^-(X) = \delta_E^+(\bar{X})$ (where $\bar{X} = V \setminus X$). An arc set $\delta_E^+(X)$ with $s \in X$ and $t \notin X$ is called an s - t cut. In the case without confusion, $\delta_E^+(X)$ and $\delta_E^-(X)$ is denoted briefly as $\delta^+(X)$ and $\delta^-(X)$, respectively. For $X \subseteq V$, denote by $E(X)$ the set of arcs with both tails and heads in X .

Let $\mathcal{P} = \{P_1, \dots, P_m\}$ be the set of all simple directed paths from s to t (called s - t paths) in D , and $E = \{e_1, e_2, \dots, e_n\}$. Let A be the path-arc incidence matrix, *i.e.*, the matrix with rows and columns indexed by paths P_1, \dots, P_m and

edges e_1, \dots, e_n , respectively, and the entry in position (P_i, e_j) is 1 or 0 according to whether P_i passes through e_j or not ($i = 1, 2, \dots, m; j = 1, 2, \dots, n$). Define f_i as the value of flow along path P_i ($i = 1, 2, \dots, m$), and $f = (f_1, f_2, \dots, f_m)$. Then the maximum flow from s to t can be formulated as the following linear program:

$$\max\{f\mathbf{1} : fA \leq \mathbf{c}, f \geq 0\} \tag{2.1}$$

where $\mathbf{1}$ is the all-one column vector of dimensions m and $\mathbf{c} = (c(e_1), c(e_2), \dots, c(e_n))$. The flow through arc $e \in E$ is denoted by $f(e) = \sum_{P_i: P_i \ni e} f_i$. Thus, the constraints $fA \leq \mathbf{c}$ can be rewritten as $f(e) \leq c(e)$ for each $e \in E$. The dual of the above maximum flow problem (2.1) is

$$\min\{\mathbf{c}z : Az \geq \mathbf{1}, z \geq 0\} \tag{2.2}$$

Let γ^* denote the maximum flow value in D , which equals optimal value of dual program (2.2). Set

$$\begin{aligned} Q_D &= \{z = (z_1, z_2, \dots, z_n) : Az \geq \mathbf{1} \text{ and } \mathbf{c}z = \gamma^*\}, \\ Q_D^c &= \{(c_1z_1, c_2z_2, \dots, c_nz_n) : z = (z_1, z_2, \dots, z_n) \in Q_D\}. \end{aligned} \tag{2.3}$$

The flow game associated with network $D = (V, E; c; s, t)$ is denoted by Γ_D . We assume that each player controls one arc in E , *i.e.*, we can identify the set of arcs with the set of players. The flow game $\Gamma_D = (E, \gamma)$ is defined formally as follows:

- (i) The player set is E ;
- (ii) $\forall S \subseteq E$, the characteristic function value, denoted by $\gamma(S)$, is the value of the maximum flow from s to t in the subnetwork induced by the arc set S .

Theorem 2.1 [4,7,8]. *Let Γ_D be the flow game associated with network $D = (V, E; c; s, t)$. Then $Q_D^c \subseteq C(\Gamma_D)$.*

The above theorem shows that the core of a flow game is always non-empty and an element in the core can be found in polynomial time [4,7,8]. However, the problem of testing whether an imputation is in the core is *co-NP*-complete [6]. Although in general, Q_D^c is a proper subset of $C(\Gamma_D)$, when the network D is simple, Q_D^c coincides with $C(\Gamma_D)$. The following characterization is due to Kalai and Zemel [8] and Deng *et al.* [4].

Theorem 2.2 [4,8]. *Let Γ_D be the flow game associated with simple network D . Then $C(\Gamma_D) = Q_D^c$, and moreover, $C(\Gamma_D)$ is exactly the convex hull of the indicator vectors of the minimum s - t cuts of D .*

2.2 Restricted Core Stability

Given a cooperative game $\Gamma = (N, \gamma)$, the core $C(\Gamma)$ is stable means that for any $y \in I(\Gamma) \setminus C(\Gamma)$, there exists a core element $x \in C(\Gamma)$ and a nonempty coalition $S \subset N$ such that

$$(a) \ x_i > y_i \text{ for all } i \in S; \quad (b) \ x(S) = \gamma(S). \tag{2.4}$$

The main purpose of this paper is to propose a new kind of core stability, called restricted core stability, by adding two restrictions to the original core stability on a general flow game Γ_D : (i) $C(\Gamma_D)$ is restricted to Q_D^c , and (ii) the nonempty coalition $S \subset E$ is restricted to an s - t path. More formally, the restricted version of the core stability of flow games is defined as follows.

Definition 2.1. Let $\Gamma_D = (E, \gamma)$ be the flow game associated with network $D = (V, E; c; s, t)$. Q_D^c is called the restricted core of Γ_D ; moreover, Q_D^c is called restricted stable if for every $y \in I(\Gamma_D) \setminus Q_D^c$, there exist an $x \in Q_D^c$ and an s - t path P in D such that

$$(a^*) x(e) > y(e) \text{ for all } e \in P; \quad (b^*) x(P) = \gamma(P). \tag{2.5}$$

We also say that x P -dominates y if (a^*) holds.

The reasons why we impose the restrictions on the core stability are as follows:

First, testing whether $x \in I(\Gamma_D) \setminus C(\Gamma_D)$ itself is NP-hard [6], while on the contrary, testing whether $x' \in I(\Gamma_D) \setminus Q_D^c$ is polynomial-time solvable. Second, for flow game Γ_D , a nonempty coalition $S \subset E$ satisfying (2.4)(a) contains an s - t path satisfying (2.5)(a*), conversely (2.5)(a*) is a special case of (2.4)(a). Third, when D is a simple network, two stabilities coincide, i.e., $Q_D^c = C(\Gamma_D)$ and the nonempty coalition $S \subset E$ can be restricted to an s - t path [12].

3 Restricted Core Characterization

In what follows, a flow means a flow from s to t and a minimum s - t cut means an s - t cut with minimum capacity. Without loss of generality, we may assume that every arc is on some s - t path, moreover, the capacity $c(e) > 0$ for each $e \in E$ since only the maximum flow problem is concerned. Recall the definition of Q_D and Q_D^c in (2.3), we rewrite Q_D^c as follows:

$$Q_D = \{z : Az \geq \mathbf{1} \text{ and } cz = \gamma^*\}, \quad Q_D^c = \{Mz : z \in Q_D\},$$

where M is the square diagonal matrix with the components of \mathbf{c} on its main diagonal. In this section, we shall give a structural characterization of the restricted core Q_D^c . From the complementary slackness of LP, we first give the following lemma.

Lemma 3.1. Let $z \in Q_D$. If $z(e) > 0$, then there is an s - t path $P \in \mathcal{P}$ such that $e \in P$ and $z(P) = 1$.

Theorem 3.2. The set Q_D is exactly the convex hull of the indicator vectors of the minimum s - t cuts of D .

Proof. It is easy to see that the indicator vector of any minimum s - t cut of D is in Q_D , so is the convex combination of these indicator vectors.

Conversely, suppose $z \in Q_D$. Consider z as a length function on E , and let $d(s, v)$ be the length of the shortest path from s to $v \in V$ w.r.t. length z . Then we first have the claim:

Claim. $d(s, t) = 1$ and $d(s, v) \leq 1$ for every $v \in V$.

In fact, the complementary slackness of LP yields directly $d(s, t) = 1$. For $v \in V$, let P be an s - v path in D and $e = (u', u'')$ be the last arc on P with $z(e) > 0$. Then by Lemma 3.1, there is an s - t path P_i passing through e with $z(P_i) = 1$, which implies that $d(s, v) \leq d(s, u'') \leq z(P_i) = 1$.

Based on the claim, we may rearrange the different values of distances as $0 = p_0 < p_1 < \dots < p_k = 1$, and consequently, partition the vertex set V into k subsets:

$$V_i = \{u \in V \mid d(s, u) = p_i\}, \quad i = 0, 1, \dots, k. \tag{3.1}$$

From the partition (3.1), we obtain k s - t cuts of D :

$$X_i = \bigcup_{j=0}^{i-1} V_j, \quad C_i = \delta^+(X_i), \quad i = 1, \dots, k. \tag{3.2}$$

Now let us show each C_i ($i = 1, \dots, k$) in (3.2) to be a minimum s - t cut. Clearly, $z(e) > 0$ for each $e \in C_i$, implying that $f(e) = c(e)$ for any maximum flow f . We also have $f(e') = 0$ for any $e' = (u', u'') \in \delta^-(X_i)$ and any maximum flow f . For otherwise, by Lemma 3.1, there is an s - t path P_j passing through e' with $z(P_j) = 1$. Denote \hat{P}' the sub-path of P_j from s to u'' , then $z(\hat{P}') \geq d(s, u') > d(s, u'')$. Replacing \hat{P}' by a shortest s - u'' path, we may obtain an s - t path with length less than 1, a contradiction.

Let χ_i denote the indicator vector of C_i ($i = 1, \dots, k$), then it can be shown that z is a convex combination of $\chi_1, \chi_2, \dots, \chi_k$, furthermore,

$$z = \sum_{i=1}^k (p_i - p_{i-1}) \chi_i.$$

In fact, if $z(e) = 0$, then $e \notin \cup_{i=1}^k C_i$; if $z(e) > 0$ for $e = (u, v)$, then by Lemma 3.1, there is an s - t path P_i passing through e with $z(P_i) = 1$, which implies that P_i is a shortest path and $d(s, v) = d(s, u) + z(e)$. Suppose $u \in V_i$ and $v \in V_j$, then $j > i$ and e is covered only by C_{i+1}, \dots, C_j , and hence, $z(e) = p_j - p_i = \sum_{\ell=i}^{j-1} (p_{\ell+1} - p_\ell)$.

The proof is completed. ■

When the network D is simple, $Q_D^c = Q_D$, Theorem 2.2 is a direct corollary of Theorems 3.2.

4 Restricted Core Stability

In this section, we discuss sufficient and necessary conditions on the restricted core stability for flow games. Given the network $D = (V, E; c; s, t)$, define

$$\begin{aligned} E_1 &= \{e \in E : e \in \delta^+(X) \text{ for some minimum } s\text{-}t \text{ cut } \delta_E^+(X) \text{ in } D\}, \\ \overline{E}_1 &= E \setminus E_1. \end{aligned}$$

It can be seen that for any maximum flow f ,

$$f(e) = c(e), \quad \forall e \in E_1. \tag{4.1}$$

Lemma 4.1. *There exists a minimal laminar collection $\mathbb{C}^* = \{\delta^+(X_1), \dots, \delta^+(X_q)\}$ of minimum s - t cuts such that*

- (a) $X_1 \subset X_2 \subset \dots \subset X_q$ and $\bigcup_{i=1}^q \delta^+(X_i) = E_1$;
- (b) For each $v \in X_i$ ($i = 1, 2, \dots, q$), there exist an s - v path with all vertices in X_i .

Proof. (a) Since every $e \in E_1$ is covered by a minimum s - t cut, there exists a collection \mathbb{C} of minimum s - t cuts covering E_1 . Among all such covering collections, we choose a collection such that

$$\sum_{\delta^+(X) \in \mathbb{C}} |X| |\overline{X}| \tag{4.2}$$

is minimized, and denote it by $\mathbb{C}^* = \{\delta^+(X_1), \dots, \delta^+(X_q)\}$.

We show that in \mathbb{C}^* , either $X_i \subset X_j$ or $X_j \subset X_i$ for $i \neq j$. Suppose to the contrary that, there is a pair $1 \leq i < j \leq q$ such that $X_i \setminus X_j \neq \emptyset$ and $X_j \setminus X_i \neq \emptyset$. Then it is easy to see that both $\delta^+(X_i \cap X_j)$ and $\delta^+(X_i \cup X_j)$ are minimum s - t cuts, furthermore,

$$|X_i \cap X_j| |\overline{X_i \cap X_j}| + |X_i \cup X_j| |\overline{X_i \cup X_j}| < |X_i| |\overline{X_i}| + |X_j| |\overline{X_j}|,$$

$$\delta^+(X_i \cap X_j) \cup \delta^+(X_i \cup X_j) \supseteq \delta^+(X_i) \cup \delta^+(X_j).$$

Thus, replacing $\delta^+(X_i)$ and $\delta^+(X_j)$ by $\delta^+(X_i \cap X_j)$ and $\delta^+(X_i \cup X_j)$ in \mathbb{C}^* , we obtain a new collection covering E_1 with the sum (4.2) decreased, contradicting to the assumption of \mathbb{C}^* . We can relabel the members of \mathbb{C}^* such that $X_1 \subset \dots \subset X_q$ holds. The minimality of \mathbb{C}^* derives directly from the minimization of (4.2).

Let C_i denote $\delta^+(X_i)$ in \mathbb{C}^* , $i = 1, 2, \dots, q$. To prove (b), let j be the minimum index such that for some $v \in X_j$, there is no s - v path with all vertices in X_j , and let V'_j denote the set of such vertices in X_j .

We claim that $\delta^+(v) \cap C_j = \emptyset$ for every $v \in V'_j$. For otherwise, say $e = (v, w) \in \delta^+(v) \cap C_j$, then $f(e) = c(e)$ for any maximum flow f . On the other hand, every s - w path P' must contains backward arc of some minimum cut in \mathbb{C}^* , implying that $f(e) = 0$, a contradiction.

Modify the sets X_1, X_2, \dots, X_q obtained in (a) as follows:

$$X'_i = \begin{cases} X_j \setminus V'_j & \text{if } i = j \\ X_{j+1} \cup V'_j & \text{if } i = j + 1 \\ X_i & \text{otherwise.} \end{cases}$$

It is easy to see that $X'_1 \subset X'_2 \subset \dots \subset X'_q$, $\bigcup_{i=1}^q \delta^+(X'_i) = E_1$ and for each $v \in X'_i$ ($i \leq j$), there exists an s - v path with vertices all in X'_i . Repeat this process until the collection \mathbb{C}^* has the property (b). ■

Lemma 4.1 plays an important role in our discussion. In the rest of this paper, X_1, X_2, \dots, X_q are fixed to be the sets defined in Lemma 4.1, denote $C_i = \delta^+(X_i)$ ($i = 1, 2, \dots, q$) and $\mathbb{C}^* = \{C_1, C_2, \dots, C_q\}$. From Lemma 4.1, it can be seen that for any maximum flow f ,

$$f(e) = 0, \quad \forall e \in \delta^-(X_i). \tag{4.3}$$

Before giving the main result of this section, we first present some properties on P -restricted domination.

Lemma 4.2. *If each $y \in I(\Gamma_D) \setminus Q_D^c$ can be P -dominated by an element in Q_D^c , then*

$$c(\delta^-(v) \cap E_1) = c(\delta^+(v) \cap E_1), \quad \forall v \in V \setminus \{s, t\}. \tag{4.4}$$

Proof. Suppose $v \in X_i \setminus X_{i-1}$ and $c(\delta^-(v) \cap E_1) > c(\delta^+(v) \cap E_1)$, i.e., $c(\delta^-(v) \cap C_{i-1}) > c(\delta^+(v) \cap C_i)$. Then for any maximum flow f ,

$$\begin{aligned} f(\delta^+(v) \setminus C_i) &= f(\delta^-(v)) - f(\delta^+(v) \cap C_i) \\ &\geq c(\delta^-(v) \cap C_{i-1}) - c(\delta^+(v) \cap C_i) > 0. \end{aligned}$$

By Lemma 4.1 and formula (4.3), we have $\emptyset \neq E(v, X_i \setminus X_{i-1}) \subseteq \delta^+(v) \setminus C_i \subseteq \overline{E_1}$. Hence, $\delta^+(X_{i-1} \cup \{v\})$ is not a minimum s - t cut, implying that $c(\delta^-(v) \cap C_{i-1}) < c(\delta^+(v) \cap C_i) + c(E(v, X_i \setminus X_{i-1}))$. Let

$$\begin{aligned} \lambda &= \frac{c(\delta^-(v) \cap C_{i-1}) - c(\delta^+(v) \cap C_i)}{c(E(v, X_i \setminus X_{i-1}))}, \\ y(e) &= \begin{cases} c(e) & \text{if } e \in (\delta^+(v) \cap C_i) \cup (C_{i-1} \setminus \delta^-(v)) \\ \lambda c(e) & \text{if } e \in E(v, X_i \setminus X_{i-1}) \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

From the above analysis, we have $0 < \lambda < 1$. And it is easy to check that $y \in I(\Gamma_D) \setminus Q_D^c$, but for all $x \in Q_D^c$ and any s - t path P , $x(e) \leq y(e)$ for $e \in P \cap \delta^+(X_{i-1} \cup \{v\})$. It follows that y can not be P -dominated by any element in Q_D^c , a contradiction. With similar arguments, it is also shown that $c(\delta^-(v) \cap E_1) < c(\delta^+(v) \cap E_1)$ is not true. ■

Lemma 4.3. *If each $y \in I(\Gamma_D) \setminus Q_D^c$ can be P -dominated by an element in Q_D^c and $e = (u, v) \notin E_1$, then*

- either (a) $u, v \in X_i \setminus X_{i-1}$ for some $2 \leq i \leq q$;
 - or (b) $u \in X_j \setminus X_{j-1}$ and $v \in X_i \setminus X_{i-1}$ for some $1 \leq i < j \leq q$.
- Moreover, in case (a), $\delta^-(v) \cap C_{i-1} = \delta^+(v) \cap C_i = \emptyset$.

Proof. We first show that $u \neq s$ and $v \neq t$. Indeed, if $u = s$, i.e., $e \in \delta^+(s)$, then $c(\delta^+(s)) > \gamma^* \geq c(\delta^+(s) \cap E_1)$. Set

$$y(e) = \begin{cases} c(e) & \text{if } e \in \delta^+(s) \cap E_1 \\ \frac{\gamma^* - c(\delta^+(s) \cap E_1)}{c(\delta^+(s) \setminus E_1)} c(e) & \text{if } e \in \delta^+(s) \setminus E_1 \\ 0 & \text{otherwise.} \end{cases}$$

Obviously, $y \in I(\Gamma_D) \setminus Q_D^c$, and it can not be P -dominated by any $x \in Q_D^c$. Similarly, $v \neq t$ can be proved. Thus, in the subsets X_1, X_2, \dots, X_q obtained in Lemma 4.1, $X_1 = \{s\}$ and $X_q = V \setminus \{t\}$, implying the cases (a) and (b).

Now we show $\delta^-(v) \cap C_{i-1} = \delta^+(v) \cap C_i = \emptyset$. Otherwise, set

$$\hat{x}(e) = \begin{cases} c(e) & \text{if } e \in (\delta^-(v) \cap C_{i-1}) \cup (C_i \setminus \delta^+(v)) \\ 0 & \text{otherwise.} \end{cases}$$

Notice (4.4), with the similar analysis as in the proof of Lemma 4.2, it can be checked that $\hat{x} \in I(\Gamma_D) \setminus Q_D^c$ and it cannot be P -dominated by any element in Q_D^c , a contradiction. ■

Let $V_0 = \{v \in V : \delta^+(v) \cup \delta^-(v) \subseteq \overline{E_1}\}$,
 $V_1 = \{v \in V : \delta^+(v) \cap E_1 \neq \emptyset, \delta^-(v) \cap E_1 \neq \emptyset\}$.

Lemma 4.3 shows that when Q_D^c is P -dominated, for any $e = (u, v) \notin E_1$ with $u, v \in X_i \setminus X_{i-1}$ for some $2 \leq i \leq q$, the head v must belong to V_0 ; and for any $e = (u, v) \in E_1$, both the tail u and the head v belong to V_1 .

Lemma 4.4. *If each $y \in I(\Gamma_D) \setminus Q_D^c$ can be P -dominated by an element in Q_D^c and $e = (u, v) \notin E_1$, then when $u \in X_j \setminus X_{j-1}$, there is a path $P = v_1 v_2 \cdots v_l$ passing through e satisfying that*

- (i) $v_1 \in X_j \setminus X_{j-1}, v_l \in X_i \setminus X_{i-1}$ for some $i < j$;
- (ii) $v_1, v_l \in V_1$ and $v_2, \dots, v_{l-1} \in V_0$.

Proof. By Lemma 4.1 (b), there is an s - u path contained in X_j , say $P_1 = s \cdots v_0 v_1 \cdots v_k (= u)$, where $v_0 \in X_{j-1}$ and $v_1, v_2, \dots, v_k \in X_j \setminus X_{j-1}$. Also by Lemma 4.1 and 4.3, we have $v_1 \in V_1$ and $v_2, \dots, v_k \in V_0$.

On the other hand, there is an s - t path traversing e , then its u - t subpath $P'_1 = v_k v_{k+1} \cdots v_l \cdots t$ ($v_k = u, v_{k+1} = v$) traverses internally some vertex in V_1 . Suppose v_l is the first such vertex traversed and $v_l \in X_i \setminus X_{i-1}$. Then $i < j$ and $v_{k+1}, \dots, v_{l-1} \in V_0$.

Clearly, the coalesced path $P = v_1 v_2 \cdots v_l$ (by P_1 and P'_1) is as required. ■

Now we are in the position to present the sufficient and necessary condition on the restricted stability of Q_D^c . We first need to identify a special kind of arcs in $\overline{E_1}$, denoted by E_0 .

- Let E_0 be the set of arcs of paths P satisfying that all internal vertices of P belong to V_0 , and moreover, there is a path $P' \subseteq E_1$ from the end vertex to the starting vertex of P .

Theorem 4.5. *Let Γ_D be the flow game associated with network $D=(V, E; c; s, t)$. Then Q_D^c is restricted stable if and only if*

- (1) $E = E_1 \cup E_0$; (4.5)
- (2) for any s - t path $P \subseteq E_1$, the arcs in P have the same capacity.

Proof. The proof depends on the laminar collection \mathbb{C}^* of minimum s - t cuts in Lemma 4.1. Denote $c_P = \min\{c(e) : e \in P\}$ for any s - t path P .

Necessity. Suppose that $E \setminus (E_1 \cup E_0) \neq \emptyset$. Then for each $e \in E \setminus (E_1 \cup E_0)$, there is a path $P_e \subseteq \overline{E_1}$ satisfying conditions (i) and (ii) of Lemma 4.4, and moreover, there is no v_l - v_1 path contained in E_1 . Among all such paths, choose

$P^* = v_1 v_2 \cdots v_l \subseteq \overline{E}_1$ with the smallest index $j = j^*$ ($v_l \in X_{j^*} \setminus X_{j^*-1}$) firstly and the largest index $i = i^*$ ($v_l \in X_{i^*} \setminus X_{i^*-1}$) secondly. Set

$$\begin{aligned} R_1 &= \{v \in X_{j^*} \setminus X_{i^*} : \text{there is a } v\text{-}v_1 \text{ path contained in } E_1\}, \\ R_2 &= \{v \in X_{j^*} \setminus X_{i^*} : \text{there are } u \in R_1 \text{ and a } u\text{-}v \text{ path contained in } \overline{E}_1\}; \\ R &= R_1 \cup R_2. \end{aligned}$$

It is not difficult to show that

- (a) $R_1 \subseteq V_1$;
- (b) $R_2 \cap V_1 \subseteq R_1$, and furthermore, $R \setminus R_1 \subseteq V_0$.

Claim 1. Let $\tilde{X} = X_{i^*} \cup R$, then $\delta^+(\tilde{X})$ is also a minimum s - t cut in D .

In fact, $E(V \setminus X_{j^*}, R) \cap E_1 = \emptyset$; and $E(X_{j^*} \setminus \tilde{X}, R) \cap E_1 = \emptyset$, since $R \setminus R_1 \subseteq V_0$. Hence

$$\delta^-(R) \cap E_1 = E(X_{i^*}, R) \subseteq C_{i^*}. \tag{4.6}$$

Also by the definition of R , we have $E(R, X_{j^*} \setminus \tilde{X}) \subseteq E_1$, yielding that

$$\delta^+(R) \cap E_1 = E(R, V \setminus \tilde{X}). \tag{4.7}$$

Combining formulas (4.6),(4.7) and $c(\delta^-(R) \cap E_1) = c(\delta^+(R) \cap E_1)$ (4.4), we have

$$c(\delta^+(\tilde{X})) = c(\delta^+(X_{i^*})) - c(E(X_{i^*}, R)) + c(E(R, V \setminus \tilde{X})) = c(\delta^+(X_{i^*})).$$

That is, $\delta^+(\tilde{X})$ is a minimum s - t cut, the claim is proved.

Based on Claim 1, we set

$$y(e) = \begin{cases} c(e) & \text{if } e \in \delta^+(\tilde{X} \setminus \{v_l\}) \cap E_1 \\ 0 & \text{otherwise.} \end{cases} \tag{4.8}$$

Let us first show $y \in I(\Gamma_D)$. Since there is no v_l - v_1 path in E_1 by the choice of P^* , hence, $E(v_l, R) \cap E_1 = \emptyset$. By Lemma 4.2, $c(\delta^+(\tilde{X} \setminus \{v_l\}) \cap E_1) = c(\delta^+(\tilde{X})) = \gamma^*$. Next, we show that $y \notin Q_D^c$. Notice that $v_1, v_l \in V_1$, so there are paths P_1 from s to v_1 and P_2 from v_l to t in E_1 . Furthermore, P_1 and P_2 are vertex-disjoint as there is no v_l - v_1 path in E_1 and $\delta^+(\tilde{X})$ is a minimum cut. Hence, the coalesced path $P = P_1 P^* P_2$ is an s - t path with $y(P) = 0$, implying that $y \notin Q_D^c$.

It is easy to verify that for all $x \in Q_D^c$ and any s - t path P , $x(e') \leq y(e')$ if $e' \in \delta^+(\tilde{X} \setminus \{v_l\}) \cap P$, i.e., y cannot be P -dominated by any element in Q_D^c . Therefore, $E \setminus (E_1 \cup E_0) = \emptyset$.

Finally, we show condition (2) holds. Let $z \in Q_D$ and $x = Mz$. Let P be an arbitrary s - t path with $P \subseteq E_1$. Since $x \in Q_D^c \subseteq C(\Gamma_D)$,

$$x(E \setminus P) \geq \gamma(E \setminus P) = \gamma^* - c_P, \text{ i.e., } x(P) \leq c_P. \tag{4.9}$$

Followed from $P \subseteq E_1$, we also have $z(e) > 0$ for each $e \in P$ and $z(P) = 1$, yielding that

$$x(P) = \sum_{e \in P} x(e) = \sum_{e \in P} c(e)z(e) \geq c_P, \tag{4.10}$$

where the equality holds if and only if $c(e) = c_P$ for each $e \in P$. Condition (2) follows directly from (4.9) and (4.10).

Sufficiency. Given $y \in I(\Gamma_D) \setminus Q_D^c$, let $z = M^{-1}y$, it follows that z is not a feasible solution to (2.2), that is, there is some s - t path P satisfying $z(P) < 1$. Let $D_r = D \setminus E_0$, \hat{z} and \hat{c} be the restrictions of z and c to E_1 , respectively. Obviously, the maximum flow value in D_r is the same as that in D .

Claim 2. There is an s - t path P^* in D_r with $z_r(P^*) < 1$.

Otherwise, z_r is an optimal solution to (2.2) restricted to D_r , that is, $z_r(P') = 1$ for all s - t path P' in D_r . By the induction on the number k of subpaths of P connecting two vertices of V_1 and with all internal vertices in V_0 , we show that $z \in Q_D$, i.e., $z(P) \geq 1$ for all s - t paths P in D .

Suppose that $P = P_u P_{uv} P_v$, where P_u is an s - u path, P_v is a v - t path and P_{uv} is a u - v path with $u, v \in V_1$ and all internal vertices in V_0 . As $u, v \in V_1$, there exists an s - v path P'_v and a u - t path P'_u in D_r . For the connection between u and v , since $E \setminus E_1 = E_0$, there exists a v - u path P'_{vu} in D_r . Therefore, the coalesced path $P'_v P'_{vu} P'_u$ is an s - t path in D_r and $z(P'_v P'_{vu} P'_u) = z_r(P'_v P'_{vu} P'_u) = 1$. By induction, $z(P_u P'_u) \geq 1$, $z(P'_v P_v) \geq 1$, which implies that

$$z(P) \geq z(P_u) + z(P_v) \geq z(P'_v P'_{vu}) + z(P'_{vu} P'_u) \geq 1.$$

Thus, $z \in Q_D$, which contradicts to the assumption $y \notin Q_D^c$. The claim is proved.

Now we proceed to define a required $z' \in Q_D$ based on P^* . For each $e \in P^*$, let $k(e)$ be the number of s - t cuts in \mathbb{C}^* covering e . Obviously, $\sum_{e \in P^*} k(e) = q$. For each $C_i \in \mathbb{C}^*$ with $C_i \cap P^* = \{e\}$ ($i = 1, 2, \dots, q$), set

$$\lambda_i = \frac{z(e)}{k(e)} + \frac{1 - z(P^*)}{q},$$

$$z'(e) = \begin{cases} \sum \{\lambda_i : e \in C_i \in \mathbb{C}^*\} & \text{if } e \in E_1 \\ 0 & \text{otherwise.} \end{cases}$$

That is, $z' = \sum_{i=1}^q \lambda_i \chi_i$, where χ_i is the indicator vector of C_i ($i = 1, 2, \dots, q$). It is easy to check that

- $0 \leq \lambda_i \leq 1$ ($i = 1, 2, \dots, q$) and $z'(P^*) = \sum_{i=1}^q \lambda_i = 1$.
- $z'(e) = z(e) + \frac{k(e)(1 - z(P^*))}{q} > z(e)$ for each $e \in P^*$.
- $z'(P) \geq \sum_{i=1}^q \lambda_i = 1$ for any s - t path P in D .

By Theorem 3.2, $z' \in Q_D$. Let $x = Mz'$, then $x \in Q_D^c$ and $x(e) = c(e)z'(e) > c(e)z(e) = y(e)$ for each $e \in P^*$. Also by condition (2), all arcs in P^* have the same capacity c_{P^*} , we have

$$x(P^*) = \sum_{e \in P^*} c(e)z'(e) = c_{P^*} \sum_{e \in P^*} z'(e) = c_{P^*} = \gamma(P^*).$$

Therefore, Q_D^c is restricted stable. ■

Corollary 4.6. *Let Γ_D be the flow game associated with simple network $D = (V, E; s, t)$. Then the core $C(\Gamma_D)$ is stable if and only if $E = E_1 \cup E_0$.*

Now we consider the computational complexity for detecting the restricted core stability. It is well known that there is a polynomial time algorithm for seeking out E_1 (see, e.g., [II]). On the other hand, the vertex sets V_1 and V_0 can be identified easily knowing the set E_1 . Furthermore, for each pair of vertices $u, v \in V_1$, by checking whether there is not only a u - v path with all internal vertices in V_0 but also a v - u path using all arcs in E_1 , we can identify all the arcs in E_0 . Therefore, we have

Theorem 4.7. *Testing the restricted core stability of a flow game can be done in polynomial time.*

References

1. Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: Network Flows. Prentice Hall, Inc., Englewood Cliffs (1993)
2. Bietenhader, T., Okamoto, Y.: Core Stability of Minimum Coloring Games. In: Hromkovič, J., Nagl, M., Westfechtel, B. (eds.) WG 2004. LNCS, vol. 3353, pp. 389–401. Springer, Heidelberg (2004)
3. Deng, X., Fang, Q., Sun, X.: Finding nucleolus of flow games. In: SODA 2006, pp. 124–131 (2006)
4. Deng, X., Ibaraki, T., Nagamochi, H.: Algorithmic aspects of the core of combinatorial optimization games. *Mathematics of Operations Research* 24, 751–766 (1999)
5. Deng, X., Papadimitriou, C.H.: On The Complexity of Cooperative Solution Concepts. *Mathematics of Operations Research* 19, 257–266 (1994)
6. Fang, Q., Zhu, S., Cai, M., Deng, X.: Membership for core of LP games and other games. In: Wang, J. (ed.) COCOON 2001. LNCS, vol. 2108, pp. 247–256. Springer, Heidelberg (2001)
7. Kalai, E., Zorn, E.: Totally balanced games and games of flow. *Mathematics of Operations Research* 7, 476–478 (1982)
8. Kalai, E., Zorn, E.: Generalized network problems yielding totally balanced games. *Operations Research* 30, 498–1008 (1982)
9. Potters, J., Reijnierse, H., Biswas, A.: The nucleolus of balanced simple flow networks. *Games Econom. Behav.* 54, 205–225 (2006)
10. Shapley, L.S.: Cores and convex games. *International Journal of Game Theory* 1, 11–26 (1971)
11. Solymosi, T., Raghavan, T.E.S.: Assignment Games with Stable Cores. *International Journal of Game Theory* 30, 177–185 (2001)
12. Sun, X., Fang, Q.: Core stability of flow games. In: Akiyama, J., Chen, W.Y.C., Kano, M., Li, X., Yu, Q. (eds.) CJCDGCGT 2005. LNCS, vol. 4381, pp. 189–199. Springer, Heidelberg (2007)
13. von Neumann, J., Morgenstern, O.: *Theory of Games and Economic Behaviour*. Princeton University Press, Princeton (1944)

Three Selfish Spanning Tree Games

Laurent Gourvès^{1,2} and Jérôme Monnot^{1,2}

¹ CNRS, UMR 7024, F-75775 Paris, France

² Université de Paris-Dauphine, LAMSADE, F-75775 Paris, France
{laurent.gourves,monnot}@lamsade.dauphine.fr

Abstract. We study a problem in a network. The input is an edge-weighted graph $G = (V, E)$ such that V contains a specific source node r . Every $v \in V \setminus \{r\}$ is an entity which wants to be connected to r either directly or via other entities. The main question is how do the entities deviate from a socially optimal network if they are not monitored by a central authority. We provide theoretical bounds on the (strong) price of anarchy of this game. In particular, three variants – each of them being motivated by a practical situation – are studied.

1 Introduction

How (in)efficiently self-interested agents make use of a common resource? This question is central in computer science because today's major platform – networks – is often operated by spontaneous and selfish users. In recent works including [1,2,3,8] the situation is modelled as a *strategic game*. A strategic game is a tuple $\langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ where N is the set of players while S_i and u_i are respectively the set of strategies of player i and her utility function. Players are supposed to be selfish and rational. Each of them chooses a strategy in order to maximize her own utility.

In this paper, we study three versions of a *selfish spanning tree game* (SST-game in short) respectively denoted by min SST-game, max SST-game and bottleneck SST-game. The SST-game is defined upon a complete graph $G = (V, E)$ where every edge $e \in E$ has a positive weight $w(e)$. Each node except a specific *source* node r is controlled by a self-interested player. Players want to be connected to r either directly or via another player, herself connected to r . For the min SST-game (resp. the max SST-game), the utility of a player is the negative of the weight (resp. the weight) of the first edge of the (unique) path between her and r . For the bottleneck SST-game, a player's utility is the minimum weight of the edges of the path between her and r .

The SST-game is motivated by two situations: when the source node r sends a message to all player nodes (one-to-all communications) and when each player node sends a request to r (all-to-one communications). For the min SST-game, the weight of an edge can represent its cost. If each node pays the price of its upstream edge then the cost of the whole network is covered. Hence each player wishes to receive the message at the lowest price. Another application of the min SST-game arises in power control for static wireless networks. The weight

of an edge represents its length and the power needed to send a message to a remote node increases with the distance. Hence each player wishes to consume as little power as possible to send its request. For the max SST-game, the weight of an edge represents its reliability: the higher the weight of an edge, the less a message passing through it is vitiated by errors. Then it is in a player’s interest to choose the most reliable link from which data arrive in order to spend as little time as possible in repairing it (it is assumed that a message is always sent or forwarded without errors). For the bottleneck SST-game, the weight of an edge represents its free bandwidth. For one-to-all communications, the flow of data at a node depends on the flow at each intermediate node, i.e. it is limited by the edge with smallest free bandwidth. Therefore each player of the bottleneck SST-game wants to maximize her own flow.

The players of the SST-game may not spontaneously reach a social optimum as uncoordinated decisions and conflicting interests often lead to suboptimal performances. The main concern of this paper is to say how (in)efficiently players of the SST-game make use of the graph if they act selfishly. To do so, we study the *price of anarchy* (PoA) [12] and the *strong price of anarchy* (SPoA) [2] of the SST-game. These two measures resort to two important concepts in game theory – the *Nash equilibrium* and the *strong equilibrium* – to quantify the performance’s deterioration due to self-interested behaviors.

2 Definitions and Notations

The SST-game. We are given a complete graph $G = (V, E)$ on $n + 1$ vertices where V contains a specific node r and a weight function $w : E \rightarrow \mathbb{R}_+$. The SST-game is such that $N = V \setminus \{r\}$ (each node except r is a player) and $S_i = V$ for all $i \in N$ (each player selects a node). The number of players is finite and denoted by n . A *strategy profile* (or state) $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ is an element of $S = S_1 \times S_2 \times \dots \times S_n$. We say that σ connects $i \in N$ to the source iff there is a path between i and r in the graph $(V, \{[u, v] : \sigma_u = v \text{ or } \sigma_v = u\})$. For the min SST-game, the utility of a player i is defined as $u_i(\sigma) = -w([i, \sigma_i])$ if σ connects i to r , otherwise $u_i(\sigma) = -\infty$. For the max SST-game, the utility is defined as $u_i(\sigma) = w([i, \sigma_i])$ if σ connects i to r , otherwise $u_i(\sigma) = -\infty$. For the bottleneck SST-game, the utility is defined recursively as

$$u_i(\sigma) = \begin{cases} w([i, \sigma_i]) & \text{if } \sigma_i = r \\ \min\{w([i, \sigma_i]), u_{\sigma_i}(\sigma)\} & \text{if } \sigma \text{ connects } i \text{ to } r \\ -\infty & \text{otherwise} \end{cases}$$

Every player i selects her strategy σ_i so that $u_i(\sigma)$ is maximum. The graph being complete, any player can select her direct link to r and have a utility different from $-\infty$. Therefore the players always build a spanning tree of G if none of them can unilaterally improve her utility.

Throughout this article, the spanning tree induced by a strategy profile σ is denoted by T^σ and defined as $\{[i, j] \in E : \sigma_i = j \text{ or } \sigma_j = i\}$. The strategy profile

induced by a spanning tree T is denoted by σ^T and defined as $\sigma_i^T = j$ for all $i \in N$ where $[i, j]$ is the first edge of the unique path between i and r in T .

Equilibria and the price of anarchy. Given $\sigma \in S$, let us denote by $(\sigma \mid_i s)$ the strategy profile where σ_i is replaced by s in σ whereas the strategy of the other players is unchanged, i.e. $(\sigma_1, \dots, \sigma_{i-1}, s, \sigma_{i+1}, \dots, \sigma_n)$. A Nash equilibrium is a stable state in which no player has an incentive to unilaterally move away. That is, $\sigma \in S$ is a Nash equilibrium if $\forall i \in N, \forall s \in S_i, u_i(\sigma \mid_i s) \leq u_i(\sigma)$.

We speak about *pure* Nash equilibria when every player i deterministically chooses a strategy $\sigma_i \in S_i$. The SST-game always has a pure Nash equilibrium since such a state can be computed as follows. Start with any pure strategy profile (e.g. $\forall i \in N, \sigma_i = r$) and change the strategy of a player as far as it is profitable (i.e. her utility strictly increases). If such a deviation is not possible then the current state is a Nash equilibrium. Otherwise the utility of at least one player has increased and the utility of all the others has not decreased. Hence if σ and σ' respectively denote the strategy profile before and after the deviation, we have $\sum_{i \in N} u_i(\sigma') > \sum_{i \in N} u_i(\sigma)$. One can repeat the process until a pure strategy Nash equilibrium, i.e. a local optimum for $\sum_{i \in N} u_i(\sigma)$, is reached.

We only consider pure strategies so we omit from now on the adjective "pure". The price of anarchy (PoA) [12] is the value of the worst Nash equilibrium relative to the social optimum. As usual the social welfare is defined as the sum of the players' utility. We use a function $\mathcal{S} : S \rightarrow \mathbb{R}$ defined as $\mathcal{S}(\sigma) = \sum_{i \in N} u_i(\sigma)$.

Let σ be a strategy profile at Nash equilibrium whereas σ^* denotes a strategy profile induced by a social optimum [1]. The PoA of the SST-game is the maximum value of $\mathcal{S}(\sigma)/\mathcal{S}(\sigma^*)$ over all instances.

One particular weakness of the Nash equilibrium is its vulnerability to deviations by coalitions of players. Selfishness does not avoid cooperation as long as it is profitable. Aumann introduced the notion of *strong equilibrium* [4]. It is a Nash equilibrium where no set of players can cooperatively deviate in a way that benefits all its members, taking the strategies of the players outside the coalition as given. Given two strategy profiles σ, β and a set $N' \subseteq N$, $(\sigma \mid_{N'} \beta)$ denotes the strategy profile α where $\alpha_i = \sigma_i$ if $i \in N \setminus N'$ and $\alpha_i = \beta_i$ if $i \in N'$. Then $\bar{\sigma}$ is a strong equilibrium if for all $N' \subseteq N$, there is no $\beta \in S$ such that $u_i(\bar{\sigma} \mid_{N'} \beta) > u_i(\bar{\sigma})$ for all $i \in N'$. Andelman, Feldman and Mansour [2] proposed the notion of *strong price of anarchy* (SPoA) which is the PoA restricted to strong equilibria. The SPoA of the SST-game is the maximum value of $\mathcal{S}(\bar{\sigma})/\mathcal{S}(\sigma^*)$ over all instances.

3 Contribution and Related Work

Contribution. We first show that any Nash equilibrium of the min SST-game is a strong equilibrium (actually the proof also works for the max SST-game).

¹ The strategy profile induced by a minimum (resp. maximum) weight spanning tree is a social optimum for the min (resp. max) SST-game. A simple greedy algorithm can build a social optimum for the bottleneck SST-game.

Since $\text{PoA}=\text{SPoA}$ in this situation we only study the PoA of the min SST-game. We first observe that the PoA of the min SST-game is in general unbounded so we restrict ourselves to instances satisfying the triangle inequality. Under this standard assumption it is proved in Section 4 that $\text{PoA} \in \Theta(\log n)$ and $\text{PoA} \leq d^*$ where n is the number of players (i.e., $n = |N| = |V| - 1$) and d^* is the *depth* (maximum number of edges from a leaf to the root) of the tree induced by a social optimum. In addition we bound the PoA with respect to a modified social welfare which is motivated by energy consumption in wireless communications. The PoA of the max SST-game and the bottleneck SST-game are studied in Section 5 (the triangle inequality is not imposed anymore). For the former we prove that $\text{PoA} = 1/\Delta^*$ where Δ^* is the maximum degree of a player node (i.e. all nodes except r) in the tree induced by a social optimum. For the bottleneck SST-game we show that any Nash equilibrium is a social optimum ($\text{PoA} = 1$), implying that any Nash equilibrium is a strong equilibrium. Some concluding remarks are given in Section 6. Due to space limitations, some proofs and tight examples are omitted.

Related work. In 1973 Claus & Kleitman [6] introduced the problem of allocating the cost of a spanning tree. Subsequently Bird [5] and Granot & Huberman [9,10] studied the problem as a cooperative game: the players must agree on the structure (the tree) and how its cost is shared. Bird's cost allocation [5] consists in computing a minimum cost spanning tree and each player-node pays the price of the first edge of her path to the source. This allocation belongs to the core [2] of the game; further results on the core are provided by Granot & Huberman [9,10]. Then the min SST-game (resp. max SST-game) is a non cooperative version of Claus & Kleitman's problem with a Bird like allocation rule.

A central network problem viewed as a strategic game is the one introduced by Anshelevich *et al.* [3] (subsequently studied by Albers [1]). We are given an edge-weighted graph $G = (V, E)$ and pairs $(s_i, t_i) \in V \times V$. Each s_i wants to connect to t_i so the strategies of i are the paths from s_i to t_i in G . The weight $w(e)$ of an edge e represents a cost that is evenly shared among its users. Namely e costs $w(e)/\nu(e)$ to each of its users where $\nu(e)$ denotes the number of e 's users (this cost sharing method is called the *shapley cost* in [3,1]). The total cost of a player is the sum of the prices she is charged over all edges in her path. Taking the social cost as the sum of all player's individual cost, Anshelevich *et al.* show that there always exists a Nash equilibrium with social cost at most $O(\log |V|)$ times the social optimum. In the SST-game all t_i 's are the same vertex but the main difference resides in the way an agent's utility is defined. In contrast with the shapley cost where a player pays a fraction of each edge of her path, a player's utility in the min SST-game depends only on the first edge. Actually the three variants of the SST-game studied in this paper model situations where the agents' utility is ruled by the topology since quantities like electric power, time (to repair a message) or bandwidth are, unlike money, hard to transfer.

² Roughly speaking, the core is the set of all allocations such that no group of players is "mistreated".

In [8] Fabrikant *et al.* study the price of anarchy of a game where each player is the node of a graph. A player’s strategy set consists of all subsets of edges incident to her. Each player’s cost is the number of edges times a positive constant plus the sum of the distances from the node to all others.

The min SST-game is strongly related to two topology control problems (the *Connectivity Game* and the *Strong Connectivity Game*) in static ad hoc networks studied by Eidenbenz *et al.* [7]. The input of these problems is a graph $H = (V, E)$ where each edge e has a weight $w(e) \geq 0$ (w satisfies the triangle inequality).

In the *Connectivity Game* [7] we are given pairs $(s_i, t_i) \in V \times V$ such that s_i needs to connect to t_i (possibly over several intermediate vertices). Each s_i has to choose a *radius*. The *radius function* ρ is a mapping from V to \mathbb{R}_+ ($\rho_v \geq 0$ denotes the radius of v). It induces a directed graph $G_\rho = (V, A)$ where $A = \{(u, v) : [u, v] \in E \text{ and } \rho_u \geq w([u, v])\}$. G_ρ is called the *transmission graph* [13], i.e. s connects to t if there is a directed path from s to t in G_ρ . The utility of s_i is defined as $-M$ if s_i does not connect to t_i in G_ρ (M being some very large number), otherwise it is $-(\rho_{s_i})^\alpha$. Here $\alpha \geq 1$ is a constant known as the *power gradient* [13]. So, each s_i is a selfish agent whose best strategy is the minimum radius that connects her to t_i . The social cost is defined as $\sum_{v \in V} (\rho_v)^\alpha$. Deciding whether an instance of the Connectivity Game has a pure NE is an NP-complete problem [7]. The *Strong Connectivity Game* [7] is a special case of the Connectivity Game where each vertex needs to connect with every other vertex. It always has a pure Nash equilibrium and its PoA is $\Theta(n^\alpha)$.

The min SST-game is a particular case of the Connectivity Game where $t_i = r$ for all $i = 1, \dots, k$, $\{s_i : i = 1, \dots, k\} = V \setminus \{r\}$ and $\alpha = 1$. It can be also viewed as a relaxation of the Strong Connectivity Game where each vertex needs to connect to r (or if we consider the transmission graph as an unoriented graph, each vertex is connected to every other vertex).

4 The Min SST-Game

Strong-Nash equivalence. A strong equilibrium is a Nash equilibrium but a Nash equilibrium is not necessarily strong (e.g. the prisoner’s dilemma). We show that the min SST-game has a particular structure since any Nash equilibrium is strong. This property is known as the strong-Nash equivalence [11].

Theorem 1. *Any Nash equilibrium of the min SST-game is a strong equilibrium.*

Proof. Given $i \in N$ and a Nash equilibrium σ , let $DEP(i, \sigma)$ be the players who depend on i to be connected to r : $DEP(i, \sigma) = \{j \in N \setminus \{i\} : u_i(\sigma \mid_i j) = -\infty\}$. Suppose that σ is a Nash equilibrium but not a strong equilibrium. Then there exists a coalition $N' \subseteq N$ and a strategy profile β such that $u_i(\sigma \mid_{N'} \beta) > u_i(\sigma)$ for all $i \in N'$. W.l.o.g., we suppose that $(\sigma \mid_{N'} \beta) = \beta$.

If there exists $i' \in N'$ such that $\beta_{i'} \notin DEP(i', \sigma)$ then $u_{i'}(\sigma \mid_{i'} \beta_{i'}) = u_{i'}(\sigma \mid_{N'} \beta) > u_{i'}(\sigma)$. Since i' can unilaterally change her strategy and improve her utility, σ is not a Nash equilibrium, contradiction.

If $\beta_i \in DEP(i, \sigma)$ for all $i \in N'$ then $(\sigma \upharpoonright_{N'} \beta)$ does not connect any $i \in N'$ to r . To see this, let $L' = \bigcup_{i \in N'} DEP(i, \sigma) \setminus N'$ and $U' = N \setminus (N' \cup L')$. Actually L', N' and U' form a partition of N . The strategy of a node in $L' \cup N'$ is to select a node in $L' \cup N'$. The strategy of a node in U' is to select a node in $U' \cup \{r\}$. Therefore $u_i(\sigma \upharpoonright_{N'} \beta) = -\infty < u_i(\sigma)$ for all $i \in N'$, contradiction. \square

Then we only consider the price of anarchy of the min SST-game but our results also hold for the strong price of anarchy.

Definitions and properties. The utility of each player (to be maximized) is the negative of a cost (to be minimized). Given $i \in N$ and $\sigma \in S$, the cost of i is denoted by $c_i(\sigma)$ and defined as $-u_i(\sigma)$. For the sake of convenience, we often manipulate the cost of a player (which is a non negative value) instead of her utility (which is a non positive value).

We introduce several definitions which will be useful in the next proofs and give two basic properties. In the following, σ denotes a strategy profile at Nash equilibrium. Meanwhile σ^* denotes a social optimum, that is $\sum_{i \in N} c_i(\sigma^*)$ is minimum.

Definition 1. Given $(i, j) \in V \times V$, let $path^*(i, j)$ be the path (i.e. set of edges) between i and j in T^{σ^*} .

Let $weight^*(i, j) = \sum_{e \in path^*(i, j)} w(e)$ and $visited^*(i, j) \subseteq V$ be all nodes of $path^*(i, j)$.

Let $root^*(i, j)$ be the vertex of $visited^*(i, j)$ which plays the role of the source, i.e. $root^*(i, j) = r$ if $r \in visited^*(i, j)$, otherwise $root^*(i, j) = \ell$ where $\ell \in visited^*(i, j)$ and $\sigma_\ell^* \notin visited^*(i, j)$.

Let $path, weight, visited$ and $root$ be defined similarly w.r.t. the tree induced by σ .

Property 1. If σ is a Nash equilibrium then $c_i(\sigma) \leq w([i, r])$ holds for all $i \in N$.

Property 2. If σ is a Nash equilibrium then $\min\{c_i(\sigma), c_j(\sigma)\} \leq w([i, j])$ holds for all $(i, j) \in N \times N$ such that $i \neq j$.

The PoA according to the number of players. We start with a disappointing observation: the players of the min SST-game can arbitrarily deviate from the social optimum. Consider an instance with two players 1 and 2. The weight function is defined as $w([1, 2]) = w([1, r]) = 1$ and $w([2, r]) = X$ where $X > 1$. The strategy profile σ^* where $\sigma_1^* = r$ and $\sigma_2^* = 1$ is a social optimum. The strategy profile σ where $\sigma_1 = 2$ and $\sigma_2 = r$ is a Nash equilibrium since 1 has no incentive to change her strategy. Therefore the price of anarchy $\mathcal{S}(\sigma)/\mathcal{S}(\sigma^*) = (1 + X)/2$ tends towards ∞ when X tends towards ∞ . This instance does not satisfy the standard hypothesis which says that the cost of a link depends on its length. From now on, we assume that weights satisfy the triangle inequality, i.e. $w([x, y]) + w([y, z]) \geq w([x, z])$ for all triple of nodes x, y and z .

Let us give a lower bound on the PoA of the min SST-game.

Proposition 1. *When the triangle inequality holds, the PoA of the min SST-game is at least $1 + (\log n)/2$.*

Proof. We consider an instance of the game with 2^k players for all positive integer k . Players are numbered from 1 to 2^k . We assume that the source node r has number 0. The weight function is defined as $w([i, j]) = |i - j|$ for all pair of nodes (i, j) . It is not difficult to see that the triangle inequality holds. The minimum weighted tree (an optimal solution), denoted by T_k^* , consists of all edges $[i, i + 1]$ ($i = 0, \dots, 2^k - 1$). Its total weight is 2^k .

We are going to describe a spanning tree T_k such that the strategy profile induced by T_k is a Nash equilibrium for the min SST-game. T_k consists of the edges $[2^{k-1}, 2^k]$, $[0, 2^k]$ and a subtree \mathcal{T}_k which is a complete binary tree on vertices $1, \dots, 2^k - 1$. To construct \mathcal{T}_k , we start from a complete binary tree on $2^k - 1$ vertices which are subsequently numbered.

Each leaf of \mathcal{T}_k receives an odd integer between 1 and $2^k - 1$. The numbering follows the DFS order. Next the number of a non-leaf vertex is defined as the average number of its two children. For instance the root of \mathcal{T}_k has number 2^{k-1} and it is denoted by R_k in the following. It suffices to make the union of $\{[2^{k-1}, 2^k], [0, 2^k]\}$ and \mathcal{T}_k to get T_k .

Given $i \in \{1, \dots, k\}$, let N_i be the set of vertices of \mathcal{T}_k whose distance (i.e. number of edges) to R_k in \mathcal{T}_k is $i - 1$. In particular $N_1 = \{R_k\}$, N_k contains all leaves of \mathcal{T}_k and $|N_i| = 2^{i-1}$. The following properties can be easily shown by induction:

- (i) The numbering of the vertices of N_i is $\{2^{k-i}(2j - 1) : j = 1, \dots, 2^{i-1}\}$.
- (ii) Denote by $p(v)$ the father of v in \mathcal{T}_k . If $v \in N_i$ with $i \in \{2, \dots, k\}$ then $w([v, p(v)]) = 2^{k-i}$.
- (iii) Given a tree T and a node v of T , we denote by $Sub(T, v)$ the subtree of T rooted at v . If $v_0 \in N_{i_0}$ with $i_0 \neq k$ has the number $2^{k-i_0}(2j_0 - 1)$ for some $j_0 \in \{1, \dots, 2^{i_0-1}\}$ then the number of the vertices of $Sub(\mathcal{T}_k, v_0)$ are the integers between $2^{k-i_0+1}(j_0 - 1) + 1$ and $2^{k-i_0+1}j_0 - 1$, that is $\{2^{k-i_0+1}(j_0 - 1) + 1, \dots, 2^{k-i_0+1}j_0 - 1\}$.

Let us prove that the strategy profile induced by T_k is a Nash equilibrium. By construction a vertex v of \mathcal{T}_k can be linked to any vertex except those of $Sub(\mathcal{T}_k, v)$ (otherwise v is not connected to r anymore). So, assume that $v \in N_{i_0}$ has number $2^{k-i_0}(2j_0 - 1)$ for some $i_0 \in \{1, \dots, k\}$ and $j_0 \in \{1, \dots, 2^{i_0-1}\}$. By Property (iii), we know that the two vertices v_0 and v_1 closest to v are numbered $2^{k-i_0+1}(j_0 - 1)$ and $2^{k-i_0+1}j_0$ respectively. Thus, we deduce that $w([v, v_0]) = |2^{k-i_0}(2j_0 - 1) - 2^{k-i_0+1}(j_0 - 1)| = 2^{k-i_0}$ and $w([v, v_1]) = |2^{k-i_0}(2j_0 - 1) - 2^{k-i_0+1}j_0| = 2^{k-i_0}$. By Property (ii), these distances are equals to $w([v, p(v)])$. Finally, by observing that the distance of R_k (numbered 2^{k-1}) to r is equal to the distance of R_k to 2^k , we conclude that T_k is a Nash equilibrium. Now, using Properties (i) and (ii), we deduce that the total weight of T_k is $2^k + 2^{k-1} + \sum_{i=1}^k |N_i|2^{k-i} = 2^k + 2^{k-1} + (k - 1)2^{k-1} = (k + 2)2^{k-1}$. Then the price of anarchy of the min SST-game is at least $((k + 2)2^{k-1})/2^k = (k + 2)/2 = 1 + \frac{\log n}{2}$ since $n = 2^k$. \square

Before giving an upper bound on the PoA of the min SST-game we prove an intermediate result which is useful in the subsequent proof.

Lemma 1. *Given an oriented tree $T = (V, E)$ rooted at r on $n + 1$ vertices, and for any subset $V' = \{v_1, \dots, v_{2p+1}\} \subseteq V \setminus \{r\}$ of $2p + 1$ vertices (hence, $2p + 1 \leq n$), one can exhibit $p + 1$ edge disjoint unoriented paths $path_i$ for $i = 1, \dots, p + 1$ in T such that one path, say $path_{p+1}$, is a path from v_j to r for $j \leq 2p + 1$ whereas the endpoints set of the p other paths is exactly $V' \setminus \{v_j\}$.*

Let us return to the PoA of the min SST-game.

Theorem 2. *When the triangle inequality holds, the PoA of the min SST-game is at most $\lfloor \log n \rfloor + 1 + \frac{2}{n}$.*

Proof. Let σ be a Nash equilibrium while σ^* denotes a social optimum. We suppose that the players are sorted by decreasing cost, i.e. $c_1(\sigma) \geq \dots \geq c_n(\sigma)$ where $n = |N|$. In order to keep simple, we assume that n is odd. We mainly prove the following property:

$$\forall p = 0, \dots, \frac{n-1}{2}, \quad \sum_{i=p+1}^{2p+1} c_i(\sigma) \leq \sum_{i=1}^n c_i(\sigma^*) \tag{1}$$

Suppose that (I) holds and replace p by $2^j - 1$. One has $\sum_{i=2^j}^{2^{j+1}-1} c_i(\sigma) \leq \sum_{i=1}^n c_i(\sigma^*)$ and the result follows:

$$\sum_{i=1}^n c_i(\sigma) \leq \sum_{j=0}^{\lfloor \log n \rfloor - 1} \sum_{i=2^j}^{2^{j+1}-1} c_i(\sigma) + \sum_{i=\frac{n+1}{2}}^n c_i(\sigma) \leq (\lfloor \log n \rfloor + 1) \sum_{i=1}^n c_i(\sigma^*)$$

If n is even then we conduct the same proof on the $n - 1$ first players and we use the fact that $\frac{n}{2}c_n(\sigma) \leq \sum_{i=\frac{n}{2}}^{n-1} c_i(\sigma) \leq \sum_{i=1}^n c_i(\sigma^*)$. Then it suffices to prove inequality (II).

Given $p \in \{0, \dots, \frac{n-1}{2}\}$, let V_p be the $2p + 1$ players with largest costs, i.e. $V_p = \{1, \dots, 2p + 1\}$. Let T^* be the tree induced by σ^* . The total weight of T^* is denoted by OPT ; it is equal to $\sum_{i=1}^n c_i(\sigma^*)$. We know from Lemma I with input V_p that one can find $p + 1$ edge-disjoint paths in T^* . W.l.o.g., assume that these paths are $path^*(\pi(2i - 1), \pi(2i))$ for $i = 1, \dots, p$ and $path^*(\pi(2p + 1), r)$ where π is a permutation of V_p . We deduce:

$$\sum_{i=1}^{p+1} weight^*(\pi(2i - 1), \pi(2i)) \leq OPT. \tag{2}$$

By the triangle inequality, we also get:

$$w([\pi(2i - 1), \pi(2i)]) \leq weight^*(\pi(2i - 1), \pi(2i)) \tag{3}$$

for $i = 1, \dots, p$ and

$$w([\pi(2p + 1), r]) \leq weight^*(\pi(2p + 1), r). \tag{4}$$

Using Property 2 we have:

$$\min\{c_{\pi(2i-1)}(\sigma), c_{\pi(2i)}(\sigma)\} \leq w([\pi(2i-1), \pi(2i)]) \text{ for } i = 1, \dots, p \quad (5)$$

Using Property 1, the cost of player $\pi(2p+1)$ is less than $w([\pi(2p+1), r])$ if she is connected to r via another player in σ . Therefore

$$c_{\pi(2p+1)}(\sigma) \leq w([\pi(2p+1), r]) \quad (6)$$

Since players are numbered by decreasing cost, we have

$$\sum_{i=p+1}^{2p+1} c_i(\sigma) \leq \sum_{i=1}^p \min\{c_{\pi(2i-1)}(\sigma), c_{\pi(2i)}(\sigma)\} + c_{\pi(2p+1)}(\sigma). \quad (7)$$

Summing up inequalities (5) for $i = 1, \dots, p$ and adding inequality (6) we get:

$$\sum_{i=1}^p \min\{c_{\pi(2i-1)}(\sigma), c_{\pi(2i)}(\sigma)\} + c_{\pi(2p+1)}(\sigma) \leq \sum_{i=1}^p w([\pi(2i-1), \pi(2i)]) + w([\pi(2p+1), r]) \quad (8)$$

Finally, using inequalities (7), (8), (2), (3) and (4), we obtain inequality (11), that is $\sum_{i=p+1}^{2p+1} c_i(\sigma) \leq OPT$. \square

Therefore the PoA of the min SST-game belongs to $\Theta(\log n)$. One can generalize the min SST-game to the case where the social cost is defined as $\sum_{i=1}^n (c_i(\sigma))^\alpha$ (problem denoted by min SST-game (α) in the following). Here α is a constant greater than 1 such that d^α is the minimum power to transmit a message at distance d [13].

Theorem 3. *For any $\alpha \geq 1$, the PoA of the min SST-game (α) is upper bounded by $n^{\alpha-1}(\lceil \log n \rceil + 1 + \frac{2}{n})$.*

The PoA according to the depth of an optimum tree. We propose an alternative bound on the PoA of the min SST-game. We need some notations and two intermediate results to show the result.

Let \mathcal{T} be an oriented tree rooted at r in the following definitions. Given $i \in N$, let \mathcal{T}_i be the subtree of \mathcal{T} rooted at i . Let $V(\mathcal{T})$ be the set of vertices spanned by \mathcal{T} . The weight of \mathcal{T} , denoted by $w(\mathcal{T})$, is defined as $\sum_{e \in \mathcal{T}} w(e)$. The depth of \mathcal{T} , denoted by $depth(\mathcal{T})$, is defined as $\max_{i \in N} \{|path_{\mathcal{T}}(i, r)|\}$ where $path_{\mathcal{T}}(i, r)$ is the path (i.e. set of edges) between i and r in \mathcal{T} . The level of $i \in N$ in \mathcal{T} , denoted by $\ell(\mathcal{T}, i)$, is defined as $|path_{\mathcal{T}}(i, r)|$. Given $N' \subseteq N$, let $f(\mathcal{T}, N')$ be a node of N' with minimum level in \mathcal{T} , i.e. $f(\mathcal{T}, N') = argmin_{i \in N'} \{\ell(\mathcal{T}, i)\}$.

Lemma 2. *We are given a spanning tree \mathcal{T} of $G = (V, E)$ rooted at r , a spanning tree T induced by a Nash equilibrium σ , a player $i \in N$ and $H \subseteq V(\mathcal{T}_i) \setminus \{i\}$ such that $\forall a, b \in H, a \notin \mathcal{T}_b$ and $b \notin \mathcal{T}_a$. Let i^* be the element of $H \cup \{i\}$ with minimum level in T , i.e. $i^* = f(T, H \cup \{i\})$. If the triangle inequality holds then*

$$\left(\sum_{j \in H \cup \{i\}} c_j(\sigma) \right) - c_{i^*}(\sigma) \leq 2 \left(\sum_{j \in H} w([i, j]) \right) - w([i, i^*]).$$

Proposition 2. *We are given a spanning tree \mathcal{T} of $G = (V, E)$ rooted at r , a spanning tree T induced by a Nash equilibrium σ and $i \in N$ (i.e., $i \neq r$). Let $i^* = f(T, V(\mathcal{T}_i))$. If the triangle inequality holds then*

$$\sum_{j \in V(\mathcal{T}_i)} c_j(\sigma) - c_{i^*}(\sigma) \leq (\text{depth}(\mathcal{T}_i) + 1)w(\mathcal{T}_i) - w([i, i^*]).$$

Now we are able to prove that the PoA is upper bounded by the depth of a minimum weight spanning tree rooted at r (denoted by d^* in the following). For instance the case $d^* = 1$ is a direct consequence of Property \square .

Theorem 4. *When the triangle inequality holds, the PoA of the min SST-game is at most d^* .*

Proof. Let T^* be a minimum weight spanning tree rooted at r with depth d^* . T^* is induced by a social optimum σ^* . Let T be the tree induced by a Nash equilibrium σ . Denote by $1, \dots, k$ the sons of r in T^* , i.e. $\{i \in N : [i, r] \in T^*\} = \{1, \dots, k\}$. Proposition \square for any $i \in \{1, \dots, k\}$ gives

$$\sum_{j \in V(T_i^*)} c_j(\sigma) \leq (\text{depth}(T_i^*) + 1)w(T_i^*) - w([i, i^*]) + c_{i^*}(\sigma) \tag{9}$$

where $i^* = f(T, V(T_i^*))$. Using Property \square and the triangle inequality we obtain $c_{i^*}(\sigma) \leq w([r, i^*]) \leq w([r, i]) + w([i, i^*])$. Then inequality (9) becomes:

$$\sum_{j \in V(T_i^*)} c_j(\sigma) \leq (\text{depth}(T_i^*) + 1)(w(T_i^*) + w([r, i])) \tag{10}$$

Using $\text{depth}(T_i^*) + 1 \leq \text{depth}(T^*)$ and summing up inequality (10) for $i = 1$ to k we get:

$$\begin{aligned} \sum_{i=1}^k \sum_{j \in V(T_i^*)} c_j(\sigma) &\leq \text{depth}(T^*) \sum_{i=1}^k (w(T_i^*) + w([r, i])) \\ \sum_{j \in N} c_j(\sigma) &\leq \text{depth}(T^*)w(T^*) = \text{depth}(T^*) \sum_{j \in N} c_j(\sigma^*). \quad \square \end{aligned}$$

5 The Max SST-Game and the Bottleneck SST-Game

In contrast with the previous section we do not suppose that the triangle inequality holds. Let us begin with the max SST-game. It is noteworthy that the proof of Theorem \square also works for the max SST-game. Then any Nash equilibrium of the max SST-game is strong. Again we only study the PoA since $\text{PoA} = \text{SPoA}$.

Theorem 5. *The PoA of the max SST-game is $1/\Delta^*$ where Δ^* is the maximum degree of a player node in a maximum cost spanning tree.*

Proof. We know that $w([i, j]) \leq \max\{u_i(\sigma), u_j(\sigma)\}$ for any pair of players (i, j) (the property is similar to Property **2**). Moreover $w([i, r]) \leq u_i(\sigma)$ holds for all $i \in N$ (the property is similar to Property **1**). We use these inequalities for each edge of the optimal tree as follows.

$$\begin{aligned} \mathcal{S}(\sigma^*) &:= \sum_{\{i \in N : \sigma_i^* \neq r\}} w([i, \sigma_i^*]) + \sum_{\{i \in N : \sigma_i^* = r\}} w([i, r]) \\ \mathcal{S}(\sigma^*) &\leq \sum_{\{i \in N : \sigma_i^* \neq r\}} \max\{u_i(\sigma), u_{\sigma_i^*}(\sigma)\} + \sum_{\{i \in N : \sigma_i^* = r\}} u_i(\sigma) \end{aligned}$$

A player’s utility may appear several times in the upper bound but at most Δ^* times. Hence, $\mathcal{S}(\sigma^*) \leq \Delta^* \sum_{i \in N} u_i(\sigma) = \Delta^* \mathcal{S}(\sigma)$ and $\text{PoA} = \mathcal{S}(\sigma) / \mathcal{S}(\sigma^*) \geq 1 / \Delta^*$. Tight examples (with and without the triangle inequality) exist. \square

Theorem 6. *The PoA of the bottleneck SST-game is 1.*

Proof. Let σ^* (resp. σ) be a strategy profile which maximizes the social welfare (resp. a Nash equilibrium). Let $B = \{i \in N : u_i(\sigma) < u_i(\sigma^*)\}$ be the players whose utility is suboptimal. Denote by $h(i)$ the number of edges which compose the unique path between node i and the source r along the tree induced by σ^* . Let $i^* = \text{argmin}\{h(i) : i \in B\}$ (ties are arbitrarily broken).

If $h(i^*) = 1$ then $u_{i^*}(\sigma \mid_{i^*} r) = u_{i^*}(\sigma^*)$. However $u_{i^*}(\sigma) \geq u_{i^*}(\sigma \mid_{i^*} r)$ always holds because σ is a Nash equilibrium. We get $u_{i^*}(\sigma) \geq u_{i^*}(\sigma^*)$ which is in contradiction with $i^* \in B$. Therefore $h(i^*) > 1$. Let $j^* \in N$ be such that $\sigma_{i^*}^* = j^*$. By definition, we have

$$u_{i^*}(\sigma^*) = \min\{u_{j^*}(\sigma^*), w([i^*, j^*])\} \tag{11}$$

We know that $j^* \notin B$ because $h(j^*) < h(i^*)$. We deduce

$$u_{j^*}(\sigma^*) \leq u_{j^*}(\sigma) \tag{12}$$

Moreover $u_{i^*}(\sigma \mid_{i^*} j^*) = -\infty$ since by **(11)** and **(12)** we have

$$u_{i^*}(\sigma^*) \leq \min\{u_{j^*}(\sigma), w([i^*, j^*])\} = u_{i^*}(\sigma \mid_{i^*} j^*).$$

Now, since $i^* \in B$, we deduce $u_{i^*}(\sigma \mid_{i^*} j^*) \geq u_{i^*}(\sigma^*) > u_{i^*}(\sigma)$ would contradict the fact that σ is a Nash equilibrium. Then $i^* \in \text{visited}(j^*, r)$ implies

$$u_{j^*}(\sigma) \leq u_{i^*}(\sigma) \tag{13}$$

Using inequalities **(11)**, **(12)** and **(13)**, we get $u_{i^*}(\sigma^*) \leq u_{i^*}(\sigma)$ which is in contradiction with $i^* \in B$. Then B must be empty and $u_i(\sigma) \geq u_i(\sigma^*)$ holds for all $i \in N$. The result follows since $\mathcal{S}(\sigma^*) \geq \mathcal{S}(\sigma) = \sum_{i \in N} u_i(\sigma) \geq \sum_{i \in N} u_i(\sigma^*) = \mathcal{S}(\sigma^*)$. \square

Corollary 1. *The SPoA of the bottleneck SST-game is 1 and any Nash equilibrium is a strong equilibrium.*

6 Concluding Remarks

We prove that a Nash equilibrium is a strong equilibrium for all games under consideration. Then, is cooperation inefficient? On the one hand cooperating with a player who depends on you to be connected (i.e. you are on her path to r) can be profitable. On the other hand cooperating with a player who does not depend on you to be connected (i.e. you are not on her path to r) is never profitable. Since two players cannot mutually depend on the other to connect to r , cooperation is unlikely.

We study three variants of a strategic spanning tree game but the focus is on the min SST-game. The gap between the lower bound $(1 + (\log n)/2)$ and the upper bound $(\lceil \log n \rceil + 1 + 2/n)$ is quite narrow but it is certainly possible to give a better analysis of the upper bound. The same goes for Theorem 3 (when the social cost of the min SST-game is $\sum_{i \in N} (c_i(\sigma))^\alpha$), we believe that a more accurate upper bound can be derived. We were not able to find tight examples for Theorem 4 (except when $d^* = 2$) and believe that the upper bound d^* can be improved.

References

1. Albers, S.: On the value of coordination in network design. In: Proc. of SODA 2008, pp. 294–303 (2008)
2. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: Proc. of SODA 2007, pp. 189–198 (2007)
3. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The Price of Stability for Network Design with Fair Cost Allocation. In: Proc. of FOCS 2004, pp. 295–304 (2004)
4. Aumann, R.: Acceptable points in general cooperative n-person games. In: Contribution to the Theory of Games. Annals of Mathematics Studies 40, vol. IV, pp. 287–324 (1959)
5. Bird, C.G.: On cost allocation for a spanning tree: A game theory approach. Networks 6, 335–350 (1976)
6. Claus, A., Kleitman, D.: Cost allocation for a spanning tree. Networks 3, 289–304 (1973)
7. Eidenbenz, S., Kumar, V.S.A., Zust, S.: Equilibria in topology control games for ad hoc networks. In: Proc. of DIALM-POMC, pp. 2–11 (2003)
8. Fabrikant, A., Luthra, A., Maneva, E.N., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: Proc. of PODC 2003, pp. 347–351 (2003)
9. Granot, D., Huberman, G.: Minimum cost spanning tree games. Mathematical Programming 21, 1–18 (1981)
10. Granot, D., Huberman, G.: On the core and nucleolus of minimum cost spanning tree games. Mathematical Programming 29, 323–347 (1984)
11. Holzman, R., Law-Yone, N.: Strong Equilibrium in Congestion Games. Games and economic behavior 21, 85–101 (1997)
12. Koutsoupias, E., Papadimitriou, C.H.: Worst Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
13. Rajaraman, R.: Topology control and routing in ad hoc networks: a survey. SIGACT News 33, 60–73 (2002)

Stochastic Submodular Maximization

Arash Asadpour, Hamid Nazerzadeh, and Amin Saberi

Stanford University, Stanford, CA
{asadpour,hamidnz,saberi}@stanford.edu

Abstract. We study stochastic submodular maximization problem with respect to a cardinality constraint. Our model can capture the effect of uncertainty in different problems, such as cascade effects in social networks, capital budgeting, sensor placement, etc. We study non-adaptive and adaptive policies and give optimal constant approximation algorithms for both cases. We also bound the adaptivity gap of the problem between 1.21 and 1.59.

1 Introduction

The problem of maximizing submodular functions with respect to known constraints is a very well-studied problem in operations research and computer science. A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is submodular if for all $x, y \in \mathbb{R}^n$:

$$f(x \vee y) + f(x \wedge y) \leq f(x) + f(y)$$

where $x \vee y$ and $x \wedge y$ denote the component-wise maximum and the component-wise minimum of x and y , respectively. If f is twice differentiable, then submodularity is equivalent to the condition $\frac{\partial^2 f}{\partial x_i \partial x_j} \leq 0$, where x_i and x_j are any two coordinates of x [17]. One may imagine the function f on the domain of 0/1-vectors as a set function where $f(S) = f(x)$, $x_i = 1$ whenever $i \in S$, and $x_i = 0$ otherwise. In other words, a set function $f : 2^N \rightarrow \mathbb{R}$ is submodular if for any two subsets $S, T \subseteq N$:

$$f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$$

A wide range of optimization problems that arise in the real world can be modeled as maximizing submodular functions with respect to some (usually cardinality) constraints. One instance is the problem of viral marketing and maximizing influence through the network [9,14], where the goal is to choose an initial “active” set of people, so as to maximize the spread of an innovation or behavior in a social network. It is well-known that under many models of influence propagation in networks (e.g. decreasing cascade model [9]), the expected size of the final cascade is a submodular function of the set of initially activated people. Also, due to some budget limitations the number of people that we can activate in the beginning is bounded. Hence, the maximizing influence problem can be seen as a maximizing submodular function problem subject to some cardinality constraint.

Another example is the capital budgeting problem which is to find the optimal investment of capital among different projects with a limited budget. There is a set of projects and one wants to invest on a group of them that maximizes his expected profit while not exceeding his budget. This problem has been studied extensively under various assumptions on the utility of the investor and dependencies among the projects [20,21,13,2]. Naturally, the utility functions are non-negative and monotone. Also, the risk-averse investors are characterized by their submodular utility functions. Therefore, such investors need to solve a submodular maximization problem to find their best bet.

Yet another example is the problem of optimal placement of sensors for environmental monitoring [11,12] where the objective is to place a limited number of sensors in the environment in order to most effectively reduce uncertainty in observations. It is known that the efficiency of a subset of sensors is a submodular set function.

For the problem of maximizing submodular set functions subject to a cardinality constraint, the celebrated result of Cornuejols et al. [4] and Nemhauser et al. [15] shows that for nonnegative monotone submodular functions the greedy algorithm that at each step chooses an element with the maximum marginal value gives a $(1 - \frac{1}{e} \approx 0.632)$ -approximation of the optimal solution. This problem has also been studied for more complicated domains. In particular, for maximizing a submodular function over a matroid (note that the cardinality constraint is a special case of this) a recent result by Vondrak [19] shows that it is still possible to get a $1 - \frac{1}{e}$ -approximation.

However, in practice one must deal with the stochasticity caused by the uncertain nature of the problem, the incomplete information about the environment, etc. For instance, in viral marketing some people in the initial set might not adopt the behavior. Another example is the capital budgeting problem where some projects taken by an investor may fail (in the beginning) due to some unexpected events in the market. Also, in environmental monitoring some sensors might not work properly because of bad weather or inconsistent connections.

All these possibilities motivate the problem of stochastic submodular maximization. In the stochastic setting, the outcome of the elements in the selected set are not known in advance; when they are picked, with some known probability they might remain in the set or not. One may think of probability p_i corresponding to each element i , and then the expected value of selecting a set S will be the expected value of function f over set \hat{S} derived from S by removing each element i independently with probability $1 - p_i$. In fact, in this paper we will consider a more general version of this problem in which the stochasticity of the problem converts the set S into a vector in the continuous space via some known probability distributions. For the exact definitions see Section 2.

The main difference between the non-stochastic and stochastic problems is that the latter can benefit from adaptivity. An *adaptive* policy can use the outcome of the steps taken so far to optimize the decisions it is going to make in the future. On the other hand, the actions chosen by a *non-adaptive* policy are independent of the outcome of the other actions. Therefore, a non-adaptive policy

is equivalent to a predetermined subset of elements. Although non-adaptive policies may not perform as well as adaptive ones, they are particularly useful when it is difficult or time-consuming to discover the outcome of an action.

Our Results: In Section 3 we first show that the expected value of a submodular set function in the stochastic setting is still a submodular set function. This immediately leads to a $(1 - \frac{1}{e})$ -approximation of the optimal non-adaptive policy. Then, we consider the adaptive policy that at each step chooses an element with the maximum *expected* marginal value, conditioned on the outcome of the previous elements. We show that the approximation ratio of this greedy policy with respect to the optimal adaptive policy is $1 - \frac{1}{e}$ 4. We also give a lower bound and an upper bound on the adaptivity gap of the problem. The *adaptivity gap* is defined as the maximum ratio between the expected value of the optimal adaptive and non-adaptive policies 5. As a lower bound, we prove that the adaptivity gap of stochastic submodular maximization problem is at least 1.21 (see Section 2.1). On the flip side, in Section 4 we show that the adaptivity gap is bounded from above by $\frac{e}{e-1} \approx 1.59$, i.e. there exists a non-adaptive policy which achieves at least $\frac{1}{1.59}$ fraction of the value of best adaptive policy. We also show that a non-adaptive policy within a factor of $(\frac{e-1}{e})^2 \approx \frac{1}{2.51}$ of the optimum adaptive policy can be found in polynomial time. In order to prove this bound, we generalize some of the techniques developed by Vondrák 18. These extensions could be of independent interest.

1.1 Related Work

We first briefly overview some parts of the literature on (non-stochastic) submodular optimization. Then, we explain some of the works that study stochastic settings similar to ours.

Cornuejols et al. 4 proved that a simple greedy algorithm gives a $(1 - \frac{1}{e})$ -approximation for the problem of maximizing monotone submodular set functions subject to capacity constraints. Later, Feige 6 proved that it is not possible to improve this ratio unless $NP \subset TIME(n^{O(\log \log n)})$. For maximizing non-monotone submodular, recently Feige et al. 7 gave a constant approximation algorithm. Another well-studied submodular maximization problem is the problem of allocating resources to agents with submodular utilities, for which several interesting approximation algorithms have been developed, see 18. In this paper, we use some of these techniques to bound the adaptivity gap.

Goemans and Vondrák 8 consider the problem of stochastic covering. In this problem the goal is to cover all elements of a target set using minimum number of subsets. The subsets are random variables and their probability distributions are given. They propose adaptive and non-adaptive policies for the problem. They also observe that the adaptivity gap is not constant.

Chan and Farias 3 study a generalization of the stochastic maximum k -cover problem where the sequence of elements arrive according to a stochastic process

¹ This is also independently observed by Chan and Farias 3.

and utility functions may vary over time. They show that under some conditions, a myopic policy is a 2-approximation of the optimal adaptive policy.

In a recent work, Streeter and Golovin [16] study an online job scheduling problem in a setting that the cost of jobs is given by a submodular function. The goal is to cover as many jobs as possible subject to a budget constraints. They take the regret minimization approach and present approximate optimal policy.

2 Problem Definition

We define the following abstraction for the stochastic submodular maximization problem. A set $\mathcal{A} = \{X_1, \dots, X_n\}$ of independent random variables is given. After choosing X_i , its actual value (outcome of an element), denoted by x_i , is discovered. We assume that $x_i \in [0, 1]$. Let $S \subseteq \mathcal{A}$ be a subset of variables. Also, let vector $s = \langle \hat{x}_1, \dots, \hat{x}_n \rangle$ denote the *realization* of set S , where $\hat{x}_i = x_i$ for $i \in S$ and $\hat{x}_i = 0$ for $i \notin S$. The value obtained by choosing the set S after the realization is equal to $f(s)$, where $f : [0, 1]^n \rightarrow \mathbb{R}^+$ is a submodular function.

Let g_i be the probability distribution of random variable X_i . For every subset $S \subseteq \mathcal{A}$, it defines a probability measure $g_S : [0, 1]^n \rightarrow \mathbb{R}$, which represents the probability density function of observing s while selecting S :

$$g_S(ds) = \int_{x \in ds} \prod_{i \in S} g_i(dx_i)$$

Also, $g_S(ds)$ is defined to be zero if there exist i such that $s_i \neq 0, i \notin S$. Now define function $F_g : [0, 1]^n \rightarrow \mathbb{R}^+$ as the expected value obtained by choosing set S , i.e.,

$$F_g(S) = \int_{s \in [0, 1]^n} f(s)g_S(ds) \tag{1}$$

Our goal is to choose a set S of size at most k which maximizes $F_g(S)$.

$$\max_{S \subseteq \mathcal{A}: |S| \leq k} F_g(S).$$

For simplicity, we assume one cannot choose an element of \mathcal{A} multiple times.² For this problem, we study two types of policies: adaptive and non-adaptive. A *non-adaptive policy* is represented by a fixed subset of \mathcal{A} . An *adaptive policy* uses the realized value of the previously chosen elements to determine the next element in the subset. In order to compare the value of these optimal policies, we study the adaptivity gap of the problem. The *adaptivity gap* is defined as the ratio between the expected value of optimal adaptive and non-adaptive policies.

² This assumption is not necessary for our results, and is made for sake of simplicity. One can create k independent copies of each random variable to simulate multiplicity. This is in contrast with the set cover problem where allowing to chose multiple copies of an element significantly reduces the adaptivity gap [8].

One issue that arises in the algorithmic discussions of this paper and many of the related works is computing the value of functions similar to F_g . We assume that we are given an oracle which computes these values up to a desired degree of accuracy. In fact, it can be shown that in many interesting cases such an oracle can be built efficiently. Two most important cases are when the probability distribution functions (g_j 's) are constant Lipschitz continuous, or when their support is a polynomial size set of discrete values. Therefore, from now on, all of our results will involve an arbitrary small error term of ϵ that we will not mention explicitly.

In the next section, we illustrate the problem by giving an example. We also present a non-adaptive and an adaptive policy for this example.

2.1 An Example: Stochastic Maximum k -Cover

A special case of stochastic submodular maximization is the stochastic maximum k -cover problem. Given a collection \mathcal{F} of subsets of $\{1, 2, \dots, n\}$, the max k -cover problem is defined as finding k subsets from \mathcal{F} such that their union has the maximum cardinality [6]. In the stochastic version, the subset that an element of \mathcal{F} would cover becomes known after choosing the element. In this section, we define an instance of this problem. We also use this example to give a lower bound on the adaptivity gap.

The instance we consider in this section is as follows: A ground set $G = \{1, 2, \dots, 2n\}$ and a collection $F = \{C_1, C_2, \dots, C_{2n}\}$ of its subsets are given. For $1 \leq i \leq n$, $C_i = \{1, 2, \dots, n\}$ with probability $\frac{1}{n}$ and is the empty set with probability $1 - \frac{1}{n}$. For $n + 1 \leq i \leq 2n$, $C_i = \{i\}$ with probability $\frac{1}{e}$ and is the empty set with the remaining probability. The goal is to cover the maximum number of elements in G by selecting at most n subsets in C .

Lemma 1. *For large enough values of n , the optimal non-adaptive policy is to select $S = \{C_1, C_2, \dots, C_n\}$. Also, the expected value of this policy is $n(1 - \frac{1}{e})$.*

Proof. Consider a subset S' selected by a non-adaptive policy. Let q the fraction of elements of S' that are in $\{C_1, C_2, \dots, C_n\}$, i.e., $q = |S' \cap \{C_1, C_2, \dots, C_n\}|/n$. Such a policy covers the elements of $\{1, 2, \dots, n\}$ with probability $1 - (1 - \frac{1}{n})^{nq}$. Also, in expectation, S' covers at most $\frac{n}{e}(1 - q)$ elements from set $\{n + 1, \dots, 2n\}$. Therefore, the expected number of covered elements is

$$n(1 - (1 - \frac{1}{n})^{nq}) + \frac{n}{e}(1 - q)$$

We can approximate the expression above by $n(1 + \frac{1}{e} - (\frac{1}{e})^q - \frac{q}{e})$ with arbitrary high precision for large enough n . This expression is increasing in q . Therefore, for the optimum non-adaptive policy we have $q = 1$ or equivalently $S = \{C_1, C_2, \dots, C_n\}$.

Now consider the following adaptive policy that at each step chooses the elements with maximum marginal value: At step i , $1 \leq i \leq n$, choose set C_i until one of

these sets covers $\{1, \dots, n\}$. After that, pick a set from $\{C_{n+1}, \dots, C_{2n}\}$ until the number of chosen sets reaches n .

The following lemma gives a lower bound on the number of elements covered by the adaptive policy.

Lemma 2. *For large enough n , the expected number of elements covered by the adaptive policy describe above is $n(1 - \frac{1}{e} + \frac{1}{e^2})$.*

Proof. The probability that C_i covers the first n elements is $\frac{1}{n}(1 - \frac{1}{n})^{i-1}$. If C_i covers the first n elements, the policy will choose $n-i$ subsets from C_{n+1}, \dots, C_{2n} , each covers a single element with probability $\frac{1}{e}$. Therefore, the expected number of covered elements is:

$$\begin{aligned} \sum_{i=1}^n [\frac{1}{n}(1 - \frac{1}{n})^{i-1} \times (n + (n - i)\frac{1}{e})] &= (1 + \frac{1}{e}) \sum_{i=1}^n (1 - \frac{1}{n})^{i-1} - \frac{1}{en} \sum_{i=1}^n i(1 - \frac{1}{n})^{i-1} \\ &\approx (1 + \frac{1}{e})n(1 - \frac{1}{e}) - \frac{1}{en} \left[n^2(1 - (1 - \frac{1}{n})^n) - n^2(1 - \frac{1}{n})^n \right] \\ &\approx n \left[(1 - \frac{1}{e^2}) - \frac{1}{e}(1 - \frac{2}{e}) \right] \\ &= n(1 - \frac{1}{e} + \frac{1}{e^2}). \end{aligned}$$

which completes the proof of the lemma.

By combining the results of Lemmas 1 and 2 we have the corollary below.

Corollary 3. *The adaptivity gap of stochastic maximum k -cover is at least:*

$$\frac{1 - e^{-1} + e^{-2}}{1 - e^{-1}} > 1.21$$

3 Near-Optimal Non-adaptive and Adaptive Policies

In this section we first present non-adaptive policy for the stochastic submodular maximization problem. Later, we give an adaptive policy. A non-adaptive policy is represented by a fixed subset $S \subseteq \mathcal{A}$. The expected value of the policy is equal to $F_g(S)$. Therefore, finding the optimal non-adaptive policy is equivalent to finding set S which maximizes $F_g(S)$. Note that the maximum k -cover problem is a special case of our problem. Therefore, it is not possible to find an approximation ratio better than $1 - \frac{1}{e}$ for the optimal adaptive policy unless $NP \subset TIME(n^{O(\log \log n)})$ [6]. In this section, we show that there exists a policy that is implementable in polynomial time and its value is within a $1 - \frac{1}{e}$ ratio of the optimal non-adaptive policy. For the ease of notation, when it is clear from the context, we use $F(S)$ instead of $F_g(S)$. Note that $F(S)$ is a convex combination of a set of monotone submodular functions. Therefore, we have the following lemma.

Lemma 4. *The function $F(S)$ is monotone and submodular in S .*

Submodularity of F , immediately leads to the following result [4,15].

Corollary 5. *Consider the non-adaptive greedy policy that at each step chooses the element with maximum marginal increase in value. The approximation ratio of this policy with respect to the optimal non-adaptive policy is at least $1 - \frac{1}{e}$.*

Now we present an adaptive greedy policy with approximation ratio $1 - \frac{1}{e}$, with respect to the optimal adaptive policy. It is easy to see that finding maximum k -cover can be reduced to designing an adaptive policy. Therefore, it is not possible to improve this ratio unless $NP \subset TIME(n^{O(\log \log n)})$.

Theorem 6. *Consider the adaptive greedy policy that at each step selects an element with the maximum marginal value, conditioned on the realized value of the previously chosen elements. The approximation ratio of the adaptive greedy policy with respect to the optimal adaptive policy is $1 - \frac{1}{e}$.*

Before stating the proof, we describe some notations. For $1 \leq i \leq k$, let S_i be the set of elements chosen by the greedy adaptive policy up to (and including) step i . Define S_0 to be the empty set. Also, let s_i denote the realization of S_i . The adaptive greedy policy at each step i chooses an element in

$$\operatorname{argmax}_{j \in \mathcal{A} \setminus S_{i-1}} E[F(S_{i-1} \cup j) | s_{i-1}]$$

Proof. The proof presented here is similar to the proof of Kleinberg et al. [10] for submodular set functions. Let T_j be the set chosen by the optimal adaptive policy up to step j . Also, denote the expected marginal value of the i^{th} element chosen by the greedy policy by Δ_i , i.e.,

$$\Delta_i = E[F(S_i) | s_{i-1}] - f(s_{i-1}) = E[F(S_i) - F(S_{i-1}) | s_{i-1}]$$

Consider a realization s_i of S_i . Because the realization of each element of T_j is independent from other elements, and f is submodular, we can write $F(T_j \cup S_i | s_i)$ as the sum of a set of monotone submodular functions. Therefore, $F(T_j \cup S_i | s_i)$ is monotone submodular with respect to j . Hence, for $T = T_k$ we have:

$$E[F(T) | s_i] \leq E[F(T \cup S_i) | s_i] \leq E[F(S_i) + k(F(T_1 \cup S_i) - F(S_i)) | s_i]$$

Because $\Delta_i \geq E[F(T_1 \cup S_i) - F(S_i) | s_i]$ we get,

$$E[F(T) | s_i] \leq E[F(S_i) + k\Delta_{i+1} | s_i]$$

Since the inequality above holds for every history, adding up all such inequalities, for all i , $0 \leq i \leq k - 1$, we have:

$$\begin{aligned} E[F(T)] &\leq E[F(S_i)] + kE[\Delta_{i+1}] \\ &= E[\Delta_1 + \dots + \Delta_i] + kE[\Delta_{i+1}] \end{aligned}$$

We multiply the i^{th} inequality, $0 \leq i \leq k - 1$, by $(1 - \frac{1}{k})^{k-1-i}$, and add them all up. The sum of the coefficients of $E[F(T)]$ is equal to

$$\sum_{i=0}^{k-1} (1 - \frac{1}{k})^{k-1-i} = \sum_{i=0}^{k-1} (1 - \frac{1}{k})^i = \frac{1 - (1 - \frac{1}{k})^k}{1 - (1 - \frac{1}{k})} = k(1 - (1 - \frac{1}{k})^k) \quad (2)$$

One the right hand side, the sum of the coefficient of $E[\Delta_i]$, $1 \leq i \leq k$, is equal to

$$\begin{aligned}
 k\left(1 - \frac{1}{k}\right)^{k-i} + \sum_{j=i}^{k-1} \left(1 - \frac{1}{k}\right)^{k-1-j} &= k\left(1 - \frac{1}{k}\right)^{k-i} + \sum_{j=0}^{k-i-1} \left(1 - \frac{1}{k}\right)^j \\
 &= k\left(1 - \frac{1}{k}\right)^{k-i} + k\left(1 - \left(1 - \frac{1}{k}\right)^{k-i}\right) \\
 &= k
 \end{aligned} \tag{3}$$

Therefore, by inequalities (2) and (3) we get

$$E[F(T)] \leq \left(1 - \left(1 - \frac{1}{k}\right)^k\right) \sum_{i=1}^k E[\Delta_i] = \left(1 - \left(1 - \frac{1}{k}\right)^k\right) E[F(S_k)]$$

Therefore, the approximation ratio of the greedy policy is at least $1 - \frac{1}{e}$.

It is easy to see from the proof above that if at every step, a policy chooses an element which is an α approximation of the maximum marginal value, then it achieves approximation ratio $1 - \left(\frac{1}{e}\right)^\alpha$.

4 Adaptivity Gap: An Upper Bound of 1.59

Our concern in this section will be to set an upper bound for the adaptivity gap. In other words, we want to have a lower bound on the approximation ratio of non-adaptive policies against the best adaptive policy. We establish such a bound through the following theorem:

Theorem 7. *There exists a non-adaptive policy that approximates the optimal adaptive policy within a factor of $\frac{e-1}{e} \approx \frac{1}{1.59}$. Moreover, There exists a polynomial time non-adaptive policy with the approximation ratio at least $\left(\frac{e-1}{e}\right)^2 \approx \frac{1}{2.51}$.*

The proof of the above theorem is inspired by the techniques in Section 3.5 of [18]. For the sake of consistency, we will use the same notation as [18] wherever possible. We generalize these techniques by extending the domain of the function F_g to real vectors. We will define a function f^+ which sets an upper bound on the performance of all adaptive policies and also lies within a constant factor (at most $\frac{e}{e-1}$) of the maximum value of F_g . As we will see, this implies that for every adaptive policy ADAPT there exists a non-adaptive policy which gains at least a fraction of $\frac{e-1}{e}$ of the expected value gained by ADAPT. Also, Corollary 5 shows that the greedy non-adaptive policy approximates the optimal non-adaptive by a factor of $\frac{e-1}{e}$. Hence, it will be within a factor of $\left(\frac{e-1}{e}\right)^2$ of the best adaptive policy.

Here comes our basic observation about adaptive policies. Consider an arbitrary adaptive policy ADAPT. Any such policy decides to choose a sequence of

elements, where the decision about which element to choose at any step might depend on the realized values of the previously chosen elements. Therefore, for any realization of outcomes (i.e. realized values of elements) a distribution on the sequence of elements will be implied by ADAPT³. This distribution corresponds to what ADAPT does if it observes that specific realization. Any adaptive policy can be described by a (possibly randomized) decision tree in which at each step an element is being added to the current selection. Due to the constraints of the problem, the height of the tree is k . Each path from the root to a leaf of this tree corresponds to a subset with k elements and occurs with some certain probability. Clearly, these probabilities sum up to one. Let y_i be the probability that element i is chosen by ADAPT. Also, let β_s be the probability density function for the outcome s . Then, we have the following properties:

1. $\int_s \beta_s = 1.$
2. $\forall s : \beta_s \geq 0.$
3. $\forall i, dx_i : \int_{s, s_i \in dx_i} \beta_s ds = y_i g_i(x_i) dx_i.$

The first two properties hold because β defines a probability measure on the space of all outcomes. The third property is due to the fact that the left hand side is in fact computing the probability that the element i is chosen and its observed value is x_i .

Now, we are ready to define the function $f^+ : [0, 1]^n \rightarrow \mathbb{R}$ which establishes an upper bound on the performance of any adaptive policy. The definition of f^+ is motivated by the above observation about all the possible outcomes of an arbitrary adaptive policy. It can be seen as the generalization of function f^+ in [18] to the continuous domain. For any vector $y \in [0, 1]^n$ we define $f^+(y)$ as follows:

$$\sup_{\alpha} \left\{ \int_s \alpha_s f(s) \right\}$$

where the supremum is taken over all probability measures α defined on $[0, 1]^n$ for which

$$\forall i, dx_i : \int_{s, s_i \in dx_i} \alpha_s ds = y_i g_i(x_i) dx_i.$$

Now, we can bound the performance of ADAPT using function f^+ . Consider all possible realizations of elements under ADAPT. Let y_i and β_s be defined as before. Then, the expected value of ADAPT is $\int_s \beta_s f(s)$. On the other hand, by the construction of β , it is one of the possibilities that can be used as α in the “sup” term in the definition of $f^+(y)$. Therefore, the performance of the policy is bounded by $f^+(y)$ and we have the following lemma.

Lemma 8. *The expected value of the adaptive policy ADAPT is at most $f^+(y)$.*

³ The reason that we mentioned a distribution (and not just a specific sequence) is that ADAPT may be a randomized policy by itself. But as we will see, it does not affect our arguments.

Now, let P_k be the matroid polytope $\{v : 1 \cdot v \leq k, v \geq 0\}$. Any valid policy is limited to select at most k elements, i.e., $y \in P_k$. Therefore, the desired upper bound on the optimal adaptive policy can be obtained as a corollary of Lemma 8.

Corollary 9. *The expected value of the optimal adaptive policy is bounded from above by $\max_{y \in P_k} \{f^+(y)\}$.*

We also define the extension $F_g : [0, 1]^n \rightarrow \mathbb{R}$. For any vector $y \in [0, 1]^n$, $F_g(y)$ is the expected value of its realized outcome \hat{y} when y_i is set to be 1 with probability y_i and 0 otherwise. More formally, if $N = \{1, 2, \dots, n\}$ then

$$F_g(y) = \mathbb{E}[F_g(\hat{y})] = \sum_{R \in N} \prod_{i \in R} y_i \prod_{i \notin R} (1 - y_i) F_g(R).$$

To complete the proof of Theorem 7 we need to prove that there exists a 0/1 vector $w \in P_k$ such that the value of $F_g(w)$ is a good approximation of the optimum value of f^+ . We will do that in two steps. First, we show that for any vector y , the values of $F_g(y)$ and $f^+(y)$ are within a constant of each other. Then, in Lemma 11 we will show that there exists a proper 0/1 vector w so that $F_g(w) \geq F_g(y)$. Remember that in fact $F_g(w) = F_g(S)$ for the subset S corresponding to entries equal to 1 in w . It proves that the ratio of the best non-adaptive and adaptive policy is at least $1 - \frac{1}{e}$. For the second part of the result we show that a 0/1 vector w' such that $F_g(w') \geq (1 - \frac{1}{e})F_g(w)$ can be found in polynomial time which provides an efficient way to find a non-adaptive policy within a factor $(1 - \frac{1}{e})^2$ of the optimal adaptive policy.

The following lemma proves that the defined extension f^+ cannot be too far from F .

Lemma 10. *For any monotone submodular function f and any vector y , $F_g(y) \geq (1 - \frac{1}{e})f^+(y)$.*

Proof. The proof might be viewed as of a generalization of the proof of Lemmas 3.7 and 3.8 in [18] to the continuous space. We define an auxiliary function $f^* : [0, 1]^n \rightarrow \mathbb{R}$ as the following:

$$f^*(y) = \inf_z \{f(z) + \sum_{j \in N} \int_{s_j > z_j} g_j(s)(f(z_s(j)) - f(z)) ds_j,$$

where $z_s(j)$ is the vector z with its j -th entry changed to s_j whenever $s_j > z_j$.

We prove that for any vector y , $F_g(y) \leq f^+(y) \leq f^*(y)$. The first inequality follows directly from the definition of F_g and f^+ . To prove the second inequality, note that for any feasible measure α and any vector z ,

$$\begin{aligned} \int_s \alpha_s f(s) ds &\leq \int_s \alpha_s [f(z) + \sum_{j \in N} (f(z_s(j)) - f(z))] ds \\ &\leq f(z) + \sum_{j \in N} \int_{s_j > z_j} g_j(s)(f(z_s(j)) - f(z)) ds_j. \end{aligned}$$

The first inequality above holds due to submodularity of f and the second one is a consequence of the definition of α . Also, observe that by plugging $\alpha = \beta$ in this inequality we have $f^+(y) \leq f^*(y)$.

Now, it is enough to prove that for all y , $F_g(y) \geq (1 - \frac{1}{e})f^*(y)$. Similar to the proof of Lemma 3.8 [18], for each j we define a Poisson clock \mathcal{C}_j with rate y_j . We start with a vector $z = 0$. Once clock \mathcal{C}_j sends a signal, a random variable x is produced from distribution g_j . Then, if $z_j < x$, the value of z_j will change to x . By abuse of notation, we denote this new value by $z_x(j)$ and the value of vector z at time t by $z(t)$. One can observe that $E[f(z(1))] \leq F_g(y)$, using monotonicity of f . On the other hand,

$$E[f(z(t + dt)) - f(z(t)) | z(t) = z] = \sum_j y_j dt \int_{x > z_j} g_j(x) (f(z_x(j)) - f(z)) dx.$$

But the R.H.S. is at least $(f^*(y) - f(z))dt$, by the definition of f^* . Therefore, the following bound can be derived on the derivative of $E[f(z(t))]$:

$$\begin{aligned} \frac{1}{dt} E[f(z(t + dt)) - f(z) | z(t) = z] &\geq (f^*(y) - f(z))dt \\ \Rightarrow \frac{d}{dt} E[f(z(t))] &\geq (f^*(y) - E[f(z(t))])dt. \end{aligned}$$

Solving the differential equation above, shows that $E[f(z(t))] \geq (1 - e^{-t})f^*(y)$. Combining this with the fact that $f^+(y) \leq f^*(y)$ and also that $E[f(z(1))] \leq F_g(y)$ completes the proof of lemma.

The next lemma shows how to round the vector y to a proper 0/1 vector w .

Lemma 11. *There exists a 0/1 vector $w \in P_k$ such that $\forall y : F_g(w) \geq F_g(y)$.*

Proof. The essential rounding tool for the proof is pipage rounding introduced by [1]. In order to be able to use pipage rounding we need to prove some convexity property on F_g . Define $F_{ij}^y = F_g(y_{ij}(\lambda))$ where $y_{ij}(\lambda)$ is a vector obtained by adding λ to y_i , subtracting λ from y_j and leaving all other entries of y unchanged. First, we show that F_{ij}^y is a convex function of λ . For any y , the function $F_g(y)$ can be written as below.

$$\begin{aligned} F_g(y) = \sum_{R \in N \setminus \{i,j\}} \prod_{k \in R} y_k \prod_{k \notin R \cup \{i,j\}} (1 - y_k) \times & [(1 - y_i)(1 - y_j)F_g(R) + \\ & (1 - y_i)y_j F_g(R + j) + y_i(1 - y_j)F_g(R + i) + y_i y_j F_g(R + i + j)]. \end{aligned}$$

Hence, we can write the second derivative of F_{ij}^y in an explicit form:

$$\begin{aligned} \frac{\partial^2 F_{ij}^y}{\partial \lambda^2} = \sum_{R \in N \setminus \{i,j\}} \prod_{k \in R} y_k \prod_{k \notin R \cup \{i,j\}} (1 - y_k) \times \\ [-F_g(R) + F_g(R + i) + F_g(R + j) - F_g(R + i + j)], \end{aligned}$$

which is clearly non-negative due to submodularity of F_g .

As a result of convexity of F_{ij}^y , for any vector $y \in P_k$, the main result of [1] ensures that Pipage rounding yields a 0/1 vector w inside P_k such that $F_g(w) \geq F_g(y)$. Hence, there exists such a vector w for which $F_g(w) \geq F_g(y)$ holds for all $y \in P_k$.

Now, we are ready to prove Theorem 7.

Proof. [Theorem 7]. Lemma 8 shows that $\text{OPT} = \max_{y \in P_k} f^+(y)$ is an upper-bound on the performance of the best adaptive policy. But from Lemma 10 we know that there exists a vector y^* such that $F_g(y^*) \geq (1 - \frac{1}{e})\text{OPT}$. On the other hand, Lemma 11 implies that there exists a 0/1 vector $w \in P$ such that $F_g(w) \geq F_g(y^*)$ and hence, $F_g(w) \geq (1 - \frac{1}{e})\text{OPT}$. Notice that $F_g(w)$ is in fact the expected value gained by a non-adaptive policy that selects the set S corresponding to the vector w . Also, due to Corollary 5 greedy non-adaptive policy obtains a value at least $(1 - \frac{1}{e})F_g(w)$ that will be at least $(1 - \frac{1}{e})^2\text{OPT}$.

References

1. Ageev, A.A., Sviridenko, M.: Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *J. Comb. Optim.* 8(3), 307–328 (2004)
2. Ahmed, S., Atamtürk, A.: Maximizing a class of submodular utility functions. Research Report BCOL.08.02, IEOR, University of California-Berkeley (March 2008)
3. Chan, C., Farias, V.: Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems (manuscript, 2008)
4. Cornuejols, G., Fisher, M., Nemhauser, G.: Location of bank accounts to optimize float. *Management Science* 23, 789–810 (1977)
5. Dean, B.C., Goemans, M.X., Vondrák, J.: Approximating the stochastic knapsack problem: The benefit of adaptivity. In: *FOCS*, pp. 208–217 (2004)
6. Feige, U.: A threshold of \ln for approximating set cover. *J. ACM* 45(4), 634–652 (1998)
7. Feige, U., Mirrokni, V.S., Vondrák, J.: Maximizing non-monotone submodular functions. In: *FOCS*, pp. 461–471 (2007)
8. Goemans, M.X., Vondrák, J.: Stochastic covering and adaptivity. In: Correa, J.R., Hevia, A., Kiwi, M. (eds.) *LATIN 2006*. LNCS, vol. 3887, pp. 532–543. Springer, Heidelberg (2006)
9. Kempe, D., Kleinberg, J.M., Tardos, É.: Maximizing the spread of influence through a social network. In: Getoor, L., Senator, T.E., Domingos, P., Faloutsos, C. (eds.) *KDD*, pp. 137–146. ACM, New York (2003)
10. Kleinberg, J.M., Papadimitriou, C.H., Raghavan, P.: Segmentation problems. *J. ACM* 51(2), 263–280 (2004)
11. Krause, A., Guestrin, C.: Near-optimal nonmyopic value of information in models. In: *AAAI*, pp. 324–331 (2005)
12. Krause, A., Guestrin, C.: Near-optimal observation selection using submodular functions. In: *AAAI*, pp. 1650–1654. AAAI Press, Menlo Park (2007)
13. Mehrez, A., Sinuany-Stern, Z.: Resource allocation to interrelated risky projects using a multiattribute utility function. *Management Science* (29), 439–490 (1983)
14. Mossel, E., Roch, S.: On the submodularity of influence in social networks. In: *Approx*, pp. 128–134 (2007)

15. Nemhauser, G., Wolsey, L., Fisher, M.: An analysis of the approximations for maximizing submodular set functions. *Mathematical Programming* 14, 265–294 (1978)
16. Streeter, M., Golovin, D.: An online algorithm for maximizing submodular functions. Tech Report CMU-CS-07-171 (2008)
17. Topkis, D.M.: Minimizing a submodular function on a lattice. *Operations Research* 26(2), 305–321 (1978)
18. Vondrák, J.: Submodularity in combinatorial optimization. PhD thesis. Charles University, Prague (2007)
19. Vondrák, J.: Optimal approximation for the submodular welfare problem in the value oracle model. In: *STOC* (2008)
20. Weingartner, H.: *Mathematical Programming and the Analysis of Capital Budgeting Problems*. Prentice-Hall, Englewood Cliffs (1963)
21. Weingartner, H.: Capital budgeting of interrelated projects: Survey and synthesis. *Management Science* (1966)

On Pure and (Approximate) Strong Equilibria of Facility Location Games^{*}

Thomas Dueholm Hansen and Orestis A. Telelis

Department of Computer Science, University of Aarhus, Denmark
{tdh,telelis}@cs.au.dk

Abstract. We study social cost losses in Facility Location games, where n selfish agents install facilities over a network and connect to them, so as to forward their local demand (expressed by a non-negative weight per agent). Agents using the same facility share fairly its installation cost, but every agent pays individually a (weighted) connection cost to the chosen location. We study the Price of Stability (PoS) of *pure* Nash equilibria and the Price of Anarchy of *strong* equilibria (SPoA), that generalize pure equilibria by being resilient to coalitional deviations. For unweighted agents on metric networks we prove upper and lower bounds on *PoS*, while an $O(\ln n)$ upper bound implied by previous work is tight for non-metric networks. We also prove a constant upper bound for the SPoA of metric networks when strong equilibria exist. For the weighted game on general networks we prove existence of e -approximate ($e = 2.718\dots$) strong equilibria and an upper bound of $O(\ln W)$ on SPoA (W is the sum of agents' weights), which becomes tight $\Theta(\ln n)$ for unweighted agents.

1 Introduction

We study *Facility Location* games played by n selfish agents residing on the nodes of a network. Each agent chooses strategically a certain network location to connect and forward its local demand to (expressed by a non-negative weight w_i for agent i), so as to minimize its individual facility installation and (weighted) connection costs to the chosen location. We use Shapley (fair) cost-sharing [1] for facility installation costs; agents connecting to the same location v share facility installation cost at v , so that each pays an amount proportional to the fraction of total demand that it forwards to v . This game models *Content Distribution Network* creation, and distributed *selfish caching* [2]. We study the social cost (sum of individual agents' costs) of *stable* network infrastructures, represented by *pure* Nash equilibria and *strong* equilibria of the game. Strong equilibria - introduced by Aumann in [3] - extend pure equilibria by being resilient to *pure* coalitional deviations: no subset of agents can deviate so that all of its members are better off. We prove bounds on the *Price of Stability* (PoS) of pure equilibria, i.e. the cost of the cheapest equilibrium relative to the socially optimum cost [1], and on the *Price of Anarchy* of strong equilibria (SPoA), the cost of the most expensive strong equilibrium relative to the socially optimum cost [4].

^{*} Center for Algorithmic Game Theory, funded by the Carlsberg Foundation, Denmark.

Anshelevich *et al.* [1] first studied the Price of Stability for network design games with fair cost sharing. In these games n agents wish to connect node pairs in a network, by sharing fairly installation costs of links and paying individually link delays. The authors showed that for unweighted agents these games are *potential* games (see [5]), hence they have pure equilibria. They proved logarithmic (in n) upper bounds on the *PoS*. Research thereafter was focused on games *without non-shareable delays*. For weighted agents it was shown in [6] that pure equilibria do not always exist. The authors studied approximate equilibria. Albers [7] recently considered *strong* equilibria: though they do not always exist, she showed that $O(\ln W)$ -approximate equilibria do exist (W is the sum of the agents' weights). She proved polylogarithmic upper and lower bounds on the *PoS* and the *SPoA* in the weighted and the unweighted case. However, strong equilibria in the context of (single-sink) unweighted network design games with fair cost sharing were first studied in [8]. The authors gave topological characterizations for the existence of strong equilibria and proved that $SPoA = \Theta(\log n)$. The Facility Location game is a special case of the model studied in [1], that is interesting on its own right: it finds numerous applications and exhibits intriguing characteristics. It embodies non-shareable delays explicitly and in a sense specializes single-sink network design considered in [6,8]: augment the network with a node t and set links (v, t) to have fairly shareable cost equal to the facility opening cost at v . The original network links have a delay cost only. Then every agent needs to choose at most two edges from the node it resides on, to t .

Results. For unweighted agents on metric networks we prove constant upper and lower bounds on the *PoS*, by analyzing the social cost increase caused by an iterative best response procedure. Strong equilibria do not always exist, but their *SPoA* is constant upper-bounded when they do. For weighted agents on general networks we prove that α -approximate strong equilibria exist for $\alpha \geq e = 2.718\dots$ (no subset deviation causes factor α improvement to all of its members), and that their *SPoA* is at most $\alpha(1 + \ln W)$. This becomes $\Theta(\ln n)$ for unweighted agents on general networks. See [9] for additional results, omitted proofs and technical details. Refer to [2,10,11] for related work on facility location game models. [2] is similar to ours, but does not incorporate fair cost-sharing of facility costs. [11] specializes the model of [1], but does not incorporate delays.

Definitions. The network will be a complete graph $G(V, E)$, having each edge (u, v) associated to a non-negative cost $d(u, v)$. We consider a set A of n agents; agent i resides on $u_i \in V$ and has a non-negative demand weight w_i . The strategy space of agent i is V : i chooses a location $v \in V$ to receive service from. Opening a facility at node $v \in V$ costs β_v . Denote a strategy profile (configuration) by $s = (s_1, \dots, s_n)$, $s_i \in V$. We define $W_s(v) = \sum_{i:s_i=v} w_i$. The cost $c_i(s)$ experienced by agent i in s is $c_i(s) = w_i \left(d(u_i, s_i) + \frac{\beta_{s_i}}{W_s(s_i)} \right)$. Agent i pays a fraction $\frac{w_i}{W_s(s_i)}$ of the facility installation cost at s_i . We denote facility locations specified in s by $F_s \subseteq V$. The *social cost* $c(s)$ is:

$$c(s) = \sum_i c_i(s) = \sum_i w_i d(u_i, s_i) + \sum_i \frac{w_i \beta_{s_i}}{W_s(s_i)} = \sum_i w_i d(u_i, s_i) + \sum_{v \in F_s} \beta_v$$

We use $W(I)$ for the sum of weights of agents in the set I . $c_I(s)$ is the social cost of agents in I , and $c_v(s)$ is the social cost of agents connected to $v \in V$ under s .

The *unweighted* Facility Location game is a potential game [5] specializing the network design games of [1], and therefore has pure Nash equilibria reachable by iterative best response performed by the players. For unweighted agents the PoA of pure equilibria can be n , while the PoS upper bound of $H(n)$ (n -th harmonic number) from [1] is tight for non-metric networks [9].

Definition 1. For $\alpha \geq 1$, a strategy profile s is an α -approximate strong equilibrium if no subset of agents can perform a pure deviation, and each of its members be better off by a factor more than α . If $\alpha = 1$, s is a strong equilibrium [3,8].

2 Unweighted Agents on Metric Networks

We analyze evolution of an equilibrium through iterative best response executed by the agents, when the initial configuration is the social optimum. The following lemma charges any specific agent i a bounded amount of social cost increase during the algorithm’s execution.

Lemma 1. Let $A_{s^*}(v)$ be the subset of agents that are connected to v in s^* . For any $i \in A_{s^*}(v)$ that deviates from v during iterative best response let $A_{s^*}^i(v) \subseteq A_{s^*}(v)$ be the subset of agents that have not yet deviated from v exactly before the first deviation of i . Then we can charge i with a total increase contribution to the social cost at most $\beta_v / |A_{s^*}^i(v)|$, throughout the algorithm’s execution.

Proof. For simplicity let $|A_{s^*}^i(v)| = k_i(v)$. Clearly $i \in A_{s^*}^i(v)$. Let us analyze contribution of i to social cost increase during its first deviation. By deviating i reduces its individual cost from $c_i(v) = x_i(v) + \frac{\beta_v}{k_i(v)}$ to $c_i(v') = x_i(v') + \frac{\beta_{v'}}{\lambda_i(v')}$ by joining another facility node v' . $x_i(v)$ and $x_i(v')$ is the connection cost payed by i before and after its first deviation. $\lambda_i(v')$ is the number of agents sharing facility cost at v' , including i . Since $c_i(v') < c_i(v)$, we get $x_i(v') - x_i(v) \leq \frac{\beta_v}{k_i(v)} - \frac{\beta_{v'}}{\lambda_i(v')}$. Let $\Delta sc_i(v)$ be the social cost difference caused by i . There are four cases:

1. $k_i(v) > 1, \lambda_i(v') > 1$: Then $\Delta sc_i(v) = x_i(v') - x_i(v) \leq \frac{\beta_v}{k_i(v)} - \frac{\beta_{v'}}{\lambda_i(v')}$.
2. $k_i(v) = 1, \lambda_i(v') > 1$: Then $\Delta sc_i(v) = -\beta_v + x_i(v') - x_i(v) \leq -\frac{\beta_{v'}}{\lambda_i(v')}$.
3. $k_i(v) > 1, \lambda_i(v') = 1$: Then $\Delta sc_i(v) = \beta_{v'} + x_i(v') - x_i(v) \leq \frac{\beta_v}{k_i(v)}$.
4. $k_i(v) = 1, \lambda_i(v') = 1$: Then $\Delta sc_i(v) = \beta_{v'} - \beta_v + x_i(v') - x_i(v) \leq 0$.

Clearly the above hold in general for any agent deviating from any node v to any node v' . Now we implement a charging procedure along with iterative best response. Give all agents an initial label $l(i) = i$, before executing iterative best response. The current label $l(i)$ of i will denote the agent to which an increase

caused by i is charged. Initialize $\Delta sc_{l(i)} = 0$. For every facility node $v \in F_{s^*}$ and every $i \in A_{s^*}(v)$ initialize $\lambda_{l(i)} = \lambda_i$ to a distinct value from $\{1, 2, \dots, |A_{s^*}(v)|\}$. Charging is implemented by relabeling deviating agents in the following manner. For an agent i that deviates from node v to node v' set $k_{l(i)}(v)$ to the number of agents connected to v exactly before deviation of i . If $k_{l(i)}(v) = \lambda_{l(i)}(v)$ no relabeling is needed. Otherwise there must be some $j \neq i$ connected to v such that $\lambda_{l(j)}(v) = k_{l(i)}(v)$. In this case exchange labels of i and j . Subsequently add the increase caused by deviation of i to $\Delta sc_{l(i)}(v)$. Finally, set $\lambda_{l(i)}(v')$ equal to the number of agents connected to v' right after i has joined v' .

By the previous definitions it follows that if $k_{l(i)}(v) \neq \lambda_{l(i)}(v)$, then it is always $k_{l(i)}(v) > \lambda_{l(i)}(v)$, i.e. i has joined v before some agent j with $\lambda_{l(j)}(v) = k_{l(i)}(v)$, but leaves v before i leaves. By exchanging labels of i, j we add the increase caused by i to the agent that previously labeled j . Possible increases in 1.,2.,3.,4., imply that any agent is charged by the end of iterative best response at most $\frac{\beta_v}{|A_{s^*}^i(v)|}$ for some i . Initializing $\lambda_i(v) = k_i(v)$ in s^* charges exactly i . \square

Note: In the following we assume an order of agents, so that agents of the same facility in the initial configuration best-respond consecutively.

Theorem 1. *The Price of Stability for the unweighted metric Facility Location game is upper bounded by a constant, strictly less than 2.36.*

Proof. Let Δsc_i denote the increase contributed by agent i to the social optimum $c(s^*)$, during iterative best response initialized at s^* . Assume an order of agents, such that agents $i \in A_{s^*}(v)$ “best-respond” consecutively, for each $v \in F_{s^*}$. Define $c_v(s^*) = \beta_v + \sum_{i:s_i^*=v} d(u_i, v)$. Then $c(s^*) = \sum_{v \in F_{s^*}} c_v(s^*)$. We will upper bound the *POs* by $\max_{v \in F_{s^*}} \frac{c_v(s^*) + \sum_{i:s_i^*=v} \Delta sc_i}{c_v(s^*)}$. We focus on the *first* deviation of $i \in A_{s^*}(v)$, for any facility $v \in F_{s^*}$. Let v' be the node that i deviates to, and $\delta x_i^* = d(u_i, v') - d(u_i, v)$. We also use $x_i^* = d(u_i, v)$ for convenience. Let λ_i be the number of agents serviced at v' right after deviation of i . The new cost of i right after its first deviation is: $d(u_i, v) + \delta x_i^* + \frac{\beta_{v'}}{\lambda_i}$. For a second agent $j \in A_{s^*}(v)$ deviating from v to some node v'' after i , we have:

$$d(u_j, v) + \delta x_j^* + \frac{\beta_{v''}}{\lambda_j} \leq d(u_j, v') + \frac{\beta_{v'}}{\lambda_i} \tag{1}$$

Substitute $d(u_j, v')$ in (1) by triangle inequality: $d(u_j, v') \leq d(u_j, v) + d(u_i, v) + d(u_i, v')$. Also, by lemma 1 $\delta x_j^* + \frac{\beta_{v''}}{\lambda_j} \leq \frac{\beta_{v'}}{k_i^*}$, where $k_i^* = |A_{s^*}^i(v)|$ ($A_{s^*}^i(v)$ is defined as in lemma 1). Thus:

$$d(u_i, v) \geq \frac{1}{2} \left(\delta x_j^* - \delta x_i^* + \frac{\beta_{v''}}{\lambda_j} - \frac{\beta_{v'}}{\lambda_i} \right) \geq \frac{1}{2} \left(\delta x_j^* + \frac{\beta_{v''}}{\lambda_j} - \frac{\beta_v}{k_i^*} \right) \tag{2}$$

The latter has to hold for every pair of distinct agents $i, j \in A_{s^*}(v)$, hence:

$$d(u_i, v) \geq \max \left\{ 0, \frac{1}{2} \left(\max_{j:s_j^*=v} \left(\delta x_j^* + \frac{\beta_{v''}}{\lambda_j} \right) - \frac{\beta_v}{k_i^*} \right) \right\} \tag{3}$$

We use (3) for the connection cost of agents in $A_{s^*}(v)$ under s^* , and consider two complementary cases: either some agents never deviate from v , or all of them do. Let $n_{s^*}(v) = |A_{s^*}(v)|$. We only analyze the first case here (the second is similar - see [9]). If r agents never deviate from v , then trivially in (3) we set $v'' = v$, $\delta x_j^* = 0$, and $\lambda_j = r$, whereas $k_i^* \geq r + 1$. The cost $c_v(s^*)$ is:

$$c_v(s^*) \geq \beta_v + \frac{\beta_v}{2} \sum_{k=r+1}^{n_{s^*}(v)} \left(\frac{1}{r} - \frac{1}{k} \right) = \beta_v + \frac{\beta_v}{2} \left(\frac{n_{s^*}(v) - r}{r} - H(n_{s^*}(v)) + H(r) \right) \quad (4)$$

By lemma [1] it is $\sum_{i:s_i^*=v} \Delta s c_i \leq H(n_{s^*}(v)) - H(r)$. Using equality in (4) for $c_v(s^*)$, and $c_v(s) = c_v(s^*) + \sum_{i:s_i^*=v} \Delta s c_i$, we obtain the following ratio. Simplify using $\gamma + \ln m \leq H(m) \leq 1 + \ln m$ ($\gamma > 0.5$ is Euler’s constant):

$$PoS \leq \frac{1 + \frac{1}{2} \left(\frac{n_{s^*}(v) - r}{r} + H(n_{s^*}(v)) - H(r) \right)}{1 + \frac{1}{2} \left(\frac{n_{s^*}(v) - r}{r} - H(n_{s^*}(v)) + H(r) \right)} \leq \frac{1.5 + \frac{n_{s^*}(v)}{r} + \ln \frac{n_{s^*}(v)}{r}}{0.5 + \frac{n_{s^*}(v)}{r} - \ln \frac{n_{s^*}(v)}{r}}$$

Let $y = \frac{n_{s^*}(v)}{r}$. The upper bound can be numerically maximized to < 2.36 . \square

Lower Bound. Take $2n$ agents; n on a single node v , the rest on a separate node each (black nodes in Fig. [1](a)). Facility costs are 1. In the social optimum s^* , n agents on v are serviced by v . The rest are equipartitioned to v_l^* , $l = 1 \dots k$, $k = \sqrt{n}$. We analyze a single facility v_l^* , henceforth denoted by v^* (same for the rest). By abusing notation, $c_{v^*}(s)$ is the cost of v^* -agents at equilibrium. Then:

$$PoS = \lim_{n \rightarrow \infty} \frac{1 + k c_{v^*}(s)}{1 + k c_{v^*}(s^*)} \geq \lim_{n \rightarrow \infty} \frac{c_{v^*}(s)}{(1/\sqrt{n} + c_{v^*}(s^*))} \quad (5)$$

In the least expensive equilibrium s , agents from each facility v^* of s^* miss-connect to v in s . For some constant $p \in (0, 1)$, only $r = \lceil (1 - p)k \rceil$ of these agents increase the social cost significantly, by increasing their connection cost. Follow iterative best response of these r agents starting from s^* . Assume $\lambda \geq n$ agents “play” v before r agents of v^* deviate to v . Set the i -th deviating agent to increase its connection cost x_i^* by $\delta x_i^* = \frac{1}{k-i+1} - \frac{1}{\lambda+i} - \epsilon$, $i = 1 \dots r$; it decreases c_i by ϵ . By (1), (2), and because all r agents deviate to v (hence $\frac{\beta_{v''}}{\lambda_j} - \frac{\beta_{v'}}{\lambda_i} = 0$ in (2)), it is $x_i^* = d(u_i, v^*) = \max\{0, \frac{1}{2}(\max \delta x_j^* - \delta x_i^*)\}$, $\max \delta x_j^* = \delta x_r^*$. For the rest $k - r$ agents set $x_j^* = 0$, $d(u_j, v) = \frac{1}{k-r+1}$. Summing up as in (4) yields:

$$c_{v^*}(s^*) = 1 + \frac{1}{2} \left[\frac{r}{k-r+1} - \Delta H(k, k-r) \right] - \frac{1}{2} \frac{r}{n+r} + \frac{1}{2} \Delta H(\lambda+r, \lambda) \quad (6)$$

where $\Delta H(n, m) = H(n) - H(m)$. Then $\sum_{i=1}^r \delta x_i^* = \Delta H(k, k-r) - \Delta H(\lambda+r, \lambda)$, and we can set $c_{v^*}(s) = c_{v^*}(s^*) - 1 + \frac{k-r}{k-r+1} + \sum_{i=1}^r \delta x_i^*$ in (5). Simplify ΔH by logarithmic bounds and substitute $r = \lceil (1 - p)k \rceil$ appropriately. Then limits of numerator and denominator in (5) exist (see [9] for details); the resulting simplified fraction can be maximized numerically to > 1.45 for $p \simeq 0.18$. Experimental




Fig. 1. Lower bounds: (a) unweighted metric PoS, (b) unweighted non-metric SPOA

evidence showed that $PoS > 1.77$. It is easy to verify that any configuration other than s^* and s is more expensive [9].

Strong equilibria do not always exist, even for unweighted agents on metric networks. We prove the following (see [9] for the proof, and for existence of 2.36-approximate strong equilibria with constant strong Price of Anarchy):

Theorem 2. *When strong equilibria exist in the unweighted metric Facility Location game, their Price of Anarchy is at most a constant.*

3 Approximate Strong Equilibria for Weighted Agents

The existence of pure equilibria for weighted agents is an open issue. We reduce the logarithmic approximation factor known for general network design [7,6] to a constant. Our result is more general, as it concerns *strong* equilibria. We make use of the following remark.

Remark 1. If an instance of the Facility Location game does not have strong equilibria, then there is at least one cycle of deviations of particular coalitions that results in a circular sequence of configurations $\{s^j\}_{j=1}^k$ with $s^1 = s^k$.

Given such a sequence $\{s^j\}_{j=1}^k$, we denote the coalition that deviates from s^j to form s^{j+1} by I_j . Such a deviation causes a cost decrease of agents in I_j and possibly a cost increase of agents in $A \setminus I_j$. Recall that A is the set of all agents. We define two quantities, the *weighted improvement* $\text{impr}(I_j)$ for agents in I_j and the *weighted damage* $\text{dam}(I_j)$ caused by agents in I_j respectively:

$$\text{impr}(I_j) = \prod_{i \in I_j} \left(\frac{c_i(s^j)}{c_i(s^{j+1})} \right)^{w_i} \quad \text{dam}(I_j) = \prod_{i \in A \setminus I_j} \left(\frac{c_i(s^{j+1})}{c_i(s^j)} \right)^{w_i}$$

We derive an approximation factor that eliminates cycles.

Lemma 2. *Let $\{s^j\}_{j=1}^k$ with $s^1 = s^k$ be a cycle of configurations in a Facility Location game instance, caused by consecutive deviations of coalitions. The game*

instance has an α -approximate strong equilibrium if for all such sequences $\alpha \geq \text{dam}_{\max}(\{s^j\}_{j=1}^k)$, where $\text{dam}_{\max}(\{s^j\}_{j=1}^k) = \max_{j=1 \dots k-1} \text{dam}(I_j)^{1/W(I_j)}$.

Proof. If there is no α -approximate strong equilibrium we know that there is at least one cycle $\{s^j\}_{j=1}^k$ such that $\forall j \in \{1, \dots, k-1\} \forall i \in I_j : \frac{c_i(s^j)}{c_i(s^{j+1})} > \alpha$.

Because $s^1 = s^k$ we have that $\prod_{j=1}^{k-1} \frac{c_i(s^j)}{c_i(s^{j+1})} = 1$ for every agent i . Then:

$$1 = \prod_{i=1}^n \left(\prod_{j=1}^{k-1} \frac{c_i(s^j)}{c_i(s^{j+1})} \right)^{w_i} = \prod_{j=1}^{k-1} \frac{\text{impr}(I_j)}{\text{dam}(I_j)} > \prod_{j=1}^{k-1} \frac{\alpha^{W(I_j)}}{\left(\text{dam}(I_j)^{1/W(I_j)}\right)^{W(I_j)}}$$

It follows that $\text{dam}_{\max}(\{s^j\}_{j=1}^k) > \alpha$. The lemma follows by contradiction. \square

We derive an approximation factor as an upper bound of $\text{dam}_{\max}(\{s^j\}_{j=1}^k)$ for any cycle.

Theorem 3. *For every $\alpha \geq e$ there exist α -approximate strong equilibria in the Facility Location game.*

Proof. We prove that $\text{dam}_{\max}(\{s^j\}_{j=1}^k) < e$ for every cycle $\{s^j\}_{j=1}^k$ of configurations and the result follows from Lemma 2. Let $I_j(v)$ be the set of agents going to v in s^j , but not in s^{j+1} , and $A_j(v)$ be the set of agents going to v in both s^j and s^{j+1} . Note that $I_j = \bigcup_{v \in V} I_j(v)$, therefore:

$$\text{dam}_{\max}(\{s^j\}_{j=1}^k) = \max_j \left(\prod_{v \in V} \left(\prod_{i \in A_j(v)} \left(\frac{c_i(s^{j+1})}{c_i(s^j)} \right)^{w_i} \right)^{\frac{W(I_j(v))}{W(I_j)}} \right)^{\frac{1}{W(I_j)}} \Rightarrow$$

$$\text{dam}_{\max}(\{s^j\}_{j=1}^k) \leq \max_{j,v} \left(\prod_{i \in A_j(v)} \left(\frac{c_i(s^{j+1})}{c_i(s^j)} \right)^{w_i} \right)^{\frac{1}{W(I_j(v))}}$$

Hence, we need only consider what happens at the worst case node. For an agent i in $A_j(v)$ we get that:

$$\frac{c_i(s^{j+1})}{c_i(s^j)} = \frac{w_i \left(d(u_i, v) + \frac{\beta_v}{W_{s^{j+1}}(v)} \right)}{w_i \left(d(u_i, v) + \frac{\beta_v}{W(I_j(v)) + W(A_j(v))} \right)} \leq 1 + \frac{W(I_j(v))}{W(A_j(v))}$$

It follows that:

$$\text{dam}_{\max}(\{s^j\}_{j=1}^k) \leq \max_{j,v} \left(1 + \frac{W(I_j(v))}{W(A_j(v))} \right)^{\frac{W(A_j(v))}{W(I_j(v))}} < \lim_{r \rightarrow \infty} \left(1 + \frac{1}{r} \right)^r = e \quad \square$$

Corollary 1. *The Facility Location game with non-uniform agent demands has α -approximate pure strategy Nash equilibria for every $\alpha \geq e$.*

For the SPoA of α -approximate strong equilibria we show [9]:

Theorem 4. *The Price of Anarchy of α -approximate strong equilibria, for the Facility Location game is upper bounded tightly by $\alpha H(n)$ for unweighted and by $\alpha(1 + \ln W)$ for weighted agents, where W is the sum of weights.*

Fig. 1(b) shows a tight (non-metric) example for $w_i = 1$. Facility opening costs are 1 and agents reside on v_{opt} . A single facility at v_{eq} is the most expensive α -approximate strong equilibrium, of cost $\alpha H(n)$: no coalition has incentive to deviate to v_{opt} , the sole optimum facility location of social cost 1, for $\epsilon = n^{-2} \rightarrow 0$

Open Problems. Existence of pure equilibria for the weighted game merits further investigation. Extending our (unweighted) metric analysis of the PoS for weights (or proving a non-constant lower bound) appears to be quite challenging. This seems to apply for the lower bounding of the weighted SPoA on general networks as well. Lower bounding techniques of [7] do not seem applicable.

Acknowledgements. We thank participants of CAGT *Open Problems Jam Sessions*: Daniel Andersson, Gudmund S. Frandsen, Kristoffer A. Hansen, Peter Bro Miltersen, Rocio S. Rodriguez, Troels B. Sørensen, Nikos Triandopoulos.

References

1. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, E., Wexler, T., Roughgarden, T.: The Price of Stability for Network Design with Fair Cost Allocation. In: Proc. IEEE Symp. on Foundations of Computer Science (FOCS), pp. 295–304 (2004)
2. Chun, B., Chaudhuri, K., Wee, H., Barreno, M., Papadimitriou, C.H., Kubiataowicz, J.: Selfish caching in distributed systems: a game-theoretic analysis. In: Proc. ACM Symposium on Principles of Distributed Computing, pp. 21–30 (2004)
3. Aumann, Y.: Acceptable Points in General Cooperative n-Person Games IV. In: Tucker, A.W., Luce, R.D. (eds.) Contributions to the Theory of Games. Annals of Mathematics Study 40, pp. 287–324 (1959)
4. Koutsoupias, E., Papadimitriou, C.H.: Worst-case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
5. Monderer, D., Shapley, L.S.: Potential games. Games and Economic Behavior 14, 124–143 (1996)
6. Chen, H., Roughgarden, T.: Network Design with Weighted Players. In: Proc. ACM Symposium on Parallel Algorithms and Architectures (SPAA), pp. 29–38 (2006)
7. Albers, S.: On the value of coordination in network design. In: Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 294–303 (2008)
8. Epstein, A., Feldman, M., Mansour, Y.: Strong Equilibrium in Cost Sharing Connection Games. In: Proc. ACM Conf. on Electronic Commerce (EC), pp. 84–92 (2007)
9. Hansen, T.D., Telelis, O.A.: On Pure and (approximate) Strong Equilibria of Facility Location Games. arXiv:0809.4792 [cs.GT] (2008)
10. Vetta, A.: Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions. In: Proc. IEEE-FOCS, p. 416 (2002)
11. Chekuri, C., Chuzhoy, J., Lewin-Eytan, L., Naor, J., Orda, A.: Non-cooperative multicast and facility location games. In: Proc. ACM-EC, pp. 72–81 (2006)

Efficiency, Fairness and Competitiveness in Nash Bargaining Games

Deeparnab Chakrabarty¹, Gagan Goel², Vijay V. Vazirani^{2,*},
Lei Wang², and Changyuan Yu^{3,**}

¹ Department of Combinatorics and Optimization, University of Waterloo, Waterloo

² College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280

³ Institute for Theoretical Computer Science, Tsinghua University, Beijing, China

Abstract. Recently, [8] defined the class of Linear Nash Bargaining Games (LNB) and obtained combinatorial, polynomial time algorithms for several games in this class. [8] also defines two natural subclasses within LNB, UNB and SNB, which contain a number of natural Nash bargaining games. In this paper we define three basic game theoretic properties of Nash bargaining games: price of bargaining, fairness and full competitiveness. We show that for each of these properties, a game in UNB has this property iff it is in SNB.

1 Introduction

The bargaining game was first modeled in John Nash's seminal 1950 paper [5] using the framework of game theory given a few years earlier by von Neumann and Morgenstern [9]. Since bargaining is perhaps the oldest situation of conflict of interest, and since game theory develops solution concepts for negotiating in such situations, it is not surprising that this paper led to a theory (of bargaining) that lies today at the heart of game theory (e.g., see [3,7,6]).

In a recent paper, Vazirani [8] initiated a study of Nash bargaining games via combinatorial, polynomial time algorithms. [8] defines LNB (Linear Nash Bargaining Games) – the class of games whose feasible set of utilities is defined by finitely many packing constraints. [8] also defines two natural subclasses within LNB: UNB and SNB. These classes contain a number of natural Nash bargaining games. In this paper we define three basic game theoretic properties of Nash bargaining games and show that for each of these properties, a game in UNB has this property iff it is in SNB. Below we intuitively define the classes UNB and SNB and then state the three properties; formal definitions appear in Section 2.

UNB is the subclass in which for each available resource, each agent who uses this resource uses it in the same way, i.e., the coefficients in the packing constraints are 0/1. Clearly, only $2^{|A|}$ such constraints are needed, where A is the set of agents – one for each subset of A . We can now view the right hand

* Supported by NSF Grant CCF-0728640.

** Supported by National Natural Science Foundation of China Grant 60553001, and the National Basic Research Program of China Grant 2007CB807900,2007CB807901.

sides of these constraints as being given by a set valued function over the power set of A . If this function is submodular, the game is said to be in the subclass SNB of UNB.

We define price of bargaining in a way that is analogous to the notion of price of anarchy [4], i.e., it measures the loss in efficiency in resorting to the Nash bargaining solution as compared to the most efficient solution that can be obtained in a centralized manner. We will say that a Nash bargaining game is *fully competitive* if whenever one player increases his disagreement utility, no other player’s utility can increase in the resulting Nash bargaining solution. We consider max-min and min-max fairness of the Nash bargaining solution in relation to all feasible solutions that are Pareto optimal. For all these properties, we gave a complete characterization of bargaining games in UNB. That is, any bargaining game in UNB has one of these properties if and only if it is in SNB.

2 Uniform Nash Bargaining Games

For a set of agents A , a Nash bargaining game is defined by a pair (\mathbf{c}, P) , where $P \subseteq R_+^{|A|}$ is a compact and convex set which defines the feasible set of utilities of all the agents, and $\mathbf{c} \in P$ is known as the disagreement point which defines the amount of utility each agent will get if the bargaining process fails.

Nash [5] defined the bargaining solution $\mathbf{u} \in P$ of this game to be the one which satisfies four axioms: Pareto optimality, Invariance under affine transformation, Symmetry, and Independence of irrelevant alternatives. Nash proved that there is a unique point in P which satisfies these axioms, and moreover this point $(\mathbf{u} \in P)$ is the one that maximizes $\prod_{i \in A} (u_i - c_i)$ or equivalently $\sum_{i \in A} \log(u_i - c_i)$.

The class Linear Nash Bargaining Games (LNB), defined in [8], consists of games whose feasible set P is defined by a finite number of linear packing constraints. The main focus of our paper will be on a natural subclass of LNB called Uniform Nash bargaining games (or UNB) which was also defined in [8]. In these games, the coefficients of the variables in the linear packing constraints are either 0 or 1. Clearly there can be at most $2^{|A|}$ such constraints, thus a function of the form $v : 2^A \rightarrow R^+$ uniquely encodes a feasible set in UNB games.

Now given a disagreement point \mathbf{c} , and a fixed set of agents $T \subseteq A$, the solution to a UNB game can be captured by the following convex program:

$$\begin{aligned}
 & \max \sum_{i \in T} \log(u_i - c_i) \\
 & \text{s.t. } \forall S \subset T : \sum_{i \in S} u_i \leq v(S) \\
 & \forall i \in T : u_i \geq 0
 \end{aligned} \tag{1}$$

For a fixed function $v : 2^A \rightarrow R^+$, we will define a family of games $F(v)$ to be the set of all Nash bargaining games for various choices of disagreement points \mathbf{c} and set $T \subseteq A$. An instance $(\mathbf{c}, T) \in F(v)$ will refer to a particular Nash bargaining game in $F(v)$ with a fixed set T and disagreement point \mathbf{c} . A **UNB** game is called an **SNB** if the function v is a submodular function.

We will assume that the following two natural conditions are satisfied by the function v :

1. **Non degenerate:** $v(\emptyset) = 0$.
2. **Non redundancy of sets:** $\forall S \subseteq A$, there exists a feasible utility vector u such that set S is *tight* w.r.t. u , i.e. $\sum_{i \in S} u_i = v(S)$.

We will call such functions to be *valid* functions. Note that the second condition above implies 1) *Monotonicity*: for any $Z_1 \subset Z_2 \subseteq A$, we have $v(Z_1) \leq v(Z_2)$, and 2) *Complement freeness*: $v(Z_1 \cup Z_2) \leq v(Z_1) + v(Z_2)$.

In this paper, we are interested in the following three game theoretic properties of UNB games:

Price of Bargaining: For any valid function $v : 2^A \rightarrow R^+$, we define the *Price of bargaining* of $F(v)$ to be $\min_{(\mathbf{c}, T) \in F(v)} \frac{u(\mathbf{c}, T)}{v(T)}$, where $u(\mathbf{c}, T)$ is the total utility obtained by set T of agents in the bargaining solution of the instance (\mathbf{c}, T) .

Full competitiveness: For any valid function $v : 2^A \rightarrow R^+$, we say that $F(v)$ is *fully competitive* if, for all games in $F(v)$, the following property holds: On increasing the disagreement utility c_i of an agent i , the bargaining solution doesn't increase the utility for any other agent j , where $j \neq i$.

Fairness: For any instance $I = (\mathbf{c}, T) \in F(v)$, define $core(I)$ to be the set of all feasible Pareto optimal solutions. For any vector \mathbf{u} , let \mathbf{u}_{dec} be the vector obtained by sorting the components of \mathbf{u} in decreasing order. A vector \mathbf{x} *min-max dominates* \mathbf{y} if \mathbf{x}_{dec} is lexicographically smaller than \mathbf{y}_{dec} . Also let \mathbf{u}^* be bargaining solution of instance I . Instance I is said to be *min-max fair* if the vector $\mathbf{u}^* - \mathbf{c}$ *min-max dominates* $\mathbf{y} - \mathbf{c}$ for all $\mathbf{y} \in core(I)$. $F(v)$ is said to be *min-max fair* if all the instances in $F(v)$ are min-max fair. Similarly we define the notion of *max-min fairness*.

Main results of this paper are described in theorems [1](#), [2](#), and [3](#).

3 Preliminaries

For any valid function v , we say that S is *tight* w.r.t. \mathbf{u} if $\sum_{i \in S} u_i = v(S)$. Let \mathbf{u}^* be the solution to the convex program given in [\(1\)](#). Then by KKT conditions, there must exist variables $\{p_S, \forall S \subseteq T\}$ such that:

1. $\forall S \subseteq T, p_S \geq 0$.
2. $\forall S \subseteq T, p_S > 0 \Rightarrow \mathbf{u}^*$ makes set S tight.
3. $\forall k \in T$, we have $\sum_{S:k \in S} p_S = \frac{1}{u_k^* - c_k}$.

We will call p_S to be the price of set S .

Now we give some properties of the submodular and non-submodular functions which will be used in our proofs.

Property 1. Given a valid submodular function $v : 2^A \rightarrow R_+$, and a utility vector \mathbf{u} , if $Z_1, Z_2 \subseteq A$ are tight sets w.r.t. \mathbf{u} , then $Z_1 \cup Z_2$ and $Z_1 \cap Z_2$ are also tight sets w.r.t. \mathbf{u} .

By using the uncrossing argument and the above property, we get the following corollary.

Corollary 1. *Given any SNB instance specified by v, \mathbf{c} and T , we can choose the prices for all subsets of T in the KKT conditions, such that the tight sets with positive prices form a nested set family, i.e. $T = T_1 \supset T_2 \supset \dots \supset T_k \supset T_{k+1} = \emptyset$.*

Also, we will use the following property of non-submodular functions which is similar to the one given in [1]. Proof is given in the full version.

Property 2. Given a valid non-submodular function $v : 2^A \rightarrow R_+$, there exists set $S \subseteq A, i, j \in A \setminus S, l \in S$ and a feasible utility vector \mathbf{u} such that:

1. $S \cup \{i\}, S \cup \{j\}$ are both tight w.r.t. \mathbf{u} .
2. Let $T = S \cup \{i, j\}, \mathcal{F}_k = \{Z \subseteq T : k \in Z, \text{ and } Z \text{ is tight w.r.t. } \mathbf{u}\}$. Then following holds

$$\mathcal{F}_l = \mathcal{F}_i \cup \mathcal{F}_j, \quad \mathcal{F}_i \cap \mathcal{F}_j = \emptyset.$$

3. $u_k > 0, \forall k \in T$.

4 Price of Bargaining

Theorem 1. *For any valid function $v, F(v)$ has Price of bargaining equal to 1 if and only if v is submodular.*

Proof. \Leftarrow : Suppose v is submodular. We want to show that for any disagreement point \mathbf{c} , and set $S \subseteq A$, if we restrict to the subproblem among agents in S , the Nash bargaining solution \mathbf{u}^* satisfies $\sum_{i \in S} u_i^* = v(S)$.

Since \mathbf{u}^* is the solution of Nash bargaining game, it must be Pareto optimal. Therefore every agent i is in some tight set T_i . Therefore by Property 1, we have $S = \cup_{i \in S} T_i$ is also tight, which means $\sum_{i \in S} u_i^* = v(S)$.

\Rightarrow : Suppose v is not submodular. By Property 2, there is a set $T = S \cup \{i, j\}$ and a feasible utility vector $\mathbf{u} = (u_k)_{k \in T}$ such that: (1) $u_k > 0, \forall k \in T$, (2) $S \cup i$ and $S \cup j$ are tight w.r.t. \mathbf{u} , (3) T is not tight w.r.t. \mathbf{u} . This is obtained from $\mathcal{F}_i \cap \mathcal{F}_j = \emptyset$.

Now for any $k \in T, k$ is in some tight set w.r.t \mathbf{u} , hence by the lemma 1 below (proof in full version), there exists \mathbf{c} such that \mathbf{u} is the Nash bargaining solution corresponding to \mathbf{c} .

By condition 3 above, we have $\sum_{k \in T} u_k < v(T)$, which implies the Price of bargaining is strictly less than 1.

Lemma 1. *Given any valid function v , an instance (\mathbf{c}, T) in $F(v)$, and a utility vector \mathbf{u} with $u_i > 0, \forall i \in T, \mathbf{u}$ is Pareto optimal if and only if there exists a vector \mathbf{c} , with $c_i > 0 \forall i \in T$, such that \mathbf{u} is the bargaining solution for the instance (\mathbf{c}, T) .*

5 Full Competitiveness

Theorem 2. *For any valid function v , $F(v)$ is fully competitive if and only if v is submodular.*

Proof. \Leftarrow : We first describe the algorithm for finding the optimal solution to the convex program (II) when the function v is a submodular function. Let $\mathbf{y} := \mathbf{u} - \mathbf{c}$. Then, an equivalent convex program is $\max\{\sum_i \log y_i : y(S) \leq f(S); \mathbf{y} \geq \mathbf{0}\}$, where $f(S) := v(S) - c(S)$. Call a set tight if $y(S) = f(S)$. Call an agent free if w.r.t the current \mathbf{y} it is not in any tight set. The algorithm maintains a set of tight sets \mathcal{T} initially empty. For all agents i which are free increase y_i simultaneously until some new set X gets tight. If X intersects with any set in \mathcal{T} , then since v is submodular, their union must be tight. Pick X to be the maximal (inclusion-wise) tight set and put it in \mathcal{T} . Continue till \mathcal{T} (the set of all agents) becomes tight. We have the following lemma (Proof in the full version):

Lemma 2. *The utility allocation returned by the above algorithm is an optimal solution to the convex program.*

Now we prove that SNB games are fully competitive. Suppose the disagreement of agent i goes from c_i to $c_i + \delta$. Call the new disagreement vector \mathbf{c}' . Let $f'(S) := v(S) - c'(S)$ for all S . To show full competitiveness, it suffices to show that the optimum, \mathbf{y}' of the convex program $\max\{\sum_i \log y_i : y(S) \leq f'(S); \mathbf{y} \geq \mathbf{0}\}$ is dominated by \mathbf{y} , the solution to the original convex program with $f(\cdot)$. We will use the continuous time algorithm above to prove this.

Firstly, note that $f'(S) = f(S)$ for all sets not containing i and $f'(S) = f(S) - \delta$ for all others. This implies, that there is at least one agent j with $y'_j < y_j$. Secondly, observe from the description of the algorithm that for any agent j with $y'_j < y_j$, there must be a corresponding tight set in \mathcal{T}' which contains both j and i .

We now show that if an agent j became non-free at time t in the original run (which means $y_j = t$), then by time t it must be in a tight set in the new run. We do so by showing that at time t if $y'_j = t$, then some set containing j at that time is tight (or over-tight which would imply $y'_j < t$).

Let A be the set containing j which went tight in the original run of the algorithm. Consider the set A in the new run of the algorithm at time t . Let $Q := \{j \in A : y'_j < y_j\}$. Note that if $j \in Q$, we are done. Assume $j \notin Q$. By the second observation made above and using the submodularity of v (to show union of intersecting tight sets is tight), we know there must exist a set Z which contains both Q and i , and which is tight. That is, $y'(Z) = f'(Z) = f(Z) - \delta$. We claim that $y'(Z \cup A) \geq y'(Z \cup A)$ and thus we are done.

This is because

$$\begin{aligned} y'(Z \cup A) &= y'(A \setminus Q) + y'(Z) \geq y(A \setminus Q) + f'(Z) = y(A) - y(Q) + f'(Z) \\ &\geq f(A) - f(Q) + f(Z) - \delta \geq f(A \cup Z) - \delta = f'(A \cup Z) \end{aligned}$$

The first inequality follows from definition of Q , the second from the tightness of A under y and feasibility of y and the last follows from submodularity of f .

⇒: Suppose v is not submodular, then by property [2](#) there must exist a set S and agents $i, j \in A \setminus S, l \in S$, and a feasible utility vector \mathbf{u} such that: (1) $S \cup \{i\}, S \cup \{j\}$ are both tight w.r.t. \mathbf{u} , (2) $\mathcal{F}_l = \mathcal{F}_i \cup \mathcal{F}_j, \mathcal{F}_i \cap \mathcal{F}_j = \emptyset$, (3) $u_k > 0, \forall k \in T$, where $T = S \cup \{i, j\}$.

We will now construct an instance $(c, T) \in F(v)$ which is not fully competitive. Let $\delta = \min_{k \in T} u_k > 0$. For tight sets $S \cup \{i\}, S \cup \{j\}$, we set their prices to be $p_{S,i}, p_{S,j}$ respectively, where $p_{S,i} = p_{S,j} = P = \frac{2}{\delta}$. For any other set $Z \subseteq T$, we set its price p_Z to be zero.

Let

$$\forall k \in T \quad c_k = u_k - \frac{1}{\sum_{Z \subseteq T, k \in Z} p_Z}$$

Since $S \cup \{i\}$ and $S \cup \{j\}$ are both tight, so for any $k \in T$, there exist at least one $Z \subseteq S$ such that $p_Z = P$, and we have

$$c_k \geq u_k - \frac{\delta}{2} > 0.$$

By the definition of \mathbf{c} , all the KKT conditions hold, thus \mathbf{u} is the bargaining solution w.r.t. (\mathbf{c}, T) .

We will now construct a \mathbf{c}' and its corresponding bargaining solution \mathbf{u}' , such that: (1) $\forall k \in T, k \neq j, c'_k \geq c_k$, and (2) $c'_j = c_j$ and $u'_j > u_j$.

Note that if the above conditions hold, then we can show that there exists a game in $F(v)$ which is not fully competitive. This is because \mathbf{c}' can be obtained from \mathbf{c} by increasing only the coordinates other than j . If $F(v)$ is fully competitive, then each time a coordinate of \mathbf{c} is increased utility allocated to j shouldn't increase. But if $u'_j > u_j$ is true then we get a contradiction.

Now we give the construction of \mathbf{u}' and \mathbf{c}' . Let \mathbf{u}' equals \mathbf{u} except that $u'_j = u_j + \epsilon, u'_i = u_i + \epsilon, u'_l = u_l - \epsilon$. Using arguments similar to the proof of property [2](#), one can show that there exists small enough ϵ (given below) so that \mathbf{u}' is feasible.

$$\epsilon < \min\{\epsilon_0, u_l/2\}, \text{ where } \epsilon_0 := \min_{\text{non-tight } Z \subseteq T} \frac{(v(Z) - \sum_{k \in Z} u_k)}{2}.$$

Now we construct \mathbf{c}' , so that it satisfies the condition mentioned above and the KKT conditions.

Note that for the KKT conditions, if we only assign positive price to tight sets $S \cup i, S \cup j$, say $p'_{S,i}$ and $p'_{S,j}$ respectively, then \mathbf{u}', \mathbf{c}' satisfy the KKT conditions and the above requirements iff

- $c'_i = u'_i - \frac{1}{p'_{S,i}} = u_i + \epsilon - \frac{1}{p'_{S,i}} \geq u_i - \frac{1}{p_{S,i}} = c_i$;
- $c'_j = u'_j - \frac{1}{p'_{S,j}} = u_j + \epsilon - \frac{1}{p'_{S,j}} = u_j - \frac{1}{p_{S,j}} = c_j$;
- $c'_l = u'_l - \frac{1}{p'_{S,i} + p'_{S,j}} = u_l - \epsilon - \frac{1}{p'_{S,i} + p'_{S,j}} \geq u_l - \frac{1}{p_{S,i} + p_{S,j}} = c_l$;
- $c'_k = u'_k - \frac{1}{p'_{S,i} + p'_{S,j}} = u_k - \frac{1}{p'_{S,i} + p'_{S,j}} \geq u_k - \frac{1}{p_{S,i} + p_{S,j}} = c_k, \forall k \neq l, k \in S$.

The above conditions can be reduced to the following:

$$p'_{S,j} = \frac{1}{\epsilon + \frac{1}{p_{S,j}}}, p'_{S,i} \geq \frac{1}{\frac{1}{p_{S,i}+p_{S,j}} - \epsilon} - p'_{S,j}$$

This can be satisfied as long as $\epsilon < \frac{1}{\frac{1}{p_{S,i}+p_{S,j}}} = \frac{\delta}{4}$.

To sum up, by setting $\epsilon = \min\{\epsilon_0/2, \delta/8\}$, we can find $p'_{S,i}, p'_{S,j}$ such that:

$$p'_{S,j} = \frac{1}{\epsilon + \frac{1}{p_{S,j}}}, p'_{S,i} \geq \frac{1}{\frac{1}{p_{S,i}+p_{S,j}} - \epsilon} - p'_{S,j}$$

Note that this value of ϵ is consistent with the previous mentioned upper bound on it. Therefore, we can construct \mathbf{c}' such that \mathbf{u}' is the bargaining solution w.r.t. \mathbf{c}' and $c'_k \geq c_k, \forall k \in T, c'_j = c_j$. Thus $(c, T) \in F(v)$ is not fully competitive.

6 Fairness

Theorem 3. *For any valid function v , $F(v)$ is min-max and max-min fair if and only if v is submodular.*

Proof. \Leftarrow : Suppose v is submodular. let \mathbf{u}^* be the Nash bargaining solution for (\mathbf{c}, T) where $T \subseteq A$. By corollary [1](#), we can choose the prices such that the tight sets w.r.t \mathbf{u}^* with positive price form a nested set family, $T = T_1 \supset T_2 \supset \dots \supset T_t \supset \emptyset$.

Pick any element \mathbf{g} in $core((\mathbf{c}, T))$ i.e. \mathbf{g} is Pareto optimal. If \mathbf{u}^* does not min-max dominate \mathbf{g} , then \mathbf{g} min-max dominates \mathbf{u}^* . In this case we will show that $\mathbf{g} = \mathbf{u}^*$, which leads to a contradiction.

Since \mathbf{g} is Pareto optimal therefore every agent is in some tight set w.r.t \mathbf{g} . Hence by property [1](#), T_1 is tight, i.e. $\sum_{k \in T_1} g_k = v(T_1)$. Since \mathbf{g} is feasible, we have

$$\sum_{k \in T_2} g_k \leq v(T_2)$$

T_1 and T_2 are tight sets w.r.t \mathbf{u}^* , so we have

$$\sum_{k \in T_1 - T_2} g_k \geq \sum_{k \in T_1 - T_2} u_k^* \tag{2}$$

Since each agent i in $T_1 - T_2$ has the highest $u_i - c_i$ among all the agents, if \mathbf{g} min-max dominates \mathbf{u}^* , then for any $k \in T_1 - T_2$, we have $g_k \leq u_k^*$. Then by [\(2\)](#), we have $g_k = u_k^*, \forall k \in T_1 - T_2$. Then we can use induction to show for any $1 \leq i \leq t$ and any $k \in T_i - T_{i+1}$, $g_k = u_k^*$. Hence $\mathbf{g} = \mathbf{u}^*$.

This proof also shows that \mathbf{u}^* is the unique min-max fair utility vector. By using an argument similar to [\[2\]](#), we can show that any unique min-max fair utility vector is also max-min fair.

\Rightarrow : Suppose v is not submodular, then by property [2](#), there is a set $T = S \cup \{i, j\}$ and a $\mathbf{g} = (g_k)_{k \in T}$ such that: (1) $g_k > 0, \forall k \in T$, (2) $S \cup \{i\}$ and $S \cup \{j\}$ are tight w.r.t \mathbf{g} , (3) $\mathcal{F}_l = \mathcal{F}_i \cup \mathcal{F}_j, \mathcal{F}_i \cap \mathcal{F}_j = \emptyset$.

For each $k \in T$, let $c_k = g_k - \epsilon$, where $0 < \epsilon < \min_{k \in T} \{g_k\}$. Clearly \mathbf{g} is a feasible core element corresponding to \mathbf{c} , since each k is in a tight set (either $S \cup \{i\}$ or $S \cup \{j\}$).

Let \mathbf{u}^* be the Nash bargaining solution corresponding to (\mathbf{c}, T) . Since \mathbf{g} is the unique min-max and max-min feasible solution, thus if \mathbf{u}^* min-max and max-min dominates \mathbf{g} , then \mathbf{u}^* must equal \mathbf{g} . Next we show that this is not possible.

Suppose $\mathbf{u}^* = \mathbf{g}$, by KKT conditions, we can price all the subsets of T such that:

$$\frac{1}{g_l - c_l} = \sum_{Z \in \mathcal{F}_l} p_Z = \sum_{Z \in \mathcal{F}_i} p_Z + \sum_{Z \in \mathcal{F}_j} p_Z = \frac{1}{g_i - c_i} + \frac{1}{g_j - c_j}$$

which contradicts the fact that $g_l - c_l = g_i - c_i = g_j - c_j = \epsilon$.

References

1. Chakrabarty, D., Devanur, N.: On competitiveness in uniform utility allocation markets. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 374–380. Springer, Heidelberg (2007)
2. Jain, K., Vazirani, V.V.: Equitable cost allocations via primal-dual-type algorithms. In: STOC (2002)
3. Kalai, E.: Solutions to the bargaining problem. In: Hurwicz, L., Schmeidler, D., Sonnenschein, H. (eds.) Social Goals and Social Organization, pp. 75–105. Cambridge University Press, Cambridge (1985)
4. Koutsoupias, E., Papadimitriou, C.H.: Worst case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563. Springer, Heidelberg (1999)
5. Nash, J.F.: The bargaining problem. *Econometrica* 18, 155–162 (1950)
6. Osborne, M., Rubinstein, A.: A Course in Game Theory. MIT Press, Cambridge (1994)
7. Thomson, W., Lensberg, T.: Axiomatic Theory of Bargaining With a Variable Population. Cambridge University Press, Cambridge (1989)
8. Vazirani, V.V.: Nash bargaining via flexible budget markets (manuscript, 2008)
9. von Neumann, J.L., Morgenstern, O.: Theory of games and economic behavior (1944)

Computing an Extensive-Form Correlated Equilibrium in Polynomial Time

Wan Huang and Bernhard von Stengel

Department of Mathematics, London School of Economics,
London WC2A 2AE, United Kingdom
wan.huang@gmail.com, stengel@nash.lse.ac.uk

Abstract. We present a polynomial-time algorithm for finding one *extensive form correlated equilibrium* (EFCE) for multiplayer extensive games with perfect recall. This is the first such algorithm for an equilibrium notion for games of this generality. The EFCE concept has been defined by von Stengel and Forges [1]. Our algorithm extends the constructive existence proof and polynomial-time algorithm for finding a correlated equilibrium in succinctly representable games by Papadimitriou and Roughgarden [23]. We describe the set of EFCE with a polynomial number of consistency and incentive constraints, and exponentially many variables. The algorithm employs linear programming duality, the ellipsoid algorithm, and Markov chain steady state computations. We also sketch a possible interpretation of the variables in the dual system.

1 Introduction

Extensive games with perfect recall are a fundamental model of noncooperative game theory. They are game trees where players may have imperfect information about the game state, modeled by information sets [4]. The standard rationality assumption of *perfect recall* is a condition on the information sets that asserts that a player never forgets what he knew or did earlier.

The game tree, with its information sets, moves, chance probabilities, and payoffs, is a *succinct* representation of a game. The *strategic form* of the game is in general exponentially larger because a pure strategy of a player is a tuple of moves, one for each information set, so there are exponentially many strategies per player; in the terminology of [23], this means the game is *not* of “polynomial type”. Already for zero-sum two-player games, finding an equilibrium is therefore an interesting computational problem. It is solved by the *sequence form* [5], which is a strategic description of the same size as the game tree, and allows to solve huge two-person zero-sum games, for example of poker (see [6], also for earlier references related to the sequence form).

Finding a Nash equilibrium of an extensive game with any number of players is as difficult as for a game in strategic form. For the latter, a (more general) *correlated equilibrium* (CE) [7] can be found in polynomial time. Papadimitriou and Roughgarden [23] give a polynomial-time algorithm for succinctly representable games. Applying the ellipsoid method [8] to a linear program derived from the existence proof for CE due to [9] and [10], the method generates a polynomial-sized LP. The solution to that LP gives a distribution on product distributions that describes the desired CE for the

succinctly represented game. However, this method cannot be applied to extensive games because they are not of polynomial type.

We present the *first polynomial-time algorithm for finding an equilibrium for general extensive games* that have any number of players and perfect recall. We have to adapt the concept of correlated equilibrium, but preserve the spirit of the extensive game in the sense that any “uncorrelated” such equilibrium is a Nash equilibrium (this is not the case for the *agent form*, where [23] can be applied, as discussed in [1]). We consider the new concept of *extensive form correlated equilibrium* (EFCE) [1]. A strategic-form CE can be considered as a device that selects a profile of pure strategies from a joint probability distribution and sends each player privately his pure strategy in that profile as a recommendation of what to play. In an EFCE, such a profile is also selected by the device before the game starts. However, the recommendations in an EFCE are “delayed”: rather than telling the entire strategy in advance, each recommended *move* is given only when the player reaches the information set where she can make that move.

For two-player perfect recall extensive game without chance moves, [1] give a polynomial-size system that describes the set of all EFCE. Hence, any solution to that system can be found in polynomial time. However, already in two-player games with chance moves, it is not possible to give a polynomial-sized description of the set of all EFCE, unless $P = NP$ [11,12,1].

In this paper, we use a description of the set of EFCE for general extensive games with any number of players with perfect recall, and possible chance moves. We describe the set of EFCE by a polynomial number of constraints, but with exponentially many variables, which allows us to extend the method of [23].

In an EFCE, unlike in a CE, a player has only to decide whether a recommended *move* is optimal, which involves a small number of comparisons with the other moves at the respective information set. In addition, a player who considers deviating from a recommended move has to look at additional future moves at previously unreached own information sets. We represent these by suitable incentive constraints with variables and constraints that, essentially, mimic “dynamic programming” in a single-player decision tree. The resulting system is small but somewhat involved, because one has to study carefully the propagation of payoffs for possible deviations through the game tree via dual variables.

Our blueprint, the proof by [23], uses the constructive existence proof of CE of [9,10], and employs linear programming duality, the ellipsoid algorithm, Markov chain steady state computations, as well as application specific methods for computing expectations over product distributions. We run the ellipsoid algorithm on the, more complicated, dual system. In contrast to the computation of a CE for a succinct strategic-form game [23], the dual system contains additional “consistency” information in the form of certain equalities for the dual variables. At each step, the ellipsoid algorithm finds a violated convex combination of the constraints of the dual system. These correspond, dually, to product distributions on moves at information sets. These do *not* represent steady states of a Markov chain, unlike in [23]. However, by assuming that the required distribution exists for preceding information sets, it can be considered as a steady state. For this reason, our algorithm uses Markov chain computation at each information set from the root of the game tree down to the information sets closest to the leaves of the game tree.

We give our construction first for perfect-recall extensive games without chance moves, and then consider chance as an extra player who gets no payoff and never deviates. For reasons of space, we have to omit all proofs, and, unfortunately, examples, which are available in a long version of this paper.

2 Incentive Constraints

We use the notation of [1], except that a player is typically denoted by p and N is the set of all players, often omitted. The set of all information sets of player p is H_p . For $h \in H_p$, the set of moves or choices at h is C_h . A pure strategy s_p of player p is an element of $\prod_{h \in H_p} C_h$, and a strategy profile s is an element of $\prod_p \prod_{h \in H_p} C_h$.

In an EFCE, a player receives a *move* recommendation when reaching an information set, unlike in a CE where a player gets a recommended strategy at the beginning of the game. The player then compares the expected payoffs of moves at that information set and chooses the move with maximum expected payoff.

We first consider games without chance moves. Given an extensive game with perfect recall, is there a “sequence form” to compute one EFCE? For two-player games without chance moves, the answer is yes, as described in detail in [1]. The system of consistency and incentive constraints of sequences defines the set of EFCE. This holds because the condition of perfect recall imposes strong restrictions on the players’ information sets, so that the recommended move at each information set can be generated uniquely. However, for games of more than two players, or with chance moves, the consistency constraints of the sequences on the marginal probabilities of moves that are correlated across information sets are only necessary conditions. For this reason, our system does not use the sequence form.

Therefore, instead of introducing an auxiliary variable $u(\sigma)$ to denote the expected payoff contribution of a sequence σ as in [1], we use $u(c)$ to denote the expected payoff contribution of the move c when the player follows his recommendations at all information sets he reaches. Before giving the expression of $u(c)$, the relations between moves and strategies need to be clarified.

Definition 1. *An information set $h \in H_p$ precedes another information set $k \in H_q$ if and only if $p = q$ and there are nodes $u \in h$ and $v \in k$ such that u is on the path from the root to v . Furthermore, $h \in H_p$ immediately precedes $k \in H_p$ when h precedes k and there is no information set $l \in H_p$ that succeeds h and precedes k .*

Unlike in [1], the relation “precedes” is only between two information sets of the same player.

Definition 2. *A move $c \in C_h$ agrees with a strategy profile s if and only if the information set h is reached and the move c is chosen when players play according to s . A move $c \in C_h$ of player p terminates a strategy profile s_{-p} if and only if h is reached and no further information set of player p is reached if p plays the moves leading to h at all preceding information sets and plays c at h , and the other players play s_{-p} . An information set $h \in H_p$ is reachable by a strategy profile $s \in S$ if and only if the player p reaches h at a certain stage if all players choose the strategies in s .*

Let $a^p(s)$ be the payoff to player p if all players choose the strategies in s , and let $z(s)$ be the probability according to which the correlation device selects the strategy profile s . For any $c \in C_p$, the variable $u(c)$ is given by

$$u(c) = \sum_{s \in S : c \text{ agrees with } s} a^p(s)z(s). \tag{1}$$

Thus for any move c so that no further information set of player p is reached afterwards, $u(c)$ is the expected payoff to player p if he plays the recommended move c . The following lemma shows that $u(c)$ is the expected payoff contribution also for a move c that leads to further information sets of player p .

Lemma 3. *Given a move $c \in C_h$ and the set $\{k \in H_p : h \text{ precedes } k\}$ is not empty, we have*

$$u(c) = \sum_{\substack{s \in S : c \text{ agrees with } s, \\ c \text{ terminates } s_{-p}}} a^p(s)z(s) + \sum_{\substack{l \in H_p : h \text{ immediately} \\ \text{precedes } l \text{ via } c}} \sum_{c' \in C_l} u(c').$$

The expected payoff $u(c)$ when the player chooses the recommended move c is compared with the possible payoff when the player deviates from his recommendation. Given a move $c \in h$ and an information set k such that $k = h$ or k succeeds h , we use $v(k, c)$ to denote the optimal expected payoff at k given the player is recommended move c at h . It is the maximum of the payoffs for the possible moves $d \in C_k$, which may either directly give a payoff $a^p(s_d^k)$ when d terminates s_{-p} (where s_d^k is the strategy profile that specifies moves leading to k at information sets preceding k , and d at k , and the same moves as s at all other information sets), or are obtained from subsequent optimal payoffs at later information sets. This is expressed by the following inequalities:

$$v(k, c) \geq \sum_{\substack{s \in S : c \text{ agrees with } s, \\ d \text{ terminates } s_{-p}}} a^p(s_d^k)z(s) + \sum_{\substack{l \in H_p : k \text{ immediately} \\ \text{precedes } l \text{ via } d}} v(l, c), \quad d \in C_k. \tag{2}$$

These incentive constraints are completed by

$$u(c) = v(h, c), \tag{3}$$

for any move $c \in C_h$. That is, given a recommended move c , the player does not gain by deviating from move c .

Theorem 4. *In a perfect-recall extensive game, a probability distribution z that fulfills for all players the incentive constraints (1), (2) and (3) defines an EFCE. The number of constraints that describe the set of EFCE is polynomial in the size of the game tree.*

3 Existence Proof

In the system describing the set of EFCE, (3) states that the expected payoff contribution of the recommended move must be optimal, as expressed by (1) and (2). One can

obviously substitute $v(h, c)$ with $u(c)$ when $c \in C_h$ in (2). We rewrite these simplified constraints as matrix inequalities and consider the linear program

$$\text{maximize } \sum_{s \in S} z(s), \quad \text{subject to } Az + Bv \geq 0, \quad z \geq 0. \tag{4}$$

So the entries of A are either 0 or linear terms of $a(s)$ for certain s and the entries of B , for (2) and (3), are either 0, 1 or -1 . The LP (4) is either trivial with the objective function being 0 or unbounded. When it is unbounded the normalized solution is an EFCE. Therefore by duality, to prove the existence of EFCE, it suffices to show that the dual of (4)

$$A^T y \leq -1, \quad B^T y = 0, \quad y \geq 0 \tag{5}$$

is always infeasible. We need the following lemma, analogous to [9][10][2].

Lemma 5. *If $y \geq 0$, $B^T y = 0$, then there is a product distribution z so that $z^T A^T y = 0$.*

Here, $z^T A^T y$ is a convex combination of left sides of the constraints $A^T y \leq -1$ in (5), and hence for every feasible $y \geq 0$, $B^T y = 0$, it should evaluate to something negative. Thus this lemma shows that (5) is infeasible.

The proof of Lemma 5 uses the following lemma which has a long but straightforward proof. There is one dual variable $y_{c,d}^k$ for each information set k and move $d \in k$ and $c \in h$ where $h = k$ or h precedes k . To prove that given $y \geq 0$ and $B^T y = 0$, there is a convex combination of components of $A^T y$ equal to 0, we first show how a component of $A^T y$ can be expressed in terms of the payoff a^p and the dual variable y .

Lemma 6. *Given a strategy profile $s \in S$ and $y \geq 0$ such that $B^T y = 0$,*

$$(A^T y)_s = \sum_p \sum_{k \in H_p} \sum_{\substack{h \in H_p : h = k \text{ or} \\ h \text{ precedes } k \text{ and } h \\ \text{is reachable by } s}} \sum_{d \in C_k} y_{c_s^h, d}^k [a^p(s^k) - a^p(s_d^k)] \tag{6}$$

where s^k is the strategy profile in S that specifies moves leading to k at information sets preceding k and the same moves as s at all other information sets (k may not be reached according to the moves s specifies at information sets of other players), and c_s^h is the move that s specifies at h .

The following lemma provides the main step to prove Lemma 5.

Lemma 7. *For any y such that $y \geq 0$ and $B^T y = 0$, there is a product distribution $z = \prod_{p \in N} \prod_{k \in H_p} z^k$ such that for any information set k , the probability distribution z^k on the moves d at k satisfies*

$$z^k(d)[\alpha^k(d) + \alpha^k(\emptyset)] = \sum_{c \in C_k} z^k(c)\beta_1^k(c, d) + \beta_2^k(d) \tag{7}$$

where for any $c \in C_k$,

$$\alpha^k(c) = \left[\prod_{l \in H^k} z^l(c_k^l) \right] \sum_{d \in C_k} y_{c,d}^k, \quad \alpha^k(\emptyset) = \sum_{h \in H_p : h \text{ precedes } k} \prod_{l \in H^h} z^l(c_h^l) \sum_{c \in C_h} z^h(c) \sum_{d \in C_k} y_{c,d}^k,$$

$$\beta_1^k(c,d) = \left[\prod_{l \in H^k} z^l(c_k^l) \right] y_{c,d}^k, \quad \beta_2^k(d) = \sum_{h \in H_p : h \text{ precedes } k} \prod_{l \in H^h} z^l(c_k^l) \sum_{c \in C_h} z^h(c) y_{c,d}^k.$$

Here H^k is the set of information sets l (of the same player as k) that precede k , and c_k^l is the unique move at l that leads to k .

Nau and McCardle [10] discussed “joint coherence” in noncooperative games, and thus gave a possible interpretation of the variables involved in both the primal and the dual system. Myerson [13] used this interpretation to obtain further properties of CE. For EFCE, we consider a certain *move transition matrix* T^k in order to prove Lemma 7 for each information set k . Any such move transition matrix for information set k can be interpreted as a random deviation plan for the player who will make a move at that information set. Each number $y_{c,d}^k$ in (6), where $d \in C_k$, represents the trend that player would deviate to the move d when c is recommended at this information set or some earlier stage of the game (and the player ignores any recommendation after getting c). More precisely, the trend that the player ignores a recommendation at some earlier stage and chooses move d at information set k is $\beta_2^k(d)$, and the trend that the player chooses d at information set k is $\sum_{c \in C_k} z^k(c) \beta_1^k(c,d) + \beta_2^k(d)$. On the other hand, the trend that the player would be getting recommended move d (he or she may ignore it) is $z^k(d) \alpha^k(c_k^d) + z^k(d) \alpha^k(\emptyset)$. One can then show that the deviation plan does not change the distribution on the player’s actions at the information set k .

The preceding lemmas require several pages of proofs in full detail. Lemma 5 then implies the existence of EFCE.

Theorem 8. *Every game of extensive form without chance moves has an EFCE.*

4 Algorithm for Games without and with Chance Moves

To find an EFCE in polynomial time, we follow [23] and apply the ellipsoid algorithm to the dual (5) of the system (4) that characterizes the set of EFCE. The LP (4) has polynomially many constraints and exponentially many variables. Thus for the dual (5) the opposite holds, which makes it suitable for the ellipsoid algorithm.

In each iteration of the ellipsoid algorithm, an extra step is needed to maintain the candidate solution y_i to satisfy the consistency constraints $B^\top y_i = 0$. At the initial iteration, for the system $B^\top y = 0, y = 0$, let \bar{y} be the free variables in y , and $y = \bar{B} \bar{y}$. Thus the system (5) is equivalent to

$$A^\top \bar{B} \bar{y} \leq -1, \quad \bar{y} \geq 0. \tag{8}$$

We apply the ellipsoid algorithm to the system (8).

Let $\bar{y}_0 = 0$ be the candidate for the initial iteration. Thus every constraint $A^\top \bar{B} \bar{y} \leq -1$ is violated. Any product distribution z satisfies $z^\top A^\top y = 0$. Choose any product distribution z_0 , and a violated inequality $z_0^\top A^\top y \leq -1$. At each iteration, the candidate \bar{y}_{i-1} is replaced by \bar{y}_i . Let $y_i = \bar{B} \bar{y}_i$. By Lemma 5, a product distribution z_i such that $z_i^\top A^\top y = 0$ can be found. Thus the inequality $(z_i^\top A^\top) \bar{B} \bar{y} \leq -1$ is violated. We proceed to the next step.

Since we know that (5) is infeasible, the algorithm will end up with recognizing the system as infeasible after polynomially many iterations. Thus when the algorithm halts, we have polynomially many candidate solutions y_i and for each y_i a corresponding product distribution z_i .

We now claim that a convex combination, denoted $Z^\top \xi$, of these product distributions can be found in polynomial time, such that the system $AZ^\top \xi + Bv \geq 0, \xi \geq 0$ is unbounded. When the ellipsoid algorithm is applied to (8), in each iteration the inequality $(z_i^\top A^\top \bar{B}) \bar{y} \leq -1$ is violated by \bar{y}_i . Let Z be the matrix where each row i is the product distribution z_i found by the ellipsoid algorithm. We consider the system of linear inequalities

$$[ZA^\top \bar{B}] \bar{y} \leq -1, \quad \bar{y} \geq 0. \tag{9}$$

Clearly, the number of variables of (9) is equal to that of (8), and is polynomial in the size of the game tree. Thus the ellipsoid algorithm is appropriate to (9) too. Apply it to (9). Let the initial candidate solution be $\bar{y}_0 = 0$. In each iteration i , the i th constraint of (9) $(z_i^\top A^\top \bar{B}) \bar{y} \leq -1$ is violated by the i th candidate solution \bar{y}_i . Thus the algorithm will determine that (9) is infeasible too. That is,

$$[ZA^\top] y \leq -1, \quad y = \bar{B} \bar{y}, \quad \bar{y} \geq 0$$

or equivalently

$$[ZA^\top] y \leq -1, \quad B^\top y = 0, \quad y \geq 0$$

is infeasible. The dual problem

$$\text{maximize } \sum_i (\xi_A)_i \quad \text{subject to } [AZ^\top] \xi_A + B \xi_B \geq 0, \quad \xi_A \geq 0 \tag{10}$$

is unbounded. Here (ξ_A, ξ_B) is a partition of the variable vector ξ .

For any feasible solution ξ of (10), ξ_A after normalization is a probability distribution on the set of strategy profiles. The product $Z^\top \xi_A$ is a convex combination of the rows of Z^\top , which are the product distributions that are computed at all the iterations of the ellipsoid algorithm. Thus the nonnegative constraints $\xi_A \geq 0$ are satisfied if and only if $Z^\top \xi_A \geq 0$. Let $z = Z^\top \xi_A$, and $v = \xi_B$. The system (10) becomes

$$\text{maximize } \sum_s z(s) \quad \text{subject to } Az + Bv \geq 0, \quad z \geq 0$$

which is the system that characterizes an EFCE. Therefore, $(z, v) = (Z^\top \xi_A, \xi_B)$ is a nontrivial solution to (5) when ξ is a nontrivial solution to (10). Furthermore, $z = Z^\top \xi_A$ is the desired EFCE.

So far, all the arguments and inductions are based on the assumption that there are no chance moves in the game. With chance moves, the system (4) is no longer appropriate, because a move may “agree” with more than one strategy profile. However, the impact of the chance moves on the reachability of an information set can be expressed by considering chance as an extra player 0, without any incentive constraints. The chance moves become part of a strategy profile, but their probabilities in the construction of product distributions will be constants rather than variables, with minor modifications of the algorithm for games without chance moves. We obtain the following result.

Theorem 9. *Every multi-player, perfect-recall extensive game, which may have chance moves, has an EFCE, which can be computed in polynomial time.*

The EFCE concept is crucial to limit the number of incentive constraints. It is an *open question* if one can find one (strategic-form) CE for extensive games, even with only two players, in polynomial time as well. Because of the exponential number of strategies for each player, it is not even clear if such a CE has a polynomial-sized description and certification of the equilibrium property (analogous to the NP property for a decision problem).

References

1. von Stengel, B., Forges, F.: Extensive form correlated equilibrium: Definition and computational complexity. *Math. Oper. Res.* (in press, 2008)
2. Papadimitriou, C.H.: Computing correlated equilibria in multi-player games. In: Proc. 37rd Annual ACM Symposium on Theory of Computing (STOC), pp. 49–56 (2005)
3. Papadimitriou, C.H., Roughgarden, T.: Computing correlated equilibria in multi-player games. *J. ACM* 55, 14 (2008)
4. Kuhn, H.W.: Extensive games and the problem of information. In: Contributions to the theory of games. *Annals of Mathematics Studies* 28, vol. 2, pp. 193–216. Princeton Univ. Press, Princeton (1953)
5. von Stengel, B.: Efficient computation of behavior strategies. *Games Econom. Behav.* 14, 220–246 (1996)
6. Gilpin, A., Sandholm, T.: Lossless abstraction of imperfect information games. *J. ACM* 54, 25 (2007)
7. Aumann, R.J.: Subjectivity and correlation in randomized strategies. *J. Math. Econom.* 1, 67–96 (1974)
8. Khachiyan, L.G.: A polynomial algorithm in linear programming. *Soviet Mathematics* 20, 191–194 (1979); (transl. of *Dokl. Akad. Nauk SSSR* 244, 1093–1096)
9. Hart, S., Schmeidler, D.: Existence of correlated equilibria. *Math. Oper. Res.* 14, 18–25 (1989)
10. Nau, R.F., McCardle, K.F.: Coherent behavior in noncooperative games. *J. Econom. Theory* 50, 424–444 (1990)
11. Chu, F., Halpern, J.: On the NP-completeness of finding an optimal strategy in games with common payoffs. *Internat. J. Game Theory* 30, 99–106 (2001)
12. von Stengel, B.: Computational complexity of correlated equilibria for extensive games. Research Report LSE-CDAM-2001-03, London School of Economics (2001)
13. Myerson, R.B.: Dual reduction and elementary games. *Games Econom. Behav.* 21, 183–202 (1997)

Homogeneous Interference Game in Wireless Networks

Joseph (Seffi) Naor^{1,*}, Danny Raz¹, and Gabriel Scalosub²

¹ Computer Science Department, Technion - Israel Institute of Technology, Haifa, Israel

{naor, danny}@cs.technion.ac.il

² Department of Computer Science, University of Toronto, Toronto, ON, Canada

scalosub@cs.toronto.edu

Abstract. We consider the problem of joint usage of a shared wireless channel in an interference-bound environment, and focus on a distributed setting where there is no central entity managing the various transmissions. In such systems, unlike other multiple access environments, several transmissions may succeed simultaneously, depending on spatial interferences between the different stations. We use a game theoretic view to model the problem, where the stations are selfish agents aiming at maximizing their success probability. We show that when interferences are homogeneous, system performance suffers an exponential degradation in performance at an equilibrium, due to the selfishness of the stations. However, when using a proper penalization scheme for aggressive stations, we can ensure the system's performance value is at least $1/e$ of the optimal value, while still being at equilibrium.

1 Introduction

Wireless networks often involve the joint usage of common communication channels in a multiple access environment. In most of the models capturing such settings, simultaneous transmission by more than one station results in a collision causing all transmissions at that time to fail. CSMA-type methods are usually used in such scenarios in order to deal with collisions, in an attempt to maximize the system's throughput. However, in many current wireless networks, such as mesh WiFi networks, or 802.15 clusters, simultaneous usage of the same wireless channel is possible. Consider, for example, the settings described in Fig. 1 where we outline two stations, A, B and their transmission ranges. If the clients of A and B are a and b respectively, then simultaneous transmissions will cause a collision at client a , while b can receive the message from B . However, if the clients of A and B are a' and b respectively, then simultaneous transmissions will both succeed, since they do not collide at either of the receiving ends.

In wireless networks where channel access need not be exclusive, one of the major optimization issues is the efficient use of radio resources. In this paper we consider the problem of joint usage of a common communication channel by a finite number of stations, where stations are always backlogged, i.e., always have a packet to send. We present a generalization of classic multiple access models by introducing the notion of *spatial interference parameters*, which capture the pairwise interferences between the stations contending for the common radio resource. It is important to notice that in

* Supported by ISF grant 1366/07 and BSF grant 2002276.




Fig. 1. Outline of two stations, A, B and their transmissions ranges

this model several transmissions may succeed simultaneously, and thus the commonly assumed upper bound of one on the overall throughput of the system no longer holds. The overall number of successful transmissions at any time can take any value between 0 and n , where n is the number of stations in the network. The exact value depends on the inter-station interferences.

As a preliminary step in understanding this model, we focus our attention on the case of homogeneous interferences, where every station inflicts the same amount of interference on any other station. Indeed, such homogeneous interferences will usually not provide an accurate modeling for real-life scenarios where different agents have different interference patterns. Yet, we believe that a better understanding of the restricted settings (where the system is described by a single parameter) is both interesting and serves as an important step toward providing insights into understanding more general non-homogeneous environments.

Model. We model our problem as a game played by selfish agents. We consider a system consisting of n agents using a common wireless medium. For every agent i , we let $S = [0, 1]$ be the strategy space of agent i , and let $R_i \in S$ denote a strategy chosen by agent i . We refer to R_i as the *probability* that agent i transmits. Due to interferences, the probability of a *successful* transmission also depends upon the transmission of other agents. Given a profile $\bar{R} = (R_1, \dots, R_n) \in [0, 1]^n$, we define the *success probability* of agent i 's transmission as:

$$r_i(\bar{R}) = R_i \cdot \prod_{j \neq i} (1 - \alpha_{i,j} R_j),$$

where for every $1 \leq i, j \leq n$, $\alpha_{i,j} \in [0, 1]$ is a fixed network-dependent parameter denoting the amount of interference inflicted on i upon simultaneous transmission of both i and j . One way of thinking about the $\alpha_{i,j}$ -s is by viewing them as the probability that a transmission by j will interfere with a transmission of i . Given a profile $\bar{R} = (R_1, \dots, R_n)$, the *social welfare* $\varphi(\bar{R})$ (also referred to as the *throughput* of the system) is considered to be the overall use of resources in the system, i.e.

$$\varphi(\bar{R}) = \sum_{i=1}^n r_i(\bar{R}) = \sum_{i=1}^n R_i \prod_{j \neq i} (1 - \alpha_{i,j} R_j).$$

$\varphi(\bar{R})$ can be interpreted as the expected number of successful transmissions.

Note that a-priori, $\varphi(\bar{R})$ can take any value between 0 and n . For example, if for all i, j , $\alpha_{i,j} = 0$, i.e., there are no interferences, then for the profile where $R_i = 1$ for all i

we obtain $\varphi(\bar{R}) = n$, which implies an optimal use of resources. On the other hand, if for all i, j , $\alpha_{i,j} = 1$, then the same profile obtains zero-throughput, i.e., $\varphi(\bar{R}) = 0$. In the latter case our model coincides with classic multiple access models, simultaneous transmissions result in a collision, causing all transmissions to fail.

We refer to the interferences as *homogeneous* if there exists some $\alpha \in (0, 1)$ such that for all i, j , $\alpha_{i,j} = \alpha$. By the above observations, when considering homogeneous interferences, we restrict our attention to the case where for all i, j , $\alpha_{i,j} = \alpha \in (0, 1)$. In what follows we refer to a profile \bar{R} as *uniform*, if $R_i = R_j$ for all i, j .

For every agent i , we let $U_i(\bar{R})$ be the utility function of agent i , assuming agents play profile \bar{R} . In the following sections we consider several choices for these utility functions, and discuss the system's performance where agents are selfish, and aim at maximizing their own utility, regardless of the effect their choices have on the overall social welfare. We refer to the above setting as the *homogeneous interferences multiple-access (HIMA) game*.

Given any profile $\bar{R} = (R_1, \dots, R_n)$, we let \bar{R}_{-i} denote the subprofile defined by strategies of all agents except for agent i . We further let (\bar{R}_{-i}, R'_i) be the profile where every agent other than i plays the same strategy as in \bar{R} , while agent i plays strategy R'_i . Profile \bar{R} is said to be a *Nash equilibrium* (NE) if for every i , and every $R'_i \in [0, 1]$, $U_i(\bar{R}) \geq U_i(\bar{R}_{-i}, R'_i)$. Intuitively, a profile is an NE if no agent can increase its benefit by unilaterally deviating from his choice. We let $\bar{R}_{\text{NE}}^{(n)}$ denote an NE profile for n agents, and use $\bar{R}_{\text{OPT}}^{(n)}$ to denote any profile for n agents which maximizes the social welfare. Assuming an NE exists, we use the notion of *Price of Anarchy (PoA)* [1] in order to evaluate this effect, defined by the supremum over all NEs $\bar{R}_{\text{NE}}^{(n)}$ of the ratio between $\varphi(\bar{R}_{\text{OPT}}^{(n)})$ and $\varphi(\bar{R}_{\text{NE}}^{(n)})$, capturing the performance of the worst case equilibrium. We further consider the notion of *Price of Stability (PoS)* [2,3], defined by the infimum of the above ratio over all NEs, capturing the performance of the best case equilibrium.

Our Contribution. We study the rational choices of agents in an HIMA game, and analyze the performance of NE compared to the optimal performance. We focus on the case of homogeneous interferences, and show that when the utility of an agent is its success probability, then selfishness causes the system's performance to be up to an exponential factor away from the optimal performance. Specifically, we show that for any constant α , the price of anarchy as well as the price of stability are exponential in the number of agents, i.e., *any* equilibrium suffers an exponential degradation in performance. These results appear in Sec. 2.

We then turn to explore the effect of penalization, and to what extent does such an approach provide better system performance at a state of equilibrium. We show that there exists a penalty function which is proportional to the amount of aggressiveness demonstrated by an agent, such that for the case where the utility of an agent is the sum of its success probability and its penalty, then the price of stability with regards to the resulting *coordinated equilibria* can be made to drop to at most $e \approx 2.718$, thus demonstrating that an exponential improvement is possible compared to the uncoordinated case. We further show that for interferences which are not too large, namely, for $\alpha \leq 2/e \approx 0.735$, the price of anarchy is also bounded by e , thus ensuring that the degradation in performance due to the selfishness of the agents can be guaranteed to be

made very small. These results mean that if we impose these penalties upon the agents, either in the form of payment for transmission to the network operator, or considering them as an intrinsic cost suffered by the agent due to transmission (e.g., due to power consumption), then the performance can be dramatically improved compared to the general case where the agent’s utility is merely its success probability. These results are presented in Sec. 3. We note that our results for the homogeneous settings also extend to the finite horizon repeated HIMA game [4]. Due to space constraints, some of the proofs are omitted and may be found in [5].

Issues involving selfish behavior of agents in multiple access environments have received much attention in recent years. The slotted Aloha model in Markovian settings was studied in terms of stability conditions, and convergence to equilibrium (e.g., [6, 7, 8, 9], and references therein). Additional works considered rate control games in wireless networks, (e.g., [10]), and other recent work [11, 12, 13, 14, 15] has also considered the role of introducing costs for transmissions, and pricing schemes, and their effect on the stability of the system. Other models of interferences in wireless networks in Markovian settings are discussed in [16, 17, 18].

2 General Nash Equilibria

In this section we present several analytical results as to the effect of selfishness on the performance of the network, in the theoretical case where interferences are homogeneous, i.e., for every i, j , $\alpha_{i,j} = \alpha$, for some system’s parameter $\alpha \in (0, 1)$. We first consider the simple utility function $U_i(\bar{R}) = r_i(\bar{R})$, and show that in such a case, the system’s performance can be very far from optimal. Specifically, we prove the following:

Theorem 1. *Given n stations, and any $k \in \{1, \dots, n - 1\}$,*

1. *If $\alpha \in \left[\frac{1}{k+1}, \frac{1}{k}\right)$ then $\text{PoA}^{(n)} = \text{PoS}^{(n)} = \frac{k}{n(1-\alpha)^{n-k}}$.*
2. *If $\alpha \leq \frac{1}{n}$ then $\text{PoA}^{(n)} = 1$.*

Theorem 1 implies that for any constant α we have $\text{PoA}^{(n)} = \text{PoS}^{(n)} = 2^{\Omega(n)}$. In what follows we provide the necessary elements in order to prove the above theorem. The following lemma follows immediately from the definition of the utility function:

Lemma 1. *For utility functions $U_i(\bar{R}) = r_i(\bar{R})$, the only NE solution is obtained by the uniform profile \bar{R} where every station i plays the strategy $R_i = 1$. The social welfare value of this NE is $n(1 - \alpha)^{n-1}$.*

Lemma 1 implies in particular that since there is only one NE in these settings, then the price of stability equals the price of anarchy. In order to determine this value, we now turn to analyze the value of a profile which maximizes the social welfare. We first analyze the value of the social welfare function on the boundary of the profiles domain $[0, 1]^n$, and then turn to analyze the maximum value obtained in the interior of the domain.

Since φ is symmetric, any two integral profiles $\bar{R}, \bar{R}' \in \{0, 1\}^n$ having the same number of 1’s, satisfy $\varphi(\bar{R}) = \varphi(\bar{R}')$. Let $B_k = (R_1, \dots, R_n)$ denote any profile with

exactly k 1's. It therefore follows that the value in every extreme point where k stations play the 1-strategy and $n - k$ stations play the 0-strategy, is given by $v_k = \varphi(B_k)$ where

$$v_k = \sum_{i:R_i=1} R_i \prod_{j \neq i} (1 - \alpha R_j) + \sum_{i:R_i=0} R_i \prod_{j \neq i} (1 - \alpha R_j) = k(1 - \alpha)^{k-1}.$$

The following lemma characterizes the maximum value on the border of the profiles domain, and its dependence on α (proof omitted).

Lemma 2. *If $\alpha \in [\frac{1}{k+1}, \frac{1}{k})$ then $\max_j v_j = v_k$.*

Since clearly $\varphi(\overline{R}_{\text{OPT}}^{(n)}) \geq v_k$ for all k and for all α , we therefore have $\varphi(\overline{R}_{\text{OPT}}^{(n)}) \geq \max_k v_k$ for all α . In order to show that indeed $\varphi(\overline{R}_{\text{OPT}}^{(n)}) = \max_k v_k$, we wish to show that the maximum of $\varphi(\cdot)$ is not obtained in the interior of the domain. The following lemma, whose proof is omitted, shows that there is only one possible extreme point \overline{R}_0 in the interior of the domain $(0, 1)^n$, and characterizes the value of $\varphi(\overline{R}_0)$.

Lemma 3. *For any n and $k \in \{1, \dots, n-1\}$, if $\alpha \in [\frac{1}{k+1}, \frac{1}{k})$ then the only possible extreme point of the social welfare function $\varphi(\cdot)$ in the interior of $(0, 1)^n$ is $\overline{R}_0 = (\frac{1}{\alpha n}, \dots, \frac{1}{\alpha n})$. Furthermore, $\varphi(\overline{R}_0) \leq v_k$.*

Combining lemmas 2 and 3 immediately implies the following corollary:

Corollary 1. *For any n and $k \in \{1, \dots, n-1\}$, if $\alpha \in [\frac{1}{k+1}, \frac{1}{k})$, then $\varphi(\overline{R}_{\text{OPT}}^{(n)}) = v_k$.*

Combining Lemma 1 which shows that there exists a single NE solution $\overline{R}_{\text{NE}}^{(n)} = (1, \dots, 1)$ whose value is $\varphi(\overline{R}_{\text{NE}}^{(n)}) = n(1 - \alpha)^{n-1}$, along with Corollary 1, we can conclude the proof of Theorem 1.

As a consequence of Theorem 1, we restrict our attention in the following sections to the case where $\alpha \in (1/n, 1)$, since for $\alpha \leq 1/n$, the single NE of the HIMA game is indeed optimal. The following sections present a penalization scheme which enables the system to obtain a much better throughput, while still being at equilibrium.

3 Coordinated Nash Equilibria

In this section we introduce a penalty based scheme, where every station i incurs a penalty $p_i(\cdot)$ for transmission. We consider two types of penalties. The first type depends upon the choices of all the stations in the system, i.e., $p_i(\overline{R})$, while the second type only depends upon the choice of station i , i.e., $p_i(R_i)$. We refer to the former as an *exogenous penalty*, whereas the latter is referred to as an *endogenous penalty*. The general form of the utility function of station i is therefore $U_i(\overline{R}) = r_i - p_i$ (see [14] for a similar approach in the context of power control in cellular networks). We use the notion of coordinated NE and show that for both penalty functions, the selfishness of the stations does not result in more than a constant factor degradation in performance compared to the optimal performance. This should be contrasted with the results presented in the previous section showing that the price of stability for the uncoordinated case can be exponential in the number of stations.

We first show that there exists some $q_0 \in [0, 1]$ such that the uniform profile \bar{R} where $R_i = q_0$ for all i implies a mere constant degradation in performance compared to the optimal throughput possible. Note however that such a uniform profile need not be at NE. We then show that there exist penalty functions which cause such a uniform profile to be at NE. It therefore follows that by the use of appropriate penalties, selfishness can be tamed into providing a throughput that is at most a constant factor far from the optimal throughput.

3.1 The Power of Uniform Profiles

Given any $q \in [0, 1]$, let \bar{R}^q denote the uniform profile where $R_i = q$ for all i . Note that the social welfare value of \bar{R}^q is given by the function $\psi(q) = nq(1 - \alpha q)^{n-1}$. It is easy to verify that the value of q which maximizes $\psi(\cdot)$ is $q_0 = \frac{1}{\alpha n}$. It follows that the social welfare value of \bar{R}^{q_0} is $\psi(q_0) = \frac{1}{\alpha} \left(1 - \frac{1}{n}\right)^{n-1} \geq \frac{1}{e\alpha}$, where the inequality follows from the fact that $\left(1 - \frac{1}{n}\right)^{n-1}$ is strictly monotone decreasing, and converges to $1/e$.

As we have seen, the optimal value of the social welfare function for $\alpha \in [\frac{1}{k+1}, \frac{1}{k}]$ ¹ is obtained for a profile where k stations play the 1-strategy, and $n - k$ stations play the 0-strategy, resulting in a social welfare value of $\varphi(\bar{R}_{\text{OPT}}^{(n)}) = k(1 - \alpha)^{k-1} < k \left(1 - \frac{1}{k}\right)^{k-1} \leq \frac{1}{\alpha}$.

It follows that $\frac{\varphi(\bar{R}_{\text{OPT}}^{(n)})}{\varphi(\bar{R}^{q_0})} \leq e$. In the following sections we show that we can choose a penalty function such that the profile \bar{R}^{q_0} is at NE.

3.2 Exogenous Penalties

Let $q \in [0, 1]$, and consider the following utility function:

$$U_i^q(\bar{R}) = R_i \prod_{j \neq i} (1 - \alpha R_j) (2q - R_i),$$

which can be cast as a utility function of the form $U_i^q(\bar{R}) = r_i(\bar{R}) - p_i^q(\bar{R})$, where the exogenous penalty is defined by $p_i^q(\bar{R}) = \prod_{j \neq i} (1 - \alpha R_j) ((1 - 2q)R_i + R_i^2)$. Assume all stations except for i play strategy q . It follows that

$$U_i^q(R) = R_i(1 - \alpha q)^{n-1}(2q - R_i) = (1 - \alpha q)^{n-1}(2qR_i - R_i^2). \quad (1)$$

By taking derivatives, we obtain that the maximum is obtained for $R_i = q$, i.e., the uniform profile \bar{R}^q is at NE.

It therefore follows that the price of stability is at least $\max_q \frac{\varphi(\bar{R}_{\text{OPT}}^{(n)})}{\varphi(\bar{R}^q)}$. In addition, since choosing $R_i = q$ is the best response of station i regardless of the strategy chosen by any station $j \neq i$, we can conclude that the uniform profile \bar{R}^q is the only NE solution, hence the price of anarchy is the same as the price of stability.

Combining this result with the result presented in the previous section, for $q_0 = \frac{1}{\alpha n}$, we obtain the following theorem:

¹ Equivalently, $k < \frac{1}{\alpha} \leq k + 1$.

Theorem 2. *For every station i there exists an exogenous penalty function $p_i(\overline{R})$ for which the price of anarchy, as well as the price of stability, are at most e .*

Although Theorem 2 guarantees that aggressiveness can be tamed by using exogenous penalties, this might not be completely satisfactory. Exogenous penalties incurred by a station might change even if the station does not change its strategy. This might not be considered a handicap if the penalties cannot increase if the station remains put, however in our case, other stations being less aggressive actually increases the penalty incurred by a station, even if this station does not change its strategy. We address this issue in the following section, and present an endogenous penalty scheme, in which the penalty imposed on a station depends solely on its strategy, where increased aggressiveness is matched by increased penalties.

3.3 Endogenous Penalties

In this section we use insights from Sec. 3.2 and discuss endogenous penalty functions where the penalty function of station i depends only on R_i . Specifically, given any $q \in [0, 1]$, we consider for every station i the utility function

$$U_i^q(\overline{R}) = r_i(\overline{R}) - p_i^q(R_i),$$

where $p_i^q(R_i) = (1 - \alpha q)^{n-1}((1 - 2q)R_i + R_i^2)$ is an endogenous penalty function.

Assume all stations but i play the strategy q . It follows that the utility function of station i is again defined by Eq. (1). This implies that the best strategy for station i to play is $R_i = q$, hence the uniform profile \overline{R}^q is at NE. Similarly to Theorem 2, we thus obtain the following theorem:

Theorem 3. *For every station i there exists an endogenous penalty function $p_i(R_i)$ for which the overall price of stability is at most e .*

When considering the price of anarchy, the following theorem provides the conditions for which the uniform profile \overline{R}^q is actually the *only* NE (proof omitted).

Theorem 4. *If $\alpha \leq \frac{2}{e}$, then for every station i there exists an endogenous penalty function $p_i(R_i)$ for which the overall price of anarchy is at most e .*

4 Conclusion and Open Questions

We present a generalization of the classic multiple access model, by considering spatial interferences parameter, modeling scenarios where different transmissions to succeed simultaneously. This new model captures the fact that collisions are a phenomenon experienced by the receiving end of transmissions, and it depends on the amount of interferences sensed by this receiver from the various simultaneous transmissions.

We show that for homogeneous interferences, if agents are selfish, then the system's performance at equilibrium can be up to an exponential factor far away from the optimal performance. We further introduce a penalty function to be cast on the agents, inducing a much better performance in an equilibrium, which is at most a factor of e away from the optimal performance.

Several interesting questions remain open. First, it would be interesting to obtain analytic guarantees as to the price of anarchy and the price of stability for non-homogeneous interferences. We believe that our results serve as a mere first step in understanding such interference-bound environments. Second, it is interesting to see if our model enables the design of better medium-access protocols, taking into account possible prior knowledge of inter-agents interferences.

References

1. Koutsoupias, E., Papadimitriou, C.: Worst-Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
2. Anshelevich, E., Dasgupta, A., Tardos, É., Wexler, T.: Near-optimal Network Design with Selfish Agents. In: STOC 2003, pp. 511–520 (2003)
3. Correa, J.R., Schulz, A.S., Stier, N.E.: Selfish Routing in Capacitated Networks. *Mathematics of Operations Research* 29(4), 961–976 (2004)
4. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press, Cambridge (1994)
5. Naor, J., Raz, D., Scalosub, G.: Homogeneous Interference Game in Wireless Networks. Technical Report CS-2008-11, Department of Computer Science, Technion (2008)
6. Jin, Y., Kesidis, G.: Equilibria of a Noncooperative Game for Heterogeneous Users of an ALOHA Network. *IEEE Communications Letters* 6(7), 282–284 (2002)
7. MacKenzie, A.B., Wicker, S.B.: Stability of Multipacket Slotted Aloha with Selfish Users and Perfect Information. In: INFOCOM 2003, pp. 1583–1590 (2003)
8. Altman, E., Barman, D., El-Azouzi, R., Jiménez, T.: A Game Theoretic Approach for Delay Minimization in Slotted Aloha. In: ICC 2004, pp. 3999–4003 (2004)
9. Ma, R.T.B., Misra, V., Rubenstein, D.: Modeling and Analysis of Generalized Slotted-Aloha MAC Protocols in Cooperative, Competitive and Adversarial Environments. In: ICDCS 2006, p. 62 (2006)
10. Menache, I., Shimkin, N.: Decentralized Rate Regulation in Random Access Channels. In: INFOCOM 2008, pp. 394–402 (2008)
11. Marbach, P., Pang, R.: Transmission Costs, Selfish Nodes, and Protocol Design. In: WiOpt 2005, pp. 31–40 (2005)
12. Wang, D., Comaniciu, C., Tureli, U.: A Fair and Efficient Pricing Strategy for Slotted Aloha in MPR Models. In: VTC 2006, pp. 1–5 (Fall 2006)
13. Saraydar, C.U., Mandayam, N.B., Goodman, D.J.: Efficient Power Control via Pricing in Wireless Data Networks. *IEEE Transactions on Communications* 50(2), 291–303 (2002)
14. Xiao, M., Shroff, N.B., Chong, E.K.P.: A Utility-Based Power-Control Scheme in Wireless Cellular Systems. *IEEE/ACM Transactions on Networking* 11(2), 210–221 (2003)
15. Lam, R.K., Chiu, D.M., Lui, J.C.S.: On the Access Pricing and Network Scaling Issues of Wireless Mesh Networks. *IEEE Transactions on Computers* 56(11), 1456–1469 (2007)
16. Carvalho, M.M., Garcia-Luna-Aceves, J.J.: A Scalable Model for Channel Access Protocols in Multihop Ad Hoc Networks. In: MOBICOM 2004, pp. 330–344 (2004)
17. Hui, K.H., Lau, W.C., Yue, O.C.: Characterizing and Exploiting Partial Interference in Wireless Mesh Networks. In: ICC 2007, pp. 102–108 (2007)
18. Qiu, L., Zhang, Y., Wang, F., Han, M.K., Mahajan, R.: A General Model of Wireless Interference. In: MOBICOM 2007, pp. 171–182 (2007)

A Network Coloring Game

Kamalika Chaudhuri¹, Fan Chung Graham², and Mohammad Shoaib Jamall²

¹ Information Theory and Applications Center, UC San Diego

`kamalika@soe.ucsd.edu`

² Department of Mathematics, UC San Diego

`{fan,mjamall}@math.ucsd.edu`

Abstract. We analyze a network coloring game which was first proposed by Michael Kearns and others in their experimental study of dynamics and behavior in social networks. In each round of the game, each player, as a node in a network G , uses a simple, greedy and selfish strategy by choosing randomly one of the available colors that is different from all colors played by its neighbors in the previous round. We show that the coloring game converges to its Nash equilibrium if the number of colors is at least two more than the maximum degree. Examples are given for which convergence does not happen with one fewer color. We also show that with probability at least $1 - \delta$, the number of rounds required is $O(\log(n/\delta))$.

1 Introduction

We perform a theoretical study of the following coloring game on networks. In the coloring game, there is an associated network G , and each player is associated with a vertex in this network. Each player has a set of available strategies or colors; the payoff of a player is 1 if he plays a color different from the color played by any of his neighbors in G , and 0 otherwise. The goal of each player is to maximize his payoff, and we are interested in the behaviour and strategies of the players as a function of the network structure.

Our study is motivated by the fact that graph-coloring problems arise as natural formalizations of many conflict-resolution problems in practice. An example due to [7], is a scenario where faculty members wish to schedule classes in a limited number of classrooms, and must avoid conflicts with other faculty members. This can be modelled as a coloring problem, where the faculty members represent vertices, classrooms represent colors, and any two faculty members who have classes with overlapping times are connected by an edge. However, in this scenario, typically there is no centralized agency which assigns classrooms to faculty members, players coordinate among themselves to decide on a non-conflicting assignment, and it is also unreasonable to lay down a distributed protocol and expect the players to abide by the rules of this protocol. As a result, a game-theoretic formulation is an appropriate model for this scenario. A second example of a scenario where a game-theoretic formulation is appropriate is when players are employees in an organization, and colors are skills, and employees attempt

to perfect skills that are different from the skills possessed by other employees in their department.

An experimental study of various coloring games was initiated by [7], which reports on behavioral experiments on human subjects who are incentivized to play the coloring game on specified networks. Comparisons were made among several different types of networks of moderate sizes. In addition, examples were given to illustrate the difficulties in analyzing the dynamics of large networks in which each node takes simple and selfish steps. The paper [7] has attracted much attention and pointed to the need for theoretical analysis, although there has not been any other prior work in this specific direction to the extent of our knowledge.

In this paper, we model the coloring game as a game played on a network over multiple rounds. We study the dynamics of the game when the players play a very simple, greedy strategy. At each round, each player picks a color uniformly at random from the set of colors unused by any of his neighbors in the previous round, and plays this color. We note that this is the *best response myopic strategy* [1] for the coloring game. We say that the coloring game converges, when the color played by each player is different from the color played by any of its neighbors. This is a *Nash equilibrium* of the coloring game, because no player has an incentive to change his strategy under this configuration. We are interested in the time taken by the players to converge, when each player adopts the greedy strategy.

Our main result in this paper is that for a coloring game played on a network on n vertices with maximum degree Δ , if the number of colors available to each vertex is $\Delta + 2$ or more, and if each player plays the greedy strategy, then the coloring game converges in $O(\log n)$ steps with high probability. Our result is also accompanied by a lower bound, in which we show a graph and a starting assignment of colors, such that, if the number of colors available to each vertex is $\Delta + 1$, and if each player plays the greedy strategy, then, the coloring game does not converge. Our upper bound holds even in the presence of non-participating vertices, which maintain the same color throughout the game. This indicates that if sufficient number of colors are available, then, even in the game-theoretic scenario, even when the players play a simple greedy strategy, convergence to the nash equilibrium is very rapid. In fact, this convergence bound is even comparable to the convergence bound for distributed protocols for graph coloring, which require the nodes to follow a distributed protocol and cooperate with each other (e.g., see, for example, the work of Luby [9]).

Related Work

The problem of graph coloring has a long history, and there is a lot of literature on centralized as well as distributed algorithms for this problem. The network coloring game has been studied experimentally by Kearns *et al.* however, to the best of our knowledge, this game has not been analyzed theoretically. On the experimental side, a number of behavior experiments with human subjects were conducted on the coloring game by Kearns, Suri, and Montford [7]. In their

experiment, a group of human subjects were assigned to vertices of a graph, and were asked to play the coloring game for some amount of time. Each subject has access to the colors of their neighbors, and did not have any knowledge about the structure of the rest of the graph. Three graphs were studied – a cycle, a cycle with chords, as well as a random preferential attachment graph. In the experiment, the subjects found it more difficult to color the preferential attachment graph in their allotted amount of time.

In general, finding the minimum number of colors required to color a graph, even in a centralized manner, is NP-Hard [3], as well as hard to approximate [8]. Coloring a graph when the number of available colors is more than its maximum degree can be easily done in linear time by a centralized algorithm; however, the same problem becomes more challenging when the algorithm is required to be distributed.

There has also been a line of work on distributed graph-coloring. Luby [9] provides a distributed algorithm which finds a coloring of a graph in $O(\log n)$ rounds, when the number of available colors is $\Delta + 1$ or more. Notice that here Δ is the maximum degree of any vertex in the graph. [2] also provide theoretical as well as experimental results on a simple algorithm for coloring a graph in a distributed manner in $O(\log n)$ rounds when the number of available colors is $\Delta + 1$. Both Luby’s algorithm and the algorithm of [2] require each node to communicate to its neighbors its status, which indicates whether this node has any conflicts with its neighbors or not. In contrast, our scenario is purely game-theoretic : our algorithm does not require any cooperation among the vertices, and will succeed even if some nodes do not participate in the game.

Another line of work which is relevant to ours is the literature on Markov Chains for randomly sampling colorings of a graph. In this case, the goal is to bound the mixing time of a Markov Chain on colorings of a graph G . Pioneering work in this field was done by Mark Jerrum [6], who showed a way to randomly sample colorings of a graph G with maximum degree Δ in $O(n \log n)$ time when the number of colors available is at least $2\Delta + 1$. This was later improved by Vigoda [10], who could randomly sample colorings from graphs of degree Δ , when $\frac{11\Delta}{6}$ colors were available. Hayes and Vigoda [5] showed a better bound for triangle-free graphs when the number of colors needed was $\min(\Delta + O(1), O(\log n))$. Finally, it was shown in [4] that colorings from planar graphs can be sampled in $O(n \log n)$ time when the number of colors is at least $\Delta / \log \log \Delta$.

A Summary of Our Results

We consider the coloring game played on a graph G . Before we state our results, we need to define the following concepts.

The Coloring Game. In the coloring game, each vertex in G represents a player. Each player has a set of k available strategies or *colors*; the payoff of a player is 1 if he picks a color different from the colors picked by his neighbors in G , and 0 otherwise. The game is played in *rounds*; each round, the players choose their strategies simultaneously. It is assumed that the players only have

a *local view* of G , which means that they only know who their neighbors are, and the colors picked by their neighbors, and they do not have any knowledge of the rest of the graph G .

Greedy Strategy. In this paper, we study the dynamics of the coloring game when each player has a local view and plays the *greedy strategy*. A player is said to play the *greedy strategy* if, in each round, he picks a color uniformly at random from the set of colors unused by any of his neighbors in the previous round. We note that this is the *best response myopic strategy* for the coloring game.

Convergence of the Coloring Game. The coloring game is said to *converge* if, for every vertex v in G , the color chosen by v is different from the color chosen by all its neighbors. We note that when the coloring game has converged, none of the players have any incentive to change their strategy.

Participants in the Coloring Game. We say that a vertex v in graph G is a participant in the coloring game on G if v plays according to the greedy strategy, otherwise we say that v is a *non-participant*. An instance of a coloring game on a graph G which has non-participant nodes, is said to converge, if, for every participant vertex v in G , the color chosen by v is different from the color chosen by all its neighbors.

The main result of this paper can be summarized by the following theorem.

Theorem 1. *Let G be any graph on n vertices, and let Δ be the maximum degree of any vertex in G . If the number of colors available to each vertex is at least $\Delta + 2$, and if each player plays the greedy strategy, then, for any starting assignment of colors, the coloring game on G converges after at most $O(\log(\frac{n}{\delta}))$ rounds with probability at least $1 - \delta$.*

In addition, we show that if there exists a set S of non-participant vertices, such that any vertex $v \in S$ has a fixed color throughout the game, then, the convergence time of the game is still at most $O(\log(\frac{n-|S|}{\delta}))$ round.

Corollary 1. *Let G be any graph on n vertices, Δ be the maximum degree of any vertex in G , and S be a set of non-participant vertices. If the vertices in S do not change their color throughout the game, and the number of colors available to each vertex is at least $\Delta + 2$, and if each player plays the greedy strategy, then, for any starting assignment of colors, the coloring game on G converges after at most $O(\log(\frac{n-|S|}{\delta}))$ rounds with probability at least $1 - \delta$.*

We show that when the number of colors is $\Delta + 1$, there is a graph G and a starting assignment of colors such that the greedy strategy does not converge.

Theorem 2. *There exists a graph G and a starting assignment of colors \mathcal{C} such that if the number of colors is $\Delta + 1$, and if each player plays the greedy strategy, the coloring game on G never converges.*

2 Several Lemmas

We use $c_t(u)$ to denote the color played by player u at round t , and $\mathcal{N}(u)$ to denote the set of neighbors of u in G . We say that player u at round t has a *conflict* if

$$\exists v \in \mathcal{N}(u), \quad c_t(u) = c_t(v)$$

At any time t , we use *number of conflicts* to mean the number of vertices v which have a conflict. We observe that if a vertex u has no conflict at time t , then $c_{t+1}(u) = c_t(u)$. We use k to denote the number of colors available and Δ to denote the maximum degree of the graph G over which the game is played.

We use the following lemma which is a minor variant of Markov's inequality.

Lemma 1. *Let X be a random variable such that $0 \leq X \leq M$. Then, for any a ,*

$$\Pr[X < a] \leq \frac{M - \mathbf{E}[X]}{M - a}.$$

We observe that if a vertex v has no conflict at time t , then at round $t + 1$, v does not change its color; moreover, no neighbor of v picks color $c_t(v)$, and after round $t + 1$, v still has no conflict. Thus we have the following lemma.

Lemma 2. *If a vertex v has no conflict at time t , then it has no conflict at any subsequent time.*

The main idea behind the proof of Theorem 1 is to show that if we consider any two successive rounds of the coloring game, then, the conflict at each vertex is resolved with constant probability. We note that considering the game over two successive rounds is essential; in a single round, it is possible that a player has as little as $(1 - \frac{1}{k-\Delta})^\Delta$ chance of getting its conflict resolved.

The main step in the proof of Theorem 1 is the following lemma.

Lemma 3. *Consider an instance of the coloring game played on a graph G , with maximum degree Δ , in which each node has k color choices, where $k \geq \Delta + 2$. If, after round t , some vertex v in G has a conflict, then, there exists some constant c such that*

$$\Pr[v \text{ has no conflict after round } t + 2] \geq c$$

Proof. For vertex v , Let M denote the neighbors of v which do not have a conflict after round t . We define

$$F = \cup_{u \in M} \{c_t(u)\}$$

and let $f = |F|$. Then, the number of choices available to v in rounds $t + 1$ as well as $t + 2$ is at most $k - f$. However, since t is any arbitrary round, the number of colors available to v after round t can possibly be much less.

The proof of this lemma proceeds in two steps. First, we show in Lemma 4 that after round $t + 1$, with constant probability, the vertex v has at least $\frac{k-f}{6}$ color choices. Next, we show in Lemma 5 that, given that v has $\frac{k-f}{6}$ color choices

after round $t + 1$, with constant probability, v has no conflict after round $t + 2$. Combining these two facts, we get a proof of the lemma.

After round t , for any color $i \in [k]$, we define the random variable Y_i as follows:

$$Y_i = \begin{cases} 1 & \text{if } c_{t+1}(u) \neq i, \text{ for all } u \in \mathcal{N}(v) \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

Let $Y = \sum_{i \in [k]} Y_i$; thus Y is the number of colors available to vertex v after round $t + 1$.

For any $u \in \mathcal{N}(v) \setminus M$, we let $\chi(u)$ be the set of colors available to vertex u after round t . Further, let $p_u = \frac{1}{|\chi(u)|}$. For $u \in M$, we define $\chi(u) = \{c_t(u)\}$, and $p_u = 1$. Note that if $u \in \mathcal{N}(v) \setminus M$, p_u is the probability with which vertex u picks a fixed color available to it during round $t + 1$. Also, as $k \geq \Delta + 2$, each vertex $u \in \mathcal{N}(v) \setminus M$ has at least two color choices available, and therefore $|\chi(u)| \geq 2$. We are interested in the probability that $Y \geq \frac{k-f}{6}$. We show the following lemma.

Lemma 4

$$\Pr(Y \geq \frac{k-f}{6}) \geq \frac{1}{25}$$

Proof. From the definition of Y_i , we can write that:

$$\Pr(Y_i = 1) = \prod_{\{u \in \mathcal{N}(v) | i \in \chi(u)\}} (1 - p_u)$$

The expectation of Y can therefore be estimated as:

$$\begin{aligned} \mathbf{E}(Y) &= \sum_{i \in [k]} \prod_{\{u \in \mathcal{N}(v) | i \in \chi(u)\}} (1 - p_u) \\ &\geq \sum_{i \in [k] \setminus F} \prod_{\{u \in \mathcal{N}(v) \setminus M | i \in \chi(u)\}} (1 - p_u) \end{aligned}$$

The second step follows by the definition of F and M . As, for each $u \in \mathcal{N}(v) \setminus M$, $|\chi(u)| \geq 2$, we can use the inequality $1 - x \geq e^{-\frac{3}{2}x}$ for $0 \leq x \leq \frac{1}{2}$ to write:

$$\mathbf{E}(Y) \geq \sum_{i \in [k] \setminus F} e^{-\frac{3}{2} \sum_{\{u \in \mathcal{N}(v) \setminus M | i \in \chi(u)\}} p_u}$$

Using the convexity of the exponential function,

$$\mathbf{E}[Y] \geq (k - f) e^{-\frac{3}{2} \frac{1}{k-f} \sum_{i \in [k] \setminus F} \sum_{\{u \in \mathcal{N}(v) \setminus M | i \in \chi(u)\}} p_u}$$

Observe that

$$\sum_{i \in [k] \setminus F} \sum_{\{u \in \mathcal{N}(v) \setminus M | i \in \chi(u)\}} p_u \leq \sum_{u \in \mathcal{N}(v) \setminus M} p_u \chi(u)$$

and therefore, we can write that:

$$\begin{aligned} \mathbf{E}[Y] &\geq (k - f)e^{-\frac{3}{2} \frac{1}{k-f} \sum_{u \in \mathcal{N}(v) \setminus M} p_u \chi(u)} \\ &\geq (k - f)e^{-\frac{3}{2}} \geq \frac{k - f}{5} \end{aligned}$$

The final step follows from the fact that $|\{u \in \mathcal{N}(v) \setminus M\}| \leq k - f - 2$. To complete the lemma, we now use Lemma [1](#), which is a variant of Markov's inequality. \square

Now we analyze the dynamics in round $t + 2$ given that $Y > \frac{k-f}{6}$.

For any color $i \in [k]$, we define variables \tilde{Y}_i as follows:

$$\tilde{Y}_i = \begin{cases} 1 & \text{if } c_{t+2}(u) \neq i \text{ for all } u \in \mathcal{N}(v) \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

Let $\tilde{Y} = \sum_{i \in [k]} \tilde{Y}_i$. We are interested in the probability of the event that $\tilde{Y} \geq \frac{k-f}{7e^9}$.

Lemma 5

$$\Pr(\tilde{Y} \geq \frac{k-f}{7e^9} | Y \geq \frac{k-f}{6}) \geq \frac{1}{42e^9}$$

Proof. We define \tilde{M} to be the set of vertices in $\mathcal{N}(v)$ which have no conflicts after round $t + 1$, and \tilde{F} as the set $\cup_{u \in \tilde{M}} \{c_{t+1}(u)\}$. Note, by Lemma [2](#), that $\tilde{M} \supseteq M$ and $|\tilde{F}| \geq |F|$. Further, we let $\tilde{f} = |\tilde{F}|$. Also we define $H = [k] \setminus \cup_{u \in \mathcal{N}(v)} \{c_{t+1}(u)\}$. Thus H is the number of colors available to v in round $t + 2$. Obviously, $|H| \leq k - \tilde{f}$; given that $Y \geq \frac{k-f}{6}$, $|H| \geq \frac{k-f}{6} \geq \frac{k-\tilde{f}}{6}$.

For any $u \in \mathcal{N}(v) \setminus \tilde{M}$, we define $\tilde{\chi}(u)$ as the set of colors available to vertex u after round $t + 1$, and \tilde{p}_u as $\frac{1}{|\tilde{\chi}(u)|}$ respectively. For $u \in \tilde{M}$, we define $\tilde{\chi}(u) = \{c_{t+1}(u)\}$ and $\tilde{p}_u = 1$. We can write:

$$\Pr(\tilde{Y}_i = 1) = \prod_{\{u \in \mathcal{N}(v) | i \in \tilde{\chi}(u)\}} (1 - \tilde{p}_u)$$

Similar to the proof of Lemma [4](#), we can write:

$$\begin{aligned} \mathbf{E}[\tilde{Y}] &= \sum_{i \in [k]} \prod_{\{u \in \mathcal{N}(v) | i \in \tilde{\chi}(u)\}} (1 - \tilde{p}_u) \\ &\geq \sum_{i \in H} \prod_{\{u \in \mathcal{N}(v) \setminus \tilde{M} | i \in \tilde{\chi}(u)\}} (1 - \tilde{p}_u) \end{aligned}$$

As, for each $u \in \mathcal{N}(v) \setminus \tilde{M}$, $|\tilde{\chi}(u)| \geq 2$, we can write:

$$\mathbf{E}(Y) \geq \sum_{i \in H} e^{-\frac{3}{2} \sum_{\{u \in \mathcal{N}(v) \setminus \tilde{M} | i \in \tilde{\chi}(u)\}} \tilde{p}_u}$$

Using the convexity of the exponential function,

$$\mathbf{E}[\tilde{Y}] \geq \frac{1}{6}(k - \tilde{f})e^{-\frac{3}{2}\frac{6}{k-\tilde{f}} \sum_{i \in H} \sum_{\{u \in \mathcal{N}(v) \setminus \tilde{M} \mid i \in \tilde{\chi}(u)\}} \tilde{p}_u}$$

Observe that

$$\sum_{i \in H} \sum_{\{u \in \mathcal{N}(v) \setminus \tilde{M} \mid i \in \tilde{\chi}(u)\}} \tilde{p}_u \leq \sum_{u \in \mathcal{N}(v) \setminus \tilde{M}} \tilde{p}_u \tilde{\chi}(u)$$

and therefore, we can write that:

$$\begin{aligned} \mathbf{E}[\tilde{Y}] &\geq \frac{1}{6}(k - \tilde{f})e^{-\frac{3}{2}\frac{1}{k-\tilde{f}} \sum_{u \in \mathcal{N}(v) \setminus \tilde{M}} \tilde{p}_u \tilde{\chi}(u)} \\ &\geq \frac{1}{6}(k - \tilde{f})e^{-9} \geq \frac{k - f}{6e^9} \end{aligned}$$

The lemma now follows by an application of Lemma [11](#). □

Now, given that the event $\tilde{Y} \geq \frac{k-\tilde{f}}{7e^9}$ occurs, there is a set \mathcal{C} of $\frac{k-\tilde{f}}{7e^9}$ colors, such that the conflict of v is resolved if it picks a color from \mathcal{C} in round $t + 2$. Since v picks one out of at most $k - \tilde{f}$ colors, given the event $\tilde{Y} \geq \frac{k-\tilde{f}}{7e^9}$, v has no conflict after round $t + 2$ with probability at least $\frac{1}{42e^9}$. Combining this with the probability of occurrence of the event $Y > \frac{k-f}{6}$, Lemma [3](#) follows for $c = \frac{1}{1050e^9}$. □

3 Proofs of the Main Theorems

First we prove Theorem [11](#).

Proof. (Of Theorem [11](#)) For a vertex v , and a time t , we define random variables $X_v(t)$ as follows:

$$X_v(t) = \begin{cases} 1 & \text{if vertex } v \text{ has a conflict after round } t \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

From Lemma [3](#), for any vertex v , for some constant c ,

$$\Pr(X_v(t+2) = 1 \mid X_v(t) = 1) \leq 1 - c \tag{4}$$

Using Lemma [2](#), for any τ ,

$$\begin{aligned} \Pr(X_v(2\tau) = 1 \mid X_v(0) = 1) &= \Pr\left(\bigcap_{i=1}^{\tau} X_v(2i) = 1 \mid X_v(0) = 1\right) \\ &= \prod_{i=1}^{\tau} \Pr(X_v(2i) = 1 \mid \bigcap_{j=1}^{i-1} X_v(2j) = 1) \\ &\leq (1 - c)^\tau \leq e^{-c\tau} \end{aligned}$$

Plugging in $\tau = \frac{1}{c} \log(\frac{n}{\delta})$, we get, for any vertex $v \in G$,

$$\Pr(X_v(2\tau) = 0 | X_v(0) = 1) \geq 1 - \frac{\delta}{n} \quad (5)$$

The theorem follows by applying an union bound over all vertices in G . \square

Proof. (Of Corollary [1](#)) All steps and lemmas in the proof of Theorem [1](#) remain valid except that we use $n - |S|$ in place of n . \square

It remains to prove Theorem 2.

Proof. (Of Theorem [2](#)) Let G be a cycle of length 5, and let $V = \{v_1, \dots, v_5\}$. Let $k = 3$, and suppose the initial configuration is $(c(v_1), c(v_2), c(v_3), c(v_4), c(v_5)) = (1, 0, 2, 2, 0)$. Then, only v_3 and v_4 have conflicts. If v_3 and v_4 follow the dynamics, then, in the next round, $c(v_3) = c(v_4) = 1$. This again causes them both to have conflicts, so in the following round, $c(v_3) = c(v_4) = 2$, and we have a cycle in the dynamics. \square

References

1. Arcaute, E., Johariand, R., Mannor, S.: Two stage myopic dynamics in network formation games. In: Workshop on Network and Economics(WINE) (2008)
2. Finocchi, I., Panconesi, A., Silvestri, R.: Experimental analysis of simple, distributed vertex coloring algorithms. In: SODA, pp. 606–615 (2002)
3. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W.H. Freeman, New York (1979)
4. Hayes, T.P., Vera, J.C., Vigoda, E.: Randomly coloring planar graphs with fewer colors than the maximum degree. In: STOC 2007: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing, pp. 450–458. ACM Press, New York (2007)
5. Hayes, T., Vigoda, E.: Coupling with the stationary distribution and improved sampling for colorings and independent sets. In: SODA 2005: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms, pp. 971–979. Society for Industrial and Applied Mathematics, Philadelphia (2005)
6. Jerrum, M.: A very simple algorithm for estimating the number of k -colorings of a low-degree graph. *Random Struct. Algorithms* 7(2), 157–165 (1995)
7. Kearns, M., Suri, S., Montfort, N.: An experimental study of the coloring problem on human subject networks. *Science* 313(5788), 824–827 (2006)
8. Khot, S.: Improved inapproximability results for maxclique, chromatic number and approximate graph coloring. In: FOCS, pp. 600–609 (2001)
9. Luby, M.: Removing randomness in parallel computation without a processor penalty. In: FOCS, pp. 162–173 (1988)
10. Vigoda, E.: Improved bounds for sampling colorings. In: Proc. of FOCS (1999)

Asynchronous Best-Reply Dynamics

Noam Nisan¹, Michael Schapira², and Aviv Zohar²

¹ Google Tel-Aviv and The School of Computer Science and Engineering,
The Hebrew University of Jerusalem, Israel

² The School of Computer Science and Engineering,
The Hebrew University of Jerusalem, Israel
{noam,mikesch,avivz}@cs.huji.ac.il

Abstract. In many real-world settings (e.g., interdomain routing in the Internet) strategic agents are instructed to follow best-reply dynamics in asynchronous environments. In such settings players learn of each other’s actions via update messages that can be delayed or even lost. In particular, several players might update their actions *simultaneously*, or make choices based on outdated information. In this paper we analyze the convergence of best- (and better-)reply dynamics in asynchronous environments. We provide sufficient conditions, and necessary conditions for convergence in such settings, and also study the convergence-rate of these natural dynamics.

1 Introduction

Many real-life protocols can be regarded as executions of best-reply dynamics, i.e, players (computational nodes) are instructed to repeatedly best-reply to the actions of other players. In many cases, like Internet settings, this occurs in *asynchronous* environments: Think of the players as residing in a computer network, where their best-replies are transmitted to other players and serve as the basis for the other players’ best-replies. These update messages that players send to each other may be delayed or even lost, and so players may update their actions simultaneously, and do so based on outdated information. Perhaps the most notable example for this is the Border Gateway Protocol (BGP) that handles interdomain routing in the Internet. As observed in [1], BGP can indeed be seen as an execution of best-reply dynamics in asynchronous environments.

Asynchronous best-reply dynamics. The most fundamental question regarding best-reply dynamics in asynchronous settings is “When are such dynamics guaranteed to converge?”. This will certainly not happen if a pure Nash equilibrium does not exist, but is not guaranteed even in very simple and well-structured games that have a pure Nash. We present a formal framework for the analysis of best-reply dynamics in asynchronous environments. We then exhibit a simple class of games for which convergence to a *unique* pure Nash equilibrium is guaranteed. We term this class, which contains all strictly-dominance-solvable games (games where iterated elimination of strictly dominated strategies leaves a single strategy profile [2]), “*max-solvable-games*”. We also discuss the

convergence-rate of best-reply dynamics in asynchronous settings. We propose a notion of an *asynchronous phase*, and show that for max-solvable games convergence also happens quickly.

Theorem: *Best reply-dynamics converge within $\sum_i m_i$ phases for every max-solvable game, and in every asynchronous schedule. Here m_i is the size of the strategy space of the i 'th player. In particular, this holds for all strictly-dominance-solvable games.*

This theorem shows that even though the “input” (a normal-form representation of a max-solvable game) is of exponential size (in the size of the strategy-spaces), best-reply dynamics converges in a linear number of phases.

We consider a generalization of max-solvable games, called “weakly-max-solvable games” that contains the class of weakly-dominance-solvable games (games where iterated elimination of weakly-dominated strategies leaves a single strategy profile [2]). For this class of games we show that no similar result holds; not only are best-reply dynamics not guaranteed to converge, but any procedure for finding a pure Nash equilibrium faces a severe obstacle.

Theorem: *Finding a pure Nash equilibrium in weakly-max-solvable games requires exponential communication in $\sum_i m_i$. This is even true for the more restricted class of weakly-dominance-solvable games.*

This result follows the line of research initiated by Conitzer and Sandholm [3], and further studied in the work of Hart and Mansour [4].

Asynchronous better-reply dynamics. At this point we turn our attention to better-reply dynamics. Now, players are not required to continuously best-reply to the strategies of the others, but merely to always choose strategies that are better replies than the ones they currently have. Once again, we are interested in figuring out when these dynamics converge in asynchronous settings. A natural starting point for this exploration is the well-known class of *potential games*, introduced by Monderer and Shapley [5], building on the seminal work of Rosenthal [6]. For these games, it is known that better-reply dynamics are guaranteed to converge (if players update their strategies one by one, and learn of each other’s action immediately).

We show, in contrast, that even for these games asynchrony poses serious challenges and may even lead to persistent oscillations. We consider a restricted, yet expressive, form of asynchrony – settings in which players may update strategies simultaneously (and not necessarily one by one), but update messages arrive at their destinations immediately (no delay). We call such restricted asynchronous settings “*simultaneous settings*”. We prove the following theorem:

Theorem: *If every subgame of a potential game has a unique pure Nash equilibrium then better-reply dynamics are guaranteed to converge for every simultaneous schedule. (By subgame, we mean a game that is the result of elimination of players’ strategies from the original game.)*

In fact, we show that this result is almost a characterization, in the sense that the uniqueness of pure Nash equilibria in every subgame is also a necessary condition for convergence in simultaneous settings for a large subclass of potential games.

Organization of the Paper: In Section 2 we present a model for analyzing best- and better-reply dynamics in asynchronous settings. In Section 3 we present and discuss max-solvable games. In Section 4 we explore potential games. Due to space constraints many of the proofs are omitted (see 7 for a full version).

2 Synchronous, Simultaneous, and Asynchronous Environments

We use standard game-theoretic notation: Let G be a *normal-form game* with n players $1, 2, \dots, n$. We denote by S_i the (finite) *strategy space* of the i 'th player. Let $S = S_1 \times \dots \times S_n$, and let $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$ be the cartesian product of all strategy spaces but S_i . Each player i has a *utility function* u_i that specifies i 's *payoff* for any strategy-profile of the players. For any strategy $s_i \in S_i$, and every $(n - 1)$ -tuple of strategies $s_{-i} \in S_{-i}$, we shall denote by (s_i, s_{-i}) the strategy profile in which player i plays $s_i \in S_i$ and all other players play their strategies in s_{-i} . Given $s_{-i} \in S_{-i}$, $s_i \in S_i$ is said to be a best reply to s_{-i} if $u_i(s_i, s_{-i}) = \max_{s'_i \in S_i} u_i(s'_i, s_{-i})$. Given $s_{-i} \in S_{-i}$, $s'_i \in S_i$ is said to be a better-reply of player i than $s_i \in S_i$ if $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$.

Consider the following best-reply dynamics procedure: We start with an initial strategy profile of the players $s \in S$. There is set of rounds $R = \{1, 2, \dots\}$ of *infinite* size. In each round one *or more* players are chosen to participate. Every player chosen to participate must switch to a best-reply to his *most recent information* about the strategies of the other players, and send update messages to all other players announcing his strategy (a player must announce his strategy to all other players even if it did not change).

As in 11, there is an adversarial entity called the *Scheduler* that is in charge of making the following decisions: Choosing the initial strategy profile $s \in S$. Determining which players will participate in which round (a function f from R to subsets of the players). Determining when sent update messages reach their destinations (see below). The Scheduler must be restricted not to indefinitely starve any player from best-replying (that is, each player participates in infinitely many rounds). We shall name all the choices made by the Scheduler a *schedule*. We distinguish between three types of settings:

Synchronous settings: In these settings, the Scheduler can only choose *one* player to play in each round (that is, $|f(r)| = 1$ for any $r \in R$). In addition, update messages sent by players arrive at their destinations *immediately* (that is, at the end of the round in which they were sent). Hence, players' actions are *observable* to other players. Observe, that a game is a potential game iff for each of its subgames, better-reply dynamics are guaranteed to converge to a pure Nash equilibrium for any synchronous schedule.

Simultaneous settings: In simultaneous settings, the Scheduler can choose *any* number of players to play in each round ($|f(r)|$ can be any number in $1, 2, \dots, n$ for any $r \in R$). As in synchronous settings, players' actions are observable (update messages sent by players arrive at their destinations immediately).

Asynchronous settings: As in simultaneous settings, the Scheduler can choose any number of players to play in each round. However, in asynchronous settings the Scheduler can also decide when each sent update message arrives at its destination (at the end of the round in which it was sent or in some subsequent round) subject to the limitation that messages that were sent earlier arrive before later ones. It can also decide to drop update messages. The Scheduler may not prevent all update messages of a player from reaching another player indefinitely.

Elementary examples (like the “Battle of the Sexes” game) show that even in very simple games, in which best-reply dynamics are guaranteed to converge in synchronous settings, they might not converge in simultaneous settings (and, in particular, in asynchronous settings). Similarly, it can be shown that convergence of best-reply dynamics in simultaneous settings does not imply convergence in asynchronous settings.

In an analogous way, we can now define synchronous, simultaneous, and asynchronous convergence of better-reply dynamics.

3 Max-solvable Games

In this section we present a class of games called “max-solvable games” for which best-reply dynamics are guaranteed to converge to a pure Nash equilibrium even in asynchronous settings. We then discuss a generalization of these games, that contains all dominance-solvable games (games in which the iterated removal of dominated strategies results in a single strategy profile).

3.1 Max-solvable Games – Definitions

We start by defining max-solvable games.

Definition 1. *A strategy $s_i \in S_i$ is max-dominated if for every strategy-profile of the other players $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ there is a strategy s'_i such that $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$.*

That is, a strategy of a player is max-dominated if it is not a best-reply to *any* strategy-profile of the other players. Observe, that every strictly dominated strategy is max-dominated. In fact, a strategy is max-dominated even if it is strictly dominated by a *mixed* strategy.

Informally, a max-solvable game is a game in which the iterated elimination of max-dominated strategies results in a single strategy-profile.

Definition 2. *A game G is said to be max-solvable if there is a sequence of games G_0, \dots, G_r such that:*

- $G_0 = G$
- For every $k \in \{0, \dots, r-1\}$, G_{k+1} is a subgame of G_k achieved by removing a max-dominated strategy from the strategy space of one player in G_k .
- The strategy space of each player in G_r is of size 1.

The class of max-solvable-games contains all strictly-dominance-solvable ones. We shall refer to an elimination order of max-dominated strategies, that results in a single strategy-profile as an *elimination sequence* of a max-solvable game.

3.2 Asynchronous Best-Reply Dynamics and Max-solvable Games

One of the helpful features of max-solvable games is the fact that such games always have a unique pure Nash equilibrium.

Proposition 1. *Any max-solvable game has a unique pure Nash equilibrium.*

We now show that in max-solvable games, best-reply dynamics always converge to the unique pure Nash equilibrium, even in asynchronous settings. How long does this take? Answering this question requires further clarifications as we must account for the fact that update messages can be arbitrarily delayed, and that players might be prevented from best-replying for long periods of time. We define an *asynchronous phase* to be a period of time in which *every* player is activated at least once, and every player receives at least one update message from *each* of his neighbours. We prove that, for any asynchronous schedule, best-reply dynamics converge to the unique pure Nash equilibrium in a number of asynchronous phases that is at most $\sum_i m_i$, where m_i is the size of the strategy space of the i 'th player.

Theorem 1. *In any max-solvable game, best-reply dynamics converges for every asynchronous schedule within $\sum_i m_i$ asynchronous phases.*

Proof. Consider an elimination sequence of max-dominated strategies that results in a single strategy-profile. Let strategy s_1 of some player i be the first strategy to be eliminated. Player i is activated once during the first asynchronous phase. If he is playing s_1 then he will switch to another strategy since s_1 is max-dominated. Furthermore, no best-reply of player i in the future will *ever* cause him to choose strategy s_1 . From this point onwards, the best-reply dynamics are effectively occurring in a game where s_1 does not exist. Let us now consider the next strategy in the elimination order s_2 , which belongs to some player j (that can be i , or some other player). Given that player i never plays s_1 , s_2 is now max-dominated. Player j is activated during the second asynchronous phase. If he is playing s_2 he will move to another strategy. No matter what, s_2 will never be played again. More generally, after k asynchronous phases the k 'th strategy in the elimination order will never be played again. Therefore after $\sum_i (m_i - 1)$ asynchronous phases we are bound to reach the pure Nash equilibrium, which is the remaining strategy-profile.

3.3 Weakly-Max-solvable-Games

The definition of max-dominated strategies required that, for any strategy-profile of the other players, a max-dominated strategy be *strictly* worse than another strategy. In this section we discuss the case of ties.

Definition 3. A strategy $s_i \in S_i$ is weakly-max-dominated if for every strategy-profile of the other players $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ there is another strategy s'_i such that $u_i(s'_i, s_{-i}) \geq u_i(s_i, s_{-i})$.

Now, we can define *weakly-max-solvable games* as games in which the iterative removal of weakly-max-dominated strategies results in a single strategy-profile. Observe that any weakly-dominance-solvable game is a weakly-max-solvable game. Unfortunately, as the following example demonstrates, best-reply dynamics are not guaranteed to converge even in weakly-dominance-solvable games.

Example 1. Consider the game depicted by the following matrix (the rows are player 1's strategies and the columns are player 2's strategies):

1,1	0,0
1,0	0,1
0,1	1,0

First, observe that this is indeed a weakly-dominance-solvable game. Observe that if the initial strategy-profile is the leftmost entry in the lower row (row 3) of the game-matrix, then the following best-reply dynamics is possible: Player one moves from row 3 to row 2, player 2 moves from the left column to the right one, player 1 moves from row 2 to row 3, player 2 moves from the right column to the left one, and so on.

Weakly-dominance-solvable games always have pure Nash equilibria. As we have just seen, best-reply dynamics are not guaranteed to converge to such an equilibrium. Is there a different procedure that can do so in reasonable time? We prove the following impossibility result:

Theorem 2. Finding a pure Nash equilibrium in games that are weakly-dominance-solvable requires communicating exponentially many bits (in $\sum_i m_i$).

4 Potential Games and Asynchrony

In this section we explore better-reply dynamics in the context of potential games. While it is easy to see that in potential games better-reply dynamics converge for any synchronous schedule, what happens in simultaneous and asynchronous environments? We study the structural properties of potential games for which convergence of better-reply dynamics in simultaneous settings is assured.

We prove the following theorem:

Theorem 3. *If every subgame of a potential game has a unique pure Nash equilibrium, then better-reply dynamics converge for any simultaneous schedule.*

We show that the uniqueness of pure Nash equilibria in every subgame of a potential game is almost a characterization of potential games for which better-reply dynamics always converge in simultaneous settings. We show this by proving that this is indeed also a necessary condition for a large subclass of potential games, we term “strict potential games”.

Definition 4. *A game G is strict if for any two strategy profiles $s = (s_1, \dots, s_n)$ and $s' = (s'_1, \dots, s'_n)$, such that there is some $j \in [n]$ for which $s' = (s'_j, s_{-j})$, $u_j(s) \neq u_j(s')$.*

That is, a game is strict if for any player i , for any two strategies of that player $s_i, s'_i \in S_i$, and for any strategy-profile of the other players s_{-i} , i strictly prefers one strategy over the other. A strict potential game is a potential game that is strict.

Theorem 4. *If a strict potential game is such that better-reply dynamics converge for any simultaneous schedule, then every subgame of that games has a unique pure Nash equilibrium.*

Remark 1. One might hope that any strict game in which every subgame has a unique pure Nash equilibrium is a potential game. However, in the full version [7] of the paper we give an example that shows that this is not the case.

What about asynchronous settings? We now show that the property that guarantees the convergence of best-reply dynamics in a potential game (i.e., that every one of its subgames has a unique pure Nash equilibrium) does not necessarily guarantee convergence in asynchronous schedules.

Example 2. Consider the game described by Fig. 1. The arrows describe the better-replies of players from any strategy-profile (an arrow between strategy-profiles denotes the transition caused by a best-reply update of a single player).




Fig. 1. A game in which better-reply dynamics might diverge for some asynchronous schedule

The reader can verify that this is a potential game and that every subgame has a unique Nash equilibrium. Recall, that in asynchronous settings, the Scheduler may delay messages. We shall show that better-reply dynamics may never converge in such settings. Let us show such an oscillation (messages arrive immediately unless specifically noted): We begin with state A and allow the row player to update his strategy and notify everyone, thus arriving at state C . We then activate the column player and the matrix player simultaneously and arrive at state H . However, we delay the message sent to the row player by the matrix player so that the row player in fact believes we are in state D . We then activate the row player and allow him another update. He believes he moves to state B while in fact we arrive at state F . We then release the message to the row player and invoke the column player which updates his strategy from F to E . Then, the matrix player is activated and we return to state A . Repeating this over and over gives a permanent oscillation.

Acknowledgements

The first and second authors are supported by a grant from the Israeli Academy of Sciences.

References

1. Levin, H., Schapira, M., Zohar, A.: Interdomain routing and games. In: Proceedings of STOC 2008 (2008)
2. Osborne, M.J., Rubinstein, A.: A Course in Game Theory. MIT Press, Cambridge (1994)
3. Conitzer, V., Sandholm, T.: Communication complexity as a lower bound for learning in games
4. Hart, S., Mansour, Y.: The communication complexity of uncoupled Nash equilibrium procedures. In: Proceedings of STOC 2007 (2007)
5. Monderer, D., Shapley, L.: Potential games. *Games and Economic Behavior* (14), 124–143 (1996)
6. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* (2), 65–67 (1973)
7. Nisan, N., Schapira, M., Zohar, A.: Asynchronous best-reply dynamics. Technical report, The Leibniz Center for Research in Computer Science (2008)

Fault Tolerance in Distributed Mechanism Design

Ronen Gradwohl*

Department of Computer Science and Applied Mathematics
The Weizmann Institute of Science, Rehovot, 76100 Israel
`ronen.gradwohl@weizmann.ac.il`

Abstract. We argue that in distributed mechanism design frameworks it is important to consider not only rational manipulation by players, but also malicious, faulty behavior. To this end, we show that in some instances it is possible to take a centralized mechanism and implement it in a distributed setting in a fault tolerant manner. More specifically, we examine two distinct models of distributed mechanism design – a Nash implementation with the planner as a node on the network, and an ex post Nash implementation with the planner only acting as a “bank”. For each model we show that the implementation can be made resilient to faults.

1 Introduction

In the standard Mechanism Design setting a centralized planner wishes to implement an objective function of some players’ private information (their types). The players are assumed to act rationally, and so the planner must provide incentives for the players in such a way that they are motivated to reveal their types truthfully.

Recent research in Multi-Agent Systems, Distributed Artificial Intelligence, and Distributed Algorithmic Mechanism Design, has considered the implementation problem in a distributed setting: There is no centralized planner, and instead the players reside on some network. A mechanism in this setting no longer consists of simple information-revelation to a planner, but rather involves communication between and computation by the players. The players follow some distributed protocol, and the result of the protocol’s execution provides the information necessary for the computation of the objective function. As in the centralized setting, players are assumed to follow a protocol only if it is in their best interest to do so. Thus, we desire protocols that are faithful – a player should not be able to gain by deviating from any part of the protocol. The protocol itself should be in an equilibrium.

The design of distributed protocols is largely the domain of distributed computing research. This field, however, is not at all concerned with the faithfulness of protocols – non-faulty players are assumed to follow the protocol blindly – and

* Research supported by US-Israel Binational Science Foundation Grant 2006060.

instead focuses on their fault tolerance. There are no assumptions about players' incentives in following a protocol, but rather a desire to make the protocol tolerant of arbitrary and possibly Byzantine faults.

A model that contains Byzantine players is often well-motivated, since computers that execute a protocol may crash and communication between parties might be faulty. This motivation applies equally well in a distributed mechanism design setting. Most of the work in this field, however, has modelled players only as being rational, and this precludes the possibility of faulty behavior.

So what happens if Byzantine players appear in a distributed implementation? Suppose a planner implements an objective function in a distributed setting. Players follow a faithful protocol because it serves their interests. If some player acts in an adversarial manner and deviates from the protocol then, even though he does not gain anything by this deviation (since following the protocol is optimal), other players may lose quite a bit. It is possible and even likely that, given this player's fault, some other player would have had a better strategy, and that the protocol he is supposed to follow is no longer optimal for him.

Thus, we would like the protocol to be more than just faithful – we would like it to be faithful *even in the presence of faults*. That is, following the protocol should be an optimal strategy for players, even if some of the players act maliciously.

1.1 Our Results

The main thesis of this paper is that fault tolerant faithfulness is very relevant for distributed mechanism design, and that it can be implemented. Implementability is demonstrated via two examples in different models for distributed implementation, and in each model we show that sometimes mechanisms can be implemented in a fault tolerant manner. Our results are of the following form: given an objective function f implementable in a centralized setting with some equilibrium concept E_1 (i.e. some notion of what it means for a player to play “optimally”), we show that it can be implemented in a distributed setting with an equilibrium concept E_2 . Note that E_2 can never be “stronger” than E_1 , and our goal is to have $E_1 = E_2$. In the first model, we show that if E_1 is some robust notion of Nash equilibria (more formally defined in Section 2), then we can get $E_2 = E_1$. In the second model, we examine the case where E_1 is a dominant-strategy equilibrium, and note that all previous research on this problem has considered the case where E_2 is an ex post Nash equilibrium, which is much weaker than E_1 (the reason being that it is in general impossible to obtain the latter – see more on this in Section 4). We bridge this gap by showing that E_2 can be a fault tolerant version of ex post Nash, which may be viewed as something between an ex post Nash and a dominant-strategy equilibrium.

We now describe the two models and our respective contributions.

Distributed Nash Implementation. The standard mechanism design problem studied in economics is one in which players' types are drawn from a commonly-known distribution, and the planner wishes to implement an objective function in a Nash equilibrium: Truth revelation should maximize players' expected utilities, where the expectation is over the distributions over the types. Monderer

and Tennenholtz [15] generalize this problem to a distributed setting, in which there is some network $G = (N, L)$, and the planner is a node on the network. Unlike the standard setting, the planner may not be able to communicate directly with all the players. Thus, players must pass messages to the planner through other players. This means that the entire protocol (information-revelation and message-passing) must be in a Nash equilibrium. Monderer and Tennenholtz [15] show that this can be done if the network is connected enough. In Section 3 we extend their results, and show that sometimes such a protocol can be implemented even in the presence of faulty players.

Distributed Ex Post Implementation. We examine a model in which the players are once again nodes on a network $G = (N, L)$, but this time there is no planner that participates in the computation. In fact, there is no planner at all, only a centralized “bank” responsible for executing the result of the computation, such as distributing the eventual payments. The players execute the protocol, and at the end each player sends his computation to the bank. If all players agree on the computation then the bank executes the payment distribution.

Research in algorithmic mechanism design has largely avoided making assumptions on the distribution of players’ types and the knowledge players have about this distribution. For direct-revelation mechanisms players’ types can be arbitrary, and mechanisms should be strategyproof – this means that truth revelation is an optimal strategy for every player regardless of others’ types. In the distributed setting described above, the literature has converged to the notion of ex post Nash as the correct solution concept. In an ex post Nash equilibrium it is optimal for players to follow the protocol regardless of the other players’ types, but assuming that others *do* follow the protocol. In Section 4 we show that sometimes we can get a fault tolerant ex post Nash equilibrium, in which we can drop the assumption that *all* players follow the protocol (and allow up to $n/3$ faulty players to act arbitrarily).

Related Work. Our results on Nash implementation are direct extensions to the work of Monderer and Tennenholtz [15]. Also closely related is work done on the implementation of mediators, most notably by Abraham et al. [2].

Our results on ex post implementation are motivated by the work of Shneidman and Parkes [19]. There has been quite a bit of additional work in this model, with various applications: For example, Feigenbaum et al. [7] and Shneidman and Parkes [19] studied the interdomain routing problem, and Parkes and Shneidman [17] and Petcu et al. [16] studied VCG mechanisms (see also Feigenbaum et al. [8]). There are two main differences between our work and the other work in this model. First, all previous work has concentrated on implementing an ex post Nash equilibrium, and we show that it is actually possible to obtain a fault tolerant version of this. Second, previous work has had an additional goal of minimizing the network complexity, an issue that we completely ignore in the current paper (in fact, combining fault tolerance with low network complexity is a main open question that arises from this work).

Notions of fault tolerance like those in this paper have appeared before in a few centralized implementation frameworks. For example, Eliaz [5] shows how to Nash implement a mechanism in a fault tolerant manner. Also, Aiyer et al. [1] implement a protocol for a specific objective function in an asynchronous setting, where the protocol is tolerant of Byzantine, altruistic, and rational behavior.

Finally, numerous researchers have expressed the need to incorporate faulty or malicious behavior into distributed settings with rational players (see the surveys of Halpern [12,13] and Feigenbaum and Shenker [9]).

2 Definitions

There is a set N of n players, and each player i has some private information $\theta_i \in \Theta_i$ called player i 's type. There is also a planner who wishes to implement some function $f(\theta) \in \mathcal{O}$, where $\theta = (\theta_1, \dots, \theta_n)$. We use the following standard notation: for a set $T \subseteq N$, $\theta_T = (\theta_i)_{i \in T}$. Also, $\theta_{-T} : \theta'_T = (\hat{\theta}_1, \dots, \hat{\theta}_n)$, where $\hat{\theta}_i = \theta'_i$ if $i \in T$ and $\hat{\theta}_i = \theta_i$ otherwise.

In a direct-revelation mechanism $M = (f, \Theta)$ players report types $\hat{\theta} \in \Theta = \Theta_1 \times \dots \times \Theta_n$ to the planner, who then computes the outcome $f(\hat{\theta})$. Players are modelled as rational utility maximizers, and so each player i has a utility function $u_i : \Theta_i \times \mathcal{O} \mapsto \mathbb{R}$ such that $u_i(\theta_i, o_1) > u_i(\theta_i, o_2)$ if and only if player i of type θ_i prefers outcome o_1 to outcome o_2 . Generally, incentives are provided so that players maximize their utilities by reporting $\hat{\theta} = \theta$ in equilibrium.

There are numerous types of equilibria in which we may be interested. In the following definition of a Nash equilibrium and a t -tolerant Nash equilibrium, we assume that players' types $\theta = (\theta_1, \dots, \theta_n)$ are drawn from a commonly-known prior product distribution $\Theta = \Theta_1 \times \dots \times \Theta_n$. The expectations in the definition are taken over this distribution.

Definition 1 (t -tolerant Nash equilibrium). *A direct-revelation mechanism M implements f in a t -tolerant Nash equilibrium if for every set $T \subset N$ of size $|T| \leq t$, every strategy θ'_T , every player $i \notin T$, and every $\theta_i, \theta'_i \in \Theta_i$,*

$$E [u_i(\theta_i, f(\Theta_{-(T \cup \{i\})} : \theta'_T, \theta_i))] \geq E [u_i(\theta_i, f(\Theta_{-(T \cup \{i\})} : \theta'_T, \theta'_i))].$$

M implements f in a Nash equilibrium if the above holds with $t = 0$.

There are other variants of Nash equilibria that handle faults and coalitions. For example, in a t -strong Nash equilibrium (see Aumann [3]), no coalition of players can increase the utility of all colluding players. In a t -coalitional Nash equilibrium, no single player can corrupt a coalition of at most t others to increase his own payoff. We note that our results apply to these notions as well – see details in the full version [11].

A different notion of equilibrium which does not require assumptions about the knowledge of players is that of a (weakly) dominant-strategy equilibrium, in which truth revelation is a player's optimal strategy regardless of other players' strategies.

Distributed Mechanisms. In a distributed setting, the mechanism is no longer a one-shot game in which each player simply sends his type to the center. Instead, the game may proceed in many stages, and information may also be transferred between the players. A player’s strategy is no longer simply an act of revealing some information, but rather a protocol specifying how to act in many possible situations. There are several ways of formalizing this, and much of the notation we use here is from Shneidman and Parkes [19].

Shneidman and Parkes [19] divide the kinds of moves players can make into three categories: Information-revelation actions, in which players reveal some of their own private information (which may be partial and untruthful), message-passing actions, in which players pass on messages received from other players to one or more of their neighbors, and computational actions, in which players perform some computation that affects the outcome rule.

In a centralized protocol, the only relevant actions are information-revelation. In such a setting players do not interact with one another but only with the center, and so all messages are passed directly to the planner. Additionally, the planner does all the computation and only requires knowledge of players’ types, and so the players themselves do not do any of the computation. In a distributed setting, however, the latter two types of actions affect the outcome of the protocol, and so players may desire to deviate from the protocol in an attempt to increase their utility.

For player i of type θ_i , we denote a possible strategy in such a setting as $s_i(\theta_i)$. Throughout the paper we often refer to such strategies as protocols that players are encouraged to follow. Essentially, a strategy is a recommended action (information-revelation, message-passing, or computational) for every situation (e.g. as a response to any incoming message). Given a finite sequence of actions specified by strategies $s = (s_1, \dots, s_n)$, the final choice of an outcome is given by a function $g(s(\theta)) = g(s_1(\theta_1), \dots, s_n(\theta_n)) \in \mathcal{O}$. The function g can be viewed as a protocol for the planner: he obtains some elements of the computation, applies his protocol g to the computation, and hopefully g is such that $g(s(\theta)) = f(\theta)$.

All the solution concepts above apply in a distributed setting as well. For example, strategies for players in a distributed setting are a t -tolerant Nash equilibrium if following the strategy is expected utility maximizing for every player assuming all but t players follow the strategy, and t players play arbitrarily. A different solution concept which has been widely adopted for distributed mechanisms is that of an ex post Nash equilibrium. In this solution concept we make no assumptions on players’ prior distribution on types (as in dominant-strategy equilibria), but we do assume players follow the protocol, given their true type. This distinction between dominant-strategy equilibria and ex post Nash equilibria is not so clear in direct-revelation mechanisms, since in such mechanisms the optimal strategy is to truthfully reveal the type. In distributed mechanisms, however, this is an important difference.

Definition 2 (t -tolerant ex post Nash equilibrium). *Strategies $s^* = (s_1^*, \dots, s_n^*)$ implement f in a t -tolerant ex post Nash equilibrium if there exists a function g satisfying $g(s^*(\theta)) = f(\theta)$, such that for every set $T \subset N$ of size*

$|T| \leq t$, every player $i \notin T$, every alternate strategies s'_T and s'_i , and every type profile $\theta = (\theta_1, \dots, \theta_n)$,

$$u_i(\theta_i, g(s^*(\theta)_{-T} : s'_T(\theta_T))) \geq u_i(\theta_i, g(s^*(\theta)_{-(T \cup \{i\})} : s'_T(\theta_T), s'_i(\theta_i))).$$

s^* implements f in an ex post Nash equilibrium if the above holds for $t = 0$.

We need two more definitions:

Definition 3 (network-oblivious strategy). A strategy s_i is network-oblivious if s_i does not depend on the structure of the network.

Definition 4 (k -connected network). A network $G = (N, L)$ is k -connected if for every two distinct nodes $i, j \in N$ there exist at least k node-disjoint paths from i to j in G .

3 Distributed Nash Implementation

Monderer and Tennenholtz [15] consider the following setting: there is some network $G = (N, L)$, and the planner is a node on the network. The planner wishes to implement some function in a Nash equilibrium, but may not have direct communication with all the players. Thus, players must pass messages to the planner through other players. This means that the entire protocol (information-revelation and message-passing) must be in a Nash equilibrium. Monderer and Tennenholtz [15] prove the following theorem.

Theorem 1 ([15]). Let f be some Nash implementable objective function and let $G = (N, L)$ be a 2-connected network. Then there is a protocol s^* that implements f on G in a Nash equilibrium. If G is a ring (a connected network in which each player has exactly two neighbors) and f is implementable in a strong Nash equilibrium, then there is a protocol $\overline{s^*}$ that implements f on G in a strong Nash equilibrium. The lengths of messages in s^* and $\overline{s^*}$ are equal to the number of bits needed to represent players' prior distributions, and the number of rounds is constant.

We prove the following theorem.

Theorem 2. Fix a nonnegative integer t . Let f be an objective function implementable in a t -robust Nash equilibrium, where robust $\in \{\text{tolerant, strong, coalitional}\}$. Let $G = (N, L)$ be a $(3t + 4)$ -connected network. Then there is a protocol s^* that implements f on G in a t -robust Nash equilibrium. The lengths of messages in s^* are $\text{poly}(n, \max_i \log |\Theta_i|)$, and the number of rounds is constant.

Note that for $t = 0$, we get a result on Nash implementation comparable to the first part of Theorem 1: the message lengths can be shorter – and independent of the prior probability distribution – at the cost of requiring 4-connectivity rather than 2-connectivity of the underlying network. We note that it is actually possible to reduce this to 3-connectivity if communication between the players

and the planner can be bidirectional. Also, this result with t -strong Nash is comparable to the second part of Theorem 1 we have to bound the size of coalitions, but our result holds for a much more general class of networks (and not just a ring).

Proof (Outline). Theorem 2 relies on protocols of Dolev et al. [4] for Secret Message Transmission (SMT). Consider two entities, a sender S and a receiver R , connected by w wires. S has a secret message m drawn from some finite set which he wishes to send to R via the w wires. The problem is that some wires may be controlled by various adversaries, and so the communication should satisfy the following two properties: Secrecy – for all sets L of at most σ wires, no party listening to all wires in L can learn anything about m – and Resiliency – for all sets D of at most ρ wires, R learns m even if an adversary controls and coordinates the behavior of wires in D . Dolev et al. [4] give a protocol for (σ, ρ) -SMT as long as $w \geq \sigma + 2\rho + 1$. They also show that such a bound is necessary.

We now describe the protocol s_i^* that player i should follow.

- Computation and information-revelation: Upon learning his type $\theta_i \in \Theta_i$, player i :
 1. Chooses $3t + 4$ node-disjoint paths from i to the planner.
 2. Follows the protocol of [4] for $(t + 1, t + 1)$ -SMT to transmit θ_i to the planner via the chosen paths. Each message sent along a path contains the list of nodes to be traversed along that path.
- Message-passing: Upon receiving a message m from a neighbor and a list of nodes that the message should traverse, player i passes m on to the node following i on the list.

The planner decodes the values sent to him (or assigns a default value if no decoding is possible).

To see that this protocol is expected utility maximizing for player i even in the presence of up to t faulty players, note the following. All other non-faulty players will safely transmit their values to the planner regardless of the message-passing actions of the faulty players or player i . Additionally, player i can not learn anything about the values sent by the other non-faulty players, even if the t faulty ones “accidentally” send him information. Thus, from his perspective, the distribution over non-faulty players’ types sent to the planner is as it was originally. Because f is a t -tolerant Nash equilibrium, his optimal strategy is to also send his true type to the planner. See [11] for a more detailed proof.

Dominant-Strategy Equilibria. Suppose we wish to define a protocol for a distributed t -tolerant ex post implementation of a mechanism that has a dominant-strategy implementation. Then we can actually do this in a network-oblivious manner via the following SMT protocol, whose proof is deferred to [11].

Proposition 1. *For any $(2t + 1)$ -connected network G , there is a network-oblivious protocol for communication between any two nodes with resiliency t .*

4 Distributed Ex Post Nash Implementation

In this section we examine a model studied by Parkes and Shneidman [19,17] and Petcu et al. [16]. Once again the players are nodes on a network $G = (N, L)$, but this time there is a centralized “bank” responsible for executing the result of the computation, such as distributing the eventual payments. The players follow a protocol to share information and make computations, and then each player sends his computation to the bank. If all players agree on the computation the bank executes the payments. We assume that if the players do not agree, the bank takes no action, and this is an undesirable outcome for the players. In the presence of faulty players, it is clearly impossible to get all the players to agree on the computation, so we allow the bank to execute the result of the computation if a *majority* of the players agree.

Feigenbaum et al. [8] observe that it is impossible to have a dominant-strategy equilibrium here – if all players but one send an incorrect (but identical) value to the bank, then an optimal strategy for the one player is to lie as well (and so this is not incentive compatible). Actually, the same argument shows that it is impossible to have t -tolerant ex post Nash either, for any $t \geq n/2$. It is straightforward to see that $t \geq n/3$ is also impossible, and this follows directly from lower bounds on the resiliency of Byzantine Agreement protocols (see the lower bounds of Pease et al. [18] and Karlin and Yao [14]).

We prove two theorems that provide solutions for the distributed ex post implementation problem with a nearly optimal bound on the number of faulty players. The protocol of Theorem 3 runs in fewer rounds, whereas the protocol of Theorem 4 has the advantage of being network-oblivious.

Theorem 3. *For any $t < n/3 - 1$, objective function f implementable with t faults in dominant strategies, and $(2t + 3)$ -connected network G there exists a protocol s^* that faithfully implements f in a t -tolerant ex post Nash on G . The communication of s^* is polynomial, and its expected number of rounds is $O(1)$.*

Theorem 4. *For any constant $\varepsilon > 0$, $t < n/(3 + \varepsilon)$, objective function f implementable with t faults in dominant strategies, and $(2t + 3)$ -connected network G there exists a network-oblivious protocol s^* that faithfully implements f in a t -tolerant ex post Nash on G . The communication of s^* is polynomial, and its expected number of rounds is $O(\log n)$.*

Essentially, each player uses an SMT protocol to send his type to all others, and then they all follow a Byzantine Agreement protocol (see Pease, Shostak and Lamport [18]) to reach agreement among the non-faulty players. Theorem 3 uses the SMT protocol of [4] to emulate secret communication, followed by the Byzantine Agreement protocol of Feldman and Micali [6] (which requires secret communication). Theorem 4 uses the network-oblivious SMT protocol of Proposition 1 followed by the Byzantine Agreement protocol of Goldwasser et al. [10] (which works in the full-information model, without assuming secret communication). See the full version of this paper [11] for more details.

References

1. Aiyer, A., Alvisi, L., Clement, A., Dahlin, M., Martin, J.-P., Porth, C.: BAR fault tolerance for cooperative services. In: Proceedings of 20th ACM Symposium on Operating Systems Principles, pp. 45–58 (2005)
2. Abraham, I., Dolev, D., Gonen, R., Halpern, J.: Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In: Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing, pp. 53–62 (2006)
3. Aumann, R.J.: Acceptable points in general cooperative n -person games. In: Contributions to the Theory of Games IV. Annals of Mathematical Studies, pp. 287–324 (1959)
4. Dolev, D., Dwork, C., Waarts, O., Yung, M.: Perfectly secure message transmission. *Journal of the ACM* 40(1), 17–47 (1997)
5. Eliaz, K.: Fault-tolerant implementation. *Review of Economic Studies* 69(3), 589–610 (2002)
6. Feldman, P., Micali, S.: An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM J. Comput.* 26(4), 873–933 (1997)
7. Feigenbaum, J., Papadimitriou, C., Sami, R., Shenker, S.: A BGP-based mechanism for lowest-cost routing. *Distributed Computing* 18, 61–702 (2005)
8. Feigenbaum, J., Schapira, M., Shenker, S.: Distributed algorithmic mechanism design. In: *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2008)
9. Feigenbaum, J., Shenker, S.: Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. In: 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, pp. 1–13 (2002)
10. Goldwasser, S., Pavlov, E., Vaikuntanathan, V.: Fault-tolerant distributed computing in full-information networks. In: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (2006)
11. Gradwohl, R.: Fault tolerance in distributed mechanism design (manuscript), <http://www.wisdom.weizmann.ac.il/~rgradwoh/files/ft-mech.pdf>
12. Halpern, J.: A computer scientist looks at game theory. *Games and Economic Behavior* 45(1), 114–131 (2003)
13. Halpern, J.: Computer science and game theory: A brief survey. In: *The New Palgrave Dictionary of Economics* (to appear)
14. Karlin, A., Yao, A.: Probabilistic lower bounds for byzantine agreement (manuscript, 1986)
15. Monderer, D., Tennenholtz, M.: Distributed games: from mechanisms to protocols. In: *AAAI*, pp. 32–37 (1999)
16. Petcu, A., Faltings, B., Parkes, D.: MDPOP: Faithful distributed implementation of efficient social choice problems. In: *AAMAS* (2006)
17. Parkes, D., Shneidman, J.: Distributed implementation of Vickrey-Clarke-Groves mechanisms. In: *AAMAS* (2004)
18. Pease, M., Shostak, R., Lamport, L.: Reaching agreement in the presence of faults. *Journal of the ACM* 27, 228–234 (1980)
19. Shneidman, J., Parkes, D.: Specification faithfulness in networks with rational nodes. In: *PODC 2004* (2004)

Bargaining Solutions in a Social Network

Tanmoy Chakraborty and Michael Kearns

Department of Computer and Information Science
University of Pennsylvania

Abstract. We study the concept of bargaining solutions, which has been studied extensively in two-party settings, in a generalized setting involving arbitrary number of players and bilateral trade agreements over a social network. We define bargaining solutions in this setting, and show the existence of such solutions on all networks under some natural assumptions on the utility functions of the players. We also investigate the influence of network structure on equilibrium in our model, and note that approximate solutions can be computed efficiently when the networks are trees of bounded degree and the parties have nice utility functions.

1 Introduction

Bargaining has been studied extensively by economists and sociologists, and the most studied setup consists of two parties A and B , with utility functions \mathcal{U}_A and \mathcal{U}_B , negotiating a bilateral deal. The deal, if agreed to by both parties, yields some fixed profit c . Such a scenario arises if two persons want to go into some business as partners. A and B also have *alternate options* α_A and α_B respectively, which is the amount of money they receive if the deal is not agreed upon. The negotiation involves how the profit from this deal is divided between the two parties. The final share that both parties agree to receive from the deal constitutes a *bargaining solution*.

Several bargaining solutions, which are predictions of how the profit will be shared, have been proposed by economists, the most well-known being the *Nash Bargaining Solution (NBS)* [1], which states that the bargaining solution finally adopted will be one that maximizes the product of the differential utilities of this deal to each party. The *differential utility* of A from the deal is the utility A receives by agreeing to the deal in excess of what it would receive without agreeing to the deal, that is, $\mathcal{U}_A(x) - \mathcal{U}_A(\alpha_A)$, where x is the share of profit A gets from the deal. Similarly, the differential utility of B is $\mathcal{U}_B(c - x) - \mathcal{U}_B(\alpha_b)$. NBS seeks to maximize $(\mathcal{U}_A(x) - \mathcal{U}_A(\alpha_A))(\mathcal{U}_B(c - x) - \mathcal{U}_B(\alpha_b))$.

Another extensively studied bargaining solution concept, known as the *Proportional Bargaining Solution (PBS)*, seeks to maximize the minimum of the differential utilities of the parties, that is, $\min\{\mathcal{U}_A(x) - \mathcal{U}_A(\alpha_A), \mathcal{U}_B(c - x) - \mathcal{U}_B(\alpha_b)\}$. There is a crucial axiomatic difference between the concepts of NBS and PBS – in fact, they are representatives from two broad classes of bargaining solution concepts that have been formulated and studied in literature (see Chapter 2 of

[2] for a discussion). One of the major axioms satisfied by NBS is that the bargaining solution should not be altered if the scale of the utility functions of the parties are altered by arbitrary constant factors. In other words, NBS is based on the axiom that utilities of different parties cannot be compared. However, this axiom is highly debated and several solutions that neglect this axiom, and instead choose to make *interpersonal comparison of utility*, have been proposed, and PBS is one of the most extensively studied among these solutions.

In this paper, we consider a generalization of the above two-party setting to a setting that involves arbitrary number of parties, but where the deals are still bilateral, and the alternate options are all zero. The parties shall be represented as vertices of a social network, where the edges represent bilateral deals that shall be negotiated, and weights on the edges represent the total profits from each deal. Thus the input to the problem, which we call the *network bargaining problem*, is an undirected graph with weights on edges, and an efficiently computable utility function for every vertex. Different deals may have different profits, which are represented by weights on edges of the input graph. A solution to the network bargaining problem is a prediction of how the profits on each edge is divided. We will primarily be interested in studying the effects of network topology on the solution, and so we shall often restrict our study to the case where the edges have unit weight, and all vertices have the same utility function. Effects of network topology on solutions of various network exchange models have been studied, theoretically as well as through human subject experiments ([3,4,5,6]). Our goal is to develop a bargaining solution concept for the network bargaining problem, that will have a strong intuitive justification.

Braun and Gautschi [5] studied the network bargaining problem, and proposed a solution. Their solution is a direct generalization of the *weighted Nash Bargaining Solution* for bilateral deals. They assign a numerical bargaining power to each vertex based solely on its degree and the degrees of its neighbors, and also assume linear utility functions, and then negotiate each edge independently according to the bargaining powers. Kleinberg and Tardos [6] studied a variant of the network bargaining problem, which has the same input as our problem, but the solution has a restriction that each vertex can agree to a deal with at most one of its neighbors. They define an equilibrium-based solution, which they call a balanced outcome, where the agreement on every deal is required to meet a stability condition. The stability condition used for each edge in [6] is the NBS in a two-party setting with intuitively defined alternate options and linear utility functions. This model was also studied previously by Cook and Yamagishi [7]. In contrast to the model of [5], this model allows the equilibrium conditions and network topology to naturally exhibit bargaining power, instead of directly assigning a value.

Inspired by the notion of balanced outcomes in [7,6], our solution for the network bargaining problem is also an equilibrium-based concept. We propose that the bargaining solution should be stable, and so no party should be keen on renegotiating a deal. For an edge $e = (u, v)$, we define the alternate options of u and v to be the total profits received by u and v , respectively, from deals

with their other neighbors. The differential utility of the deal on e to u and v is now intuitively clear (see Section 2 for definition). Renegotiation occurs if the current division does not satisfy a standard two-party bargaining solution. In a bargaining state (specifying the division of profits on each edge), the deal on an edge is *stable* if the division of profits satisfy the two-party bargaining solution. We say that a state is an *equilibrium* if all edges are stable, and this is our bargaining solution. Depending on the two-party bargaining solution used for renegotiating edges, we have thus proposed two bargaining solutions, the *NBS equilibrium* and *PBS equilibrium*. These bargaining solutions are formally defined in Section 2.

A question crucial to the applicability of our model is to characterize structures (networks and utility functions) in which there exists an equilibrium. In this paper, we completely characterize the PBS and NBS equilibria on every social network when all the vertices of the network have linear utility functions (functions with constant marginal utility). In this case, we show that there is a unique PBS equilibrium and a unique NBS equilibrium in every social network, and that the network topology has no influence on the solutions. We also show that on any network, there exists a PBS equilibrium if all the utility functions are increasing and continuous. Further, we show that on any network, there exists an NBS equilibrium if all the utility functions are increasing, concave and twice differentiable.

The rest of this paper is organized as follows. Section 2 contains a formal introduction to the model and some basic lemmas that are applicable to two-party settings. Section 3 characterizes equilibria in our model when the utility functions are linear. Section 4 contains the proof of existence of PBS and NBS equilibria on all networks, for broad classes of utility functions. Section 5 provides some results about the effect of network structure on NBS equilibrium. And finally, Section 6 briefly describes an efficient algorithm to compute approximate PBS equilibria on trees with bounded degree and a specific utility function $\log(1+x)$.

2 Preliminaries

The input to the *network bargaining problem* consists of an undirected graph $G(V, E)$ with n vertices and m weighted edges, where vertices represent people and edges represent possible bilateral trade deals, and a utility function U_v for each vertex v . The utility functions are all represented succinctly and are computable in polynomial time.

Let $c(e)$ be the weight of an edge e in G . Let $e_1, e_2 \dots e_m$ be an arbitrary ordering of the edges in $E(G)$. For every edge e_i , we assign it an arbitrary direction, and refer to its end-points as u_i and v_i , such that e_i is directed from u_i to v_i . A *state of the bargaining model* is described by the division of profits on each edge of the graph. Let $x(u_i, e_i)$ and $x(v_i, e_i)$ denote the profits u_i and v_i receive from the agreement on the edge e_i , respectively. Note that $x(v_i, e_i) = c(e_i) - x(u_i, e_i)$. We shall represent a state of the bargaining model as a vector $s = (s_1, s_2 \dots s_m) \in \mathbb{R}^m$ such that $s_i = x(u_i, e_i)$. Note that s uniquely determines the division of profits on all edges.

Definition 1. Let $s \in \mathbb{R}^m$ be a state of the bargaining model for a graph G . For any vertex u and any edge e incident on u , let $\gamma_s(u)$ denote the total profit of a vertex u from all its deals with its neighbors. Let $x_s(u, e)$ denote the profit u gets from the agreement on edge e . Let $\alpha_s(u, e) = \gamma_s(u) - x_s(u, e)$ be the profit u receives from all its deals except that on e .

If the current state of the bargaining model is s , and $e = (u, v)$ is renegotiated, then we say that $\alpha_s(u, e)$ and $\alpha_s(v, e)$ are the *alternate options* for u and v respectively, that is, the amount they receive if no agreement is reached on the deal on e . We shall drop the suffix s if we make a statement for any arbitrary state, or if the state is clear from the context.

Definition 2. Let s be any state of the bargaining model. Let x be the profit of u from the deal on $e = (u, v)$. Then, the *differential utility* of u from this deal is $a_s(x) = \mathcal{U}_u(\alpha_s(u, e) + x) - \mathcal{U}_u(\alpha_s(u, e))$, and the *differential utility* of v from this deal is $b_s(x) = \mathcal{U}_v(\alpha_s(v, e) + c(e) - x) - \mathcal{U}_v(\alpha_s(v, e))$.

Definition 3. Let s be any state of the bargaining model. Define $y_s(u, e)$ to be the profit u would get on the edge $e = (u, v)$ if it is renegotiated (according to some two-party solution), the divisions on all other edges remaining unchanged. Also define $update(s, e) = |x_s(u, e) - y_s(u, e)|$.

If e is renegotiated according to the *Nash Bargaining Solution* (NBS), then $y_s(u, e)$ is a value $0 \leq x \leq c(e)$ such that the NBS condition is satisfied, that is, the function $W_N(x) = a_s(x)b_s(x)$ is maximized. Instead, if e is renegotiated according to the *Proportional Bargaining Solution* (PBS), then $y_s(u, e)$ is a value $0 \leq x \leq c(e)$ such that the PBS condition is satisfied, that is, function $W_P(x) = \min\{a_s(x), b_s(x)\}$ is maximized.

The following lemmas give simpler equivalent conditions for PBS and NBS under certain assumptions about the utility functions, and are also applicable to the two-party setting.

Lemma 1. *If the utility functions of all vertices are increasing and continuous, then the PBS condition reduces to the condition $a_s(x) = b_s(x)$, and there is a unique solution x satisfying this condition.*

Lemma 2. *Let the utility functions of all vertices be increasing, concave and twice differentiable,. Moreover, let $q_s(x) = \frac{a_s(x)}{a'_s(x)}$, and let $r_s(x) = -\frac{b_s(x)}{b'_s(x)}$. Then the NBS condition reduces to $q_s(x) = r_s(x)$, and there is a unique solution x satisfying this condition.*

Definition 4. *We say that an edge e is stable in a state s if renegotiating e does not change the division of profits on e , that is, $update(s, e) = 0$. We say that a state s is an equilibrium if all edges are stable. We say that s is an ϵ -approximate equilibrium if $update(s, e) < \epsilon$ for all edges e .*

We refer to an equilibrium as an *NBS equilibrium* if the renegotiations satisfy the NBS condition. We refer to the equilibrium as a *PBS equilibrium* if the renegotiations satisfy the PBS condition.

3 Linear Utility Functions: Characterizing All Equilibria

In this section, we characterize all possible NBS and PBS equilibria when all vertices have linear increasing utility functions, for every vertex v . Braun and Gautschi [5] make this assumption in their model, and so do Kleinberg and Tardos [6].

We show that in our model, if we make this assumption, there is a unique NBS equilibrium and a unique PBS equilibrium, and network topology has no influence on the division of profits on the deals at equilibrium. The following two theorems formalise these observations. Their proofs are simple, and are omitted due to lack of space.

Theorem 1. *Suppose all vertices have linear increasing utility functions. Then there is a unique NBS equilibrium, in which the profit on every edge is divided equally between its two end-points.*

Theorem 2. *Suppose all vertices have linear increasing utility functions. Let $U_i(x) = k_i x + l_i \forall i \in V(G)$. Then there is a unique PBS equilibrium, such that for any edge $e = (u, v)$, $x_s(u, e) = c(e) \frac{k_v}{k_u + k_v}$.*

4 Existence of Equilibrium for General Utility Functions

We now turn our focus towards non-linear utility functions. In this section, we prove that PBS and NBS equilibria exist on all graphs, when the utility functions satisfy some natural conditions. The proofs use the Brouwer fixed point theorem, and is similar to the proof of existence of mixed Nash equilibrium in normal form games.

Theorem 3. *PBS equilibrium exists on any social network when all utility functions are increasing and continuous. NBS equilibrium exists on any social network when all utility functions are increasing, concave and twice differentiable.*

Essentially, it is sufficient for the utility functions to satisfy the following general condition of continuity:

Condition 1. *Let s be any state of the bargaining model, and $e = (u, v)$ be an edge. For every $\epsilon > 0$, there exists $\delta > 0$ such that for any state t such that $|\alpha_t(u, e) - \alpha_s(u, e)| < \delta$ and $|\alpha_t(v, e) - \alpha_s(v, e)| < \delta$, we have $|y_t(u, e) - y_s(u, e)| < \epsilon$.*

Note that $y_s(u, e)$ and $y_t(u, e)$ are influenced both by the utility functions as well as the two-party solution concept that is used (NBS or PBS). Thus whether Condition 1 holds will depend on whether the renegotiations follow the NBS or the PBS condition, and also on the utility functions.

Lemma 3. *If Condition 1 holds for the NBS solution concept or the PBS solution concept, then NBS or PBS equilibrium exists, respectively.*

Proof. We define a function $f : [0, 1]^m \rightarrow [0, 1]^m$ that maps every state s to another state $f(s)$. Given s , we can construct the unique solution t such that the deal on an edge $e = (u, v)$ in t is the renegotiated deal of e in s , that is, $x_t(u, e) = y_s(u, e)$. We define $f(s)$ to be t . Thus, $f(s)$ is the “best-response” vector for s .

Clearly, s is an ϵ -approximate equilibrium if and only if $\|s - f(s)\|_\infty < \epsilon$. In particular, s is an equilibrium if and only if $f(s) = s$, that is, s is a fixed point of f . Also, $[0, 1]^m$ is a closed, bounded and convex set. So if f were continuous, then we can immediately use Brouwer fixed point theorem to deduce that the equilibrium exists. Thus the following claim completes the proof. \square

Claim. f is continuous if and only if Condition \square holds.

Proof. Suppose Condition \square holds for some ϵ and δ . Thus, if $\|s - t\|_\infty < \delta/n$, then $|\alpha_t(u, e) - \alpha_s(u, e)| < \delta$ and $|\alpha_t(v, e) - \alpha_s(v, e)| < \delta$, so, by Condition \square , $|y_t(u, e) - y_s(u, e)| < \epsilon$, and thus $\|f(s) - f(t)\|_\infty < \epsilon$. Since there exists a δ for every $\epsilon > 0$, so f is continuous.

Now suppose f is continuous. Let $\epsilon > 0$. Then there exists $\delta > 0$ such that for any solution t , if $\|s - t\|_\infty < \delta$, then $\|f(s) - f(t)\|_\infty < \epsilon$, which implies that coordinatewise for every edge e , we have $|y_t(u, e) - y_s(u, e)| < \epsilon$. Since this is true for all $\epsilon > 0$, Condition \square holds. \square

Lemma 4. *Condition \square holds for all increasing, continuous utility functions when renegotiations follow the PBS condition.*

Proof. Let s be any state of the bargaining model and let $e = (u, v)$ be any edge. Here, Lemma \square is applicable. Let $h(s, x) = a_s(x) - b_s(x)$. Also, let $g_s(x) = h(s, x)$ be a function defined on a particular state s . Note that g_s is an increasing, continuous function on the domain $[0, c(e)]$, $g_s(0) < 0$ and $g_s(c(e)) > 0$. The renegotiated value $y_s(u, e)$ is the unique zero of $g_s(x)$ between 0 and $c(e)$.

Let $y = y_s(u, e)$ be the zero of g_s . Let $\eta = \max\{|g_s(y - \epsilon)|, |g_s(y + \epsilon)|\}$. Then, since g_s is increasing, $\eta > 0$, and for all $x \in [0, c(e)] \setminus (y - \epsilon, y + \epsilon)$, $|g_s(x)| \geq \eta$.

Now, observe that $h(s, x)$ is dependent on $\alpha_s(u, e)$, $\alpha_s(v, e)$ and x only, and is continuous in all three of them when the utility functions are continuous. Thus, there exists $\delta > 0$ such that for any state t where $|\alpha_s(u, e) - \alpha_t(u, e)| < \delta$ and $|\alpha_s(v, e) - \alpha_t(v, e)| < \delta$, we have $|h(s, x) - h(t, x)| < \eta \forall x \in [0, c(e)]$, that is $|g_t(x) - g_s(x)| < \eta$. This implies that $g_t(x) \neq 0$ for all $x \in [0, c(e)] \setminus (y - \epsilon, y + \epsilon)$, and so the zero of g_t , which is $y_t(u, e)$, lies in the range $(y - \epsilon, y + \epsilon)$. \square

Lemma 5. *Condition \square holds for all increasing, concave and twice differentiable utility functions when renegotiations follow the NBS condition.*

Proof. Let s be any state of the bargaining model and let $e = (u, v)$ be any edge. Here, Lemma \square is applicable. Let $h(s, x) = q_s(x) - r_s(x)$. Also, let $g_s(x) = h(s, x)$. The rest of the proof identically follows that of Lemma \square . \square

Combining Lemmas \square , \square and \square , we get Theorem \square .

5 Effect of Network Structure on NBS Equilibrium

In this section, we shall study the effect of network topology on NBS equilibrium. In the rest of this section, we shall assume here that all vertices have the *same utility function* $\mathcal{U}(x)$, and that the *deal on every edge has unit profit*, so that the network topology is solely responsible for any variation in the distribution of profits in the NBS equilibrium. We also assume some natural properties of the utility function, and the following is our main result under these assumptions.

Theorem 4. *Let $\mathcal{U}(x)$ be the utility function of every vertex, and let all edges have unit weight. Let $\mathcal{U}(x)$ be increasing, twice differentiable and concave. Also, suppose $\frac{\mathcal{U}(x)-\mathcal{U}(0)}{\mathcal{U}'(x)} < Kx \ \forall x \in [0, 1]$ for some constant K , and $|\mathcal{U}''(x)| \leq \epsilon(x)\mathcal{U}'(x)$ for some decreasing function $\epsilon(x)$. Let s be any NBS equilibrium in this network. Let $e = (u, v)$ be an edge such that u and v have degree more than $(K + 1)d + 1$ for some positive integer d . Then, $|x_s(u, e) - \frac{1}{2}| < \epsilon(d)$.*

Note that the assumptions on the utility function guarantee the existence of NBS equilibrium. Also note that the function $\mathcal{U}(x) = x^p$ for some $0 < p < 1$ satisfies the conditions of Theorem 4 with $K = p^{-1}$ and $\epsilon(x) = (1 - p)/x$. The function $\mathcal{U}(x) = \log(1 + x)$ satisfies the conditions of Theorem 4 as well, with $K = 2$, since $(1 + x)\log(1 + x) < (1 + x)x \leq 2x$ when $x \in [0, 1]$, and $\epsilon(x) = \frac{1}{1+x}$. To prove the above theorem, we will need the next two lemmas. Their proofs are technical, and are omitted due to lack of space.

Lemma 6. *Let $\mathcal{U}(x)$ be increasing, twice differentiable and concave. Also, suppose $\frac{\mathcal{U}(x)-\mathcal{U}(0)}{\mathcal{U}'(x)} < Kx \ \forall x \in [0, 1]$ Then at an NBS equilibrium s , for every edge $e = (u, v)$, $x_s(u, e) \geq \frac{1}{K+1}$ and $x_s(v, e) \geq \frac{1}{K+1}$.*

Lemma 7. *Let $\mathcal{U}(x)$ be increasing, twice differentiable and concave. Let s be an NBS equilibrium, $e = (u, v)$ be any edge, and $\epsilon > 0$. Also, let $|\mathcal{U}''(\alpha_s(u, e))| \leq \epsilon\mathcal{U}'(\alpha_s(u, e))$ and $|\mathcal{U}''(\alpha_s(v, e))| \leq \epsilon\mathcal{U}'(\alpha_s(v, e))$. Then, if u gets x on this agreement at equilibrium (and v gets $1 - x$), then $|x - \frac{1}{2}| < \epsilon$.*

Proof (of Theorem 4). There are $(K + 1)d$ edges incident on each vertex u and v excluding (u, v) , so Lemma 6 implies that at an NBS equilibrium, $\alpha_s(u, e) > \frac{1}{K+1}(K + 1)d = d$ and $\alpha_s(v, e) > \frac{1}{K+1}(K + 1)d = d$. Since $|\mathcal{U}''(x)| \leq \epsilon(x)\mathcal{U}'(x)$ and $\epsilon(x)$ is decreasing, we put $\epsilon = \epsilon(d) < \min\{\epsilon(\alpha_s(u, e)), \epsilon(\alpha_s(v, e))\}$ in Lemma 7 to obtain our result. □

6 Computing Approximate PBS Equilibria on Trees of Bounded Degree

In this section, as a first step towards settling the computational complexity of finding an equilibrium in our model, we note that approximate PBS equilibria can be computed efficiently when the networks are trees of bounded degree and utility function is same for all vertices and is very specific, as follows.

Theorem 5. *Suppose that the bargaining network is a tree with n vertices and maximum degree k , and weights on all edges bounded by C , and where all vertices have the same utility function $U(x) = \log(1+x)$. There is an algorithm that computes an ϵ -approximate PBS equilibrium of this network in time $n(C\epsilon^{-1}k)^{O(k)}$.*

Since this algorithm is not central to this paper, and due to lack of space, we shall only provide its intuition and omit the details. Our algorithm is essentially a modification of the *TreeNash* algorithm of Kearns et. al. [8]. It is a dynamic programming technique on a rooted tree, where computation for the root u of a subtree can be easily completed if the same computation has been already completed for the children of u . The algorithm discretizes the division of profits on each edge to the multiples of some fraction $\delta = \epsilon/k$, and then computes a table for each subtree, under root u . A typical entry of the table stores whether there exists an approximate equilibrium in the subtree, given the total profit of u and its profit from the deal with its parent, and also the deals of u in at least one such equilibrium, if it exists.

Lemma 8 below is crucial for the correctness of our algorithm. It implies that the approximation factor achieved by the algorithm is proportional to the discretization factor δ . The lemma follows quite easily from Lemma 9. Lemma 9 depends heavily on the fact that the utility function is $\log(1+x)$. However, similar results hold for many other utility functions, and our algorithm can be modified to apply to any such utility function.

Lemma 8. *Let s be an exact equilibrium on any graph of maximum degree at most k , and let $U(x) = \log(1+x)$. Let t be any state with $l_\infty(s, t) = \max_{i=1}^m |s_i - s'_i| < \delta$. Then, t is a $k\delta$ -approximate equilibrium.*

Lemma 9. *Let $U(x) = \log(1+x)$. Let s and t be any two states, and $e = (u, v)$ be an edge, such that $|\alpha_t(u, e) - \alpha_s(u, e)| < \epsilon_1$ and $|\alpha_t(v, e) - \alpha_s(v, e)| < \epsilon_2$. If we use PBS for renegotiations, then $|y_t(u, e) - y_s(u, e)| < (\epsilon_1 + \epsilon_2)/2$.*

References

1. Nash, J.: The bargaining problem. *Econometrica* 18, 155–162 (1950)
2. Binmore, K.: *Game Theory and the Social Contract. Just Playing*, vol. 2. MIT Press, Cambridge (1998)
3. Kakade, S.M., Kearns, M.J., Ortiz, L.E., Pemantle, R., Suri, S.: Economic properties of social networks. In: *NIPS* (2004)
4. Judd, S., Kearns, M.: Behavioral experiments in networked trade. In: *ACM Conference on Electronic Commerce* (2008)
5. Braun, N., Gautschi, T.: A nash bargaining model for simple exchange networks. *Social Networks* 28(1), 1–23 (2006)
6. Kleinberg, J., Tardos, E.: Balanced outcomes in social exchange networks. In: *STOC* (2008)
7. Cook, K.S., Yamagishi, T.: Power in exchange networks: A power-dependence formulation. *Social Networks* 14, 245–265 (1992)
8. Kearns, M.J., Littman, M.L., Singh, S.P.: Graphical models for game theory. In: *UAI*, pp. 253–260 (2001)

Sharing Online Advertising Revenue with Consumers

Yiling Chen^{2,*}, Arpita Ghosh¹, Randolph Preston McAfee¹,
and David Pennock¹

¹ Yahoo! Research
{arpita,mcafee,pennockd}@yahoo-inc.com

² Harvard University
yiling@eecs.harvard.edu

Abstract. Online service providers generate much of their revenue by monetizing user attention through online advertising. In this paper, we investigate *revenue sharing*, where the user is rewarded with a portion of the surplus generated from the advertising transaction, in a cost-per-conversion advertising system. While revenue sharing can potentially lead to an increased user base, and correspondingly larger revenues in the long-term, we are interested in the effect of cashback in the short-term, in particular for a single auction. We capture the effect of cashback on the auction's outcome via *price-dependent conversion probabilities*, derived from a model of rational user behavior: this trades off the direct loss in per-conversion revenue against an increase in conversion rate. We analyze equilibrium behavior under two natural schemes for specifying cashback: as a fraction of the search engine's revenue per conversion, and as a fraction of the posted item price. This leads to some interesting conclusions: first, while there is an equivalence between the search engine and the advertiser providing the cashback specified as a fraction of search engine profit, this equivalence no longer holds when cashback is specified as a fraction of item price. Second, cashback can indeed lead to short-term increase in search engine revenue; however this depends strongly on the scheme used for implementing cashback *as a function* of the input. Specifically, given a particular set of input values (user parameters and advertiser posted prices), one scheme can lead to an increase in revenue for the search engine, while the others may not. Thus, an accurate model of the marketplace and the target user population is essential for implementing cashback.

1 Introduction

Advertising is the act of paying for consumers' attention: advertisers pay a publisher or service provider to display their ad to a consumer, who has already been engaged for another purpose, for example to read news, communicate, play games, or search. Consumers pay attention and receive a service, but are typically not directly involved in the advertising transaction. Revenue sharing, where

* Part of this work was done while Y. Chen was at Yahoo! Research.

the consumer receives some portion of the surplus generated from the advertising transaction, is a method of involving the user that could potentially lead to an increased user base for the service provider, albeit at the cost of a possible decrease in short-term revenue.

In May 2008, Microsoft introduced cashback in LiveSearch, where users who buy items using the LiveSearch engine receive cashback on their purchases. As in Livesearch, revenue sharing is best implemented in a *pay-per-conversion* system, where advertisers need to make a payment only when users actually purchase items—since money must change hands to trigger an advertising payment and a revenue share, the system is less susceptible to gaming by users as compared to systems based on cost-per-impression (CPM) or cost-per-click (CPC). Since revenue sharing in this setting corresponds directly to price discounts on items purchased, this gives users a direct incentive to engage with the advertisements on the page. Thus, there is in fact a potential for short-term revenue benefits to the search engine in the form of increased conversion probabilities, in addition to the possibility of an increased user base in the long term.

In this paper, we present a model to study the effect of revenue sharing on search engine revenues, and advertiser and user welfare in a single auction (specifically, we do not model long-term effects). We model the impact of cashback on the user via a price dependent conversion probability, and investigate equilibrium behavior in an auction framework. There are multiple natural schemes for revenue sharing: should cashback be specified as a fraction of the item price, or as a fraction of the search engine's profit from each transaction? Since advertisers might also potentially benefit from cashback in the form of increased sales, should the burden of providing cashback be the advertiser's or the search engine's? Since advertising slots are sold by auction, the choice of scheme (which includes the ranking and pricing functions for the auction) influences the strategic behavior of advertisers, and therefore the final outcome in terms of the winning advertiser, his payment and the final price to the user. As we will see, these different methods of revenue sharing essentially reduce to creating a means for sellers to price discriminate between online and offline consumer segments (or different online consumer segments): the difference in outcomes arises due to differences in the nature and extent of price discrimination allowed by these revenue-sharing schemes.

The analysis, while technically straightforward, leads to some interesting results, even for the simplest case of an auction for a single slot. First, search engines may earn higher advertising revenue when sharing part of that revenue with consumers rather than keeping all revenue to themselves, even ignoring the effect of the policy on overall user growth. (This is because providing cash back to consumers can increase their likelihood of purchasing items, thereby increasing the probability of an advertising payment.)

However, whether, and how much, revenue increases depends strongly on the scheme used for implementing cashback *as a function* of the input: that is, given a particular set of input values (user parameters and advertiser prices), one scheme can lead to an increase in revenue for the search engine, while the others may not.

Further, while one might expect an equivalence between cashback being provided by the search engine and the advertiser (since advertisers can choose their bids strategically in the auction), this is true when the cashback is a fraction of the search engine's profit but *not* when it is a fraction of the item price. Finally, the effect on advertiser or user welfare is also not obvious: depending on the particular scheme being used, it is possible to construct examples where the final, effective, price faced by the user might actually increase with cashback, owing to increased competition amongst advertisers. Thus the problem of cashback is not a straightforward one, and none of these schemes always dominates the others: understanding the marketplace and target user population is essential for effective implementation of revenue sharing.

Related work. The most relevant prior research is that of Jain ([5]), making the case that search engines should share the surplus generated by online advertising with users. In contrast, we take a completely neutral approach to revenue-sharing, and provide a model for analyzing its effects on search engine revenue, and user and advertiser welfare.

In some advertising systems, a portion of advertisers' payments go to consumers in the form of coupons, cash back incentives, or membership rewards, either directly from the advertiser or indirectly through an affiliate marketer or other third party lead generator. Several large online affiliate marketing aggregators, for example ebates.com, mypoints.com, and jellyfish.com, function this way, collecting from advertisers on every sale and allocating a portion of their revenue back to the consumer. The main distinction in our work is that the cashback mechanism is embedded in an auction model: advertisers are competing for a sales channel, and the search engine's revenue is determined by the ranking and pricing function used, as well as by the discount offered. We build on the work on equilibrium in sponsored search auctions [74].

Goel et al. [1] explore revenue sharing in a ranking or reputation system, describing an ingenious method to incentivize users to fix an incorrect ranking. There is a large body of empirical work on the effect of price discounts and sales on purchases of items, and the impact of different methods of specifying the discount; see, for example, [2,6]. Researchers have examined consumers' perceptions of search and shopping intentions, at different levels of discounts across two shopping environments, one online and the other offline, showing that the shopping intention of the consumers differ at varying discount levels in the two environments [3].

2 Model

We model the simplest instance of revenue sharing, where n sellers, each selling an item with posted price p_i , compete for a *single* advertising slot in a cost-per-conversion system (*i.e.*, the winning advertiser makes a payment only when a user buys the item). The search engine, which auctions off this ad slot amongst the sellers, controls the ranking and pricing functions for the auction, and can choose whether and how to include cashback in the mechanism. The key

element in our model capturing the effect of revenue sharing is a *price-dependent conversion probability*, $g_i(p)$, which is a decreasing function of p : this introduces a trade-off since decreasing the final price to the user increases the *probability* of a conversion, which may lead to higher expected revenue. This conversion probability function is derived from the following user model: a user is a rational buyer, whose value for item i , v_i , is drawn i.i.d. from a distribution with CDF $F_i(v_i)$. The user buys the item only if the price $p_i \leq v_i$, which has probability $1 - F_i(p)$. Since the user's probability of purchasing item i need not solely be determined by price (it might depend, for instance, on the reputation of seller i , or the relevance of product i to the user), we introduce a price-independent multiplier x_i ($0 < x_i \leq 1$). Thus, the final probability of purchase given price p is $g_i(p) = x_i(1 - F(p))$, which is a decreasing function of p ¹.

Associated with seller i , in addition to the posted price p_i and the conversion probability (function) $g_i(p)$, is a production cost c_i , so that a seller's net profit when he sells an item at a price p is $p - c_i$. We assume that posted prices p_i 's and the functions g_i are common knowledge to both the search engine and advertisers (this assumption is discussed later); the costs c_i are private to the advertisers. We investigate the trade-off between cashback and expected revenue to the search engine in a single auction; we clarify again that we do not model and study long-term effects of cashback on search engine revenues in this paper.

3 Schemes for Revenue Sharing

We describe and analyze four variants of natural revenue-sharing schemes that the search engine could use when selling a single advertising slot through an auction. For each scheme, we analyze the equilibrium behavior of advertisers, and where possible, state the conditions under which cashback leads to an increase in revenue for the search engine. (Our focus is on search engine revenue since decrease in revenue is the primary argument for a search engine against implementing cashback.) Finally, we compare the schemes against each other. Due to space constraints, all proofs and examples can be found in the full version of this paper.

3.1 Cashback as a Fraction of Posted Price

Specifying cashback as a fraction of the posted price of an item is most meaningful to the user, since he can now compute the exact final price of an item. We consider three natural variants, and specify their equilibria, in order to perform a revenue comparison. Note that the ranking, and therefore the winning advertiser and welfares, are a function of the variants and can also depend on the cashback fraction.

¹ For example, if the density f_i is uniform on $[0, W_i]$, $g_i(p) = x_i(1 - \frac{1}{W_i}p)$ is a linear function; if f_i is exponential with parameter λ_i , the resulting g function is exponential as well.

1. *Cashback as a fixed fraction of posted price paid by advertiser.*

We first consider the scheme where the auction mechanism also dictates the winning advertiser to pay a fixed fraction α of its posted price as cashback to users for every conversion. The fraction α is determined by the search engine ahead of time and is known to all advertisers. In such an auction, advertisers submit a bid b_i which is the maximum amount they are willing to pay the search engine per conversion. The search engine ranks advertisers by expected value per conversion *including* the effect of cashback on conversion probability, *i.e.*, by $g_i(p_i - \alpha p_i)b_i$ (note b_i is the bid and p_i is the posted price). For every conversion, the winning advertiser must pay the search engine the minimum amount he would need to bid to still win the auction; he also pays the cashback to the consumer.

In such an auction, an advertiser’s dominant strategy is to bid so that his maximum payment to the search engine plus the revenue share to the user equals his profit, in order to maximize his chance of winning the slot. The following describes the equilibrium of the auction.

Proposition 1. (*Equilibrium behavior*) *At the dominant strategy equilibrium, advertisers bid $b_i^I = \max(0, (1 - \alpha)p_i - c_i)$ and are ranked by the mechanism according to $z_i^I = \max(0, g_i(p_i(1 - \alpha))((1 - \alpha)p_i - c_i))$. Let σ_1 be the ranking of advertisers. The winning advertiser, $\sigma_1\{1\}$ pays*

$$p_c^I = \frac{z_{\sigma_1\{2\}}^I}{g_{\sigma_1\{1\}}(p_{\sigma_1\{1\}}(1 - \alpha))} \tag{1}$$

for every conversion. The search engine’s expected revenue equals the second highest expected value after cashback,

$$R^I = g_{\sigma_1\{1\}}(p_{\sigma_1\{1\}}(1 - \alpha))p_c^I = z_{\sigma_1\{2\}}^I. \tag{2}$$

Note that the ranking σ_1 is a function of α . For different values of α , different advertisers may win the auction and the advertisers’ bids also change.

Even though it is the advertiser who pays the cash-back, it is not always beneficial for the search engine to choose a non-zero fractional cashback, *i.e.*, $\alpha > 0$. We present some sufficient conditions for cashback to be (or not to be) revenue-improving in this case.

Theorem 1. *Suppose g_i is such that $(p - c_i)g_i(p)$ is continuous and differentiable with respect to p , and has a unique maximum at some price p_i^* . Let σ_0 be the ranking of advertisers when there is no cash-back. If $(p_{\sigma_0\{1\}} - c_{\sigma_0\{1\}})g_{\sigma_0\{1\}}(p_{\sigma_0\{1\}}) > (p_{\sigma_0\{2\}} - c_{\sigma_0\{2\}})g_{\sigma_0\{2\}}(p_{\sigma_0\{2\}})$ and $p_{\sigma_0\{2\}} > p_2^*$, there exists $\alpha > 0$ that increases the search engine’s revenue. Conversely, if all advertisers’ posted prices satisfy $p_i \leq p_i^*$, revenue is maximized by setting $\alpha = 0$.*

Theorem □ implies that cash-back may be beneficial to the search engine when the original product prices are “too high”, *i.e.* higher than the optimal prices. The natural question to ask is why any advertiser would want to set

a price higher than his optimal price. This relates to the assumption that each advertiser keeps a universal price across all markets (buyer segments or sales channels). Buyers in each market can have a different price sensitivity function g_i . Thus, the universal price can be the optimal price in other markets but higher than the optimal price in the market that the advertiser attempts to reach through the search engine. (It is possible, for instance, that shoppers typically look for deals online, or would want to pay lower prices online than in stores due to uncertainty in product quality or condition.) Example 1 in Appendix B in the full version of the paper illustrates the increase of expected revenue for search engine by choosing a positive α .

2. *Search engine pays cashback as a fixed fraction of posted price.*

Next we consider the scheme where the search engine pays a fixed fraction β of the winning advertiser’s posted price as cashback for every conversion. β is determined by the search engine and is known to all advertisers. Naturally, the search engine will only choose values of β so that p_c , the payment per conversion received by the search engine, is greater than or equal to βp_i . Advertisers submit bids b_i . The search engine ranks advertisers by their final (post-cash-back) conversion rate multiplied by their bid, i.e., $g_i(p_i - \beta p_i)b_i$. An advertiser’s dominant strategy is to bid so as to maximize his chances of winning the slot without incurring loss:

Proposition 2. *(Equilibrium behavior) Advertisers bid $b_i^{\text{II}} = p_i - c_i$ and are ranked by $z_i^{\text{II}} = g_i(p_i(1 - \beta))(p_i - c_i)$ at the dominant strategy equilibrium. Let σ_{II} be the ranking of advertisers. The winning advertiser, $\sigma_{\text{II}}\{1\}$, pays*

$$p_c^{\text{II}} = \frac{z_{\sigma_{\text{II}}\{2\}}^{\text{II}}}{g_{\sigma_{\text{II}}\{1\}}(p_{\sigma_{\text{II}}\{1\}}(1 - \beta))}. \tag{3}$$

for every conversion. The search engine’s expected revenue is

$$\begin{aligned} R^{\text{II}} &= g_{\sigma_{\text{II}}\{1\}}(p_{\sigma_{\text{II}}\{1\}}(1 - \beta))(p_c^{\text{II}} - \beta p_{\sigma_{\text{II}}\{1\}}) \\ &= z_{\sigma_{\text{II}}\{2\}}^{\text{II}} - \beta p_{\sigma_{\text{II}}\{1\}} g_{\sigma_{\text{II}}\{1\}}(p_{\sigma_{\text{II}}\{1\}}(1 - \beta)). \end{aligned} \tag{4}$$

In this case also, the search engine may increase its expected revenue when using this scheme. Suppose $g_i(p_i) = 1 - 0.1p_i$. Three advertisers A, B, and C participate in the auction. They have prices $p_A = 6$, $p_B = 9$, and $p_C = 10$ respectively; $c = 0$ for all advertisers. Then, by setting $\beta = 0.4737$ the search engine increases its expected revenue from 0.9 to 2.4931 and the final price faced by the user drops from 6 to 4.74.

3. *Advertiser chooses amount of cashback and pays it.*

More expressiveness is provided to the advertisers if they are allowed to bid both on the fractional discount they offer, as well as their per-conversion payment to the search engine. Both of these are then used in the ranking function. The search engine runs an auction that does not specify the fraction of revenue share required. Instead, the auction rule requires the advertiser to submit both a bid b_i and a fraction γ_i ($0 \leq \gamma_i \leq 1$). Advertisers are ranked by

conversion rate (including cashback) multiplied by bid, i.e. $g_i(p_i(1 - \gamma_i))b_i$. The payment of the winning advertiser is as follows: his net payment is $\gamma_i p_i + p_c$, where p_c is the minimum amount he needs to bid, keeping γ_i fixed, to win the auction. The dominant strategy for all advertisers is to choose γ_i to maximize their values, and for the choice of γ_i , to bid their true value after the effect of cashback.

Proposition 3. *At the dominant strategy equilibrium, advertisers select $\gamma_i^* = \arg \max_{0 \leq \gamma_i \leq 1} x_i g_i((1 - \gamma_i)p_i)((1 - \gamma_i)p_i - c_i)$, bid $b_i^{\text{III}} = (1 - \gamma_i^*)p_i - c_i$, and are ranked by $z_i^{\text{III}} = g_i(p_i(1 - \gamma_i^*))((1 - \gamma_i^*)p_i - c_i)$. Let σ_{III} be the ranking of advertisers. The winning advertiser, $\sigma_{\text{III}}\{1\}$, pays the search engine*

$$p_c^{\text{III}} = \frac{z_{\sigma_{\text{III}}\{2\}}^{\text{III}}}{g_{\sigma_{\text{III}}\{1\}}(p_{\sigma_{\text{III}}\{1\}}(1 - \gamma_{\sigma_{\text{III}}\{1\}}^*))}} \tag{5}$$

and pays the user $\gamma_{\sigma_{\text{III}}\{1\}}^* p_{\sigma_{\text{III}}\{1\}}$ per conversion. The search engine’s expected revenue is

$$R^{\text{III}} = g_{\sigma_{\text{III}}\{1\}}(p_{\sigma_{\text{III}}\{1\}}(1 - \gamma_{\sigma_{\text{III}}\{1\}}^*)) p_c^{\text{III}} = z_{\sigma_{\text{III}}\{2\}}^{\text{III}}. \tag{6}$$

Note that allowing the advertiser to choose γ_i as well as b_i essentially allows them to choose an effective new “price”. Consequently, if possible the advertiser selects γ_i so that the new price equals his optimal price. For $p_i > p_i^*$, this γ_i^* is such that $(1 - \gamma_i^*)p_i = p_i^*$, where p_i^* is the price that maximizes the function $(p - c_i)g_i(p)$. The following theorem shows that in this scheme, the search engine’s expected revenue is always weakly larger than without cashback.

Theorem 2. *Let R^0 denote search engine’s expected revenue without cashback. For the same set of advertisers, $R^{\text{III}} \geq R^0$.*

Example 2, Appendix B in the full version of this paper illustrates the increase of search engine’s expected revenue with this scheme.

3.2 Cashback as a Fraction of Search Engine Revenue

Another natural way to specify a revenue share is to describe it as a fraction α of the search engine’s revenue, i.e., the payment per conversion; this corresponds to the search engine sharing its surplus with the user, who is an essential component of the revenue generation process. Unless the search engine charges a fixed price per conversion, it is hard to include post-cashback conversion rates to determine the ranking, since the amount of cashback depends on the ranking. Thus, we use the conversion rate before cashback to rank advertisers. In this scheme, advertisers are ranked according to $g_i(p_i)b_i$, where b_i is the per-conversion bid submitted by advertiser i , and search engine pays a fixed fraction δ of its revenue per conversion as cashback. Again, it is a dominant strategy for advertisers to bid their true value:

Proposition 4. Advertisers bid $b_i^{IV} = p_i - c_i$ and are ranked by $z_i^{IV} = g_i(p_i)(p_i - c_i)$ at the dominant strategy equilibrium. Let σ_{IV} be the ranking of advertisers. The winning advertiser, $\sigma_{IV}\{1\}$, pays

$$p_c^{IV} = \frac{g_{\sigma_{IV}\{2\}}(p_{\sigma_{IV}\{2\}})(p_{\sigma_{IV}\{2\}} - c_{\sigma_{IV}\{2\}})}{g_{\sigma_{IV}\{1\}}(p_{\sigma_{IV}\{1\}})} \tag{7}$$

per conversion. The revenue of the search engine with cashback is

$$R^{IV} = g_{\sigma_{IV}\{1\}}(p_{\sigma_{IV}\{1\}} - \delta p_c^{IV}) \frac{z_{\sigma_{IV}\{2\}}}{g_{\sigma_{IV}\{1\}}(p_{\sigma_{IV}\{1\}})} p_c^{IV} (1 - \delta). \tag{8}$$

Note that this ranking is independent of the value of δ , the cashback fraction: σ_{IV} is the same as σ_0 , the ranking without cashback.

It is also possible to request the advertiser to pay the cashback that is specified as a fixed fraction of the search engine’s revenue. We show that it is equivalent to the case that the search engine pays the cashback.

Theorem 3. The scheme where search engine pays δ fraction of its revenue per conversion as cashback is equivalent to the scheme where the advertiser pays $\delta/(1 - \delta)$ fraction of the search engine’s revenue per conversion as cashback, regarding to the utilities of the user, the advertisers, and the search engine.

Note that when revenue share is specified as a fraction of search engine revenue, the search engine may choose the optimal fraction δ after advertisers submit their bids. This will not change the equilibrium bidding behavior of advertisers, in contrast to the case where advertisers pay the cashback. Since the optimal cashback δ might be 0, choosing δ after collecting bids ensures that the search engine’s revenue never decreases because of cashback.

Whether or not the search engine can increase its revenue by giving cash-back depends on the posted prices of the top two advertisers and their g functions.

Theorem 4. If there exists $\delta > 0$ such that $g_{\sigma_{IV}\{1\}}(p_{\sigma_{IV}\{1\}} - \alpha p_c)(1 - \delta) \geq g_{\sigma_{IV}\{1\}}(p_{\sigma_{IV}\{1\}})$, revenue sharing with parameter δ increases the expected revenue of the search engine. For linear $g_i = x_i(1 - kp_i)$ and $c_i = 0$, $\delta > 0$ when $p_{\sigma_{IV}\{1\}} + p_c^{IV} > 1/k$.

3.3 Comparison between Schemes

The first three schemes described above all specify revenue share as a fraction of posted price, while the fourth scheme specifies revenue share as a fraction of the search engine revenue. The following results characterize the choice of mechanism to maximize the search engine’s revenue, when revenue share is expressed as a fraction of posted price.

Theorem 5. Given a set of advertisers, $R^{III} \geq R^I$ for all α .

Theorem 6. Given a set of advertisers, $R^I \geq R^{II}$ if $\alpha = \beta$ and the ranking according to $p_i * g(p_i(1 - \beta))$ is the same as the ranking according to $(p_i - c_i) * g(p_i(1 - \beta))$.

This gives us a result on maximizing revenue when cashback is specified as a fraction of the posted prices for the special cases below.

Corollary 1. *When $c_i = 0$, or $c_i = \mu p_i$ for all i , $R^{\text{III}} \geq R^{\text{I}} \geq R^{\text{II}}$. Thus revenue is maximized when the search engine allows advertisers to choose and pay the fraction γ_i of their posted prices.*

When revenue share is expressed as a fraction of the posted price, allowing advertisers to choose the fraction of revenue share (the third scheme) can lead to the highest revenue for the search engine in many cases. Thus, we compare it with the case when revenue share is specified as a fraction of the advertising revenue (the fourth scheme). We have the following result.

Proposition 5. *Neither the revenue-maximizing cashback scheme with cashback as a fraction of posted price, nor the revenue-maximizing scheme with cashback as a fraction of search engine revenue, always dominates the other in terms of generating higher expected revenue for the search engine.*

Thus, depending on the set of posted prices, the expected revenue of the search engine in either the third or the fourth scheme can be higher. Both schemes, however, are always weakly revenue improving; in the third scheme where advertisers specify the cashback amount, the search engine needs to make no choice and, according to Theorem 2, the search engine's revenue is at least as large as that without cashback. In the fourth scheme also, the search engine can choose the optimal fraction after the bids have been submitted, ensuring that cashback never leads to loss in revenue.

We note that whether cashback can increase search engine revenue or not also depends on the revenue sharing schemes. Given a set of advertiser prices, it is possible that one scheme can increase the revenue of search engine by providing positive cashback, while the other scheme is better off not giving cashback at all. Examples 3 and 4 in Appendix B in the full version of the paper support this with two specific instances.

4 Conclusion

We model revenue sharing with users in the context of online advertising auctions in a cost-per-conversion system, in which the winning advertiser pays the search engine only in the event of a conversion. The conversion probability of a user is modeled as a decreasing function of the final product price that the user faces. Thus, sharing revenue with the user may increase the conversion probability sufficiently to lead to a short-term increase in the search engine's expected revenue, despite the fact that the per-conversion revenue decreases.

We study four schemes for a search engine to specify the revenue share in the auction setting. When the revenue share is expressed as a fraction of the winning advertiser's posted price, we have (1) advertiser pays cashback as a fixed fraction of posted price; (2) search engine pays cashback as a fixed fraction of posted price; and (3) advertiser determines and pays cashback. If the revenue

share is specified as a fraction of the advertiser's revenue per conversion, we consider (4) the search engine pays cashback as a fixed fraction of its revenue. We analyze the equilibrium of the auction for the four schemes and show that for all four schemes there are situations in which search engine can increase its short-term expected revenue by allowing revenue sharing. Scheme (3) dominates scheme (1) and (2) in many situations in terms of maximizing search engine revenue. However, neither scheme (3) nor scheme (4) are universally better for generating higher search engine revenue. We note that although revenue sharing often leads to lower final prices to users, this need not always be the case: there exist advertiser prices under which the revenue maximizing cashback fraction leads to increased final price to the user, as shown in Example 1, Appendix B in the full version of this paper.

The properties of these revenue sharing mechanisms rely strongly on the assumption that advertisers keep a universal price across all sales channels, which is often the case in reality. If advertisers can or are willing to charge channel-specific-prices, they will select an optimal price to participate in the advertising auction. In return, the search engine no longer needs to, or will not find it profitable to share revenue with the user. In fact, revenue sharing with users is an indirect way, controlled by the search engine, to achieve price discriminations across different sales channels.

References

1. Bhattacharjee, R., Goel, A.: Algorithms and incentives for robust ranking. In: ACM-SIAM Symposium on Discrete Algorithms (SODA) (2007)
2. Darkea, P., Chung, C.: Effects of pricing and promotion on consumer perceptions. *Journal of retailing* 81(1) (2005)
3. Das, N., Burman, B., Biswas, A.: Effect of discounts on search and shopping intentions: the moderating role of shopping environment. *International Journal of Electronic Marketing and Retailing (IJEMR)* 1(2) (2006)
4. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalised second price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97(1), 242–259 (2007)
5. Jain, K.: The good, the bad, and the ugly of the search business (manuscript, 2007)
6. Kim, H.M., Kramer, T.: The effect of novel discount presentation on consumer's deal perceptions. *Marketing Letters* 17(4) (2006)
7. Varian, H.R.: Position auctions. *International Journal of Industrial Organization* 25(6), 1163–1178 (2006)

Budget Constrained Bidding in Keyword Auctions and Online Knapsack Problems*

Yunhong Zhou¹, Deeparnab Chakrabarty², and Rajan Lukose³

¹ Rocket Fuel Inc, Redwood Shores, CA
yzhou@rocketfuelinc.com

² Georgia Tech, Atlanta, GA
deepc@cc.gatech.edu

³ HP Labs, Palo Alto, CA
rajan.lukose@hp.com

Abstract. We consider the budget-constrained bidding optimization problem for sponsored search auctions, and model it as an *online (multiple-choice) knapsack problem*. We design both deterministic and randomized algorithms for the online (multiple-choice) knapsack problems achieving a *provably optimal* competitive ratio. This translates back to fully automatic bidding strategies maximizing either profit or revenue for the budget-constrained advertiser. Our bidding strategy for revenue maximization is oblivious (i.e., without knowledge) of other bidders' prices and/or clickthrough-rates for those positions. We evaluate our bidding algorithms using both synthetic data and real bidding data gathered manually, and also discuss a sniping heuristic that strictly improves bidding performance. With sniping and parameter tuning enabled, our bidding algorithms can achieve a performance ratio above 90% against the optimum by the omniscient bidder.

1 Introduction

Sponsored search auction is an effective way of monetizing search query activities for search engine providers, while shifting the burden to advertisers/bidders to figure out how to automate and optimize the keyword bidding process. In this work we focus on the bid optimization problem under the budget constraint. Formally, given an advertiser with a fixed budget over a fixed time horizon, and a set of keywords that he is interested to bid on, we try to design bidding strategies to address the following problem: For each keyword and each time period, how much should the advertiser bid to obtain which position, so as to maximize return on investment (ROI) of these auctions?

Keyword Bidding Models. For simplicity, assume that the default advertiser has a budget B over a fixed time horizon, discretized into time periods $1, \dots, T$. He is interested in a single keyword with expected value-per-click V . The model can be easily extend to the multiple-keyword case.

* Work was done while the author were at HP Labs.

There are bidders $\{1, \dots, N\}$ at time t for this keyword and their bids are sorted in decreasing order $b_1(t) > \dots > b_N(t)$. There are S ad slots, and are assigned to the top- S bids as follows: bidder s gets slot s ; for each user click on his ad, bidder s is charged a price b_{s+1} , if $s < S$ or a minimum fee b_{\min} (e.g. 10¢). Each slot s has a clickthrough-rate (CTR), denoted $\alpha(s)$, which is defined as the total number of clicks on an ad divided by the total number of impressions (displays). Assuming other bidders have their bids fixed, the default advertiser can obtain slot s by bidding slightly over $b_s(t)$; for each user click, he incurs a cost of $b_s(t)$, obtains an expected revenue V and profit $V - b_s(t)$.

Online Knapsack Problems. Fix a keyword with positions $1, \dots, S$. At time t , $X(t)$ is the number of clicks at period t , while $b_s(t)$ is the maximum bid corresponding to position s . Winning position s at time t costs the advertiser $w_s(t)$ and earns him a profit of $v_s(t)$ where

$$w_s(t) \equiv b_s(t)X(t)\alpha(s), \quad v_s(t) \equiv (V - b_s(t))X(t)\alpha(s). \tag{1}$$

For revenue maximization, $v_s(t) = VX(t)\alpha(s)$. Let $N(t) = \{(w_s(t), v_s(t)) | s = 1, \dots, S\}$, then winning position s at time t corresponds to selecting item $(w_s(t), v_s(t)) \in N_t$. Since the default bidder has to decide either overbidding $b_s(t)$ or not at time t , thus keyword bidding corresponds to the *online* multiple-choice knapsack problem (Online-MCKP). The multiple-choice knapsack problem is a generalization of the classic knapsack problem, where there are multiple item-sets and you can select *at most one* item from each item-set; the multiple-choice constraint of MCKP corresponds to the sponsored search auction policy where each advertiser can select to win at most one ad slot for each keyword at each time.

Our Assumptions. We use *competitive analysis* to evaluate our bidding strategies, comparing our result with the maximum profit attainable by the *omniscient bidder* who knows the bids of all the other users ahead of time. In general, no online algorithm can achieve any non-trivial *competitive ratio* (the ratio between the output of the given algorithm and the offline optimum) for Online-KP [4]. Fortunately, in our setting, we make two reasonable assumptions on the knapsack items, which allow us to develop interesting online algorithms. These two assumptions are:

$$(i) w_s(t) \ll B; \quad (ii) L \leq \frac{v_s(t)}{w_s(t)} \leq U, \quad \forall t, \forall s. \tag{2}$$

2 Results

In this work we model budget-constrained bidding optimization as variants of online knapsack problems. In Section 3, we design a deterministic algorithm for the online knapsack problem with two assumptions given above. The algorithm has a competitive ratio $\ln(U/L) + 1$, and is robust against any adaptive adversary. We also show a *matching lower bound* in section 3.1. Therefore our algorithm is *provably optimal* in the worst-case sense. We also give a $(\ln(U/L)+2)$ -competitive

online algorithm for the multiple-choice knapsack problem (MCKP), the classic generalization of the knapsack problem which corresponds to the general bidding optimization problem with multiple slots per keyword.

In Section 4, we translate the algorithms for online knapsack problems into bidding strategies for sponsored search auctions, for both profit and revenue maximization. For single-slot auctions, the corresponding strategies are *oblivious*, and thus work even if other bidders' bids were not known. It also implies that the strategy is an *approximate dominant strategy* in the sense that it is an approximate best response to *any* bid profile of other bidders. For the multiple-slot case, we translate the algorithm for Online-MCKP to bidding strategies for both profit and revenue-maximizing bidding strategies. The profit maximizing strategy is not oblivious and requires knowledge of other players' bids and also the CTRs of all slots. The revenue-maximizing strategy remains oblivious.

In Section 5, we report experimental results evaluating our bidding strategies using both synthetic bidding data and real bidding data collected manually. We modify our strategy by adding a *sniping* heuristic, and it performs much better empirically while maintaining the same theoretical bounds. Our limited experimental evaluation also suggests that parameter tuning helps to improve the performance of our bidding algorithms. With both sniping and parameter tuning enabled, our bidding algorithms (for both profit and revenue maximization) achieve an output value which is consistently more than 90% of the optimum by the omniscient bidder.

2.1 Related Work

Due to page limit as well as the vast amount of research literature in sponsored search auctions, knapsack problems, and online algorithms, we will only discuss previous work most relevant to ours.

Keyword Bidding. Sponsored search auctions have attracted a lot of attention, for both auctioneer revenue maximization and advertiser bidding optimization. Among all these work, Mehta et al. [5] studied the auctioneer revenue maximization with budget-constrained bidders, using a trade-off function Ψ (compare it to our threshold function) to grant queries to bidders, and the technique they use is probably most similar to the threshold function we use.

Online Algorithms. Awerbuch et al. [2] studied the online call routing which generalizes the online classical knapsack problem. More recently, Buchbinder et al. [3] designed online algorithms for fractional versions of general packing problems which imply an $O(\ln(U/L))$ -competitive algorithm for the online knapsack problem.

3 Online Knapsack Problems

Consider the online version of the classic 0/1 knapsack problem. The input sequence consists of a knapsack of capacity B and a stream of T items where item t has value $v(t)$ and weight $w(t)$. We call the value-to-weight ratio $v(t)/w(t)$ of

item t its *efficiency*. The goal is to choose these items in an *online* fashion, i.e., making a decision as an item arrives and not revoking them later, so as to maximize the total value of selected items. For the online *multiple-choice* knapsack problem, at each step a set of items N_t arrives and we need to choose at most one item from each set.

We say that an online algorithm \mathcal{A} has competitive ratio γ (or equivalently is γ -competitive) if for *any* input sequence σ , we have $\text{OPT}(\sigma) \leq \gamma \cdot \mathcal{A}(\sigma)$, where $\mathcal{A}(\sigma)$ is the (expected, if \mathcal{A} is randomized) value obtained by \mathcal{A} given σ , and $\text{OPT}(\sigma)$ is the maximum value which can be obtained by any *offline* algorithm with the knowledge of σ .

We now give a deterministic algorithm for the online knapsack problem achieving the optimal bound of $\ln(U/L) + 1$. In the remainder of the paper, e denotes the base of the natural logarithm.

Algorithm. ONLINE-KP-THRESHOLD
 Let $\Psi(z) \equiv (Ue/L)^z(L/e)$.
 At time t , let $z(t)$ be the fraction of capacity filled, pick element t iff

$$\frac{v(t)}{w(t)} \geq \Psi(z(t)).$$

Observe that for $z \in [0, c]$ where $c \equiv 1/(1 + \ln(U/L))$, $\Psi(z) \leq L$, thus the algorithm will pick all items available until c fraction of the knapsack is filled. In fact, we will assume henceforth $\Psi(z) = L$ for $z \in [0, c]$. When $z = 1$, $\Psi(z) = U$, and since Ψ is strictly increasing, the algorithm will never over-fill the knapsack.

Theorem 1. ONLINE-KP-THRESHOLD has a competitive ratio of $\ln(U/L) + 1$.

Proof. Fix an input sequence σ . Let the algorithm terminate filling Z fraction of the knapsack and obtaining a value of $\mathcal{A}(\sigma)$. Let S and S^* respectively be the set of items picked by the Algorithm ONLINE-KP-THRESHOLD and the optimum. Denote the weight and the value of the common items by $W = w(S \cap S^*)$ and $P = v(S \cap S^*)$. For each item t not picked by the algorithm, its efficiency is $< \Psi(z(t)) \leq \Psi(Z)$ since $\Psi(z)$ is a monotone increasing function of z . Thus,

$$\text{OPT}(\sigma) \leq P + \Psi(Z)(B - W)$$

Since $\mathcal{A}(\sigma) = P + v(S \setminus S^*)$, the above inequality implies that

$$\frac{\text{OPT}(\sigma)}{\mathcal{A}(\sigma)} \leq \frac{P + \Psi(Z)(B - W)}{P + v(S \setminus S^*)}. \tag{3}$$

Since each item j picked in S must have efficiency at least $\Psi(z_j)$ where z_j is the fraction of the knapsack filled at that instant, we have

$$P \geq \sum_{j \in S \cap S^*} \Psi(z_j)w_j =: P_1, \tag{4}$$

$$v(S \setminus S^*) \geq \sum_{j \in S \setminus S^*} \Psi(z_j)w_j =: P_2. \tag{5}$$

Since $\text{OPT}(\sigma) \geq \mathcal{A}(\sigma)$, inequality (3) implies

$$\frac{\text{OPT}(\sigma)}{\mathcal{A}(\sigma)} \leq \frac{P + \Psi(Z)(B - W)}{P + v(S \setminus S^*)} \leq \frac{P_1 + \Psi(Z)(B - W)}{P_1 + v(S \setminus S^*)} \tag{6}$$

Since $P_1 \leq \Psi(Z)w(S \cap S^*) = \Psi(Z)W$, plugging in the values of P_1 and P_2 , we get

$$\frac{\text{OPT}(\sigma)}{\mathcal{A}(\sigma)} \leq \frac{\Psi(Z)}{\sum_{j \in S} \Psi(z_j)\Delta z_j} \tag{7}$$

where $\Delta z_j = z_{j+1} - z_j = w_j/B$ for all j .

Based on the assumption that the weights are much smaller than B , we can approximate the summation via an integration (refer to the remark following the proof). Thus,

$$\begin{aligned} \sum_{j \in S} \Psi(z_j)\Delta z_j &\approx \int_0^Z \Psi(z)dz = \int_0^c Ldz + \int_c^Z \Psi(z)dz \\ &= cL + \frac{L}{e} \frac{(Ue/L)^Z - (Ue/L)^c}{\ln(Ue/L)} \\ &= \frac{L}{e} \frac{(Ue/L)^Z}{\ln(Ue/L)} = \frac{\Psi(Z)}{\ln(U/L) + 1}. \end{aligned}$$

Along with inequality (7), this completes the proof.

Remark: We can make the approximation made above precise. Since $\Psi(z)$ is an increasing function of z , we obtain $\sum_{j \in S} \Psi(z_j)\Delta z_j \geq (1 - \epsilon_0) \int_0^Z \Psi(z)dz$ where $\epsilon_0 = (\max_j w_j)/B$ is small constant. Thus, to be precise, the competitive ratio is actually $\ln(Ue/L) \cdot \frac{1}{1 - \epsilon_0}$. For simplicity, we ignore the factor $1 - \epsilon_0$ for subsequent analysis.

Extension to Online-MCKP. One can extend the above algorithm to multiple choice knapsack problems in the following way – at each step t , let $E_t \subseteq N_t$ denote the items with efficiency at least $\Psi(z(t))$. Pick the item in E_t with the highest profit. Call this algorithm ONLINE-MCKP-THRESHOLD. The following theorem can be proved similarly as above and we omit it from the extended abstract.

Theorem 2. ONLINE-MCKP-THRESHOLD has a competitive ratio of $(\ln(U/L) + 2)$.

3.1 A Matching Lower Bound

Theorem 3. The competitive ratio of any (possibly randomized) online algorithm for the online knapsack problem is at least $\ln(U/L) + 1$.

Proof (Sketch). The proof constructs a distribution of input sequences and shows that any deterministic algorithm against this distribution achieves a competitive ratio at most $\ln(U/L) + 1$. Based on Yao’s minimax lemma [6], the lower bound is obtained. We describe the distribution here and defer the analysis to our technical report.

Fix a parameter $\eta > 0$. Let k be an integer such that $(1 + \eta)^k = U/L$, i.e., $k = \frac{\ln(U/L)}{\ln(1 + \eta)}$. The support of the input distribution consists of the instances I_0, I_1, \dots, I_k , where I_0 is a stream of B identical items each with weight 1 and value L . I_1 is I_0 followed by a stream of B identical items each with weight 1 and value $(1 + \eta)L$, and in general I_{j+1} is I_j followed by B items with weight 1 and value $(1 + \eta)^{j+1}L$. The distribution \mathcal{D} is specified by giving probability p_j to instance I_j where

$$p_k := \frac{1 + \eta}{(k + 1)\eta + 1}, \quad p_j := \frac{\eta}{(k + 1)\eta + 1}, \quad \forall 0 \leq j < k.$$

The ratio is obtained as $\eta \rightarrow 0$.

4 Bidding Strategies for Keyword Auctions

In this section, we construct bidding strategies for either profit maximization or revenue maximization. The difference in the two are in the parameter settings. For simplicity and brevity, we start with the single-slot case and extend to the multiple-slot case.

For profit maximization, recall that outbidding $b(t)$ at time t gives an efficiency of $\frac{v(t)}{w(t)} = \frac{V}{b(t)} - 1$ while for revenue maximization its $\frac{V}{b(t)}$. Thus, the parameters U and L for revenue maximization strategies are: $U_r := \frac{V}{b_{min}}$ and $L_r := 1$ respectively. For profit maximization $U_p = U_r - 1$, though L_p could be 0. To take care of this, we introduce another parameter ϵ , such that we bid only when the efficiency is bigger than ϵ . This makes $L_p = \epsilon$ but leads to an *additive loss* in the performance guarantee.

The strategies are derived from the online algorithms: Bidder 0 outbids only if the efficiency is bigger than the threshold. Since the threshold does not depend on anything other than the fraction of knapsack filled, the strategies also depend only on the fraction of budget spent. The strategies are formally stated as follows:

Bidding Strategy: PROFIT-MAXIMIZING SINGLE-SLOT
 Let $\Psi(z) \equiv (U_p e / \epsilon)^z (\epsilon / e)$.
 At time t , if fraction of budget spent is $z(t)$, then bid $b_0(t) = \frac{V}{1 + \Psi(z(t))}$.

Bidding Strategy: REVENUE-MAXIMIZING SINGLE-SLOT
 Let $\Psi(z) \equiv (U_r e)^z (1/e)$.
 At time t , if fraction of budget spent is $z(t)$, then bid $b_0(t) = \frac{V}{1 + \Psi(z(t))}$

Notice that both strategies only need the fraction of budget spent and are thus oblivious to the other parameters of the auction. We use Profit and Revenue to

denote the profit and revenue earned by the above strategies respectively, and OPT_p and OPT_r to denote the profit and revenue of an omniscient bidder. Then we have the following theorem:

Theorem 4. (i) For single-slot profit maximization, for any $\epsilon > 0$,

$$\text{OPT}_p \leq \epsilon B + \left(\ln \left(\frac{(V - b_{\min})}{\epsilon b_{\min}} \right) + 1 \right) \text{ Profit.}$$

(ii) For single-slot revenue maximization, assuming that OPT does not overbid at time t where $b(t) > V$,

$$\text{OPT} \leq \left(\ln \left(\frac{V}{b_{\min}} \right) + 1 \right) \text{ Revenue.}$$

The proof of Theorem 4(i) follows from Theorem 1 and the fact that all items with efficiency $\leq \epsilon$ has total value at most ϵB . Theorem 4 also suggests that different ϵ values give different guarantees for Profit, thus we can choose ϵ appropriately to maximize the guaranteed value of Profit. In practice, it turns out we can treat L , the lower bound of all items' efficiency, as a tunable parameter (essentially ignoring all items with efficiency less than L), and significantly improve the performance of the bidding algorithm. We will discuss this in Section 5.2. The proof of Theorem 4(ii) follows from Theorem 1 setting $L = 1$. The assumption is valid if the budget B is not exceedingly large. In practice, even if the advertiser wants to maximize revenue, rarely is he willing to buy unprofitable keyword positions.

4.1 Multiple-Slot Bidding Strategies

For multiple-slot auctions we consider both profit-maximizing and revenue-maximizing cases. At each time period, bidder 0 has to decide which slot should he outbid to win. The algorithm suggests bidding so as to get maximum profit (revenue) while having a minimum efficiency. Unfortunately, bidding to get maximum profit requires knowledge of other bidders bids. On the other hand, assuming that clickthrough rates increase as we move up the slots, bidding higher would only give a higher revenue.

The profit-maximizing bidding strategy is presented below. The parameters are the same as in the single-slot case. Notice that the bidding strategy is still oblivious of $X(t)$, however now requires knowing the bids $b_s(t)$ and also $\alpha(s)$.

Bidding Strategy. PROFIT-MAXIMIZING MULTIPLE-SLOT

Fix $\epsilon > 0$. Let $\Psi(z) \equiv (Ue/\epsilon)^z(\epsilon/e)$.

At time t , let $z(t)$ be fraction of budget spent,

$$E_t \equiv \left\{ s \mid b_s(t) \leq \frac{V}{1 + \Psi(z(t))} \right\},$$

bid $b_s(t)$ where

$$s = \arg \max_{s \in E_t} (V - b_s(t))\alpha(s).$$

For revenue maximization, we can actually find the slot s in time t to maximize the revenue. This is because, the revenue obtained on bidding $b_s(t)$ is $VX(t)\alpha(s)$. Given that $\alpha(s)$ is a decreasing function, maximizing $VX(t)\alpha(s)$ is equivalent to minimize s , i.e., to find the rank s as low as possible. Since the efficiency condition imposes that the slot we win have $b_s(t) \leq \frac{V}{\Psi(z(t))}$, our bid should be exactly that. Thus we have a bidding strategy for revenue-maximizing multiple-slot auctions which is *exactly the same* as that for single-slot auctions and has the desirable property of obliviousness.

Similar to the performance guarantee of the single-slot bidding strategies in Theorem 4, the above bidding strategies have performance guarantees, stated as the following theorem:

Theorem 5. (i) For multiple-slot profit maximization, for any $\epsilon > 0$,

$$\text{OPT}_p \leq \epsilon B + \left(\ln \left(\frac{V}{\epsilon b_{\min}} \right) + 2 \right) \cdot \text{Profit}.$$

(ii) For multiple-slot revenue maximization,

$$\text{OPT}_r \leq \left(\ln \left(\frac{V}{b_{\min}} \right) + 2 \right) \cdot \text{Revenue} .$$

5 Experimental Exploration

In this section, we evaluate our bidding algorithms using both synthetic and real-world data, and discuss two useful heuristics: sniping and parameter tuning.

5.1 Simulation and the Sniping Heuristic

We now discuss an experiment for single-slot auctions that points out a weakness of the bidding strategy. We then modify the strategy which, although having the same theoretical guarantee, performs much better empirically. As a negative, the strategy does not remain oblivious any more: it requires knowledge of $X(t)$, the traffic function and also α , the clickthrough-rate of the slot.

Figure 1 shows the performance of our algorithm in a simulation against bidders whose bids are random variables. The budget of the bidder is \$1000 and value $V = \$8.00$. Figure 1 shows our strategy obtains around 40% of that obtained by the omniscient bidder (the theoretical bound is around 13%). The advertiser stops overbidding very early, at around $t = 200$, and has an unspent budget of \$425.

At time t , suppose the fraction of budget remaining is $y(t) = 1 - z(t)$. Moreover assume we know future click traffic $X(\tau)\alpha$ for $t < \tau \leq T$. Thus the maximum number of clicks in the remaining time is $\int_t^T X(\tau)\alpha \cdot d\tau$, and bidding at most $\frac{y(t) \cdot B}{\int_t^T X(\tau)\alpha \cdot d\tau}$ from time t to T would avoid exhausting the budget. This suggests the following modified strategy which in the toy example of Figure 1 almost doubles the profit.




Fig. 1. Performance comparison of various bidding strategies in presence of one other bidder who bids a price uniform random in $[4, 6]$

Bidding Strategy: PROFIT-MAXIMIZING SINGLE-SLOT WITH SNIPING

Fix $\epsilon > 0$. Let $\Psi(z) \equiv (Ue/\epsilon)^z(\epsilon/e)$.

At time t , if fraction of budget spent is $z(t)$, bid

$$\max \left\{ \frac{V}{1 + \Psi(z(t))}, \frac{(1 - z(t)) \cdot B}{\int_t^T X(\tau)\alpha \cdot d\tau} \right\}.$$

The following theorem shows that the sniping does not affect the worst-case behavior of the strategies.

Theorem 6. *The modified bidding strategy using sniping always obtains at least as much profit as the original bidding strategy.*

The above sniping heuristic can be generalized to the multiple-slot case as well.

Bidding Strategy: MULTIPLE-SLOT WITH SNIPING

At time t , let $z(t)$ denote fraction of budget spent, $\rho = \Psi(z(t))$

For each slot s , if $\rho > \frac{v_s(t)}{w_s(t)}$ & $b_s(t) \leq \frac{(1-z(t))B}{\alpha(s) \int_t^T X(\tau)d\tau}$:

$$\rho = \frac{v_s(t)}{w_s(t)}$$

$$E_t = \{s \mid \frac{v_s(t)}{w_s(t)} \geq \rho\}$$

bid $b_s(t)$ where $s = \arg \max_{s \in E_t} v_s(t)$

5.2 Evaluation Using Real Bidding Data

Parameter Tuning. If the lower bound L in the online knapsack problem is too small, we can replace it with a larger value $L' > L$ for the threshold function Ψ . This essentially discards items with very low efficiency, and the loss is minimal if the optimal solution consists of items with relatively high efficiency. It turns out tuning the parameter L makes a significant performance improvement empirically. If we choose $L = 0.1$ for profit maximization, we get less than 50% performance without sniping and about 70% with sniping. However, with L tuned and fixed for the non-sniping case, we get much better results.

Next we report some experimental results on evaluating bidding algorithms for multiple-slot auctions using real bidding data. We scraped bidding data from the now defunct Overture webpage [1] with continuous crawling for about two weeks, for one of the most dynamic and expensive keyword “auto insurance.” There are totally $T = 1842$ distinct time periods in our collected data, and most top-5 bids are larger than \$10. For the experiments, we use $B = 1000$, and three different values $V = 8, 10, 12$. We evaluated both the profit-maximizing and revenue-maximizing strategies with and without sniping. For all these experiments, we use $U = V/b_{\min} - 1$ for profit maximization and $U = V/b_{\min}$ for revenue maximization, and $b_{\min} = 0.9$. The lower bound L is optimized for each instance without sniping, and it remains the same for the sniping version.

We summarize the experimental results in Table 5.2. For all the examples we run, sniping improves the bidding performance significantly while exhausting the budget. Table 5.2 seems to tell us, for almost all values, with parameter tuning of L , the performance ratio (ALG/OPT) is around 70%-75% without sniping, and 90%-95% with sniping.

V	OPT	ALG	ALG/ OPT	budget left	ALG (sniping)	ALG/ OPT
8	3779	2751	73%	225.5	3541	94%
10	4974	4059	82%	116.1	4607	93%
12	6169	4463	72%	240.8	5842	95%

6 Concluding Remarks

The algorithms in the paper can be extended to the general case where there are multiple keywords and each keyword has multiple positions. The competitive ratio would now have V replaced by V_{\max} , where V_{\max} is the maximum valuation for all keywords. As an open problem, there is a gap of additive constant 1 between the lower and upper bounds for the competitive ratio of Online-MCKP, and it will be nice to close the gap.

References

1. Overture view bids (February 2007), <http://uv.bidtool.overture.com/d/search/tools/bidtool>
2. Awerbuch, B., Azar, Y., Plotkin, S.A.: Throughput-competitive on-line routing. In: FOCS, pp. 32–40 (1993)
3. Buchbinder, N., Jain, K., Naor, J.S.: Online primal-dual algorithms for maximizing ad-auctions revenue. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) ESA 2007. LNCS, vol. 4698, pp. 253–264. Springer, Heidelberg (2007)
4. Marchetti-Spaccamela, A., Vercellis, C.: Stochastic on-line knapsack problems. *Mathematical Programming* 68, 73–104 (1995)
5. Mehta, A., Saberi, A., Vazirani, U.V., Vazirani, V.V.: Adwords and generalized on-line matching. In: Proc. FOCS, pp. 264–273 (2005)
6. Yao, A.C.-C.: Probabilistic computations: towards a unified measure of complexity. In: Proc. 18th IEEE FOCS, pp. 222–227 (1977)

Position Auctions with Bidder-Specific Minimum Prices

Eyal Even-Dar, Jon Feldman, Yishay Mansour*, and S. Muthukrishnan

Google Research, New York, NY
{evendar, jonfeld, mansour, muthu}@google.com

Abstract. Position auctions such as the Generalized Second Price (GSP) auction are in wide use for sponsored search, e.g., by Yahoo! and Google. We now have an understanding of the equilibria of these auctions, via game-theoretic concepts like Generalized English Auctions and the “locally envy-free” property, as well as through a relationship to the well-known, truthful Vickrey-Clarke-Groves (VCG) mechanism. In practice, however, position auctions are implemented with additional constraints, in particular, bidder-specific minimum prices are enforced by all major search engines. The minimum prices are used to control the quality of the ads that appear on the page.

We study the effect of bidder-specific minimum prices in position auctions with an emphasis on GSP. Some properties proved for standard GSP no longer hold in this setting. For example, we show that the GSP allocation is now not always efficient (in terms of advertiser value). Also, the property of “envy-locality” enjoyed by GSP—which is essential in the prior analysis of strategies and equilibria—no longer holds. Our main result is to show that despite losing envy locality, GSP with bidder-specific minimum prices still has an envy-free equilibrium. We conclude by studying the effect of bidder-specific minimum prices on VCG auctions.

1 Introduction

The Internet economy has been revolutionized by the introduction of sponsored search links. Sponsored links are a small number of advertisements (*ads*, henceforth) that the search engine displays in addition to the standard search results. These ads are arranged in *positions* top to bottom, typically on the side. Normally, the advertiser pays only when the user clicks on the link. It is a difficult task to set a fixed price for each position because the search queries vary widely and with them the value of the positions. Hence, typically, auctions are used to determine the prices, and these are called *position auctions*. A major task for the search engine is to determine the rules of the position auction, and to select, rank and price the ads that will be displayed to the user, according to that auction.

Today, both Google and Yahoo! use a position auction called the *generalized second price auction* (GSP). The GSP auction ranks the ads by the product of

* School of Computer Science, Tel Aviv University.

the advertiser's bid with a quality score, which is often abstracted as the *click-through rate* (*ctr*)—the probability that the user will click on the advertisement. Then, the ad in position i is charged based on the bid of the ad on position $i + 1$.

There is a great need to understand the behavior of these auctions since they are part of everyday life of many, with hundreds of millions being run each day, worldwide. Decades of research in economic, game and auction theories provide the tools to design and understand auctions. However, position auctions—and GSP in particular—have needed new specific methods such as the recent results of [10] and [4]. Specifically, they developed the notion of *Generalized English Auctions* to study GSP, introduced a new notion of “locally envy-free” equilibrium to characterize GSP, and related such equilibria of GSP to that of the well-known Vickrey-Clarke-Groves (VCG) mechanism (as applied to a position auction). In particular, the “locally envy-free” property captures the dynamics of advertisers trying to move up or down the list of positions and plays a crucial role in understanding the equilibrium properties of GSP. (An excellent introduction can be found in [9].)

The departure in our work from prior research begins with the observation that while Google and Yahoo! do implement GSP, they add other features. In particular, they are driven by the need to present high quality ads, and as a result, include features that encourage advertisers to make high-quality ads. One important such feature is the use of advertiser-specific factors for setting *minimum prices*.¹ Beyond the standard use of advertisers' bids and quality scores as in GSP, the search engines force the bid and price per click of advertiser i to be at least a minimum price R_i . Minimum prices are motivated by revenue concerns, but perhaps more importantly in this setting by ad quality control.² Indeed it is well-known from the work of Myerson [8] and others that minimum prices are important for revenue; however optimizing only for short-term revenue ignores the effect of low-quality ads on the quality of service felt by the user, and therefore the future revenue of the search engine. The exact formulation of this problem is very subtle (see [1] for an example). Yet most mechanisms that will try to address this subtlety are likely to enforce bidder specific minimum prices.

The immediate impact of minimum prices is not only that advertisers may pay more than what is determined by GSP, but more importantly for sponsored search, because the “heavy tail” of infrequent keywords often has only a few advertisers per query, the minimum price determines whether or not the ad will appear for that query. So, advertiser-specific minimum prices have a profound effect on advertisers, users and search engines in practice.

This motivates the question that is the focus of this paper: What are the strategic changes in the outcome of GSP—and more generally, other position auctions—in presence of advertiser-specific minimum prices? For example, while the introduction of minimum prices looks innocuous, does it affect truthfulness

¹ See <https://adwords.google.com/support/bin/answer.py?answer=49177> and <http://help.yahoo.com/help/1/us/yahoo/ysm/sps/start/overview.qualityindex.html>

² Therefore, they are not just a side effect of charging advertisers per click and maximizing revenue per impression.

or equilibrium properties of position auctions? A quick sanity check is to study VCG, and doing so reveals that a naïve post-VCG enforcement of bidder-specific minimum bid prices can break the truthfulness property. Being more careful, we can show suitably modified allocation and pricing that is a truthful variant of VCG; this modification shows the impact of minimum prices for VCG.

We primarily study GSP. Since GSP is not truthful to begin with, we study the effect of bidder-specific minimum prices on the equilibria. A simple example with just two bidders shows that minimum prices can cause a loss of efficiency. Furthermore, we see that an important property enjoyed by basic GSP no longer holds: namely, “envy locality.” This property says that if a bidder in position i is in a state where she does not envy the bidders in adjacent positions ($i - 1$ and $i + 1$), then she does not envy any other bidders either. Envy locality is a strong property on its own, as it makes equilibrium discovery simpler for the bidder [4,10]. It is also essential in the existing proofs that there is an envy-free equilibrium of GSP.

Our main result, which was also the most technically challenging one, is to show that despite losing envy locality, GSP with bidder-specific minimum prices still has an envy-free equilibrium. To derive the prices of this equilibrium, we define a specialized Tâtonnement process that takes a global view of the best-response relationship between bidders and positions. This global view was unnecessary in the basic GSP analysis such as in [4] because of envy locality. We prove that the process converges to a set of prices from which an envy-free equilibrium set of bids is derived.

Demange et al. [3] consider a multi-item auction in which every buyer is interested in bundles of size at most one; using a Tâtonnement process similar to ours they show how to compute equilibrium prices (but without bidder-specific minima which is our concern). However in the context of position auctions, their technique is somewhat appropriate for *position-specific* reserve prices, see [5].

Remark. The proofs of the theorems appear in the full version of the paper.

2 Model

A position auction is defined by a tuple (N, K, v, α, β) . The set $N = \{1, \dots, n\}$ is the set of bidders and the set $K = \{1, \dots, k\}$ is the set of positions. Each bidder $i \in N$ is associated with two values, v_i which is its valuation for a click and α_i which is its click through rate (ctr). Each position $\ell \in K$ is associated with a click through multiplier β_ℓ . As a convention, $\beta_\ell > \beta_{\ell+1}$, $\beta_1 = 1$ and $\beta_{k+1} = \dots = \beta_n = 0$ (therefore, effectively we have k positions).

We use the standard assumption that the actual click through rate of bidder i in position j is the product of the bidder’s ctr α_i and the position click through multiplier β_j , i.e., if bidder i is placed at position j then she receives a click with probability $\alpha_i \beta_j$. We assume that the value of v_i is known only to bidder i while all the other parameters are publicly known.

In a position auction mechanism, each bidder i submits a bid b_i . Given all the bids b and the public information (N, K, α, β) , the mechanism assigns each bidder i to a position $\text{loc}(i)$ and charges it a price $P_{\text{loc}(i)}(b, \alpha, \beta)$ per click. The utility of bidder i at position j is: $u_i(j) = \alpha_i \beta_j (v_i - p_j)$, where $p_j = P_j(b, \alpha, \beta)$ is the price per click. The mechanism assigns to each position exactly one bidder, and therefore we have an inverse function $\text{loc}^{-1}(j)$ that returns the bidder that was assigned to position j . We call such an assignment loc a legal assignment.

The two most studied position auction mechanisms are GSP [10,4] and VCG [11,2,6]. The VCG mechanism ranks the bidders by $b_i \alpha_i$, which can be thought of as the expected advertiser value if $b_i = v_i$. Therefore, the VCG allocation maximizes the *social welfare*, which is the sum of the bidders' expected value, i.e., $\sum_{i \in N} v_i \alpha_i \beta_{\text{loc}(i)}$. The VCG mechanism charges each bidder i the total value lost by other bidders that is caused by i 's presence in the auction; it has the property that each bidder's dominant strategy is to bid her true value, i.e., $b_i = v_i$.

The GSP mechanism ranks the bidders by $b_i \alpha_i$.³ Again, without loss of generality, assume that bidder i is assigned to position i . The price that the bidder at position i pays per click is $b_{i+1} \alpha_{i+1} / \alpha_i$. It was shown by [10,4] that for any position auction, there exists an envy-free equilibrium (defined below) such that both the allocation and the payments of the GSP and the VCG mechanism are identical.

While so far we have described the traditional theoretical model, the position auctions used in practice contain an additional important feature, namely bidder-specific minimum prices. The minimum prices imply that each bidder i has a minimum price R_i , which is *known* to the bidder. If bidder i submits a bid b_i that is smaller than R_i then it will not participate in the auction (and have a zero utility). For simplicity we will assume that for each bidder i the bid b_i is at least R_i . Therefore, the price per click of bidder i is at least R_i , i.e., its price per click at position $j = \text{loc}(i)$ is $\max\{R_i, P_j(b, \alpha, \beta)\} = \max\{R_i, b_{i+1} \alpha_{i+1} / \alpha_i\}$. The focus of our work is to study the effect of bidder-specific minimum prices. We will show that this small modification to the auction mechanism can dramatically influence the behavior of the bidders.

3 Generalized Second Price (GSP) Auctions

Since GSP is not a truthful mechanism, our main focus is to show that there exists an equilibrium. In fact we will show a stronger result, that there are envy-free prices for GSP.

The existence of envy-free prices for GSP with no minimum bids (or a uniform minimum bid which is identical for all the bidders) was shown in [4,10]. Along the way, these analyses show a few interesting properties of GSP. The first is that there are envy-free prices that result in an efficient allocation (i.e., maximize the sum of bidders valuations). The second is the fact that local envy-free prices imply (global) envy-free prices. In the full version of the paper we show that both of those properties do not hold once bidder-specific minimum prices

³ One can derive the "rank by bid" mechanism by setting $\alpha_i = 1$ for all $i \in N$.

are introduced. We also discuss in the full version the effect of bidder specific minimum prices on the revenue and extend the results to unequal CTR.

For the proof of the existence of envy-free prices we use a specific Tâtonnement process. Our process increases prices while ensuring that certain properties of the allocation are maintained. The remainder of the section will be devoted to sketching our main theorem, extending the existence of envy-free equilibria to the case of bidder-specific minimum prices:

Theorem 1. *The GSP mechanism with bidder-specific minimum prices has an envy-free prices equilibrium.*

The proof technique that we will use to show Theorem 1 is to define a specific Tâtonnement process, and show that it converges to a set of envy-free prices.

Best Response Graph. We will define the Tâtonnement process on a particular graph that models the envy relation between bidders and positions. Let $K' = \{1, \dots, k\}$ be the set of positions with non-zero multiplier, i.e., $\beta_i > 0$. Given a price vector P for any subset of bidders $B \subset N$ and subset of positions $S \subseteq K'$ we define the *best response graph*, $G(P, B, S) = (B, S, E)$. The graph $G(P, B, S)$ is a bipartite graph where $(b, s) \in E$ if and only if position $s \in S$ is a best response for bidder $b \in B$. We say that positions $i \in S$ and $j \in S$ are *connected* if there exists a path between i and j in $G(P, B, S)$. We denote by $\nu_G(v)$ the neighbors of a node v in $G = G(P, B, S)$. We use the notation $P' = (P, \epsilon, j)$ to denote a price update of position j by ϵ , i.e., $p'_i = p_i$ for every $i \neq j$ and $p'_j = p_j + \epsilon$. We also let $S_{NE}(P) \subseteq K'$ be the set of positions that are a best response for at least one bidder at the prices P ; equivalently $S_{NE}(P)$ is the set of position nodes $s \in K'$ in $G = G(P, N, K')$ with at least one incident edge, i.e., $|\nu_G(s)| \geq 1$. We say that a set $S \subseteq K'$ is *matched* in $G = G(P, N, S)$ if there is a perfect matching in $G(P, B, S)$ for some $B \subseteq N$. To simplify our notation, whenever P, B, S or G are clear from the context we might omit them. We will also assume that all α_i 's are equal and that all v_i s are different; the extension to the general case appears in the full version.

The Tâtonnement process. The Tâtonnement process is described formally in the figure, and here we provide some useful intuition. The Tâtonnement process begins with a set of prices P_1 such that the first k bidders all prefer the first position; i.e., $S_{NE}(P_1) = \{1\}$, $B_1 = \{1, \dots, k\}$ and $G(P_1, B_1, K')$ is a star graph where each bidder $i \in B_1$ has exactly one edge to the node for position 1. The Tâtonnement process gradually increases prices, increasing the price of only one position during each update. While increasing the prices the algorithm preserves two invariants: (1) At each step, with prices P_t , the set of positions $S_t = S_{NE}(P_t)$ that are the best response for some bidder can only grow; i.e., $S_t \subseteq S_{t+1}$, and (2) There is a matching of the positions S_t , such that every position in S_t can be matched to a unique bidder in $G(P_t)$.

Both invariants are preserved by maintaining the conditions of Hall's theorem 7 on every subset of $S_t = S_{NE}(P_t)$, i.e., for every subset $S' \subset S_t$ we require that $|S'| \leq |\nu(S')|$ which is a sufficient and necessary condition for a matching by Hall's theorem.

Tâtonnement process TP

Initialize P_1 such that $p_j = v_{k+1}$ for $j \leq k$ and $p_j = 0$ for $j \geq k + 1$;
 Let $t = 1$ and $S_1 = \{1\}$;
while $\exists \epsilon > 0, j \in S_t: \text{MATCH}(P_t, \epsilon, j) = \text{TRUE}$ and $S_t \neq K'$ **do**
 For each $j \in S_t$ let $\epsilon_j = \max\{\epsilon : \text{MATCH}(P_t, \epsilon, j)\}$;
 $s_t = \arg \max_{j \in S_t} \epsilon_j; \epsilon_t = \epsilon_{s_t}; P_{t+1} = (P_t, \epsilon_t, s_t); S_{t+1} = S_{\text{NE}}(P_{t+1}); t = t + 1$;
end
 Output the set of price P_t and the allocation is a matching in $G(P_t, N, K)$.

$\text{MATCH}(P, \epsilon, j) = \text{TRUE}$ iff there is a matching for $S' = S_{\text{NE}}(P') \cup \{j\}$ in $G(P', N, S')$, where $P' = (P, \epsilon, j)$.

We first show that the Tâtonnement process TP cannot loop indefinitely if all numbers are rational. This is done by showing that there exists ϵ_{\min} , which is a function of $v_1, \dots, v_n, R_1, \dots, R_n, \beta_1, \dots, \beta_k$, and $\alpha_1, \dots, \alpha_n$, where every increase will be at least ϵ_{\min} .

Lemma 1. *The Tâtonnement process TP always terminates.*

Now we would like to prove a few facts on how TP makes progress, until it terminates. Specifically, we would like to show that the set S_t increases monotonically and furthermore, each time it changes it adds the least position that is not in S_t . Therefore, initially we have $S_1 = \{1\}$, and at any time t we will have $S_t = \{1, \dots, j\}$ for some $j \in K'$. The following two lemmas establish this property.

The following observation shows the effect of a price increase at position j , which is the basic step of the Tâtonnement process TP.

Lemma 2. *Let $P' = (P, \epsilon, j)$, $(N, K', E) = G(P, N, K')$ and $(N, K', E') = G(P', N, K')$. Then every edge $(b, i) \in E - E'$ is incident to j ; i.e., $i = j$. Also, if there is an edge $(b, s) \in E' - E$ then the edge (b, j) is in E .*

The following lemma shows that we preserve the invariant that S_t monotonically grows. Furthermore, since the positions have a strict preference order which is shared by all the bidders, the sets S_t are prefixes of $[1, \dots, k]$ and can grow by at most one position each time step.

Lemma 3. *Let $S_t = S_{\text{NE}}(P_t)$ at time t . Then, for every time $t' > t$ we have $S_t \subseteq S_{t'}$. In addition, if $j \in S_t$ then any $i \leq j$ has $i \in S_t$, and if $S_t \neq S_{t+1}$ then $|S_{t+1} - S_t| = 1$.*

We have shown that the Tâtonnement process terminates, and maintains the key invariants mentioned. Since we are maintaining a matching for the set S_t , we essentially just need to show that when the Tâtonnement process TP terminates, we have $S_t = K'$. By Lemma 3 it is sufficient to show that at some time $S_t = K'$, since S_t is monotone. Thus the remaining lemmas are steps to show that if $S_t \neq K'$ then some price can be increased, and therefore the Tâtonnement process does not terminate.

First, consider the case that there is a bidder who has only one position as a best response. We show that in this case we increase the prices and thus the Tâtonnement process TP cannot terminate.

Lemma 4. *Let $G = G(P_t, N, S_t)$ be a best response graph such that $S_t \neq K'$. If for some bidder $b \in N$ we have $\nu_G(b) = \{j\}$, then there exists an $\epsilon > 0$ such that $\text{MATCH}(P_t, \epsilon, j) = \text{TRUE}$.*

We say that prices P induce equal payments if for any position $s \in S_{\text{NE}}(P)$ and any two bidders i, i' for which s is a best response, then $\max\{R_i, p_s\} = \max\{R_{i'}, p_s\}$. For the process it is important to distinguish between prices which induce equal payments and ones which do not. The next lemma claims that in certain subgraphs if P_t do not induce equal payments then they can be increased.

Lemma 5. *Let P_t be prices which do not induce equal payments, and $G = G(P_t, N, S_t)$ be a best response graph, where $S_t \neq K'$. If in G every subset S' of S_t satisfies $|S'| \leq |\nu_G(S')| - 1$, $|S_t| \geq 2$ and G is connected then there exists a position $j \in S_t$ and $\epsilon > 0$ such that $\text{MATCH}(P_t, \epsilon, j) = \text{TRUE}$.*

Next we claim that if the prices induce equal payments, then two bidder can have at most one position in the intersection of their best response sets.

Lemma 6. *Let P_t be prices which induce equal payments, and let $G = G(P_t, N, S_t)$ be a best response graph. For any bidders $i, i' \in N$ we have $|\nu_G(i) \cap \nu_G(i')| \leq 1$.*

The next lemma shows that if in a subgraph every subset of positions has some slack with respect to the Hall's theorem condition, i.e., each subset of positions has strictly more bidders connected to in the graph, then there exists a price we can increase without violating the matching constraint.

Lemma 7. *Let $G = G(P_t, B, S_t)$ be a best response graph. If in G every non-empty subset S' of S_t satisfies $|S'| \leq |\nu_{G(P)}(S')| - 1$, $|S| \geq 2$ and G is connected then there exists some $j \in S_t$ and $\epsilon > 0$ such that $\text{MATCH}(P_t, \epsilon, j) = \text{TRUE}$.*

By combining the above lemmas we prove Theorem [11](#). We remark that in the Tâtonnement process TP we do not necessarily terminate with a price vector P_t that induces equal payments. We might terminate with prices P_t that do not induce equal payments, since we already reached a state in which $S_t = K'$. Our proof technique only shows that as long as $S_t \neq K'$ we can increase the price of some position.

4 VCG Auctions

The VCG mechanism gives a general methodology to implement truthful mechanisms. The mechanism is aimed at maximizing a social welfare function which is the sum of the bidders' utilities. The basic idea of the mechanism is that each bidder pays its marginal influence on the social welfare function of other bidders. In the full version of the paper we investigate three possible modifications to the

VCG payments, with the goal of incorporating bidder-specific minimum prices while still maintaining the desirable properties of the mechanism (being truthful and efficient):

(i) Naïve implementation of VCG: The most natural implementation of VCG when minimum prices are enforced is to first compute the VCG prices, and then the price for bidder i is the maximum of its VCG price and its minimum price. Unfortunately, we show that the resulting mechanism is not truthful.

(ii) Virtual Values: Since the naïve approach to incorporating bidder-specific minimum prices fails, we would like to explore another approach. We first make the observation that if for some bidder i , every other bidder i' with $\text{loc}(i') < \text{loc}(i)$ had $b_{i'}\alpha_{i'} \geq R_i\alpha_i$, then the (unmodified) VCG price for i would be at least R_i . This observation motivates introduction of bidder-specific “virtual values”: When computing the price for bidder i we use $\max\{b_{i'}\alpha_{i'}, R_i\alpha_i\}$ as a substitute for $b_{i'}\alpha_{i'}$ for all applicable i' . This implies that the bid of a bidder i' is interpreted differently when computing prices of different bidders. We show that the resulting mechanism is efficient and truthful.

(iii) Offsetting Bid by Minimum Price: A generic approach of incorporating minimum prices is to subtract from the original bid the bidder’s minimum price, run a truthful auction, and add the minimum price at the end (see also [11]). We show that the resulting mechanism is efficient and truthful.

Acknowledgments. We thank Hal Varian and Martin Pál for useful discussions.

References

1. Abrams, Z., Schwarz, M.: Ad auction design and user experience. In: Applied Economics Research Bulletin Special Issue on Auctions (2008)
2. Clarke, E.: Multipart pricing of public goods. *Public Choice* 11, 17–33 (1971)
3. Demange, G., Gale, D., Sotomayor, M.: Multi-item auctions. *The Journal of Political Economy* 94, 863–872 (1986)
4. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97(1), 242–259 (2007)
5. Gonen, R., Vassilvitskii: Sponsored search auctions with reserve prices: Going beyond separability. In: Ad Auction Workshop (2008)
6. Groves, T.: Incentives in teams. *Econometrica* 41(4), 617–631 (1973)
7. Harary, F.: *Graph Theory*. Addison-Wesley, Reading (1994)
8. Myerson, R.: Optimal auction design. *Mathematics of Operations Research* 6, 58–73 (1981)
9. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: *Algorithmic Game Theory*. Cambridge University Press, New York (2007)
10. Varian, H.: Position auctions. *International Journal of Industrial Organization* 25(6), 1163–1178 (2007)
11. Vickrey, W.: Counterspeculation, auctions and competitive-sealed tenders. *Finance* 16(1), 8–37 (1961)

A Cascade Model for Externalities in Sponsored Search^{*}

David Kempe^{1,**} and Mohammad Mahdian²

¹ University of Southern California, Los Angeles, CA
dkempe@usc.edu

² Yahoo! Research, Santa Clara, CA
mahdian@yahoo-inc.com

Abstract. One of the most important yet insufficiently studied issues in online advertising is the externality effect among ads: the value of an ad impression on a page is affected not just by the location that the ad is placed in, but also by the set of other ads displayed on the page. For instance, a high quality competing ad can detract users from another ad, while a low quality ad could cause the viewer to abandon the page altogether.

In this paper, we propose and analyze a model for externalities in *sponsored search ads*. Our model is based on the assumption that users will visually scan the list of ads from the top to the bottom. After each ad, they make independent random decisions with ad-specific probabilities on whether to continue scanning. We then generalize the model in two ways: allowing for multiple separate blocks of ads, and allowing click probabilities to explicitly depend on ad positions as well. For the most basic model, we present a polynomial-time incentive-compatible auction mechanism for allocating and pricing ad slots. For the generalizations, we give approximation algorithms for the allocation of ads.

1 Introduction

Online advertising auctions are run with the goal of assigning advertising slots to bidders in such a way as to maximize social welfare or the revenue of the auctioneer. The common setup is as follows: k slots are available for ads, and may be assigned to (some of) n bidders. When users *click* on an advertiser's ad, this will sometimes lead to a purchase and thus revenue for the advertiser. In other words, in the type of auctions we consider here, only clicks are of interest to the bidders, as opposed to *impressions*, which would matter if the goal were to increase product awareness.

The key quantity an advertiser a is interested in with respect to slot i is the *click-through rate*, the probability that ad a , if placed in slot i , will be clicked. The

^{*} A preliminary version of this paper was presented at the 4th Workshop on Ad Auctions, held in conjunction with the 2008 ACM Conference on Electronic Commerce.

^{**} Supported in part by NSF CAREER award 0545855, and NSF grant DDDAS-TMRP 0540420.

larger the click-through rate, the larger the expected revenue of the advertiser. Hence, any auction aiming to maximize social welfare will need to be based on a model of click-through rates of combinations of ads and slots.

Traditional models are based on the assumption that the click-through rate depends solely on the relevance of the ad and the prominence of the slot on the page. In fact, the most commonly used model makes the even stronger assumption that it is the *product* of the two quantities. The model thus completely discounts the effects of *other ads* shown on the same page. Intuition suggests that a high-quality relevant ad placed more prominently can detract from another ad, or a very low-quality ad may cause the viewer to completely disregard the other ads.

In economics jargon, this effect is called an *externality* of an ad. Ghosh and Mahdian [9] initiated the study of externalities in online advertising. They proposed several models primarily in the context of lead generation advertising, i.e., when the publisher must select an unordered set of advertisers. The main model in [9] is based on a rational choice model for the advertising audience. However, for most of these models, the allocation problem is intractable.

In this paper, we focus exclusively on the case of *sponsored search ads*. Here, the publisher needs to select ads to be placed in a number of slots on a web page. We study the allocation problem and the design of incentive-compatible mechanisms under a simple and intuitive model called the *Cascade Model*. Our model generalizes the Cascade Model recently proposed by Craswell et al. [6] in the context of click-through rates of organic search results. (The same model was proposed independently and simultaneously by Aggarwal et al. in the 4th Workshop on Ad Auctions [1].) The basic Cascade Model (defined formally in Section 2) assumes that the users scan through the ads in order. For each ad a , users decide probabilistically whether to click (with some ad-specific probability q_a), as well as whether to continue the scanning process, with a possibly different ad-specific probability c_a . The probabilistic continuation allows us to model the externality of prematurely terminating the scanning process as a result of either a very irrelevant ad, or a very high-quality web site leading to a purchase.

Craswell et al. [6] considered the special case of the Cascade Model where $q_a = c_a$ for all a , in the context of organic search results. Their work is motivated by the work of Joachims et al. [11], which provides limited experimental evidence for the hypothesis that the click-through rate of a search result depends on surrounding results. Craswell et al. compare the Cascade Model with four other models, including the commonly used model of separable click-through rates. They show that the Cascade Model provides the best fit to click logs of a large search engine. Since the click-through rates in organic search results and ads appear to be of a similar nature (and so far the same models have been used for both), this provides evidence that the Cascade Model can achieve a significant improvement over the currently used model of separable click-through rates.

We show that under the Cascade Model we define, the optimum allocation can be computed in polynomial time, and priced so as to lead to an incentive-compatible mechanism. We subsequently consider several generalizations of the

Cascade Model. The first generalization concerns the placement of multiple separate columns or slates of ads. While each slate is scanned from top to bottom, different types of users have different orders among the slates. We show (in Section 4) that the allocation problem for this model admits a polynomial-time approximation scheme (PTAS). The second generalization (in Section 5) is a common generalization of the Cascade Model and the separable click-through rate model, which augments the Cascade Model by slot-specific click probabilities. For the allocation problem in this model, we give a simple 4-approximation algorithm as well as a quasi polynomial-time approximation scheme.

1.1 Related Work

Ad Auctions in general have received a lot of recent attention (see, e.g., [2,4,12,14,15]). Many of the core theoretical results (e.g., [8,16,2,13]) are based on the simplifying assumption of *separable click-through rates*. That is, the probability of ad a being clicked in position i is the product $q_a \lambda_i$ of an ad-specific term and a position-dependent one. This assumption has been made mostly for simplicity; experimental studies find that there is very little evidence that separable click-through rates constitute an adequate model [11,6,18].

As a result, several recent papers have proposed more general models for ad auctions, including some that model externalities. In addition to the work by Craswell et al. [6] on organic search results, a paper by Aggarwal et al. [1] independently and simultaneously proposes the same Cascade Model as this paper for ads. They derive the same dynamic programming algorithm to solve the allocation problem in the basic Cascade Model, and then focus on improving the running time and proving monotonicity properties of this algorithm, while our focus is on solving the allocation problem for generalizations of the Cascade Model that take position-dependent effects or multiple slates of ads into account. A paper by Das et al. [7] studies a different model of externalities. In their model, the click-through probabilities are essentially the same as in the standard separable model. However, the authors model externalities in the *conversion* of clicks, in that users will purchase from *at most one* of the sites they visited.

Athey and Ellison [3] model the consumer search behavior when the consumer is unaware of the quality of the advertisers, and uses the ranking given by the search engine as a signal of the quality. Their model is similar to ours in that the consumer starts scanning the ads from the top, and continues until her need is met or until she decides that the probability that the next ad meets her need is so low as to not be worth the cost of a click. They analyze the equilibria of the *generalized second price mechanism* under strong assumptions about the distributions of user costs and bids, whereas in this paper, we design new allocation mechanisms for *arbitrary* parameter settings to improve the efficiency.

While there has been a huge body of work in advertising on the effects of ad placement, size, etc. on readers' attention and recall, there appears to be, surprisingly, no study of the externalities between ads even in traditional media such as printed advertising or TV. Thus, a comparison between the Cascade Model and traditional models for externalities in advertising is not possible.

2 Click-through Models and Allocations

A publisher needs to choose ads from a set of n ads to display in k slots on a page, numbered sequentially from 1 to k . Each advertiser a specifies an amount b_a : how much they are willing to pay for each click on their ads. In order to optimize either his revenue or the social welfare, the publisher therefore needs to predict the probability that an ad is clicked, and take these predictions into account when allocating the slots to ads.

The *click-through rate (CTR)* of an ad is the probability that it receives a click. In principle, this probability could depend on everything on the page, including the ad itself, the position where it is placed, other ads placed in other slots, as well as seemingly less relevant other content. Since a model with so many parameters will not be useful for designing a prediction and allocation algorithm, the models currently used simplify the dependence of click-through rates on the information on the page.

The simplest model, which is currently widely used in the industry and also is the basis for most theoretical work in the area (e.g., [8,16,21,3]), is based on *separable click-through rates*. It assumes that the CTR of an ad $a \in \{1, \dots, n\}$ placed in position $i \in \{1, \dots, k\}$ is the product $q_a \lambda_i$. Here, q_a measures the intrinsic *quality* or relevance of ad a , the probability that a user, seeing ad a , will actually click on it. λ_i measures the *prominence* of slot i , and is the probability that the user will see slot i . It is commonly assumed that λ_i is monotonically non-increasing in i . The main advantage of this model is its simplicity. Among others, simply sorting the advertisers by decreasing $b_a q_a$ yields an optimal allocation of the ad space.

2.1 The Cascade Model

In the basic Cascade Model, each ad a , in addition to the intrinsic quality q_a , has a second parameter c_a , called its *continuation probability*. The model assumes that the user behaves as follows:

1. Start with the ad a_1 in slot 1.
2. When looking at the ad a_i in slot i , click on it with probability q_{a_i} .
3. Independently of whether ad a_i was clicked or not, continue to slot $i+1$ with probability c_{a_i} ; otherwise, terminate the scanning process.
4. Terminate the scanning process also once no more ads remain.

For notational convenience, we denote empty slots by \perp , with the understanding that $q_\perp = 0$ and $c_\perp = 1$. Under this model, assuming ads a_1, \dots, a_k are in slots $1, \dots, k$, the user will see a particular slot i with probability $C_i = \prod_{j=1}^{i-1} c_{a_j}$, and the click-through rate of ad a_i is therefore

$$r_{a_i} = q_{a_i} \cdot C_i = q_{a_i} \cdot \prod_{j=1}^{i-1} c_{a_j}. \quad (1)$$

Remark 1. Our model subsumes an apparently more natural generalization with two different conditional probabilities c_a^+, c_a^- . If c_a^+ is the continuation probability

if ad a is clicked, and c_a^- the continuation probability if a is not clicked, then by setting $c_a = q_a c_a^+ + (1 - q_a)c_a^-$, it is easy to see that the resulting click-through rates remain the same for each position.

As we mentioned above, a simpler version of the Cascade Model has been proposed recently by Craswell et al. [6] in the context of organic search results. They assume that $c_a = q_a$ for all ads a .

2.2 Generalized Cascade Models

Multiple Ad Slates. Many search engines present sponsored search ads in multiple different slates, e.g., some preceding the organic search results and some on the right-hand side. As a result, different users may have different orders in which they scan the ads. We will define the corresponding *Slated Cascade Model* as a special case of a very general (and likely intractable) *Permuted Cascade Model*.

In the Permuted Cascade Model, for each permutation π of $\{1, \dots, k\}$, a fraction f_π of users will scan the ads in the order $\pi(1), \dots, \pi(k)$. Then, the probability that a user with scanning order π will look at slot i is $C_i^{(\pi)} = \prod_{j=1}^{\pi^{-1}(i)-1} c_{a_{\pi(j)}}$, and the overall CTR of ad a_i in slot i under the distribution f is

$$r_{a_i}^{(f)} = \sum_{\pi} q_{a_i} f_{\pi} C_i^{(\pi)} = \sum_{\pi} q_{a_i} f_{\pi} \cdot \prod_{j=1}^{\pi^{-1}(i)-1} c_{a_{\pi(j)}}. \tag{2}$$

We are particularly interested in the special case of the *Slated Cascade Model*, in which there is a constant number s of *slates*. Slate i has k_i slots; (j, i) denotes the i^{th} slot of the j^{th} slate, and $a_{j,i}$ the ad in slot (j, i) . Each user scans each slate from top to bottom (until stopping the scan). However, different users might have different orders over the slates. Formally, the only permutations π with non-zero frequencies f_π are those defined in terms of a permutation ψ on the slates $\{1, \dots, s\}$ of the slates, which induces the permutation π placing all $k_{\psi(j)}$ slots of slate $\psi(j)$, in their natural order, before all $k_{\psi(j+1)}$ slots of slate $\psi(j+1)$, for each j .

In the Slated Cascade Model, we write $C_{j,i} = \prod_{h=1}^{i-1} c_{a_{j,h}}$ for the probability of reaching slot i , given that slate j is entered, $\overline{C}_j = \prod_{i: a_{j,i} \neq \perp} c_{a_{j,i}}$ for the probability of moving on to the next slate given that slate j is entered, and

$$\hat{C}_{j,i} = \sum_{\psi} f_{\psi} \cdot C_{j,i} \cdot \prod_{h=1}^{\psi^{-1}(j)-1} \overline{C}_h$$

for the overall probability of a random user seeing slot (j, i) .

Position-Dependent Multipliers. In the *Cascade Model with Position-Dependent Multipliers (CMPDM)*, each slot i also has a position-dependent multiplier λ_i , the slot-specific probability of reading an ad in that slot, subject to scanning all the way to the slot [1].

¹ The λ_i can model the fact that users appear to have been conditioned to assume that results listed in higher positions may be more relevant. The model of Athey and Ellison [3] provides a rational justification for this behavior.

In accordance with most of the literature on position auctions, we assume that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k$. The expression (II) for the click-through rate is now

$$r'_{a_i} = \lambda_i q_{a_i} \cdot \prod_{j=1}^{i-1} c_{a_j}. \tag{3}$$

The CMPDM is a common generalization of both the Cascade Model (with all $\lambda_i = 1$) and the separable click-through rate model (with all $c_a = 1$), and is therefore more expressive than both. We will study the allocation problem for this model in Section 5.

2.3 Slot Allocation and Incentive-Compatibility

In the basic Cascade Model, the publisher needs to solve the following optimization problem in order to maximize the value. Assume, without loss of generality, that $n \geq k$. The objective is to select $\ell \leq k$ distinct ads a_1, \dots, a_ℓ to maximize

$$\sum_{i=1}^{\ell} b_{a_i} q_{a_i} \cdot \prod_{j=1}^{i-1} c_{a_j}. \tag{4}$$

In the next section, we will describe an efficient algorithm for solving this optimization problem. However, solving the optimization problem requires knowledge of the parameters b_a, q_a, c_a for each bidder a . While the probabilities q_a and c_a can be learned from click-through histories, the willingness to pay is the valuation v_a which an advertiser assigns to clicks, and thus intrinsically private information. In particular, utility-maximizing advertisers may submit bids $b_a \neq v_a$ if doing so stands to improve their utilities.

In order to extract *truthful* bids from the advertisers, i.e., entice them to submit $b_a = v_a$, the publisher can charge them prices p_a per click, which may differ from the submitted bids. If charged p_a , an advertiser’s utility is $r_a(v_a - p_a)$, where r_a is the click-through rate for advertiser a , and may depend on the entire assignment. A *mechanism* consists of both an allocation rule and a payment rule, giving the payments p_a per click. It is *truthful* or *incentive compatible* if the best strategy of each advertiser a , independent of the strategies of other advertisers, is to bid $b_a = v_a$. We will discuss truthful mechanisms in the Cascade Model below.

3 Winner Determination in the Cascade Model

In this section, we show that the optimal allocation for the simple Cascade Model can be computed by a dynamic program. The key tool for deriving this program is a lemma showing that whichever ads are shown must follow a simple ordering. The results of this section were obtained independently and simultaneously by Aggarwal et al. [1].

Lemma 1. *Assume that the optimal solution places ad a_i in position i . Then, w.l.o.g.,*

$$\frac{b_{a_1} q_{a_1}}{1 - c_{a_1}} \geq \frac{b_{a_2} q_{a_2}}{1 - c_{a_2}} \geq \dots \geq \frac{b_{a_k} q_{a_k}}{1 - c_{a_k}}. \tag{5}$$

Proof. The proof of this lemma relies on an exchange argument similar to the types of arguments in the analysis of greedy scheduling algorithms. Assume that there is a position $i < k$ such that $\frac{b_{a_i} q_{a_i}}{1 - c_{a_i}} < \frac{b_{a_{i+1}} q_{a_{i+1}}}{1 - c_{a_{i+1}}}$. Let $a = a_i, a' = a_{i+1}$ be the two ads in those positions, and consider the alternative ordering placing a in position $i + 1$, and a' in position i , while leaving all other ads in the same slots as before. By Equation (II), the click-through rates for all positions $j \notin \{i, i + 1\}$ remain the same. Recalling that $C_i = \prod_{j=1}^{i-1} c_{a_j}$, the new click-through rates for ads a and a' are $r'_a = q_a c_{a'} \cdot C_i$ and $r'_{a'} = q_{a'} \cdot C_i$. Thus, the total change in value is

$$\begin{aligned} r'_a b_a + r'_{a'} b_{a'} - (r_a b_a + r_{a'} b_{a'}) &= C_i (b_{a'} q_{a'} (1 - c_a) - b_a q_a (1 - c_{a'})) \\ &\geq C_i (b_a q_a (1 - c_{a'}) - b_a q_a (1 - c_{a'})) = 0, \end{aligned}$$

Hence, swapping a and a' will not decrease the value, and repeating such swaps until the ads are ordered according to (5) proves the lemma. ■

We can now use dynamic programming to design a polynomial-time algorithm for the winner determination problem. First, all ads are sorted according to (5). We then fill out a dynamic programming table $A_{n \times k}$, whose entry $A[a, i]$ contains the optimum value that can be obtained from ads a, \dots, n in positions i, \dots, k , *conditioned* on the ad in slot i being read. Once this table is filled out, the solution of the problem is contained in the entry $A[1, 1]$. To fill this table, we use the following recurrence:

$$A[a, i] = \max(A[a + 1, i], b_a q_a + c_a A[a + 1, i + 1]).$$

If ad a is placed in position i , then its *conditional expected value* is $b_a q_a$, and the reader will continue to slot $i + 1$ with probability c_a . Thus, the expected conditional value obtained from slots $i + 1, \dots, k$ is $c_a A[a + 1, i + 1]$, since the ads in slots $i + 1, \dots, k$ will be chosen optimally as well. Summing up, we obtain the following theorem.

Theorem 1. *There is an algorithm with a running time of $O(n \log n + nk)$ which computes the optimal placement of n ads in k slots in the simple Cascade Model.*

3.1 Incentive-Compatible Mechanism Design

To turn the above algorithm into an incentive compatible mechanism, we can use a pricing scheme based on the classical Vickrey-Clarke-Groves (VCG) mechanism [17, 5, 10]. The VCG payment scheme charges each bidder a an amount equal to the externality this bidder imposes on other bidders. The externality can be calculated by removing a from the set of advertisers, running the algorithm again, and computing the total utility of all advertisers in the resulting solution. The VCG payment is then the difference between the value of all advertisers in this new solution, and the value of all advertisers except a in the original optimum. It is well known that this payment scheme gives an incentive-compatible, efficient mechanism for allocating and pricing ads in the simple Cascade Model. Computing the prices can be accomplished with k separate invocations of the dynamic program described above, for a total running time of $O(n \log n + nk^2)$.

4 Multiple Ad Slates

In this section, we give a polynomial-time approximation scheme (PTAS) for the allocation problem in the Slated Cascade Model. The NP-hardness of this problem remains an open question.

The algorithm first discretizes some of the parameters. It then exhaustively searches over all possible aggregate continuation probabilities in each slate, and then runs a dynamic program to assign ads to slots. By conditioning on the aggregate continuation probabilities of slates, the choices for slates become de facto independent. Throughout, we let δ be a suitably small constant whose precise value will be determined below.

4.1 Ignoring Small Probabilities

In order to show that we can ignore, at the cost of only a small loss in approximation guarantee, any ads that will be seen only with small probability, we first prove the following lemma. We prove it in a fairly general form with position dependent multipliers and multiple slates, since it will also be a key building block for our approximation algorithms in Section 5.

Lemma 2. *Let the position dependent multipliers of slate j be $\lambda_{j,1} \geq \lambda_{j,2} \geq \dots \geq \lambda_{j,k_j}$. Let ψ be any distribution over permutations of slates, and OPT the value of the optimum solution.*

For any $\delta > 0$, there is a solution (possibly leaving some slots empty) of value at least $(1 - \delta) \cdot \text{OPT}$ such that for all non-empty slots j , we have $C_{j,i} \geq \delta$. That is, each non-empty slot is reached with probability at least δ , given that its slate is entered in the first place.

Proof. Let $(a_{j,i})_{j,i}$ be an optimal solution of value OPT. For every slate j , let $r_j \leq k_j$ be the last slot containing an ad, and ℓ_j the largest index such that $C_{j,\ell_j} \geq \delta$. Consider the solution that is obtained by moving the ads $a_{j,\ell_j+1}, \dots, a_{j,r_j}$ to slots $(j, 1), \dots, (j, r_j - \ell_j)$, for all j simultaneously, while leaving the remaining slots empty.

Let $C'_{j,i}$ be the new probability of seeing slot (j, i) after the change, given that slate j is scanned. By the choice of ℓ_j , we immediately obtain that $C'_{j,i-\ell_j} \geq \frac{1}{\delta} C_{j,i}$ for all $i \geq \ell_j + 1$. That is, the ad formerly in position (j, i) for $i \geq \ell_j + 1$ is now at least $\frac{1}{\delta}$ times as likely to be reached, given that slate j is scanned. Furthermore, because we only removed ads from slates, we immediately have that $\overline{C}'_j \geq \overline{C}_j$ for all slates j , i.e., it only becomes more likely that scanning of any slate j will finish. Thus, with $\hat{C}'_{j,i}$ denoting the new overall probability of seeing slot (j, i) , we obtain that $\hat{C}'_{j,i-\ell_j} \geq \frac{1}{\delta} \hat{C}_{j,i}$, i.e., each remaining ad is at least $\frac{1}{\delta}$ times as likely to be seen after the change.

Each such ad is now in a slot (j, i) whose position dependent multiplier $\lambda_{j,i}$ is at least as large as the original one, by the sorting of the multipliers. The value of the new solution is thus

$$\sum_j \sum_{i=1}^{r_j-\ell_j} \lambda_{j,i} \hat{C}'_{j,i} q_{a_{j,i+\ell_j}} b_{a_{j,i+\ell_j}} \geq \frac{1}{\delta} \sum_j \sum_{i=\ell_j+1}^{r_j} \lambda_{j,i} \hat{C}_{j,i} q_{a_{j,i}} b_{a_{j,i}},$$

where we used the above argument as well as the fact that the $\lambda_{j,i}$ are sorted. Because the new solution cannot be better than OPT, we obtain that

$$\sum_j \sum_{i=\ell_j+1}^{r_j} \lambda_{j,i} \hat{C}_{j,i} q_{a_{j,i}} b_{a_{j,i}} \leq \delta \cdot \text{OPT},$$

and thus

$$\sum_j \sum_{i=1}^{\ell_j} \lambda_{j,i} \hat{C}_{j,i} q_{a_{j,i}} b_{a_{j,i}} \geq (1 - \delta) \cdot \text{OPT}.$$

In other words, removing the ads in positions $\ell_j + 1, \dots, r_j$ cannot decrease the value by more than $\delta \cdot \text{OPT}$, and ensures that $C_{j,i} \geq \delta$ for all non-empty slots (j, i) by construction. ■

4.2 Description of the Algorithm

We are now ready to describe and analyze the steps of the algorithm in detail:

1. **Ignoring small continuation probabilities.** First, we round any continuation probability c_a that is less than $\delta/(s + 1)$ down to zero. Effectively, this ignores the value of any ad that comes after an ad with such small c_a , in the same slate j . It also changes the probability of leaving slate j to 0. We argue that this decreases the value of the solution by at most $\delta \cdot \text{OPT}$. By Lemma 2, there is a solution of value at least $(1 - \frac{\delta}{s+1}) \cdot \text{OPT}$ such that no ad follows any such low-probability ad in the same slate. For this altered solution, changing the continuation probability of $c_{a_{j,i}}$ to 0 does not affect the value of slate j .

Now consider the impact of the modified probabilities \overline{C}_j of leaving slates $j = 1, \dots, s$. That is, for some slates, we replace $\overline{C}_j \leq \frac{\delta}{s+1}$ by $\overline{C}'_j = 0$. For each slate j , the probability of reaching it decreases by at most $\frac{\delta}{s+1}$. Let V_j be the expected value of slate j , conditioned on reaching it in the scanning process. By leaving all other slates empty, we could make sure to reach slate j with probability 1, and since the optimum solution must be at least as good, we obtain that $V_j \leq \text{OPT}$. Summing up over all slates, the total expected value decreases by at most $\sum_j \frac{\delta}{s+1} V_j \leq \frac{s\delta}{s+1} \cdot \text{OPT}$. Thus, in total, the value decreased by at most $\delta \cdot \text{OPT}$.

2. **Rounding continuation probabilities.** Next, we round down each non-zero continuation probability to the nearest power of $(1 - \delta/k)$. For any slot (j, i) , the resulting probability of reaching (j, i) is not changed by more than a factor of $(1 - \delta/k)^k \geq 1 - \delta$. Therefore, this stage decreases the value of the optimal solution by at most $\delta \cdot \text{OPT}$. Then, the *product* of the continuation probabilities of any subset of at most k ads is one of the $O(k^2)$ values

$$\{0, 1, (1 - \delta/k)^1, \dots, (1 - \delta/k)^{k \lceil \log_{(1-\delta/k)}(\delta/(s+1)) \rceil}\}.$$

Denote this set of values by \mathcal{C} .

3. **Enumerating over all slate probabilities.** The algorithm exhaustively enumerates all combinations of probabilities $\overline{C}_s \in \mathcal{C}$ for slates s . We call such an s -tuple $(\overline{C}_1, \dots, \overline{C}_s)$ a *configuration*. Notice that there are only $O(k^{2s})$,

i.e., polynomially many, configurations. So long as for each slate s , the product of the continuation probabilities is exactly \overline{C}_s , this enumeration makes it possible to evaluate precisely the click-through rate of any ad in any position. In particular, we use γ_s to denote the overall probability of entering slate s . Notice that all γ_s can be computed efficiently from the distribution f_ψ over permutations and the \overline{C}_s values.

4. **Dynamic programming solution.** For each configuration $(\overline{C}_1, \dots, \overline{C}_s)$, we find the optimal solution consistent with that configurations using dynamic programming. By Lemma [11](#) the advertisers in each slate must be ordered in decreasing order of their ratios $\frac{b_a q_a}{1 - c_a}$. The dynamic programming idea is essentially an s -dimensional KNAPSACK program. First sort the ads such that $\frac{b_1 q_1}{1 - c_1} \geq \frac{b_2 q_2}{1 - c_2} \geq \dots \geq \frac{b_n q_n}{1 - c_n}$. The dynamic programming table has entries $A[a, i_1, \dots, i_s, y_1, \dots, y_s]$ for all $1 \leq a \leq n, 0 \leq i_j \leq k_j, y_j \in \mathcal{C}$, and $\overline{C}_j \leq y_j \leq 1$ for every slate j . This entry contains the optimal total value that can be obtained from the last i_j slots of each slate $j = 1, \dots, s$, where all ads are from the set a, \dots, n , assuming that the product of the continuation probabilities of the first $k_j - i_j$ slots of slate j is y_j .

This entry of the dynamic programming table can be computed by considering all options for ad a . By Lemma [11](#) ad a is either not used at all, or is placed at the first “empty” slot of one of the slates. If ad a is not used at all, then the optimum value is just $A[a + 1, i_1, \dots, i_s, y_1, \dots, y_s]$. Otherwise, if ad a is used in slate j (which is possible only if $i_j > 0$), the optimum value is $A[a + 1, i_1, \dots, i_j - 1, \dots, i_s, y_1, \dots, y_j \cdot c_a, \dots, y_s] + \gamma_j y_j b_a q_a$. The optimum is then simply the maximum over these $s + 1$ different options, and can be computed in constant time $O(s)$. Overall, the dynamic program then takes time $O(nk^{3s})$.

The rounding stages in the above algorithm lose at most $2\delta \cdot \text{OPT}$ of the value of the optimal solution. The last two stages use $O(k^{2s})$ invocations of a dynamic program, each taking time $O(nk^{3s})$, to compute the optimal solution for the rounded instance. Hence, by taking $\delta = \epsilon/2$, we obtain

Theorem 2. *For every constant $\epsilon > 0$, there is a polynomial time algorithm for the winner determination problem in the Slated Cascade Model which always outputs a solution of value at least $(1 - \epsilon) \cdot \text{OPT}$.*

5 Position-Dependent Multipliers

The main difficulty in designing an algorithm for the winner determination problem in the CMPDM is that the equivalent of Lemma [11](#) no longer holds. Intuitively, high continuation probabilities are much less important if subsequent slots have very low slot-specific λ_i values. Thus, there can be no simple sorting criterion based solely on properties of the ads themselves. In this section, we present a simple 4-approximation algorithm based on a reduction to the KNAPSACK problem, as well as the sketch of a quasi-polynomial-time approximation scheme.

A 4-approximation algorithm. First, by applying Lemma 2 with $\delta = \frac{1}{2}$ (and just one slate), we can restrict our attention to solutions where $C_i \geq \frac{1}{2}$ for all i , so long as we are willing to lose a factor 2 in the approximation guarantee. For such solutions, the objective function (4) can be 2-approximated by $\sum_i \lambda_i b_{a_i} q_{a_i}$. Therefore, the solution to the winner determination problem can be 4-approximated by the optimal solution to the following KNAPSACK-like problem: select $\ell \leq k$ distinct ads a_1, \dots, a_ℓ to maximize $\sum_{i=1}^\ell \lambda_i b_{a_i} q_{a_i}$ subject to $\prod_{i=1}^{\ell-1} c_{a_i} \geq \frac{1}{2}$.

To solve this problem, we first exhaustively try all ads \hat{a} which will occupy the last assigned slot ℓ . The ad \hat{a} is the only one with no restrictions on its continuation probability $c_{\hat{a}}$. Since λ_i 's are non-increasing, the remaining ads must satisfy $b_{a_1} q_{a_1} \geq b_{a_2} q_{a_2} \geq \dots \geq b_{a_{\ell-1}} q_{a_{\ell-1}}$. The optimization problem is thus to select $a_1, \dots, a_{\ell-1}$ to maximize $\sum_{i=1}^{\ell-1} \lambda_i b_{a_i} q_{a_i} + \lambda_\ell b_{\hat{a}} q_{\hat{a}}$, subject to the constraints $\sum_{i=1}^{\ell-1} \log_2 \frac{1}{c_{a_i}} \leq 1$ and $\ell \leq k$. This problem can be solved using a simple generalization of the classical fully polynomial-time approximation scheme for KNAPSACK, thus giving a $(4 + \epsilon)$ approximation for the ad allocation problem, for any $\epsilon > 0$.

A quasi-PTAS. To obtain a quasi-PTAS for the winner determination problem, we can build on the ideas from Section 4. The key idea is to round the λ_i values to $O(\log k)$ powers $(1 - \delta)^i$. This results in $O(\log k)$ segments of slots with now identical position-dependent multipliers. Similar to the PTAS from Section 4, we also round continuation probabilities, and exhaustively search over all probabilities of entering each segment. Given these probabilities, the optimum solution can then be found with a Dynamic Program akin to multi-dimensional Knapsack. Due to space constraints, the details and analysis of the algorithm are deferred to the full version of this paper.

6 Discussion and Future Research Directions

In this paper, we studied a model for externalities in sponsored search based on a probabilistic model of user behavior. One shortcoming of our model is that it does not address situations where the search term is ambiguous (e.g., “Apple” which can be matched to ads on Apple computers and on the fruit). In such cases, the set of candidate ads consists of multiple classes with significantly different “inter-class” and “intra-class” externalities. However, such situations are rare and likely to become more so with the advance of new search technologies aimed at modeling intent.

As for future research, one of the important problems we are currently investigating is to develop CTR-learning algorithms that can learn the parameters of our model, and pursue an exploration-exploitation strategy that converges to the optimal solution over time. It would also be desirable to determine the complexity of computing an *exact* solution in the Slated Cascade Model or in the Cascade Model with Position-Dependent Multipliers. While we presented approximation schemes for these models, we do not currently know whether they are NP-hard to solve optimally.

Acknowledgments. We would like to thank David Pennock for suggesting the generalization to the CMPDM model, the authors of [11] and [7] for sharing preprints of their papers, and several anonymous reviewers for useful feedback.

References

1. Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M.: Sponsored search auctions with markovian users. In: Proc. 4th Workshop on Ad Auctions (2008)
2. Aggarwal, G., Goel, A., Motwani, R.: Truthful auctions for pricing search keywords. In: Proc. 8th ACM Conf. on Electronic Commerce, pp. 1–7 (2006)
3. Athey, S., Ellison, G.: Position auctions with consumer search. Working Paper (September 2007)
4. Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A., Mathieu, C., Schwarz, M.: Greedy bidding strategies for keyword auctions. In: Proc. 9th ACM Conf. on Electronic Commerce, pp. 262–271 (2007)
5. Clarke, E.: Multipart pricing of public goods. *Public Choice* 11, 17–33 (1971)
6. Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: Proc. 1st Intl. Conf. on Web Search and Data Mining (2008)
7. Das, A., Giotis, I., Karlin, A., Mathieu, C.: On the effects of competing advertisements in keyword auctions (May 2008) (unpublished manuscript)
8. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97(1), 242–259 (2007)
9. Ghosh, A., Mahdian, M.: Externalities in online advertising. In: 17th Intl. World Wide Web Conference (2008)
10. Groves, T.: Incentives in teams. *Econometrica* 41, 617–631 (1973)
11. Joachims, T., Granka, L., Pan, B., Hembrooke, H., Radlinski, F., Gay, G.: Evaluating the accuracy of implicit feedback from clicks and query reformulations in web search. *ACM Transactions on Information Systems (TOIS)* 25(2) (2007)
12. Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: Proc. 8th ACM Conf. on Electronic Commerce, pp. 218–227 (2006)
13. Lahaie, S., Pennock, D., Saberi, A., Vohra, R.: Sponsored search auctions. In: Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V. (eds.) *Algorithmic Game Theory*, ch. 28. Cambridge University Press, Cambridge (2007)
14. Mahdian, M., Nazerzadeh, H., Saberi, A.: Allocating online advertisement space with unreliable estimates. In: Proc. 9th ACM Conf. on Electronic Commerce, pp. 288–294 (2007)
15. Muthukrishnan, S., Pál, M., Svitkina, Z.: Stochastic models for budget optimization in search-based advertising. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 131–142. Springer, Heidelberg (2007)
16. Varian, H.: Position auctions. *Intl. J. of Industrial Organization* 25(6), 1163–1178 (2007)
17. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance* 16, 8–37 (1961)
18. Wortman, J., Vorobeychik, Y., Li, L., Langford, J.: Maintaining equilibria during exploration in sponsored search auctions. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 119–130. Springer, Heidelberg (2007)

Sponsored Search Auctions with Reserve Prices: Going Beyond Separability*

Rica Gonen** and Sergei Vassilvitskii

Yahoo! Research Labs
gonenr@yahoo-inc.com, sergei@yahoo-inc.com

Abstract. The original analysis of sponsored search auctions by Varian and independently by Aggarwal et al. did not take into account the notion of reserve prices, which are common across all major search engines. We investigate this further and show that the separability assumption derived by Aggarwal et al. is not sufficient for aligning the greedy allocation employed by GSP and the efficient allocation in the presence of reserve prices. We extend separability and derive the condition under which the greedy ranking allocation is an efficient truthful mechanism. We call this generalization the *extended separability condition*.

To complement the analysis of the extended separability condition we present an extension of the laddered auction in the presence of reserve prices, which we call the *bi-laddered auction*. We show that the bi-laddered auction is the unique truthful auction for advertisers that provides a price vector support for an *extended GSP SNE* scheme. Nevertheless the bi-laddered auction is shown to allow a budget deficit.

Building on our model of reserve prices we continue by depicting advertising networks as double sided sponsored search markets with advertisers on one side, syndicators on the other, and the search engine as the market maker. For the latter model we provide a truthful scheme for the seller and show that by assuming separability one can design a SNE, individually rational, and nearly efficient syndicated market that allows the market maker (search engine) to run the market with a surplus/budget balance. The uniqueness of our bi-laddered auction scheme implies that without the separability condition no truthful syndicated market can run without a deficit.

1 Introduction

Sponsored search auctions are the primary way that companies like Google and Yahoo! monetize their search engines. They allow advertisers to bid on particular queries, thereby ensuring the relevance of the advertisement to the user, and increasing the conversion rate. Sponsored search is a very large business,

* A preliminary version of this work was presented in the 4th workshop on Ad Auctions at EC'08.

** Confidentially submitted for publication in WINE on Dec 17- 20, 2008

projected to grow to many billions of dollars in the next few years; it is not surprising that the analysis of the precise way the auctions are run has generated research interest in the past few years, (i.e., [8,4,19,9,3,11,12,7]).

The auctions have a very simple framework¹. Each advertiser specifies the query she is looking to advertise on, and submits a bid, representing the maximum amount she is willing to pay. When a user enters a query, the system collects all of the advertisers bidding for the query, and runs a generalized second price auction to determine both the winners, and the prices that each would be charged. There are usually multiple winners, as there are multiple advertiser slots on the search result page, with higher slots being more valuable since they are seen by more users. Finally, the advertiser is charged *only* in the event of a user click on the ad, otherwise no money changes hands, the so-called *pay per click* scheme.

Separability. In one of the first analyses of these auctions, Aggarwal et al. [1] showed that, notwithstanding the claims made by Google, the auctions were not truthful. The authors showed that the greedy ranking employed by Google agrees with the efficient allocation *only* when the clickthrough rates are separable, that is, they are the product of the function of the advertiser quality and the position in which the advertisement appeared. The separability property has since been used as a simplifying assumption in other work i.e., [8,11,12,7]².

The separability assumption is also implicitly present in the work of Varian [20]. Varian assumes that each slot s has a click-through rate x_s but advertisers have a quality score of 1 for their ad, meaning that an advertiser a 's click-through at slot s equals $1 \cdot x_s$. Varian's work presents a mechanism in an equilibrium state; following the revelation principle it is not surprising that his Symmetric Nash Equilibrium (SNE) supporting prices are essentially the Aggarwal et al.'s truthful prices.

GSP Enhancements. While the basic principles of the Generalized Second Price (GSP) auction are now well understood, the auctions that are run in practice have evolved beyond this bare-bones model. Some of the most pertinent extensions include advertiser budgets, exploration in learning advertisers click through rates, broad (as opposed to exact) matches of keywords and reserve prices. Not surprisingly, each engineering enhancement has unintended economic consequences, and may potentially wreak havoc on the equilibrium achieved by the players.

While these extensions are now widely acknowledged (see for instance the footnote in [7]) and used in practice, for the most part their precise effect on the equilibria has not been analyzed. In this work we tackle the notion of *slot* specific reserve prices, and detail the changes that this condition brings to the auctions and the equilibria. Independently of our work, Even-Dar et al. [7] have

¹ In practice the situation is much more complex - e.g. advertisers specify maximum daily budgets, there is fuzzy matching on the queries, etc. We do not consider these problems in this work.

² [8] used separability earlier than [1].

recently investigated the effect of *bidder* specific reserve prices. Not surprisingly, they show that many naïve modifications lead to non truthful behaviors.

1.1 Our Contribution

While the results of Aggarwal et al. and Varian provide an initial analysis of the sponsored search auctions, they fail to take into account reserve prices that search engines usually set for many of the queries. In this work we explore the effect of reserve prices, and show that the separability assumption derived by [1] is not sufficient for aligning the greedy allocation produced by the ranking function and the efficient allocation. Thus, if we are to follow the allocation produced by the ranking function, Clarke-Groves prices would not result in a truthful mechanism. Instead we present the *extended separability condition* which provides the necessary and sufficient condition for a ranking function to be truthful under the VCG prices and modify the laddered auction of [1] (which we call the *bi-laddered auction*) to derive a mechanism that is truthful in the presence of reserve prices. The bi-laddered scheme is shown to be a supporting price vector for a SNE created by the new *extended generalized second price* scheme that is used to replace the GSP in the presence of reserve prices.

We then turn our attention to the sellers. In the sponsored search auction where the ads are presented alongside the search results, it is the search engine that controls the placement and the reserve prices for each of the slots. Currently the sponsored search auction is evolving into an advertising network motivated by Google's acquisition of DoubleClick and Yahoo! buying RightMedia. In the advertising network there are advertisers and publishers and the search engines' role is that of market makers. In such a network it is easy to imagine a syndicate situation, where a publisher (e.g. LinkedIn) offers to place advertisements along its content pages, but has a set of reserve prices for these slots – as it encounters cost, such as user satisfaction decrease, in placing the ads. In this case, we describe a pricing scheme that is truthful for the *seller*. Motivated to design an advertising network in SNE we show that the greedy allocation without the separability assumption and with cost on the slots has a unique truthful pricing scheme that does not allow for budget balance or surplus for the marker maker (the search engine). Thus such a design would be unreasonable in practice.

We then tackle the problem of devising a market that would be budget balanced (or carry a budget surplus), while at the same time eliciting truthful behavior from both the buyers and the sellers simultaneously. It is easy to see that in this scenario we must relax one of the conditions to avoid the impossibility result of Myerson-Satterthwaite [18]. We present one such mechanism, that sacrifices some of the efficiency and prove its properties of maintaining SNE and budget balance/surplus under the separability condition.

Independently of our work Even-Dar et al. [7] analyzed the notion of bidder specific reserve prices. Our work differs in two major aspects. First, [7] explicitly assumes that the separability condition holds. While, this assumption has been made before, it has recently been called into question. See for example the results of [6] and the model introduced in [2,15]. In our work we showcase the *Extended*

Separability Condition and prove that the condition is necessary in order to align the greedy allocation with the efficient allocation. Second, the major contribution of [7] is the proof that GSP has an envy-free equilibrium with bidder specific reserve prices. We present an extended GSP pricing scheme for the slot specific reserve prices model which is shown to maintain SNE. Furthermore, while we provide a truthful auction in the presence of reserve prices, a major point of our work is the exploration of the strategies of the sellers (publishers in case of syndication), and the wider question of a SNE in double sided markets that carry a budget surplus.

Finally, let us define the notation that we will use for the rest of the exposition. In general there are n advertisers $= \{i_1, \dots, i_n\}$, who bid for k slots $\{j_1, \dots, j_k\}$. Denote by b_i the bid of advertiser (buyer i), by c_j the reserve price (cost) of slot j , and finally, by $\lambda_{i,j}$ the clickthrough rate of advertiser i when she is placed in slot j .

2 The Extended Separability Condition

In this section we present and analyze the extended separability condition that is required to guarantee the existence of a truthful efficient mechanism for the greedy allocation ranking auction of sponsored search with reserve prices.

Aggarwal et al. [1] showed that for sponsored search auctions with no reserve prices on slots, for all possible ranking functions $R = (w_1, \dots, w_n)$ a truthful solution exists only if for every two buyers i, i' and slots $j, j + 1$ the condition $\frac{\lambda_{i,j}}{\lambda_{i,j+1}} = \frac{\lambda_{i',j}}{\lambda_{i',j+1}}$ holds. Below we show that not only must the clickthrough rates be separable, but in any case where the same two buyers can be matched to two different slots, the costs of those slots must be equal. In other words, the only time when the allocation provided by the ranking function agrees with the socially efficient allocation (maximizing the total gain from trade), is when either such an allocation is straightforward, or many of the reserve costs are identical. We now state this formally:

Theorem 1. *For the sponsored search auction with reserve prices a truthful efficient mechanism for the greedy allocation ranking auction exists only if the following condition holds on the click through rates of the buyers and the cost of the slots. For every rank function $R = (w_1, \dots, w_n)$ and any two buyers i and i' ranked to slots j and j' in R there exists a set of VCG [27, 5, 13] weights that always yield the same ranking as R only if:*

$$\frac{\lambda_{i,j}}{\lambda_{i,j+1}} = \frac{\lambda_{i',j}}{\lambda_{i',j+1}}$$

and either $c_j = c_{j+1}$ or $\lambda_{i,j+1} = \lambda_{i',j+1}$.

Proof. The proof follows the basic structure of the proof of the separability condition in [1]. See [14] for full details.

It is important to note that unlike the Aggarwal et al. model where every advertiser can potentially be allocated in every possible slot, our model might limit the

allowed ranking functions as an advertiser i cannot be allocated to slot j where $b_i < c_j$. In other words if for some ranking function R i and i' are allocated to j and j' respectively it is not necessarily true that there exist a ranking function R' in which i and i' are allocated to j' and j respectively; and it is precisely in these situations that a seller can charge different reserve prices for different slots, and still have the allocation specified by R to be efficient.

3 Auctions with Reserve Prices

Aggarwal et al. [1] presented an auction that preserves and realizes an allocation consistent with a ranking function $R = (w_1, \dots, w_n)$ with a laddered pricing scheme; and proved that the laddered price auction is truthful. Following a similar assumption of separability Varian proved that there exists a SNE in an auction that maintains a ranking function $R = (w_1, \dots, w_n)$ where all $w_i = 1$. The laddered pricing scheme in Aggarwal et al. is the lower bound supporting price vector for the Varian SNE. Both authors' results do not take into account reserve prices (although they often occur in practice) and therefore do not consider equilibria in the presence of reserve prices. In the previous section we focused our attention on the condition required to maintain a truthful mechanism that is aligned with the efficient allocation in a ranking based auction with costs assigned to the slots. In this section we provide a pricing scheme to support a ranking based allocation with costs assigned to slots and prove that the pricing scheme is truthful for the advertisers. Our pricing scheme also provides a support price vector for a SNE in this extended model.

3.1 The Modified Laddered Auction

Recall, we are given n advertisers (buyers) with bid vector $b = (b_1, \dots, b_n)$, a ranking function $R = (w_1, \dots, w_n)$, and k slots (sellers) with associated costs $c = (c_1 \geq c_2 \geq \dots \geq c_k)$. To assign buyers to sellers we first rank the buyers by the product $b_i w_i$. For the sake of exposition, reindex the buyers so that $b_1 w_1 \geq b_2 w_2 \geq \dots \geq b_n w_n$. If the first buyer can afford the top slot, assign her to that slot, and repeat. Otherwise, leave slot the top slot unassigned, and recurse on the remaining slots. Observe that this allocation rule maintains the following two invariants:

1. If buyer i is assigned to a slot j and i' is assigned to slot j' with $j < j'$ then $b_i w_i \geq b_{i'} w_{i'}$.
2. If buyer i is assigned to slot j then $b_i \geq c_j$.

It remains to describe the prices charged. We proceed similar to the laddered auction [1] but add the effect of the reserve prices.

$$p_i = \sum_{j=i}^k \left(\frac{\lambda_{i,j} - \lambda_{i,j+1}}{\lambda_{i,i}} \right) \max \left(\frac{w_{j+1} b_{j+1}}{w_i}, c_j \right)$$

Theorem 2. *The auction presented above is truthful.*

Proof. Consider a buyer i , similar to the proof of [1], let x be the position of i in the allocation above, and r be the closest preferred position of i , holding everybody else’s bids constant. We show that there exists a rank closer to than r to x , establishing a contradiction. Let p_y be the price charged to i if she ends up in in position y . If $r > x$, i.e. the merchant prefers to be lower, then the total change in her utility by moving to rank $r - 1$ is:

$$\begin{aligned} \lambda_{i,r-1}(v_i - p_{r-1}) - \lambda_{i,r}(v_i - p_r) &= v_i(\lambda_{i,r-1} - \lambda_{i,r}) - (\lambda_{i,r-1}p_{r-1} - \lambda_{i,r}p_r) \\ &= v_i(\lambda_{i,r-1} - \lambda_{i,r}) - (\lambda_{i,r-1} - \lambda_{i,r}) \max\left(\frac{w_r b_r}{w_i}, c_j\right) \\ &= (\lambda_{i,r-1} - \lambda_{i,r})(v_i - \max\left(\frac{w_r b_r}{w_i}, c_j\right)) \\ &\geq 0, \end{aligned}$$

Where the last line follows from the two invariants we demonstrated above and the fact that click through rates decrease with position. Since the utility gain is non-negative, r cannot be the closest preferred rank. In the case that $r < x$ the proof is similar.

While the auction above is truthful, the reserve prices are not always met: it is easy to construct examples where the average per click price paid by the buyer is lower than the reserve price for that particular slot. In effect, the pricing scheme ensures that the buyer pays at least the reserve price for slot j *only* for those clicks that she is getting at j that she would not get at $j - 1$. While a limitation, as the following theorem shows this is a direct consequence of the ranking used by this widely employed mechanism. (In Section 6 we will explore budget balanced/surplus mechanisms for this problem.)

Theorem 3. *The auction defined above is the unique truthful auction that ranks buyers according to decreasing $w_i b_i$.*

Proof. The proof parallels the uniqueness proof shown in [1]. We omit it here for space reasons.

4 The Symmetric Nash Equilibrium with Reserve Prices

In the previous sections we showed that the bi-laddered pricing scheme conducts a truthful sponsored search auction with reserve prices. In this section we present a new extended Generalized Second Price auction for sponsored search auction with reserve prices that conducts a Symmetric Nash Equilibrium. As expected a price vector that is shown to support the SNE presented is the bi-laddered prices.

For simplicity of presentation and similarly to Varian’s paper [20] we will show the SNE of the sponsored search auction with reserve prices assuming that the weight w_i for all advertiser i is 1. Denote by p_j the price charged for a click at slot j . recall that we assume $\alpha \leq 1$.

Definition 1. ([20]) *A symmetric Nash equilibrium set of prices satisfies*

$$(v_i - p_j)\lambda_{i,j} \geq (v_i - p'_j)\lambda_{i,j'}$$

for all i, j and j' .

Let $b_{m(j+1)}$ be the bid of the advertiser in slot $j + 1$, where $m(j + 1)$ is the index of the advertiser placed in slot $j + 1$.

The GSP pricing scheme in [20] is defined to be $p_j = b_{m(j+1)}$. Consider the following extended GSP scheme for sponsored search auction with reserve prices as $p_j = \max(b_{m(j+1)}, c_j)$. Though our greedy ranking scheme may not maintain efficiency (some slots may be left unallocated) our allocation still maintains the same key properties as were shown in [20] that allow for a SNE to exist with the presence of reserve prices. The allocation is individual rational (it has non-negative surplus), monotone in values and prices, the SNE is included in the NE ($SNE \subset NE$) and a local SNE implies global SNE. These facts allow us to provide an explicit characterization of equilibrium prices and bids. For the proof below we make one more technical assumption, namely that in the greedy ranking among all advertisers that can afford a particular slot, the one with higher click-through rate will be ranked higher. Formally, we assume that $\frac{\lambda_{m(j+1),j}}{\lambda_{m(j),j}} = \alpha \leq 1$.

Since advertiser i in slot j (indexed $m(j)$) does not want to move down one slot it follows that

$$(v_i - p_j)\lambda_{m(j),j} \geq (v_i - p_{j+1})\lambda_{m(j),j+1}$$

or, equivalently

$$v_{m(j)}(\lambda_{m(j),j} - \lambda_{m(j),j+1}) + p_{j+1}\lambda_{m(j),j+1} \geq p_j\lambda_{m(j),j}.$$

Similarly since advertiser z in slot $j + 1$ does not want to move up one slot it follows that

$$(v_z - p_{j+1})\lambda_{m(j+1),j+1} \geq (v_z - p_j)\lambda_{m(j+1),j}$$

or that

$$p_j\lambda_{m(j+1),j} \geq v_{m(j+1)}(\lambda_{m(j+1),j} - \lambda_{m(j+1),j+1}) + p_{j+1}\lambda_{m(j+1),j+1}.$$

As we have assumed that $\frac{\lambda_{m(j+1),j}}{\lambda_{m(j),j}} = \alpha \leq 1$, it follows that when combining the above two formulas we get:

$$\begin{aligned} v_{m(j)}(\lambda_{m(j),j} - \lambda_{m(j),j+1}) + p_{j+1}\lambda_{m(j),j+1} &\geq p_j\lambda_{m(j),j} \geq \\ &\geq p_j\lambda_{m(j+1),j} \geq v_{m(j+1)}(\lambda_{m(j+1),j} - \lambda_{m(j+1),j+1}) + p_{j+1}\lambda_{m(j+1),j+1}. \end{aligned}$$

Since $p_j = \max(b_{m(j+1)}, c_j)$ it follows that

$$\begin{aligned} v_{m(j-1)}(\lambda_{m(j-1),j-1} - \lambda_{m(j-1),j}) + \max(b_{m(j+1)}, c_j)\lambda_{m(j-1),j} &\geq \\ &\geq \max(b_{m(j)}, c_{j-1})\lambda_{m(j-1),j-1} \\ &\geq \max(b_{m(j)}, c_{j-1})\lambda_{m(j),j-1} \\ &\geq v_{m(j)}(\lambda_{m(j),j-1} - \lambda_{m(j),j}) + \max(b_{m(j+1)}, c_j)\lambda_{m(j),j} \end{aligned}$$

We can then write down the upper and lower bounds on the bids:

$$\begin{aligned} \max(b_{m(j)}^U, c_{j-1})\lambda_{m(j-1),j-1} &= \\ &v_{m(j-1)}(\lambda_{m(j-1),j-1} - \lambda_{m(j-1),j}) + \max(b_{m(j+1)}, c_j)\lambda_{m(j-1),j} \\ \max(b_{m(j)}^L, c_{j-1})\lambda_{m(j-1),j-1} &= \\ &v_{m(j)}(\lambda_{m(j),j-1} - \lambda_{m(j),j}) + \max(b_{m(j+1)}, c_j)\lambda_{m(j),j} \end{aligned}$$

The solution to the recursions is:

$$\begin{aligned} b_{m(j)}^L \lambda_{m(j-1),j-1} &= \sum_{t \geq j} v_{m(t)} (\lambda_{m(t),t-1} - \lambda_{m(t),t}) \\ b_{m(j)}^U \lambda_{m(j-1),j-1} &= \sum_{t \geq j} v_{m(t-1)} (\lambda_{m(t-1),t-1} - \lambda_{m(t-1),t}) \end{aligned}$$

Note that since $c_{j-1} \leq v_{m(j-1)}$ even in the slots where $\max(b_{m(j)}, c_{j-1}) = c_{j-1}$ our bi-laddered scheme is bounded by $b_{m(j)}^U$ from above. Since $\max(b_{m(j)}, c_{j-1}) \geq b_{m(j)}$ our bi-laddered scheme is bounded by $b_{m(j)}^L$ from below.

5 A Truthful Scheme for the Seller

In the previous section we discussed the truthful scheme for the advertisers, i.e., the buyers. In this section we will present a pricing scheme for the seller of k slots and prove that the scheme is truthful for the seller.

One can imagine that the advertising space might be managed by a third party such as a syndicator (i.e., LinkedIn) and therefore we would like to design a pricing scheme that motivates the slot seller to report his true costs for the advertising slots.

First let us consider the sponsored search setting in which the syndicator, i.e., the seller is interested in selling k slots of advertising. Each slot of advertising it has an associated cost, stemming from the impact of advertising on users, opportunity cost of utilizing the space for other results (or content), etc. We will assume that the higher the advertising slot, the higher the cost inflicted on the publisher.

The seller has to determine a set of reserve prices, reporting a cost c_j for every slot j to the market maker (search engine in this case). By our assumption, the costs are decreasing in j , i.e. $c_j > c_{j+1}$.

The pricing scheme for the seller is the double-sided auction extension of the laddered scheme provided in Aggarwal et al. [1] for the buyer. Unlike the buyers' side, the seller's side of our double-sided auction has only a single seller. The above fact simplifies significantly the formula for the seller's pricing scheme. Let A be the set of all the allocated slots' indexes in the optimal efficient allocation, and denote by $m(j)$ the buyer assigned to slot j . Then the pricing scheme is as follows: For all $j \in A$, set the price per click paid to the seller for allocated slot j , to $b_{m(j)}\lambda_{m(j),j}$. So overall,

$$p_{seller} = \sum_{j \in A} \frac{b_{m(j)}\lambda_{m(j),j}}{\lambda_{j,j}} \tag{1}$$

The intuition behind the pricing scheme is simple as there is only a single seller, in every slot there will be no trade without him.

Theorem 4. *The pricing scheme in equation (7) is truthful with respect to the seller.*

Proof. We omit the proof here for space reasons. The full proof can be found in [14].

6 The Syndicated Sponsored Search

In the previous sections we showed two pricing schemes one for the seller and one for the buyers. One scheme allows the buyers to be truthful and the other allows the seller to be truthful. As the sponsored search market evolves into networks of advertisers and syndicators a natural question arises: is it possible to conduct a market where all parties, i.e., the buyers and the seller are motivated to tell the truth simultaneously. We can consider the results from the previous sections. We already saw that the bi-laddered auction, while truthful for the buyers, leads us to charge prices that may be below the reserve prices for particular slots (and this is inevitable given the ranking function). A similar problem plagues us in the seller's case - since the allocation is efficient, following the Myerson-Satterthwaite result [18], the market maker will potentially sustain a budget deficit. As the market maker in the desired syndicated market is the search engine it is unreasonable to expect the market maker to carry a loss. One way to overcome the impossibility of Myerson-Satterthwaite is to give up some of the efficiency and maintain the other properties of individual rationality (no player losses by participating in the market), truthfulness and budget balance/surplus.

Of course other properties, e.g. truthfulness, can be sacrificed in order to avoid the budget deficit. Nevertheless as we investigate the design of a market in equilibrium (such as the SNE) if the requirement for a truthful market is relaxed then following the revelation principle there will not exist a SNE in the designed market.

The question that this section tries to answer is under what condition is it possible to create a syndicated sponsored search market that operates without a loss (i.e. budget balanced/surplus), while maintaining the desired properties of SNE, individual rationality and minimal loss of efficiency.

Interestingly the question of maintaining budget balance in a truthful syndicated sponsored search market ties back to the separability condition and the extended separability condition. While under separability it is possible to create a truthful budget balanced/surplus syndicated sponsored search market, and a SNE budget balanced/surplus syndicated sponsored search market, we could not show a similar result for the case where separability does not hold. Moreover the uniqueness of the bi-laddered truthful pricing scheme and the fact that it does not maintain budget balance indicates that no truthful budget balanced/surplus syndicated sponsored search market exists for the general case.

6.1 Separable Budget Balanced Syndicated Market

Consider a syndicator sponsored search double sided auction with buyers (advertisers) and seller of multiple slots (syndicator) and denote that auction by S . We assume that S is separable: for every buyer i and slot j , $\lambda_{i,j} = x_i \cdot y_j$ where x_i is advertiser i 's click through rate and y_j is slot j 's click through rate. Now let D be a double sided auction where every buyer with a valuation v_i in S is represented by a buyer with valuation $v_i \cdot x_i$ in D ; and every slot with cost c_j in S is a slot with cost $c_j \cdot y_j$ in D . The efficient allocation that maximizes gain from trade in D is the following standard double sided auction allocation: order the buyers according to decreasing valuations and the slots according to increasing costs and consider for allocation all the buyer-slot pairs that have positive gain from trade. The difference between the efficient allocation of S and the efficient allocation of D is that after determining which buyers and slots made it into the allocation (D 's or S 's), S orders the slots in the allocation in descending order and matches to the buyers that made it into the allocation also ordered in descending order.

Since D is a standard double sided auction it is well known (e.g. [17,10]) that there exist a budget balanced solution which is truthful, individually rational, and gives up at most one trade of the efficient allocation. Thus it is reasonable to expect that there should exist a mechanism that gives up at most one buyer and one slot, such that S is budget balanced, has a SNE and individually rational. And indeed the following mechanism produces the desirable properties:

The BB SS Syndicator double-sided auction mechanism:

1. Let A be the efficient allocation for a Syndicator double-sided auction.
2. remove from A buyer $i \in A$ such that $v_i \cdot x_i$ is minimal
3. unmatch in A slot $j \in A$ such that $c_j \cdot y_j$ is maximal
4. rematch buyers and slots that remain in A according to buyers ordered in descending order and slots ordered in descending order.
5. charge buyer i $p_i = \max\{b_{i+1}, c_{e(i)}\}$ where $c_{e(i)}$ is the cost of the slot matched with buyer i in the efficient allocation.
6. charge the seller for slot j $p_j = b_{e(j)}$ where $b_{e(j)}$ is the bid of the buyer matched with slot j in the efficient allocation.

Lemma 1. *The BB SS Syndicator double-sided auction mechanism is budget balanced, SNE, and individually rational.*

Proof. The mechanism is individually rational for every buyer as his price is always less than the buyer's bid, i.e., $\max\{b_{i+1}, c_{e(i)}\} \leq b_i$. It is also individually rational for the seller as $b_{e(j)} \geq c_j$ for every j .

The mechanism is a SNE as the pricing scheme for every allocated buyer is identical to the efficient mechanism pricing scheme in SNE (see section 4) and since all of the allocated buyers after the trade reduction did not change their relative ranking, no one has a new incentive to move in the rankings. Similarly, for every slot that is allocated, the seller is payed the truthful price from the efficient allocation. Thus he also has no incentive to deviate from the truthful strategy.

The mechanism is budget balanced as the trade reduced allocation shifted every buyer by one slot down and therefore for every buyer i that is matched with seller j in the efficient allocation, buyer i 's price in the trade reduced allocation is $p_i = \max\{b_{i+1}, c_j\}$ and the seller expects to be paid for every allocated slot j in the trade reduced allocation $p_j = b_{i+1}$. Thus $p_i \leq p_{j+1}$, ensuring budget balance.

Acknowledgments

We would like to thank David Pennock for insightful comments on an earlier draft of this work.

References

1. Aggarwal, G., Goel, A., Motwani, R.: Truthful Auctions for Pricing Search Keywords. In: Proceeding of EC 2006 (2006)
2. Aggarwal, G., Feldman, J., Muthukrishnan, S., Pal, M.: Sponsored Search Auctions with Markovian Users. In: Ad Auctions Workshop (2008); In: WINE short papers 2008 (to appear)
3. Abrams, Z., Mendelevitch, O., Tomlin, J.: Optimal Delivery of Sponsored Search Advertisements Subject to Budget Constraints. In: Proc. EC 2007 (2007)
4. Borgs, C., Chayes, J., Immorlica, N., Mahdian, M., Saberi, A.: Multi-unit auctions with budget-constrained bidders. In: Proc. 6th ACM Conference on Electronic Commerce, pp. 44–51 (2005)
5. Clarke, E.H.: Multipart Pricing of Public Goods. *Journal Public Choice* 2, 17–33 (1971)
6. Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: ACM WSDM 2008 (2008)
7. Even-Dar, E., Feldman, J., Mansour, Y., Muthukrishnan, S.: Position Auctions with Bidder Specific Minimum Prices. In: Ad Auctions Workshop (2008); In: WINE short papers 2008 (to appear)
8. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review* 97 (2007)
9. Feng, J., Bhargava, H., Pennock, D.: Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *Inform Journal on Computing*
10. Gonen, M., Gonen, R., Pavlov, E.: Generalized Trade Reductions: The Role of Competition in Designing Budget Balanced Mechanisms. In: Proc. of ACM EC 2007 (2007)
11. Gonen, R., Pavlov, E.: An Incentive-Compatible Multi Armed Bandit Mechanism. In: Third Workshop on Sponsored Search Auctions WWW 2007, PODC 2007 (2007)
12. Gonen, R., Pavlov, E.: An Adaptive Sponsored Search Mechanism δ -Gain Truthful in Valuation, Time, and Budget. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 341–346. Springer, Heidelberg (2007)
13. Groves, T.: Incentives in teams. *Journal Econometrica* 41, 617–631 (1973)

14. Gonen, R., Vassilvitskii, S.: Sponsored Search Auctions With Reserve Prices: Going Beyond Separability. Yahoo! Research Technical Report 2008-004, <http://www.ricagonen.com/>
<http://theory.stanford.edu/~sergei/index.html>
15. Kempe, D., Mahdian, M.: A Cascade Model for Externalities in Sponsored Search. In: Ad Auctions Workshop (2008); In: WINE 2008 (to appear)
16. Lahaie, S.: An Analysis of Alternative Slot Auction Designs for Sponsored Search. In: Proceedings of the 7th ACM Conference on Electronic Commerce. ACM Press, New York (2006)
17. McAfee, P.: A Dominant Strategy Double Auction. *Journal of Economic Theory* 56, 434–450
18. Myerson, R., Satterthwaite, M.: Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory* 29, 265–281
19. Mehta, A., Saberi, A., Vazirani, U., Vazirani, V.: Adwords and the Generalized Bipartite Matching Problem. In: Proceedings of the Symposium on the Foundations of Computer Science, pp. 264–273 (2005)
20. Varian, H.: Position Auctions. *International Journal of Industrial Organization* (October 2006)
21. Vickrey, W.: Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance* 16, 8–37 (1961)

Auctions for Share-Averse Bidders

Mahyar Salek and David Kempe

Department of Computer Science, University of Southern California,
CA 90089-0781, USA
{salek,dkempe}@usc.edu

Abstract. We introduce and study *share-averse auctions*, a class of auctions with allocation externalities, in which items can be allocated to arbitrarily many bidders, but the valuation of each individual bidder decreases as the items get allocated to more other bidders. For single-item auctions where players have incomplete information about each others' valuation, we characterize the truthful mechanism that maximizes the auctioneer's revenue, and analyze it for some interesting cases.

We then move beyond single-item auctions, and analyze single-minded combinatorial auctions. We derive sufficient conditions for a truthful allocation in this setting. We also obtain a \sqrt{m} -approximation algorithm for maximizing social welfare, which is essentially tight unless P=NP.

1 Introduction

Consider the problem of selling a piece of technological or financial advice. In principle, such information can be sold to all participating bidders at no marginal cost to the seller. However, in reality, the value of the information to each individual bidder decreases the more other bidders receive the information, since the winner(s) of the auction will not obtain as strong a technological or financial advantage over the losers as they would have otherwise. A similar scenario can arise for physical items: for instance, the value a shared network infrastructure, a road, or a park decreases in the number of others who have access to it.

The preceding examples motivate the study of auctions for *share-averse buyers* (*share-averse auctions* for brevity): auctions in which items can in principle be allocated to arbitrarily many bidders, but the valuation of each individual bidder decreases as the items get allocated to more other bidders. Share-averse auctions fall broadly in the scenario of *auctions with allocation externalities* [10,2,8]. They differ in that the externalities take on a simpler form: we assume that agents care only about the number of other players sharing items with them, but not about their identity, or the allocation of items which the player does not share. Furthermore, traditional models [10,2,8] for externalities use additive terms, whereas in the share-averseness model, multiplicative decreases appear more appropriate.

Our main results in this paper are twofold. First, we extend the seminal work of Myerson [14] and characterize optimal auctions for share-averse bidders if only a single item is auctioned off. We then focus on the case where all bidders have the

same share-averseness response function f , and derive a partial characterization of the optimal allocation rule in those cases. As a special case, we recover a result by Maskin and Riley [13] on revenue maximization for multi-unit auctions if bidders are unit-demand, i.e., they need at most one copy of an item.

Second, we consider the case of allocating bundles to share-averse single-minded bidders in a combinatorial auction. In our model, the value ascribed to the bundle by such a bidder depends on the *maximum* number of other bidders she shares any item with. For this problem, we characterize sufficient conditions for a truthful mechanism in the spirit of [12], and provide a (tight) \sqrt{m} approximation mechanism.

2 Related Work

Share-averseness is a negative allocation externality among the winners. Auctions with externalities [9,8,10,2] are often studied in economics both for revenue maximization and efficiency. Many of these scenarios have externalities affecting the *loser* of auctions, whereas our results are based in a reduction in utility for the *winners*. Jehiel et al. [9,8] study both informational and allocative type-independent externalities. Brocas [2] looks at the extension where externalities depend on the types of both the winner and loser of the good. Recently, Ghosh et al. [4] looked at the computational challenges of allocation with externalities and showed inapproximability results for general case. In their result, the utility depends not only on the number of bidders sharing the item, but also on the identity of the winner set. This makes the problem significantly more complex.

Our work also relates to the problem of allocating public goods or clubs subject to congestion [3]. Public goods are defined as being shared by more than one agent. *Congestion* describes the decrease in utility to the individuals as a result of the sharing. Much of the work on clubs and public goods focuses on the issues of cost sharing and incentive compatibility (see, e.g., [7]). While there has been some work on equilibria in games between different clubs trying to maximize profits (e.g., [15]), these tend to focus on the competition between multiple clubs vying for customers rather than an optimal auction for membership in one club with given size.

Approximation algorithms and truthful mechanisms for combinatorial auctions [16] have recently received a lot of attention. Much of the focus has been on the single-minded case. With m denoting the number of items, Lehman et al. [12] were the first to show that a simple greedy algorithm gives a \sqrt{m} -approximation (which is best possible unless $P=NP$). Gonen et al. [6,5] use linear programming to extend the results to the more general case of *Packing Integer Programs* (PIPs), where multiple copies of each item are available. Later, Briest et al. [1] improve their result to a truthful $m^{1/b}$ approximation algorithm, where b is the minimum of the multiplicities of all items. We will use both algorithms as a black box in deriving our approximation result for single-minded bidders.

3 Single-Item Auctions with a Prior

In this section, we focus on the special case of selling a single item to risk-neutral bidders. The set of all n bidders is denoted by $N := \{1, \dots, n\}$. Each bidder i has a private valuation v_i , if she is allocated the item exclusively. If she shares the item with k other bidders, her valuation decreases to $v_i \cdot f_i(k)$. We call f_i the *share-averseness function* of bidder i , and require that $f_i(0) = 1$ and f_i is monotonically non-increasing. We assume that share-averseness functions are common knowledge, as opposed to valuations, which are private.

Following standard convention (e.g., [14]), we assume that each bidder’s valuation is drawn independently from some distribution $g_i : [\ell_i, r_i] \rightarrow \mathbb{R}^+$ over a finite interval $[\ell_i, r_i]$, and that all bidders share this belief. We denote the cumulative distribution function (CDF) of g_i by $G_i(v) = \int_{\ell_i}^v g_i(t)dt$. We let $V = [\ell_1, r_1] \times \dots \times [\ell_n, r_n]$ denote the set of all possible combinations of bidders’ values, and $V_{-i} = \times_{j \in N, j \neq i} [\ell_j, r_j]$ the set of possible values of bidders other than bidder i . The *joint distribution* on valuation vectors $\mathbf{v} = (v_1, \dots, v_n)$ is $g(\mathbf{v}) = \prod_{i \in N} g_i(v_i)$. Likewise, we define $\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ and $g_{-i}(\mathbf{v}_{-i}) = \prod_{j \in N, j \neq i} g_j(v_j)$.

In this setting, we want to derive a truthful mechanism maximizing the auctioneer’s revenue. Such a mechanism can be described by functions $a_i : V \rightarrow [0, 1]$ and $p_i : V \rightarrow \mathbb{R}^+$ for each agent i . We denote the vector of all of these functions by \mathbf{a} and \mathbf{p} , respectively. For each vector \mathbf{v} of valuations, $a_i(\mathbf{v})$ is the fraction of the item assigned to bidder i (corresponding to the *visitation rate* in club good theory [3]), and $p_i(\mathbf{v})$ the expected payment of agent i . Given that we allow fractional assignments of items, we need to define the notion of share-averseness more precisely. We set $w_i(\mathbf{v}) = \sum_{j \in N, j \neq i} a_j(\mathbf{v})$ to be the total fractional sharing of bidder i . One way to interpret $w_i(\mathbf{v})$ is as the expected number of bidders that i is sharing with if each bidder j receives the item with probability $a_j(\mathbf{v})$.

We also need to extend the share-averse function f_i to fractional values now. We define $f'_i(t)$ as the convex combination $(\lceil t \rceil - t) \cdot f_i(\lceil t \rceil) + (1 - (\lceil t \rceil - t)) \cdot f_i(\lfloor t \rfloor)$. This definition ensures monotonicity of f'_i . Henceforth, whenever the distinction is clear from the context, we will use f_i to refer to the extension of the function to fractional values.

Remark 1. It may appear natural to explicitly consider the $a_j(\mathbf{v})$ values as *probabilities*, and assign the item to each bidder *independently* with probability $a_j(\mathbf{v})$. The disadvantage of this approach is that the expected utility of bidder i now depends not only on $w_i(\mathbf{v})$, but also on the *exact fractional assignments of each other agent*, violating our framework of share-averseness.

There is a natural interpretation of the function f'_i defined above. If the item is shared over a period of time, then standard network flow techniques can be used to efficiently find an assignment over time in which *each* bidder i shares the item with $\lceil w \rceil$ other bidders for a $1 - (\lceil w \rceil - w)$ fraction of time, and with $\lfloor w \rfloor$ other bidders for the remaining $\lceil w \rceil - w$ fraction. (Here, $w = w_i(\mathbf{v})$.) From this flow argument, we can also derive an actual distribution letting us interpret the $a_i(\mathbf{v})$ as probabilities. After finding a period of time with corresponding assignments,

simply define a distribution over allocations by drawing a uniformly random point in time, and then taking the allocation at that time.

The utility of player i under valuations \mathbf{v} is then $a_i(\mathbf{v}) \cdot v_i \cdot f_i(w_i(\mathbf{v})) - p_i(\mathbf{v})$. Therefore, the *expected* utility of player i with valuation v_i is

$$u_i(v_i) = \int_{V_{-i}} (a_i(\mathbf{v}) \cdot v_i \cdot f_i(w_i(\mathbf{v})) - p_i(\mathbf{v})) \cdot g_{-i}(\mathbf{v}_{-i}) d\mathbf{v}_{-i},$$

where $\mathbf{v} = (v_i, \mathbf{v}_{-i})$, and $d\mathbf{v}_{-i} = dv_1 \cdots dv_{i-1} dv_{i+1} \cdots dv_n$. The expected utility of the seller from this auction is

$$\hat{u} = \int_V \sum_{i \in N} p_i(\mathbf{v}) g(\mathbf{v}) d\mathbf{v}.$$

In order to ensure that the auction mechanism is feasible and truthful, the payments and allocated fractions will have to satisfy voluntary participation and incentive compatibility (truthfulness), as captured by the following two conditions for each bidder i :

$$u_i(v_i) \geq 0 \tag{1}$$

$$u_i(v_i) \geq \int_{V_{-i}} (a_i(\hat{v}, \mathbf{v}_{-i}) \cdot v_i \cdot f_i(w_i(\hat{v}, \mathbf{v}_{-i})) - p_i(\hat{v}, \mathbf{v}_{-i})) \cdot g_{-i}(\mathbf{v}_{-i}) d\mathbf{v}_{-i} \forall \hat{v}. \tag{2}$$

An auction mechanism is specified by the functions determining the (fractional) assignments and the payments of each bidder i , i.e., by the pair (\mathbf{a}, \mathbf{p}) . To simplify subsequent notation, we define

$$Q_i(v) = \int_{V_{-i}} a_i(v, \mathbf{v}_{-i}) \cdot f_i(w_i(v, \mathbf{v}_{-i})) \cdot g_{-i}(\mathbf{v}_{-i}) d\mathbf{v}_{-i}$$

to be the expected conditional fraction of the original valuation bidder i expects having valuation v . We are now ready to characterize feasible truthful and individually rational mechanisms (\mathbf{a}, \mathbf{p}) .

Lemma 1. *The following conditions are necessary and sufficient for (\mathbf{a}, \mathbf{p}) to be feasible, truthful, and individually rational.*

1. *Monotonicity:* For each bidder i , if $v \leq v'$, then $Q_i(v) \leq Q_i(v')$.
2. *Individual Rationality:* For each bidder i and valuation v , $u_i(v) \geq 0$.
3. *(Extended) Incentive Compatibility (EIC):* The expected utility function of each bidder i satisfies $u_i(v) = u_i(\ell_i) + \int_{\ell_i}^v Q_i(t) dt$.

Proof. The proof is similar to Lemma 2 from [14]. We only sketch it here due to space constraints. The utility of agent i with true valuation v_i , but reporting a different valuation \hat{v} , is

$$\begin{aligned} & \int_{V_{-i}} (a_i(\hat{v}, \mathbf{v}_{-i}) \cdot v_i \cdot f_i(w_i(\hat{v}, \mathbf{v}_{-i})) - p_i(\hat{v}, \mathbf{v}_{-i})) \cdot g_{-i}(\mathbf{v}_{-i}) d\mathbf{v}_{-i} \\ &= \int_{V_{-i}} (a_i(\hat{v}, \mathbf{v}_{-i}) \cdot (\hat{v} + (v_i - \hat{v})) \cdot f_i(w_i(\hat{v}, \mathbf{v}_{-i})) - p_i(\hat{v}, \mathbf{v}_{-i})) \cdot g_{-i}(\mathbf{v}_{-i}) d\mathbf{v}_{-i} \\ &= u_i(\hat{v}) + (v_i - \hat{v}) Q_i(\hat{v}). \end{aligned}$$

Thus, incentive compatibility for bidder i is equivalent to requiring that

$$u_i(v_i) \geq u_i(\hat{v}) + (v_i - \hat{v})Q_i(\hat{v}), \tag{3}$$

for all $v_i, \hat{v} \in [\ell_i, r_i]$. The rest of the proof is nearly identical to [14]. ■

The next theorem captures the notion that a mechanism is truthful if and only if the allocation rule for each bidder is monotone, and the prices are defined appropriately. The proof is similar to the proof of Lemma 3 by Myerson [14], and due to space constraints, we defer it to the full version of this paper.

Theorem 1. *Given the allocation functions a_1, \dots, a_n , let payment functions \hat{p}_i be defined as $\hat{p}_i(\mathbf{v}) = a_i(\mathbf{v})v_i f_i(w_i(\mathbf{v})) - \int_{\ell_i}^{v_i} a_i(t, \mathbf{v}_{-i}) f_i(w_i(t, \mathbf{v}_{-i})) dt$ for valuations $\mathbf{v} = (v_1, \dots, v_n)$. Then, a share-averse auction is truthful if and only if the allocation functions a_i satisfy the monotonicity condition $Q_i(v) \leq Q_i(v')$ when $v \leq v'$. Furthermore, the revenue-maximizing auction maximizes*

$$\int_V \sum_{i \in N} \left(v_i - \frac{1 - G_i(v_i)}{g_i(v_i)} \right) \cdot a_i(\mathbf{v}) f(w_i(\mathbf{v})) g(\mathbf{v}) d\mathbf{v}. \tag{4}$$

Remark 2. Note that we focus here only on the revenue maximization problem. The problem of maximizing *social welfare* is much simpler in the single-item case. The following mechanism can be easily seen to maximize social welfare and be truthful.

For each k , let S_k be a set of at most k elements maximizing $\sum_{i \in S} v_i f_i(k - 1)$, with an arbitrary tie breaking rule consistent over all k . Given k , one can compute S_k with simple sorting. The welfare maximizing mechanism simply picks the set S_k with the largest social welfare. The proof uses a simple exchange argument both to show that the optimum uses an integral allocation and picks S_k . To make the mechanism truthful, one can simply charge each bidder the VCG payments, and standard arguments prove incentive compatibility.

3.1 Regular Auctions

The term $v_i - \frac{1 - G_i(v_i)}{g_i(v_i)}$ in Theorem 1 is traditionally called the *virtual valuation* (or *effective bid*) of agent i . The function $c_i(x) = x - \frac{1 - G_i(x)}{g_i(x)}$ is called the *virtual valuation function*. If all virtual valuation functions c_i are strictly increasing, the auction is called *regular*. Regularity is a standard assumption in auction theory, and for the rest of this section, we will focus on regular auctions. We also make the further assumption that all bidders have the same share-averseness function, i.e., $f_i = f$ for all i .

In terms of the virtual valuations, we can state the auctioneer’s objective as maximizing $\int_V \sum_{i \in N} c_i(v_i) \cdot a_i(\mathbf{v}) f(w_i(\mathbf{v})) g(\mathbf{v}) d\mathbf{v}$. Without loss of generality, we will assume for the remainder of this section that bidders are sorted by virtual valuations, i.e., $c_1(v_1) \geq c_2(v_2) \geq \dots \geq c_n(v_n)$.

Theorem 2. *Without loss of generality, the optimal mechanism for general share-averse bidders has the following properties:*

1. The allocations are monotone non-increasing, i.e., $a_i(\mathbf{v}) \geq a_{i+1}(\mathbf{v})$ for all i .
2. If $c_i(v_i) < 0$, then $a_i(\mathbf{v}) = 0$.
3. For every index i with $a_{i+1}(\mathbf{v}) > 0$, we have $a_i(\mathbf{v}) + a_{i+1}(\mathbf{v}) \geq 1$.

Notice that the theorem implies that there can be at most one bidder i with allocation $0 < a_i(\mathbf{v}) < \frac{1}{2}$.

Proof (Sketch). An easy calculation using the monotonicity of f shows that swapping the allocations of j and $j + 1$ cannot decrease the utility of the auctioneer. Therefore, the optimal allocation is monotone non-increasing using a simple exchange argument. (Details are deferred to the full version.)

If there is a bidder with negative virtual valuation who has a (fractional) allocation, it is easy to see that the auctioneer’s revenue strictly increases by taking away that bidder’s allocation.

If there is a j such that $a_j(\mathbf{v}) + a_{j+1}(\mathbf{v}) \leq 1$, then consider the new assignment giving bidder j an allocation of $a_j(\mathbf{v}) + a_{j+1}(\mathbf{v})$, and bidder $j + 1$ an allocation of 0. We obtain that

$$\begin{aligned} \hat{u}^{\text{OPT}} &= \int_V \sum_{i \neq j, j+1} c_i(v_i) a_i(\mathbf{v}) f(w_i(\mathbf{v})) \\ &= \int_V c_j(v_j) a_j(\mathbf{v}) f(w_j(\mathbf{v})) + c_{j+1}(v_{j+1}) a_{j+1}(\mathbf{v}) f(w_{j+1}(\mathbf{v})) g(\mathbf{v}) d\mathbf{v} \\ &\leq \int_V c_j(v_j) \cdot (a_j(\mathbf{v}) + a_{j+1}(\mathbf{v})) \cdot f(w_j(\mathbf{v}) - a_{j+1}(\mathbf{v})) g(\mathbf{v}) d\mathbf{v} \\ &= \hat{u}^{\text{OPT}'} - \int_V \sum_{i \neq j, j+1} c_i(v_i) a_i(\mathbf{v}) f(w_i(\mathbf{v})), \end{aligned}$$

and can therefore repeatedly perform such alterations until the third condition is satisfied. ■

3.2 Convex Share-Averseness Functions

A very natural further restriction on f is that it is convex over its entire support. Intuitively, this corresponds to bidders losing their sensitivity to more and more other bidders sharing the item: the addition of the 100th bidder causes less marginal loss in utility than the addition of the second bidder. If f is convex, we can derive stronger conditions on the allocated fractions than Theorem 2.

Theorem 3. *Under the optimal mechanism for convex share-averse bidders, at most one bidder j will obtain a fractional allocation $0 < a_j(\mathbf{v}) < 1$.*

Proof. By Theorem 2, the allocations in OPT are sorted, and no bidder with negative virtual valuation obtains an allocation. Suppose that in OPT, there is a j such that $1 > a_j(\mathbf{v}) \geq a_{j+1}(\mathbf{v}) > 0$. By Theorem 2, we know that $a_j(\mathbf{v}) + a_{j+1}(\mathbf{v}) \geq 1$. We construct an alternate solution, where bidder j ’s new allocation is 1, and bidder $(j + 1)$ ’s is $a_j(\mathbf{v}) + a_{j+1}(\mathbf{v}) - 1$.

Define $W := \sum_{i \neq j, j+1} a_i(\mathbf{v})$. We can then see that $a_j(\mathbf{v}) \geq a_{j+1}(\mathbf{v})$ implies

$$w_j(\mathbf{v}) = W + a_{j+1}(\mathbf{v}) \leq W + a_j(\mathbf{v}) = w_{j+1}(\mathbf{v}),$$

and because f is monotone non-increasing, $f(w_j(\mathbf{v})) \geq f(w_{j+1}(\mathbf{v}))$. Therefore, the convexity of f implies that $f(w_j(\mathbf{v}) - \delta) - f(w_j(\mathbf{v})) \geq f(w_{j+1}(\mathbf{v}) - f(w_{j+1}(\mathbf{v}) + \delta)$, for any $\delta \geq 0$. In other words, since agent j is currently not sharing as much as agent $j + 1$, reducing her load by δ gives a larger increase than the decrease of agent $j + 1$ by increasing her load by δ . Setting $\delta = 1 - a_j(\mathbf{v})$, we can now use the above reasoning to derive

$$\begin{aligned} \hat{u}^{\text{OPT}} & - \int_V \sum_{i \neq j, j+1} c_i(v_i) a_i(\mathbf{v}) f(w_i(\mathbf{v})) \\ & = \int_V c_j(v_j) a_j(\mathbf{v}) f(w_j(\mathbf{v})) + c_{j+1}(v_{j+1}) a_{j+1}(\mathbf{v}) f(w_{j+1}(\mathbf{v})) g(\mathbf{v}) d\mathbf{v} \\ & \leq \int_V c_j(v_j) a_j(\mathbf{v}) f(w_j(\mathbf{v}) - \delta) + c_{j+1}(v_{j+1}) a_{j+1}(\mathbf{v}) f(w_{j+1}(\mathbf{v}) + \delta) g(\mathbf{v}) d\mathbf{v} \\ & < \int_V c_j(v_j) f(w_j(\mathbf{v}) - \delta) + c_{j+1}(v_{j+1}) \cdot (a_{j+1}(\mathbf{v}) - \delta) \cdot f(w_{j+1}(\mathbf{v}) + \delta) g(\mathbf{v}) d\mathbf{v} \\ & = \hat{u}^{\text{OPT}'} - \int_V \sum_{i \neq j, j+1} c_i(v_i) a_i(\mathbf{v}) f(w_i(\mathbf{v})). \end{aligned}$$

The first inequality used the convexity observation along with the fact that $c_j(v_j) a_j(\mathbf{v}) \geq c_{j+1}(v_{j+1}) a_{j+1}(\mathbf{v})$ by the sorting. The second inequality used monotonicity of f and the sorting $c_j(v_j) > c_{j+1}(v_{j+1})$. By repeating such reallocations, we derive an allocation with at most one fractional $a_j(\mathbf{v})$. ■

Remark 3. A proof similar to Theorem 3 derives the optimal mechanism for multi-unit auctions with regular virtual valuations. It thus recovers Proposition 4 of Maskin and Riley [13]. A multi-unit auction with k items can be modeled by bidders with the share-averseness function $f(x) = 1$ for $x \leq k - 1$, and $f(x) = 0$ for $x > k - 1$. An easy calculation using Theorem 2 shows that the optimal mechanism assigns the item fully to the first $\min(j, k)$ bidders, and not at all to the remaining ones, where $j \leq n$ is the largest index such that $c_j(v_j) \geq 0$. The payment of a winning agent i can then be easily derived to be the *threshold bid*, the lowest bid with which agent i could have been assigned the item.

4 Single-Minded Combinatorial Auctions

In this section, we extend the study of share-averse auctions to the combinatorial setting in which there is more than one item. The set of all items is $M := \{1, \dots, m\}$. Bidders are *single-minded*. That is, for each bidder i , there exists a set S_i such that $v_i^*(S) = v_i^*(S_i)$ for all $S \supseteq S_i$, and $v_i^*(S) = 0$ otherwise. These are the bidders' valuations if they do not share any items in S .

An assignment $\mathbf{B} = (B_1, B_2, \dots, B_n)$ of items $B_i \subseteq M$ to bidders need not have disjoint bundles. (However, we restrict our focus to mechanisms that assign items only integrally.) We now use $w_i(\mathbf{B}) = \max_{j \in B_i} (n_{i,j})$ to denote the *maximum* number of other bidders that i shares any of her items with, where $n_{i,j} = |\{i' \neq i \mid j \in B_{i'}\}|$ is the number of users sharing item j with bidder i . The valuation of bidder i is $v_i(\mathbf{B}) = v_i^*(B_i) \cdot f(w_i(\mathbf{B}))$. Notice that we assume in this section that all bidders have the same share-averseness function f .

Remark 4. Naturally, the *maximum* number of other bidders is not the only possible measure of sharing. One could instead consider a (weighted) average,

for example. The maximum number appears natural in settings where the items are combined in a physical way, and limited access to any single item causes a bottleneck. An investigation of other aggregations is left for future work.

A bid b_i comprises a pair (S, v) . Since both S_i and $v_i^*(S_i)$ are private information, bidders can be strategic both about the set and the valuation they declare. The vector of the bids of all bidders is denoted by \mathbf{b} . The bids of all bidders except bidder i are denoted by \mathbf{b}_{-i} .

4.1 A Sufficient Condition for a Truthful Mechanism

Lehmann et al. [12] proved that an allocation rule for single-minded combinatorial auctions gives rise to a truthful mechanism if the allocation rule is monotone and exact, in the sense that each bidder i is either allocated her desired set S_i or the empty set, and increasing one’s bid can never result in moving from receiving S_i to receiving the empty set. We extend these conditions as follows:

1. *Exactness*: For each bidder i , either $B_i(\mathbf{b}) = S_i$ or $B_i(\mathbf{b}) = \emptyset$.
2. *Allocation Monotonicity*: If $S' \subseteq S$ and $v' \geq v$, and $B_i(\mathbf{b}) \neq \emptyset$ with $b_i = (S, v)$, then $w_i(\mathbf{b}) \geq w_i((S', v'), \mathbf{b}_{-i})$. The is, by requesting a smaller set and bidding higher, a bidder can only share with fewer other bidders.

Given arbitrary (but fixed) bids \mathbf{b}_{-i} by all bidders except i , and a fixed set S , we define the j^{th} critical bid τ_j^i of bidder i to be the infimum of all v such that $B_i((S, v), \mathbf{b}_{-i}) = S$ and $w_i((S, v)) \leq j$. It then follows immediately from allocation monotonicity that $\tau_1^i \geq \tau_2^i \geq \dots \geq \tau_{m-1}^i$, and that if bidder i bids less than τ_{m-1}^i , she does not receive any items. Based on the critical values, we define the following payment structure:

$$\pi_j^i = \begin{cases} 0 & \text{if } j = m \\ f(m-1)\tau_{m-1}^i & \text{if } j = m-1 \\ (f(j) - f(j+1))\tau_j^i + \pi_{j+1}^i & \text{if } j < m-1. \end{cases}$$

Expanding the recursive formula gives $\pi_j^i = f(j)\tau_j^i + \sum_{k=j+1}^{m-1} f(k)(\tau_k - \tau_{k-1})$. Given an allocation scheme, we will charge bidder i the amount π_j^i for the unique index j such that $v_i \in (\tau_j^i, \tau_{j-1}^i]$. (If $v_i > \tau_{m-1}^i$, then we define $j = m$.) Note that this payment does not depend on the amount of the agent’s bid, but only on the interval which the bid falls into. In the sequel, we assume that the bidder i is fixed, and omit it from the notation where it is clear. The following proposition follows fairly directly from the definition of the payment scheme:

Proposition 1. *If bidder i ’s bid is denied, her utility is 0. If bidder i bids truthfully, her utility is non-negative.*

The main result of this section is the following theorem:

Theorem 4. *If the allocation rule satisfies Exactness and Monotonicity, then the payment scheme π_j^i yields a truthful implementation.*

Proof. Assume that bidder i desires set S with valuation v , and submits a bid $b' = (S', v')$. By Proposition [1](#), b' must lead to winning S' , and $S' \supseteq S$. By Lemma [3](#) below, bidding (S, v') gives at least the same utility. In turn, by Lemma [2](#), the utility of bidding (S, v) is at least that of bidding (S, v') . Hence, it is a dominant strategy to declare (S, v) . ■

Lemma 2. *If bidder i desires set S with valuation v , declaring (S, v) dominates declaring (S, v') for all v' .*

Proof (Sketch). Due to space constraints, the proof is deferred to the full version of the paper. The idea is to distinguish several cases. The easy cases are when either (S, v) or (S, v') are losing bids. In those cases, it is easy to show that the utility of the truthful bid dominates the other one. If both bids lead to receiving the set S , then we distinguish whether v or v' leads to more sharing. In both cases, somewhat involved calculations show that the non-truthful declaration cannot lead to higher utility. ■

Lemma 3. *If bidder i desires set S with valuation v , declaring (S, v) dominates declaring (S', v) for all S' .*

Proof. If $S' \not\supseteq S$, then bidder i can obtain valuation at most 0. Let $\nu = w_i((S, v), \mathbf{b}_{-i})$ and $\nu' = w_i((S', v), \mathbf{b}_{-i})$. Note that $u_i((S, v), \mathbf{b}_{-i}) = f(\nu)(v - \tau_\nu) + \sum_{k=\nu+1}^{m-1} f(k)(\tau_k - \tau_{k-1})$, and $u_i((S', v), \mathbf{b}_{-i}) = f(\nu')(v - \tau_{\nu'}) + \sum_{k=\nu'+1}^{m-1} f(k)(\tau_k - \tau_{k-1})$, where $\nu \leq \nu'$ by monotonicity. Define $\phi(S, v) = f(j)$ if $\tau_j < v < \tau_{j-1}$ for $j < m$ and 0 otherwise. Note that by definition $u_i((S, v), \mathbf{b}_{-i}) = \int_0^v \phi(S, v)$. Similarly, $u_i((S', v), \mathbf{b}_{-i}) = \int_0^v \phi(S', v)$.

We show that $\phi(S', x) \leq \phi(S, x)$ for all $x \in [0, v]$. This immediately implies $u_i((S', v), \mathbf{b}_{-i}) \leq u_i((S, v), \mathbf{b}_{-i})$. If (S, x) is a losing bid, then $u_i((S', v), \mathbf{b}_{-i}) = u_i((S, v), \mathbf{b}_{-i}) = 0$. Otherwise, (S, x) is a winning bid sharing with n_x other winners. By monotonicity, the bid (S', x) would not have been granted with $n'_x < n_x$ bidders, so $x \leq \tau_{n_x-1}$, and $\phi(S', v) \leq f(n_x) = \phi(S, v)$. ■

4.2 A \sqrt{m} -Approximation Algorithm

In this section, we present a mechanism approximating the social welfare of the assignment A , $V(A) = \sum_{i \in N} v_i(B_i(\mathbf{b}))$. We will achieve a \sqrt{m} approximation. Since the problem contains the SET PACKING problem as a special case (with share averseness function $f(x) = 1$ for $x = 0$ and $f(x) = 0$ for $x > 0$), it is NP-hard to approximate the problem to within a factor $m^{\frac{1}{2}-\epsilon}$ for any $\epsilon > 0$.

Consider the modified problem (which we denote by $P_{a,b}$), in which we modify the share-averseness function as follows: $f_{a,b}(x) = f(a)$ for $x \leq a$, $f_{a,b}(x) = f(x)$ for $a \leq x \leq b$, and $f_{a,b}(x) = 0$ for $x > b$. Then, in the optimum solution, w.l.o.g., each item is shared between a and b times. This modified problem $P_{a,b}$ can be reasonably well approximated, so long as a and b are chosen “close enough”.

Lemma 4. *$P_{a,b}$ can be approximated within a factor $\min(\sqrt{m}, m^{\frac{1}{b}}) \frac{f(b)}{f(a)}$.*

Proof. Consider the following set packing problem. The sets are exactly the desired sets S_i , with valuations v_i . We impose a hard constraint that each item can be shared at most b times. Thus, we obtain a packing problem with a uniform upper bound of b on the number of times each item can be included.

Briest et al. [1] show that such a packing problem can be approximated to within $\min(\sqrt{m}, m^{\frac{1}{b}})$. Let OPT be the optimum solution to $P_{a,b}$. For each allocated set S_i , the optimum, by definition, obtains valuation at most $f(a)v_i$. Therefore, the optimum solution to the packing problem has valuation at least $\frac{1}{f(a)}V(OPT)$. Thus, the approximate solution returned by the algorithm of [1] has valuation at least $\frac{1}{f(a) \cdot \min(\sqrt{m}, m^{\frac{1}{b}})}$ times that of OPT . For each allocated set S_i , our solution obtains valuation at least $f(b)v_i$, completing the proof. ■

Next, we show that so long as we are willing to incur a constant factor loss in the approximation guarantee, we can restrict our attention to solutions in which the share-averseness function does not take on too small values.

Lemma 5. *Let \bar{n} be the largest j such that $f(j) \geq \frac{1}{n}$. Then, there is a solution OPT' whose value is within a factor 2 of that of OPT , such that no item is shared more than \bar{n} times in OPT' .*

Proof. Let I be the set of all bidders sharing items with more than \bar{n} other bidders. If the total valuation obtained from bidders in I is at most $\frac{1}{2}V(OPT)$, then we allocate each bidder in I the empty bundle. This ensures the condition, while reducing the total value of the solution by at most a factor of 2.

On the other hand, if the total valuation of bidders in I is at least half the optimum valuation, then let $i^* \in I$ be the bidder with highest valuation for her set S_i . If we allocate i^* her bundle and everyone else the empty bundle, we again ensure that no one shares with more than \bar{n} other bidders. Furthermore, the choice of i^* gives us that

$$v_{i^*} \geq \frac{1}{|I|} \cdot \sum_{i \in I} v_i \geq \frac{1}{n} \cdot \sum_{i \in I} (nf(\bar{n} + 1))v_i \geq \frac{1}{2}V(OPT).$$

Here, we used that $|I| \leq n$, and $nf(\bar{n} + 1) \leq 1$ by definition of \bar{n} . ■

The idea of our approximation algorithm is to solve several problems of the form $P_{a,b}$, and keep the best of the solutions. In order not to lose too large a factor $f(a)/f(b)$, we ensure that each interval has $f(a)$ and $f(b)$ “reasonably close” together. Formally, we define a sequence d_1, \dots, d_{k+1} by $d_1 = 0$, and $d_{i+1} = \max\{j \mid f(j - 1) \geq \frac{1}{2}f(d_i)\}$. We stop when $d_i \geq \bar{n}$, and let $k + 1$ be the total length. Notice that $k = O(\log n)$ by Lemma 5. The algorithm solves each of the problems $P_{d_i, d_{i+1}-1}$ using the algorithm of Briest et al. [1], and returns the best of the solutions found. This clearly takes polynomial time.

Theorem 5. *This algorithm gives a solution within a factor $\Omega(\sqrt{m})$ of OPT .*

Proof. Let OPT be the optimum solution. For each $i = 1, \dots, k$, let O_j denote the set of bidders j who were assigned their set S_j sharing with d other bidders,

for $d_i \leq d \leq d_{i+1} - 1$. Because in each solution O_j , there is potentially less sharing than in OPT, we obtain that $\sum_{i=1}^k V(O_i) \geq V(\text{OPT})$.

Each assignment O_i is a feasible solution to $P_{d_i, d_{i+1}-1}$. Therefore, by Lemma 4, the solution A_i found by the algorithm in the i^{th} iteration satisfies

$$V(A_i) \geq \frac{f(d_i)}{\min(\sqrt{m}, m^{1/(d_{i+1}-1)}) \cdot f(d_{i+1}-1)} \cdot V(O_i) \geq \frac{1}{2(\min(\sqrt{m}, m^{1/(d_{i+1}-1)})} V(O_i),$$

by the definition of the d_i . Now, consider 2 cases:

1. If $V(O_1) + V(O_2) \geq \frac{1}{2} \cdot V(\text{OPT})$ (i.e., sharing very little can give within a constant factor of the optimum total welfare), then

$$V(A_1) + V(A_2) \geq \frac{1}{2\sqrt{m}} \cdot (V(O_1) + V(O_2)) = \Omega\left(\frac{1}{2\sqrt{m}}\right) \cdot V(\text{OPT}).$$

Therefore, at least one of A_1, A_2 gives an $\Omega\left(\frac{1}{2\sqrt{m}}\right)$ approximation.

2. If, on the other hand, $V(O_1) + V(O_2) < \frac{1}{2} V(\text{OPT})$, then $\sum_{i=3}^k V(O_i) \geq \frac{1}{2} V(\text{OPT})$. Because $d_4 \geq 4$, and thus each item can be allocated at least three times in $P_{d_i, d_{i+1}-1}$ for $i \geq 3$, we know that the algorithm of Briest et al. 11 gives an $\Omega(m^{1/3})$ approximation for each such subproblem. The best of the solutions A_i for $i \geq 3$ is at least as good as the average, i.e., at least

$$\frac{1}{2(k-2)m^{1/3}} V(\text{OPT}) \geq \frac{1}{2m^{1/3} \log(n)} V(\text{OPT}) \geq \frac{1}{2\sqrt{m}} V(\text{OPT}),$$

so long as m and n are polynomially related. ■

5 Further Directions

In the context of single-item share-averse auctions, a promising direction for future work is to characterize the optimum mechanism more specifically when bidders have different share-averseness functions. Perhaps, stronger assumptions on the distributions could help here. It would also be interesting to draw further connections to the literature on club goods, and consider the effects of multiple competing auctioneers.

In the context of share-averse combinatorial auctions, many directions remain open. It would be desirable to obtain approximation guarantees (nearly) matching those for regular combinatorial auctions, e.g., with submodular valuations. For the single-minded case, our algorithm gives an essentially best-possible approximation guarantee. However, it does not satisfy the monotonicity condition in Theorem 4. A simple randomized variation gives monotonicity in the declared values. However, the more difficult problem is that the algorithm of Briest et al. 11 is not monotone in terms of the number of sharing agents. Modifying the algorithms of 11 to achieve monotonicity in the amount of sharing is an interesting direction for future work.

Another challenge is to obtain exact or tight approximate solutions when bidders have different share-averseness functions. Our algorithms can be generalized

to this case, but current results on approximations of PIPs [11] are not quite strong enough to give the tight \sqrt{m} approximation. Finally, we have not yet covered the case where the share-averseness functions are not public knowledge. Designing mechanisms that are also incentive compatible with regard to revealing share-averseness, or mechanisms that learn share-averseness from past bids, is an interesting direction for future work.

Acknowledgments. We thank Isabelle Brocas for useful discussions.

References

1. Briest, P., Krysta, P., Vöcking, B.: Approximation techniques for utilitarian mechanism design. In: *The 37th ACM Symposium on Theory of Computing*, pp. 39–48 (2005)
2. Brocas, I.: Auctions with type dependent and negative externalities: the optimal mechanism, IEPW working paper (November 2007)
3. Cornes, R., Sandler, T.: *The Theory of Externalities, Public Goods, and Club Goods*. Cambridge University Press, Cambridge (1996)
4. Ghosh, A., Mahdian, M.: Externalities in online advertising. In: *17th international conference on World Wide Web* (2008)
5. Gonen, R., Lehmann, D.: Linear programming helps solving large multi-unit combinatorial auctions. In: *Electronic Market Design Workshop* (2001)
6. Gonen, R., Lehmann, D.J.: Optimal solutions for multi-unit combinatorial auctions: branch and bound heuristics. In: *ACM Conference on Electronic Commerce*, pp. 13–20 (2000)
7. Jackson, M., Nicoló, A.: The strategy-proof provision of public goods under congestion and crowding preferences. *Journal of Economic Theory* 115(2), 278–308 (2004)
8. Jehiel, P., Moldovanu, B.: Efficient design with interdependent valuations. *Econometrica* 69(5), 1237–1259 (2001)
9. Jehiel, P., Moldovanu, B., Stacchetti, E.: Multidimensional mechanism design for auctions with externalities. *Journal of economic theory* 85(2), 258–294 (1999)
10. Kamien, M., Oren, S., Tauman, Y.: Optimal licensing of cost reducing innovation. *Journal of Mathematical Economics* 21, 483–508 (1992)
11. Krysta, P.: Greedy approximation via duality for packing, combinatorial auctions and routing. In: *Jedrzejowicz, J., Szepietowski, A. (eds.) MFCS 2005. LNCS, vol. 3618*, pp. 615–627. Springer, Heidelberg (2005)
12. Lehmann, D., O’Callaghan, L.I., Shoham, Y.: Truth revelation in approximately efficient combinatorial auctions. In: *ACM Conference on Electronic Commerce*, pp. 96–102. ACM Press, New York (1999)
13. Maskin, E., Riley, J.: Optimal multi-unit auction. *The economics of missing markets, information and games*, 312–335 (1989)
14. Myerson, R.: Optimal auction design. *Mathematics of Operations Research* 6, 58–73 (1981)
15. Scotchmer, S.: Profit maximizing clubs. *J. of Public Economics* 27, 25–45 (1985)
16. Smith, V., Crampton, P., Shoham, Y., Steinberg, R. (eds.): *Combinatorial Auctions*. MIT Press, Cambridge (2006)

Sponsored Search Auctions with Markovian Users

Gagan Aggarwal, Jon Feldman, S. Muthukrishnan, and Martin Pál

Google, Inc.

76 Ninth Avenue, 4th Floor, New York, NY, 10011

1600 Amphitheatre Pkwy, Mountain View, CA, 94043

{gagana, jonfeld, muthu, mpal}@google.com

Abstract. Sponsored search involves running an auction among advertisers who bid in order to have their ad shown next to search results for specific keywords. The most popular auction for sponsored search is the “Generalized Second Price” (GSP) auction where advertisers are assigned to slots in the decreasing order of their *score*, which is defined as the product of their bid and click-through rate. One of the main advantages of this simple ranking is that bidding strategy is intuitive: to move up to a more prominent slot on the results page, bid more. This makes it simple for advertisers to strategize. However this ranking only maximizes efficiency under the assumption that the probability of a user clicking on an ad is independent of the other ads shown on the page. We study a Markovian user model that does not make this assumption. Under this model, the most efficient assignment is no longer a simple ranking function as in GSP. We show that the optimal assignment can be found efficiently (even in near-linear time). As a result of the more sophisticated structure of the optimal assignment, bidding dynamics become more complex: indeed it is no longer clear that bidding more moves one higher on the page. Our main technical result is that despite the added complexity of the bidding dynamics, the optimal assignment has the property that ad position is still monotone in bid. Thus even in this richer user model, our mechanism retains the core bidding dynamics of the GSP auction that make it useful for advertisers.

1 Introduction

Targeted advertisements on search queries is an increasingly important advertising medium, attracting large numbers of advertisers and users. When a user poses a query, the search engine returns search results together with advertisements that are placed into positions, usually arranged linearly down the page, top to bottom. On most major search engines, the assignment of ads to positions is determined by an auction among all advertisers who placed a bid on a keyword that matches the query. The user might click on one or more of the ads, in which case (in the pay-per-click model) the advertiser receiving the click pays the search engine a price determined by the auction.

In the past few years, the sponsored search model has been highly successful commercially, and the research community is attempting to understand the underlying dynamics, explain the behavior of the market and to improve the auction algorithms. The most common auction being run today is the *Generalized Second Price* (GSP) auction: Each bidder i submits a bid b_i stating the maximum amount they are willing to pay for a click, and the bidders are placed in descending order of $b_i p_i$, where p_i is what is called the *click-through-rate* of advertiser i ; i.e., the probability that a user will click on the ad, given that the user looks at it. Much of previous research on sponsored search auctions has fixed this sort order, and focused on understanding the implications of different pricing schemes, assuming strategic behavior on the part of the advertisers. We now know something about GSP's equilibrium properties [8][7][3], alternative pricing that will make it truthful [3], and to some extent, impact on the revenue in principle [8] and via simulations [16].

However, by fixing this sort order, and assuming that an ad's clicks are independent of the other ads on the page, prior work exogenizes an important third party in sponsored search: the *search engine user*. Unfortunately, there is very little guidance on this in the literature, even though the user's behavior is the essential ingredient that defines *the commodity* the advertisers are bidding on, and its value.

We suggest a different framework for principled understanding of sponsored search auctions:

- Define a suitable probabilistic model for search engine user behavior upon being presented the ads.
- Once this model is fixed, ask the traditional mechanism design questions of how to assign the ads to slots, and how to price them.
- Analyze the given mechanism from the perspective of the bidders (e.g., strategies) and the search engine (e.g., user satisfaction, efficiency, revenue).

There are certain well-accepted observations about the user's interaction with the sponsored search ads that should inform the model:

- The higher the ad is on the page, the more users see it and thus click on it.
- The “better” the ad is, the more users click on it, where the ad's “goodness” depends on its inherent quality, and how well it matches the user query.

These properties govern not only how the auction is run but also how advertisers think about their bidding strategy: they prefer to appear higher and get more clicks. Indeed even though GSP is not truthful under a private value model (where each bidder has some inherent private value v_i for a click), its ranking function makes bidding strategy simple: to obtain a higher (more prominent) slot, bid higher and/or make your ad better. This simplicity is very important, since an advertiser may not have a precise notion of click value on which to base their bid. Indeed some advertisers are hoping to generate traffic just to attract attention to their brand, and the value of this attention is less clear than for an advertiser who is making direct sales through the Internet.

In this paper, we propose a natural Markov model for user clicks that retains the properties above, and no longer assumes that the number of clicks an ad

receives is independent of the other ads on the page. We show that simple ranking no longer finds the optimal assignment of ads to slots under this model, but the optimal assignment can still be found efficiently (in near-linear time). Our main technical result is to show that under this new assignment algorithm, the auction mechanism still has a simple bidding strategy; i.e., we prove that a bidder's ad slot is monotone in her bid.

1.1 Our Contributions

Modeling the Search Engine User. Most previous work on sponsored search has (implicitly) modeled the user using two types of parameters: ad-specific click-through rates p_i and position-specific visibility factors α_j . There are some intuitive user behavior models that express overall click-through probabilities in terms of these parameters. One possibility is “for each position j *independently*, the user looks at the ad i in that position with probability α_j then clicks on the ad with probability p_i .” Alternatively: “The user picks a *single* position according to the distribution implied by the α_j 's, and then clicks on the ad i in that position with probability p_i .” Under both these models, it follows that the probability of an ad i in position j receiving a click is equal to $p_i\alpha_j$, which is the so-called *separability* assumption [3]. From separability it follows that GSP ordering of ads will be suitable, because GSP ordering maximizes the total advertiser value on the page.

In both these models there is no reason *a priori* that the position factors α_j should decrease; this is simply imposed because it makes sense, and it is verifiable empirically. Also, both suggested models assume that the probability of an ad getting clicked is independent of *other ads* on the page, an assumption made without much justification. It is hard to imagine that seeing an ad, perhaps followed by a click, has no effect on the subsequent behavior of the user.

In designing a user model, we would like the monotonicity of click-through rate in position to arise naturally. Also, each ad should have parameters dictating their effect on the user both in terms of clicking on that ad, as well as looking at other ads. We propose a model based on a user who starts to scan the list of ads from the top, and makes decisions (about whether to click, continue scanning, or give up altogether) based on what he sees. More specifically, we model the user as the following Markov process: “Begin scanning the ads from the top down. When position j is reached, click on the ad i with probability p_i . Continue scanning with probability q_i .” In this model, if we try to write the click probability of an ad i in position j as $p_i\alpha_j$, we get that $\alpha_j = \prod_{i' \in A} q_{i'}$, where A is the set of ads placed above position j . Thus the “position factor” in the click probability decreases with position, and does so naturally from the model. Also note that we do not have separability anymore, since α_j depends on which ads are above position j . Consequently, it can be shown that GSP assignment of ads is no longer the most efficient.

¹ Throughout the paper, we will often refer to a position or an ad being “higher” or “above” another position or ad; this means that it is earlier on the list, and is looked at first by the user.

Auction with Markovian users. Given this new user model, we can now ask what the best assignment is of ads to slots. We will study the most efficient assignment; i.e., the one that maximizes total advertiser value derived from user clicks. It turns out that the structure of this assignment is different from that of GSP, and indeed is more sophisticated than any simple ranking. The presence of the q_i 's requires a tradeoff between the click probability of an ad and its effect on the slots below it. In this paper, we identify certain structural properties of the optimal assignment and use them to find such an optimal assignment efficiently, not only in polynomial time, but in near-linear time. Given this algorithm, a natural candidate for pricing is VCG [18,6,11], which is truthful in this setting under a private value model.

Monotonicity of Bidding. Now that we have defined a more sophisticated assignment function, even though VCG pricing makes the mechanism truthful under a private click-value model, the auction may not still admit the intuitive bidding strategies that are so important under GSP—especially for advertisers without a precise notion of click value. Our main technical result is to show that in our Markov model, if a mechanism uses the most efficient assignment, indeed position and click probabilities are monotonic in an ad's bid (with all other bids fixed), thus preserving this important property. Monotonicity of click probability follows from the general result of Archer and Tardos [4] on single-parameter mechanisms—for completeness we provide a proof from first principles. In contrast, position monotonicity turns out to be rather involved to prove, requiring some detailed combinatorial arguments, and insights into the optimal substructure of bidder assignments.

1.2 Related Work

Sponsored search has been an active area of research in the last several years after the early papers explored the foundational models [8,3,17,15]. In general, the motivation for the work that followed is that sponsored search in practice is much more complex than as described by the first models; see [9] for a discussion.

Only very recently are alternate user models that break the separability assumption starting to receive some attention. Ghosh and Mahdian [10] study a very general model and show hardness results for the allocation (winner determination) problem; they also give algorithms for several special cases, but none of those imply the algorithms discussed in this work. Craswell et al. [7] give an empirical study of several user click models. The “cascade” model, which was found to fit the data the best, is a special case of the model we study here (with $p_i = 1 - q_i$ and the events being mutually exclusive). Gunawardana and Meek [12] performed an empirical study of ad aggregators with the goal of detecting the affect of an ad on the other ads on the page. Their findings were consistent with our model; i.e., the presence of an ad can have a significant affect on the ads below it. Athey and Ellison [5] present a model where users have an inherent need, and click until that need is filled (or there is little chance of it getting filled). They analyze user behavior, advertiser bidding strategies and Bayesian equilibria in their model.

Independently of our work (which also appeared in [12]), Mahdian and Kempe [13] study the same model we do here. They also provide an $O(n \log n + nk)$ dynamic program for allocation; however at that point they generalize to the case of position-dependent continuation probabilities and provide an approximation algorithm for this case, whereas we go on to study deeper structural and incentive properties in the original model.

2 Markov User Click Model

We consider an auction with n bidders $\mathcal{B} = \{1, \dots, n\}$ and k positions. We will also refer to “ad i ,” meaning the advertisement submitted by bidder i . Each bidder $i \in \mathcal{B}$ has two parameters, p_i and q_i . The click-through-rate p_i is the probability that a user will click on ad i , given that they *look* at it. The continuation probability q_i is the probability that a user will look at the next ad in a list, given that they look at ad i .

Each bidder submits a bid b_i to the auction, representing the amount that they value a click. The quantity $p_i b_i$ then represents the value of an “impression,” i.e., how much they value a user looking at their ad. This is commonly referred to as their “ecpm.”² Throughout, we use the notation $e_i = p_i b_i$ for convenience.

Our model is as follows. Given an assignment (x_1, \dots, x_k) of bidders to the k positions, the user looks at the first ad x_1 , clicks on it with probability p_{x_1} , and then continues looking with probability q_{x_1} .³ This is repeated with the second ad, etc., until the last ad is reached, or some continuation test has failed. Thus the overall expected value of the assignment to the bidders is $e_{x_1} + q_{x_1}(e_{x_2} + q_{x_2}(e_{x_3} + q_{x_3}(\dots q_{x_{k-1}}(e_{x_k}))))$.

The goal of the auctioneer is to compute an assignment of ads to positions that maximizes overall expected value. Given this assignment, prices can be computed using VCG [18, 6, 11]: for each assigned bidder we compute the change in others’ value if that bidder were to disappear. This assures truthful reporting of bids under a profit-maximizing utility.

3 Properties of the Optimal Assignment

In this section, we establish several properties of optimal assignments in this Markov user model, including our main technical result that position and click probability will be monotone in bid and match our intuition. We also give our algorithm for finding an optimal assignment, which gives the truthful auction via VCG pricing. All proofs can be found in [1].

² The acronym ecpm stands for “expected cost per thousand” impressions, where M is the roman numeral for one thousand. We will drop the factor of one thousand and refer to $p_i b_i$ as the “ecpm.”

³ The user could also have some fixed probability of looking at the first ad, which can be trivially incorporated into our results, and we leave this out for clarity. Also, the click event and the continuation event could in principle have some correlation, and all our results will still hold. However since we only consider expected value, we never use this correlation explicitly in our analysis.

Adjusted ECPM. It turns out that the quantity $e_i/(1 - q_i)$, which we will refer to as the “adjusted ecpm (a-ecpm),” plays a central role in this model. Intuitively, this quantity is the impression value adjusted by the negative effect this bid has on the ads below it. We use $a_i = e_i/(1 - q_i)$ for convenience. The following lemma tells us how to assign a set of k selected ads to the k positions:

Lemma 1. □ *In the most efficient assignment, the ads that are placed are sorted in decreasing order of adjusted ecpm $a_i = e_i/(1 - q_i)$.*

While this result tells us how to sort the ads selected, it does not tell us *which* k ads to select. One is tempted to say that choosing the top k ads by a-ecpm would do the trick; however the following example proves otherwise.

Example 1. Suppose we have two slots and three bidders as follows:

Bidder	e_i	q_i	$a_i = e_i/(1 - q_i)$
1	\$1	.75	4
2	\$2	.2	2.5
3	\$0.85	.8	4.25

Let’s consider some possible assignments and their efficiency. If we use simple ranking by ecpm e_i , we get the assignment (2, 1), which has efficiency $\$2 + .2(\$1) = \$2.20$. If we use simple ranking by a-ecpm a_i we get the assignment (3, 1) with efficiency $\$0.85 + .8(\$1) = \$1.65$. It turns out that the optimal assignment is (1, 2) with efficiency $\$1 + .75(\$2) = \$2.50$. The assigned bidders are ordered by a-ecpm in the assignment, but are not the top 2 bidders by a-ecpm.

Now suppose we have the same set of bidders, but now we have three slots. The optimal assignment in this case is (3, 1, 2); note how bidder 3 goes from being unassigned to being assigned the first position.

Bidder Dominance. In classical sponsored search with simple ranking, a bidder j can dominate another bidder i by having higher ecpm; i.e., bidder j will always appear whenever i does, and in a higher position. Example □ above shows that having a higher ecpm (or a-ecpm) does not allow a bidder to dominate another bidder in our new model. However, we show that if she has higher ecpm *and* a-ecpm, then this does suffice. This is not only interesting in its own right, it is essential for proving our deeper structural properties.

Lemma 2. *For all bidders i in an optimal assignment, if some bidder j is not in the assignment, and $a_j \geq a_i$ and $e_j \geq e_i$, then we may substitute j for i , and the assignment is no worse.*

Subset Substructure in Optimal Assignments. We show some subset structure between optimal assignments to different numbers of slots. This is used to prove position monotonicity, and is an essential ingredient of our algorithm. Let $\text{OPT}(C, j)$ denote the set of all optimal solutions for filling j positions with bidders from the set C .

Theorem 1. *Let $j \in \{1, \dots, k\}$ be some number of positions, and let C be an arbitrary set of bidders. Then, for all $S \in \text{OPT}(C, j - 1)$, there is some $S' \in \text{OPT}(C, j)$ where $S' \supset S$.*

⁴ Interestingly, this result also essentially follows from a prior work on optimizing database queries □[\[14\]](#).

Monotonicity of Position and Click Probability. Our main theorem regarding the structure of the optimal assignments in the Markovian click model is that position and click probability are monotonic in a bidder's bid, with all other bids fixed. This is a fundamental property that makes the bidder's interaction with the system intuitive, and allows the bidder to adjust her bid intelligently without global knowledge of the other bids.

Theorem 2. *As a bidder increases her bid (keeping all other bids fixed):*

- (a) *the probability of her receiving a click in the optimal solution does not decrease, and*
- (b) *her position in the optimal solution does not go down.*

Computing the Optimal Assignment. A dynamic program gives an $O(n \log n + nk)$ time algorithm for computing the optimal assignment of bidders to positions. The algorithm proceeds as follows. First, sort the ads in decreasing order of a-ecpm in time $O(n \log n)$. Then, let $F(i, j)$ be the efficiency obtained (given that you reach slot j) by filling slots (j, \dots, k) with bidders from the set $\{i, \dots, n\}$. We get the following recurrence: $F(i, j) = \max(F(i+1, j+1)q_i + e_i, F(i+1, j))$. Solving this recurrence for $F(1, 1)$ yields the optimal assignment, and can be done in $O(nk)$ time. In fact, insights from the previous sections give an $O(n \log n + k^2 \log^2 n)$ time algorithm which is faster when k is large with respect to $\log n$:

Theorem 3. *Consider the auction with n Markovian bidders and k slots. There is an optimal assignment which can be determined in $O(n \log n + k^2 \log^2 n)$ time.*

4 Concluding Remarks

We approached sponsored search auctions as a three party process by modeling the behavior of users first and then designing suitable mechanisms to affect the game theory between the advertiser and the search engine. This formal approach shows an intricate connection between the user models and the mechanisms.

There are some interesting open issues to understand about our model and mechanism. For example, in order to implement our mechanism, the search engine needs to devise methods to estimate the parameters of our model, in particular, q_i 's. This is a challenging statistical and machine learning problem. Also, we could ask how much improvement in efficiency and/or revenue is gained by using our model as opposed to VCG without using our model.

More powerful models will also be of great interest. One small extension of our model is to make the continuation probability q_i a function of location as well, which makes the optimization problem more difficult; Mahdian and Kempe [13] have given an approximation algorithm for this case, and so it is natural to ask if position monotonicity is preserved in their algorithm. One can also generalize the Markov model to handle arbitrary configurations of ads on a web page (not necessarily a search results page), or to allow various other user states (such as navigating a landing page). Finally, since page layout can be performed dynamically, one could ask what would happen if the layout of a web page were a part of the mechanism; i.e., a function of the bids.

Acknowledgments. We thank the anonymous reviews of a previous version of this paper for their comments and for pointing us to several references.

References

1. Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M.: Sponsored search auctions with Markovian users (May 2008), <http://arxiv.org/abs/0805.0766>
2. Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M.: Sponsored search auctions with Markovian users. In: Workshop on Ad Auctions (July 2008)
3. Aggarwal, G., Goel, A., Motwani, R.: Truthful auctions for pricing search keywords. In: ACM Conference on Electronic Commerce (EC) (2006)
4. Archer, A., Tardos, E.: Truthful mechanisms for one-parameter agents. In: Proc. of the 42nd IEEE Symposium on Foundations of Computer Science (2001)
5. Athey, S., Ellison, G.: Position auctions with consumer search. Levine's Bibliography 12224700000001633, UCLA Department of Economics (October 2007)
6. Clarke, E.: Multipart pricing of public goods. *Public Choice* 11, 17–33 (1971)
7. Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: WSDM 2008 (2008)
8. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In: Second workshop on sponsored search auctions (2006)
9. Feldman, J., Muthukrishnan, S.: Algorithmic Methods for Sponsored Search Advertising, pp. 91–124. Springer, Heidelberg (2008)
10. Goel, G., Mehta, A.: Online budgeted matching in random input models with applications to adwords. In: SODA (2008)
11. Groves, T.: Incentives in teams. *Econometrica* 41(4), 617–631 (1973)
12. Gunawardana, A., Meek, C.: Aggregators and contextual effects in search ad markets. In: Workshop on Targeting and Ranking for Online Advertising (2008)
13. Kempe, D., Mahdian, M.: Cascade model for externalities in sponsored search. In: Workshop on Internet Ad Auctions (2008)
14. Krishnamurthy, R., Boral, H., Zaniolo, C.: Optimization of nonrecursive queries. In: VLDB 1986, pp. 128–137 (1986)
15. Lahie, S., Pennock, D., Saberi, A., Vohra, R.: Sponsored Search Auctions. In: Algorithmic Game Theory, pp. 699–716. Cambridge University Press, Cambridge (2007)
16. Szymanski, B.K., Lee, J.-S.: Impact of ROI on bidding and revenue in sponsored search advertisement auctions. In: 2nd Wkshp Sponsored Search Auctions (2006)
17. Varian, H.: Position auctions. *International Journal of Industrial Organization* 25(6), 1163–1178 (2007)
18. Vickrey, W.: Counterspeculation, auctions and competitive-sealed tenders. *Finance* 16(1), 8–37 (1961)

On the Equilibria and Efficiency of the GSP Mechanism in Keyword Auctions with Externalities

Ioannis Giotis and Anna R. Karlin

University of Washington, Seattle WA 98195, USA
{giotis,karlin}@cs.washington.edu

Abstract. In the increasingly important market of online search advertising, a multitude of parameters affect the performance of advertising campaigns and their ability to attract users' attention enough to produce clicks. Thus far, the majority of the relevant literature assumed an advertisement's probability of receiving a click to be dependent on the advertisement's quality and its position in the sponsored search list, but independent of the other advertisements shown on the same webpage.

We examine a promising new model [1, 16] that incorporates the *externalities* effect based on the probabilistic behavior of a typical user. We focus on the *Generalized Second Price* mechanism used in practice and examine the Nash equilibria of the model. We also investigate the performance of this mechanism under the new model by comparing the efficiency of its equilibria to the optimal efficiency.

1 Introduction

Online search engine advertising is an appealing approach to highly targeted advertising, and is *the* major source of revenue for modern web search engines such as Google, Yahoo! and MSN. The most common setup is as follows: when a user performs a query at a search engine, she is shown a collection of organic search results that contains the links the search engine has deemed relevant to the search, together with a list of *sponsored links*, i.e., paid advertisements. If the user actually clicks on a sponsored link, she will be transferred to the advertiser's web site. For each such click, in which the advertiser receives a potential customer, the advertiser pays the search engine.

Keyword auctions determine which ads get assigned to which keywords (search terms) and how much each advertiser pays. Because of the explosive growth of online advertising and the rising economic importance of ad auctions, a great deal of recent research has focused on developing mathematical models of these systems, with an eye towards understanding their equilibria, dynamics and other properties from the perspective of users, advertisers and search engines [20, 2, 11, 19, 6].

Most keyword auction models assume that each advertisement shown has an inherent click-through rate that depends only on the slot allocated to that advertisement and on the advertisement itself, regardless of the other advertisements

that are shown. However, this does not take into account *externalities*: the success of an advertisement depends to a significant extent on which other advertisements are shown alongside it. This is because high-quality competitive ads shown at the same time may detract from each other. Moreover, low-quality ads may deter the viewer from continuing to examine ads shown on the same page.

The importance of the phenomena of externalities has motivated a number of recent papers [12, 1, 16, 9]. Building on work of Craswell et al [8], independently, Kempe and Mahdian [16], and Aggarwal, Feldman, Muthukrishnan and Pal [1] have defined a new Markovian user model. Their model postulates that users scan through the ads in order. For each ad a , users decide probabilistically whether to click, with some ad-specific probability r_a , as well as whether to continue the scanning process, that depends on a different ad-specific probability q_a (as well as the slot the ad is in). This probabilistic continuation models the externality of prematurely terminating the scanning process as a result of either a frustrating irrelevant ad, or a very high-quality web site leading to a purchase. The papers of Kempe et al and Aggarwal et al focus on the problem of computing the efficient allocation given this model, and the efficient (in terms of computational complexity) implementability of an incentive-compatible mechanism (the VCG mechanism).

In this paper, we consider the impact of the new Markovian model on equilibria under the Generalized Second Price (GSP) mechanism. This is important because GSP is the standard mechanism used in practice. Moreover, it is highly unlikely that even if the Markovian model is accurate that the search engines will switch to VCG. Thus, it is of great interest to understand the effects of the new model when GSP is used.

1.1 Results

The focus of our research is on understanding the equilibria of GSP under the new user model. Our main results are the following:

We show that in the new model, in contrast to the most important result about the standard model [20, 11], GSP does not necessarily have an equilibrium in which efficiency is maximized. This raises two key questions: First, does GSP have pure equilibria? And second, how bad can these equilibria be in terms of their efficiency?

The answer to the first question is yes. In Section 3, we give a general construction showing that no matter what the parameters of the system, GSP does have pure Nash equilibria in the new model.

We then turn to the study of the efficiency of GSP equilibria. Our main result here, in Section 4, is that the efficiency of the *worst* Nash equilibrium under GSP can be a factor of k smaller than optimal but no worse, where k is the number of slots in the system. Thus, the so-called *price of anarchy* [18] of GSP with respect to efficiency is k . This latter result depends on the assumption that no advertiser ever bids more than their value. On the other hand, when advertisers can bid more than their value, the efficiency of the worst Nash equilibrium under GSP can be arbitrarily smaller than optimal.

Finally, we show that there are instances where the efficiency of the *best* Nash equilibrium under GSP has efficiency which is a factor of k smaller than optimal. Thus, even the *price of stability* [4] of GSP is k .

2 Model

We consider a model for sponsored search auctions with n participating players (bidders or advertisers) $\{1, \dots, n\}$, bidding for k advertising slots $\{1, \dots, k\}$. Each player $i \in \{1, \dots, n\}$ has three associated values: The first, $0 \leq r_i \leq 1$, represents the *position-independent click-through rate*, which is the probability that the user will click on the ad, given that they look at it. It is a measure of the *relevance* of the ad to the query as well as the general *quality* of the ad, and can also perhaps be thought of as the probability that the user will click on the ad if it is placed in the top slot. The second quantity, $0 \leq q_i \leq 1$ represents the *continuation probability*, the probability that a user will look at the next ad in the list, given that they look at ad i . The third quantity, $V_i \geq 0$ represents the *expected value or profit* of the advertiser given that the user clicks on his ad.

Each advertising slot $s \in \{1, \dots, k\}$ has an associated fixed constant θ_s representing the ad-independent probability that a user continues scanning advertisements after the s -th slot, given that she scans the s -th slot.

Each player submits a bid and depending on the mechanism used (see discussion below), the search engine produces an allocation of the k slots $\pi()$ such that advertiser $\pi(s)$ is assigned to slot s . An associated list of prices p is also produced such that each time user $\pi(s)$ receives a click he is charged a price of p_s .

We model the behavior of the end-user when presented with the sponsored search results as follows. The user begins scanning the results list with some probability θ_0 which for simplicity we normalize to 1. The first slot is scanned and the user clicks on the ad with probability $r_{\pi(1)}$. Independently of whether the user clicked on the first ad, she proceeds to scan the second slot with probability $\theta_1 \cdot q_{\pi(1)}$, where she clicks on that ad with probability $r_{\pi(2)}$. On the other hand, with probability $1 - \theta_1 \cdot q_{\pi(1)}$, the user stops scanning ads after the first and quits the whole process. Given that the user scanned the second slot, she proceeds to scan the third slot with probability $\theta_2 \cdot q_{\pi(2)}$ and so on.

Our main measure for evaluating the performance of the system will be the system's *efficiency*. The efficiency of the system for a given ranking of the players π is defined as the sum of the expected utilities of all the players.

$$\begin{aligned}
 \text{efficiency} &= r_{\pi(1)}V_{\pi(1)} + \theta_1q_{\pi(1)} \left(r_{\pi(2)}V_{\pi(2)} + \theta_2q_{\pi(2)} \left(\dots \left(V_{\pi(2)} \right) \right) \right) = \\
 &= \sum_{j=1}^k \left(\left(\prod_{i=1}^{j-1} \theta_i q_{\pi(i)} \right) \cdot r_{\pi(j)} V_{\pi(j)} \right). \tag{1}
 \end{aligned}$$

Discussion

In this model, the probability that the user proceeds to scan the ad in slot $s + 1$ given that she scanned the ad in slot s is dependent on both the slot s and the

the quality $r_{\pi(s)}$ of the ad in slot s . The dependence on slot has been documented in eye-tracking studies that show that the probability that a user looks at an ad decays with the slot number [13, 15]. This description of the user’s eye movement and clicking behavior has been studied under the term “directional market” in several economics papers [5, 3]. As stated, in [3], “the directionality arises due to cognitive burden as it is cognitively ‘costlier’ for a typical consumer to visit sellers at the bottom of the list before visiting the sellers at the top of the listing”.

The dependence of continued scanning on the quality $r_{\pi(s)}$ of the ad in slot s is the combination of two phenomena. First, if the click on slot s results in a conversion, the user is unlikely to continue scanning. Second, if the quality of the ad in slot s is low, the user may be more likely to give up in “disgust”. These factors and undoubtedly many others combine to give some ad-dependent probability of continuing to scan. This feature of the model captures the externalities inherent in this setting.

2.1 Mechanisms

The VCG Mechanism. One of the mechanisms under examination and our main comparison point is the celebrated Vickrey-Clarke-Groves (VCG) [21, 7, 14] mechanism. The VCG mechanism is a truthful mechanism which allocates the slots such that efficiency, as defined in [1], is maximized.

Recall that under the VCG mechanism, the expected payment charged to player $\pi'(j)$ at slot j is determined by $OPT_{-\pi'(j)} - (OPT - v_{\pi'(j)})$ where

$$v_{\pi'(j)} = \left(\prod_{i=1}^{j-1} \theta_i q_{\pi(i)} \right) \cdot r_{\pi(j)} V_{\pi'(j)}$$

is the expected utility of this player, OPT is the optimal efficiency with all the players and $OPT_{-\pi'(j)}$ is optimal efficiency without player $\pi'(j)$. Since the most commonly used charging scheme, both in literature and in practice, is on a per click basis, the pay per click price for VCG is defined as

$$p_j = \frac{OPT_{-\pi'(j)} - (OPT - v_{\pi'(j)})}{\left(\prod_{i=1}^{j-1} \theta_i q_{\pi(i)} \right) \cdot r_{\pi(j)}}$$

The GSP Mechanism. Our main focus in this study will be the mechanism most widely used in practice, the *Generalized Second Price* mechanism (GSP):

Definition 1. GSP mechanism

Each player i submits a bid b_i representing the maximum amount they are willing to pay for a click. The GSP mechanism ranks the players in decreasing order of $b_i \cdot r_i$. For the resulting ranking $\pi()$, the price per click of slot j is

$$p_j = b_{\pi(j+1)} \frac{r_{\pi(i+1)}}{r_{\pi(i)}}$$

Player	V	r	q	VCG ranking	VCG price(expected)
1	1	1	0.75	1	0.7
2	2	1	0.2	2	0.6
3	0.8	1	0.7	3	0

Fig. 1. Counterexample for the existence of the VCG equilibrium

The expected utility $U(\pi(j))$ of player $\pi(j)$ occupying slot j is

$$U(\pi(j)) = \left(\prod_{i=1}^{j-1} \theta_i q_{\pi(i)} \right) (r_{\pi(j)} V_{\pi(j)} - b_{\pi(j+1)} r_{\pi(j+1)}).$$

We note that in the standard model (where $q_i = 1$ for all i) the GSP ranking maximizes efficiency with respect to the declared bids¹. In our model, this is not the case: if we were to rank by declared efficiency, players with lower $b_i r_i$ might be placed in higher slots than players with higher $b_i r_i$, which would be considered unfair.

3 Nash Equilibria in the GSP Mechanism

Of particular interest in the literature on ad auctions [20, 10, 11] has been the equilibrium that yields the same allocation and prices as the VCG mechanism under the standard user model². It is of course very appealing to be able to show that GSP has an equilibrium in which optimal efficiency and several other appealing properties of the VCG equilibrium (such as envy-freeness) hold.

Unfortunately in our model the VCG equilibrium does not always exist. We present a counterexample inspired by a similar counterexample for a different purpose in [1]. Suppose we have 3 bidders and 2 slots with $\theta_1 = 1$. Given the parameters defined in figure 1 it is easy to check that the ranking and VCG prices are as stated in the figure. Notice that the prices in the figure are prices per round or in expectation, therefore the pay per click prices would have to be $p_1 = 0.7$ and $p_2 = \frac{0.6}{\theta_1 q_1} = \frac{0.6}{0.75} = 0.8$ and thus $b_2 = 0.7$ and $b_3 = 0.8$ which cannot result in the desired ranking in the GSP mechanism. Despite the fact the VCG equilibrium might not be achievable, we are able to prove that a pure equilibrium always exists:

Theorem 1. GSP Equilibria Existence

We assume the players are labeled in decreasing order of $r_i \cdot V_i$. If the players' bids are such that

$$b_s r_s = \begin{cases} V_1 r_1 & \text{for } s = 1, \\ \sum_{j=s-1}^{k+1} \left(\prod_{i=s}^j \theta_{i-1} q_i \right) V_j r_j (1 - \theta_j q_{j+1}) & \text{for } 1 > s \geq k \\ V_s r_s & \text{for } k > s, \end{cases} \quad (2)$$

or alternatively by the following recursive definition

¹ Since this mechanism is not truthful, the actual efficiency of the system is not guaranteed to be optimal as in VCG.

² We will refer to this equilibrium as the VCG equilibrium.

$$b_s r_s = \begin{cases} V_s r_s & \text{for } k > s, \\ (1 - \theta_{s-1} q_s) V_{s-1} r_{s-1} + \theta_{s-1} q_s b_{s+1} r_{s+1} & \text{for } 1 > s \geq k \\ V_1 r_1 & \text{for } s = 1, \end{cases} \tag{3}$$

then the resulting allocation and prices of the GSP mechanism is a Nash equilibrium in the new model.

The proof of this theorem is presented in the Appendix.

4 The Efficiency of GSP Equilibria

In light of the fact that the equilibria of GSP may not maximize efficiency, it is interesting to ask how low the relative efficiency (and other properties) of these equilibria can go. We do this using *price of anarchy* and *price of stability* style of analysis³. For the price of anarchy analysis we focus on the least efficient GSP equilibrium and compare it against the VCG allocation and the most efficient GSP equilibrium, while, for the price of stability, we compare the most efficient GSP equilibrium against the VCG allocation.

We will also distinguish between two cases. In the first case, the players bid in an unrestricted fashion while in the second case the players can only bid as high as their value. While in reality it is possible for players to bid above their values, it seems unlikely that such bidding behavior can be sustained in practice as the players risk paying a price higher than their value. We therefore expect the restricted case to be more interesting in practice. We will show that the price of anarchy for efficiency can be bounded as per the following theorem.

Theorem 2. Price of Anarchy

The price of anarchy of GSP equilibria both against VCG and the best GSP equilibrium is k (the number of slots) in the restricted case, and infinite in the unrestricted case.

Proof. We first look at the efficiency of GSP equilibria in the restricted case. Fix ϵ and δ arbitrarily small positive constants and consider the following setting. We have $n = k + 1$ players $\{1, 2, \dots, k, k + 1\}$ bidding for k slots with $\theta_i = 1$ for all $1 \leq i \leq k$. The players’ parameters are illustrated in figure 2.

It is easy to check that, for large enough X , the most efficient ranking is $[k + 1, k, \dots, 4, 3, 1]$, with total efficiency $kX - (k - 1)(1 + \delta) \geq kX - \epsilon X$. Although this ranking is not achievable under GSP, an equilibrium with the ranking $[2, 3, \dots, k, 1]$ can be achieved if all players bid their values except player 1 who bids $X - 1 - \delta$. The efficiency of this equilibrium is $X - \delta + \frac{1}{1+\delta} ((k - 1)X - (k - 2)(1 + \delta)) \geq \frac{1}{1+\delta} kX - \epsilon X$ for large enough X .

On the other hand, consider the case under GSP where the players are bidding their values except player 2 who bids $X - 1$. The resulting allocation is $[1, 2, \dots, k - 1, k]$ and we can verify that this results in an equilibrium. Clearly,

³ The price of anarchy was originally introduced by Koutsoupias and Papadimitriou in [18] (see also [17] for a survey) as a measure of the performance degradation by selfish autonomous users in the absence of a coordination mechanism.

Player	1	2	3	...	$k - 1$	k	$k + 1$
V	X	$X - \delta$	$X - 1 - \delta$	\dots	$X - 1 - \delta$	$X - 1 - \delta$	$X - 1 - \delta$
q	0	$\frac{1}{1+\delta}$	1	\dots	1	1	1
r	1	1	1	\dots	1	1	1

Fig. 2. Example for the price of anarchy regarding efficiency in the restricted case

all players are getting zero utility without being able to improve it. Player 1 is getting utility $X - (X - 1) = 1$ while if he were to bid lower to obtain slot j he would still get $1/(1 + \delta)(X - (X - 1 - \delta)) = 1$. The efficiency of this equilibrium is just X . We conclude that for both against VCG and over all GSP equilibria, the price of anarchy regarding efficiency can be bounded by

$$\frac{\frac{1}{1+\delta}kX - \varepsilon X}{X} \geq \frac{1}{1 + \delta}k - \varepsilon.$$

We are also able to show that this bound is tight. Indeed, assume we have an arbitrary system of players and slots and consider the least efficient equilibrium of GSP. Focusing on a player x for which $r_x V_x = \max_i r_i V_i$, we will show that the efficiency of this equilibrium is at least $r_x V_x$. Indeed, consider the case where x is not awarded the top slot. In this case some other player y gets the first slot while player x is at slot j with probability of the user getting to that slot ϕ_j . From the equilibrium conditions regarding player x 's "desire" to get the first slot by bidding higher than y 's bid b_y , we have

$$r_x(V_x - b_y \frac{r_y}{r_x}) \leq \phi_j r_x(V_x - p_j)$$

and using our bidding restriction, we can bound both sides of the inequality.

$$\begin{aligned} r_x V_x - r_y V_y &\leq r_x(V_x - b_y \frac{r_y}{r_x}) \leq \phi_j r_x(V_x - p_j) \leq \phi_j r_x V_x \\ r_x V_x - r_y V_y &\leq \phi_j r_x V_x \\ r_x V_x &\leq \phi_j r_x V_x + r_y V_y \end{aligned}$$

The efficiency of the equilibrium is at least $r_y V_y + \phi_j r_x V_x \geq r_x V_x$. But both the VCG mechanism and most efficient GSP equilibrium cannot have efficiency more than $k \cdot r_x V_x$, hence the price of anarchy is at most k .

For the unrestricted case, consider a setting of 1 slot and two players such that $r_1 V_1 = 0, r_2 V_2 = X$. It is easy to see then when $b_1 > X$ and $b_2 = 0$ we have an equilibrium of 0 efficiency. On the other hand, the VCG or optimal GSP equilibrium allocations yield efficiency X . We conclude that for both of these cases the price of anarchy is unbounded. ■

We next turn our attention to the *price of stability* of GSP equilibria relative to the VCG mechanism. Here our goal is to understand how the *best* GSP equilibrium in the worst case compares in performance to the VCG outcome.

Theorem 3. Price of Stability

The price of stability of GSP equilibria against the VCG mechanism is k in the restricted case, and between $k/2$ and k in the unrestricted case.

The proof of this Theorem is similar in spirit to the proof of Theorem 1 and is omitted from this short version of this paper.

5 Conclusions

We have examined a simple and elegant model for keyword auctions introduced in a series of papers [8, 11, 16] that is able to capture effects that appear in practice but are not considered by the standard model. This model incorporates externalities by modeling the effects advertisements have on the probability that a typical user will scan or click on *other* ads.

Our model makes use of player parameters that are considered a priori determined by the search engine. The use of the click-through probability r_i is generally considered acceptable and these values are probably computed by search engines by sampling the click performance of an ad when the listing is placed randomly in different slots. However, it is not clear if similar techniques can be used to estimate the new parameters q_i . Although determining q_i is not necessary to run the GSP mechanism, if it can be computed efficiently it would certainly open up possibilities for more efficient ranking and pricing schemes.

We have shown that the GSP mechanism always has a pure Nash equilibrium. On the other hand, unlike the standard model, it may not have a Nash equilibrium which maximizes efficiency. We thus attempted to quantify the difference in efficiency between GSP and VCG by examining the price of anarchy and stability. Although the derived bounds appear to make a strong statement in favor of the VCG mechanism, it remains undetermined how these two mechanisms would compare in practice.

An empirical study with real or simulated auction data would potentially reveal more practical results on the performance of GSP and it would be extremely interesting to evaluate its performance against alternative mechanisms that take advantage of the extended information of this model.

References

- [1] Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M.: Sponsored search auction with markovian users. In: Fourth Workshop on Ad Auctions, ACM Electronic Commerce (2008)
- [2] Aggarwal, G., Goel, A., Motwani, R.: Truthful auctions for pricing search keywords. In: EC 2006: Proceedings of the 7th ACM conference on Electronic commerce, pp. 1–7. ACM, New York (2006)
- [3] Animesh, A., Viswanathan, S., Agarwal, R.: Competing “creatively”. In: online markets: Evidence from sponsored search. Working Paper (2007)

- [4] Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: FOCS 2004: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, Washington, DC, USA, pp. 295–304. IEEE Computer Society Press, Los Alamitos (2004)
- [5] Arbatskaya, M.: Ordered search. Working Paper (2004)
- [6] Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A., Mathieu, C., Schwarz, M.: Greedy bidding strategies for keyword auctions. In: EC 2007: Proceedings of the 8th ACM conference on Electronic commerce, pp. 262–271. ACM, New York (2007)
- [7] Clarke, E.H.: Multipart pricing of public goods. *Public Choice* 11, 17–33 (1971)
- [8] Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: 1st International Conference on Web Search and Data Mining (2008)
- [9] Das, A., Giotis, I., Karlin, A., Mathieu, C.: On the effects of competing advertisements in keyword auctions. Working Paper (2008)
- [10] Edelman, B., Ostrovsky, M.: Strategic bidder behavior in sponsored search auctions. In: First Workshop on Sponsored Search Auctions, ACM Electronic Commerce (2005)
- [11] Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review* (to appear, 2006)
- [12] Ghosh, A., Mahdian, M.: Externalities in online advertising. In: Proceedings of the 17th International World Wide Web Conference (2008)
- [13] Granka, L. Eye-r: Eye tracking analysis of user behavior in online search. Masters Thesis. Cornell University Library Press (2004)
- [14] Groves, T.: Incentives in teams. *Econometrica* 41, 617–631 (1973)
- [15] Hotchkiss, G.: Enquiro eye tracking report ii: Google, msn and yahoo! compared (November 2006)
- [16] Kempe, K., Mahdian, M.: A cascade model for externalities in sponsored search. In: Fourth Workshop on Sponsored Search Auctions, ACM Electronic Commerce (2008)
- [17] Koutsoupias, E.: Coordination mechanisms for congestion games. *Sigact News* 35, 58–71 (2004)
- [18] Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: 16th Annual Symposium on Theoretical Aspects of Computer Science, pp. 404–413 (1999)
- [19] Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: EC 2006: Proceedings of the 7th ACM conference on Electronic commerce, pp. 218–227. ACM Press, New York (2006)
- [20] Varian, H.: Position auctions. *International Journal of Industrial Organization* (to appear, 2006)
- [21] Vickrey, W.: Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 8–37 (1961)

Appendix

A Proof of Theorem 1

Proof. It is easy to check that the two definitions are equivalent. Also, it can be easily seen from the recursive definition that the resulting bids are ordered in the

right way as each $b_s r_s$ is a linear combination of $b_{s+1} r_{s+1}$ and the $s-1$ 'th player's expected valuation. By the initial conditions of the recursive definition it follows that the bids are correctly ordered and the following equilibrium conditions are satisfied

- $\left(\prod_{i=1}^{s-1} \theta_i q_i\right) (V_s r_s - b_{s+1} r_{s+1}) \geq 0$ for all $s \leq k$,
- $\left(\prod_{i=1}^{j-1} \theta_i q_i\right) (r_s V_s - b_j r_j) \leq 0$ for all $j \geq k > s$.

or in other words, all the winning players have greater than or zero utility and the losing players cannot get positive utility by bidding higher.

It remains to show the remaining equilibrium conditions, or that the winning players do not have an incentive to alter their bid so as to get a different slot. Assume an arbitrary winning player s . We need to show that

- For all slots $j < s$, $\left(\prod_{i=j}^{s-1} \theta_i q_i\right) (V_s r_s - b_{s+1} r_{s+1}) \geq V_s r_s - b_j r_j$.
- For all slots $k \geq j > s$, $V_s r_s - b_{s+1} r_{s+1} \geq \left(\prod_{i=s}^{j-1} \theta_i q_{i+1}\right) (V_s r_s - b_{j+1} r_{j+1})$.

To prove the first case, we proceed as follows.

$$\begin{aligned} V_s r_s - b_j r_j &= V_s r_s - ((1 - \theta_{j-1} q_j) V_{j-1} r_{j-1} + \theta_{j-1} q_j b_{j+1} r_{j+1}) \\ &\quad \text{and since } V_s r_s \leq V_{j-1} r_{j-1} \\ &\leq V_s r_s - ((1 - \theta_{j-1} q_j) V_s r_s + \theta_{j-1} q_j b_{j+1} r_{j+1}) \\ &= \theta_{j-1} q_j (V_s r_s - b_{j+1} r_{j+1}) \\ &\leq \dots \text{ (similarly substituting using the recursive definition)} \\ &\leq \theta_{j-1} q_j \theta_j \cdots q_{s-1} \theta_{s-1} q_s (V_s r_s - b_{s+1} r_{s+1}) \\ &\leq \left(\prod_{i=j}^{s-1} \theta_i q_i\right) (V_s r_s - b_{s+1} r_{s+1}), \text{ since } \theta_{j-1} q_s \leq 1. \end{aligned}$$

To prove the second case, for $j > s$, we proceed similarly.

$$\begin{aligned} V_s r_s - b_{s+1} r_{s+1} &= V_s r_s - ((1 - \theta_s q_{s+1}) V_s r_s + \theta_s q_{s+1} b_{s+2} r_{s+2}) \\ &= \theta_s q_{s+1} (V_s r_s - b_{s+2} r_{s+2}) \\ &= \theta_s q_{s+1} (V_s r_s - (1 - \theta_{s+1} q_{s+2}) V_{s+1} r_{s+1} - \theta_{s+1} q_{s+2} b_{s+3} r_{s+3}) \\ &\geq \theta_s q_{s+1} (V_s r_s - ((1 - \theta_{s+1} q_{s+2}) V_s r_s + \theta_{s+1} q_{s+2} b_{s+3} r_{s+3})) \\ &= \theta_s q_{s+1} \theta_{s+1} q_{s+2} (V_s r_s - b_{s+3} r_{s+3}) \\ &\geq \dots \\ &\geq \theta_s q_{s+1} \cdots \theta_{j-1} q_j (V_s r_s - b_{j+1} r_{j+1}) \\ &\geq \left(\prod_{i=s}^{j-1} \theta_i q_{i+1}\right) (V_s r_s - b_{j+1} r_{j+1}). \end{aligned} \quad \blacksquare$$

Biased Voting and the Democratic Primary Problem

Michael Kearns and Jinsong Tan

Department of Computer and Information Science
University of Pennsylvania

{mkearns@cis, jinsong@seas}.upenn.edu

Abstract. Inspired by the recent Democratic National Primary, we consider settings in which the members of a distributed population must balance their individual preferences over candidates with a requirement to quickly achieve collective unity. We formalize such settings as the “Democratic Primary Problem” (DPP) over an undirected graph, whose local structure models the social influences acting on individual voters.

After contrasting our model with the extensive literature on diffusion in social networks (in which a force towards collective unity is usually absent), we present the following series of technical results:

- An impossibility result establishing exponential convergence time for the DPP for a broad class of local stochastic updating rules, which includes natural generalizations of the well-studied “voter model” from the diffusion literature (and which is known to converge in polynomial time in the absence of differing individual preferences).
- A new simple and local stochastic updating protocol whose convergence time is provably polynomial on any instance of the DPP. This new protocol allows voters to declare themselves “undecided”, and has a temporal structure reminiscent of periodic polling or primaries.
- An extension of the new protocol that we prove is an approximate Nash equilibrium for a game-theoretic version of the DPP.

1 Introduction

The recent Democratic National Primary race highlighted a tension that is common in collective decision-making processes. On the one hand, individual voters clearly held (sometimes strong) preferences between the two main candidates, Barack Obama and Hillary Clinton, and these preferences appeared approximately balanced across the population. On the other hand, as the race progressed there were frequent and increasingly urgent calls for Democrats to “unify” the party — that is, quickly determine the winner and then all rally around that candidate in advance of the general election [13]. There was thus a balancing act between determining the overall (average) preference of voters, and reaching global consensus or unity.

Inspired by these events, we consider settings in which the members of a distributed population must balance their individual preferences over candidates

with a requirement to achieve collective unity. We formalize such settings as the “Democratic Primary Problem” (DPP) over an undirected graph, whose local structure models the social influences acting on individual voters. In this model, each voter i is represented by a vertex in the network and a real-valued weight $w_i \in [0, 1]$ expressing their preference for one of two candidates or choices that we shall abstractly call *red* and *blue*. Here $w_i = \frac{1}{2}$ is viewed as indifference between the two colors, while $w_i = 0$ (red) and $w_i = 1$ (blue) are “extremist” preferences for one or the other color.

Our overarching goal is to investigate distributed algorithms in which three criteria are met:

1. *Convergence to the Global Preference*: If the global average W of the w_i is even slightly bounded away from $\frac{1}{2}$ (indifference), then *all* members of the population should eventually settle on the globally preferred choice (i.e. all red if $W < \frac{1}{2}$, all blue if $W > \frac{1}{2}$), even if it conflicts with their own preferences (party unity).
2. *Speed of Convergence*: Convergence should occur in time polynomial in the size of the network.
3. *Simplicity and Locality*: Voters should employ “simple” algorithms in which they communicate only *locally* in the network via (stochastic) updates to their color choices. These updates should be “natural” in that they plausibly integrate a voter’s individual preferences with the current choices of their neighbors, and do not attempt to encode detailed information, send “signals” to neighbors, etc.

The first two of these criteria are obviously formally precise. While it might be possible to formalize the third as well, we choose not to do so here for the sake of brevity and exposition. However, we are explicitly *not* interested in algorithms in which (for instance) voters attempt to encode and broadcast their underlying preferences w_i as a series of binary choices, or similarly unnatural and complex schemes. In particular, in our main protocol it will be very clear that voters are always updating their current choices in a way that naturally integrates their own preferences and the statistics of current choices in their local neighborhood.

We note that the formalization above clearly omits many important features of “real” elections. Foremost among these is the fact that real elections typically have strong global coordination and communication mechanisms such as polling, while we require that all communication between participants be entirely local in the network. On the other hand, our framework does allow for the presence of high-degree individuals, including ones that are indifferent to the outcome ($w_i = 1/2$) and can thus act as “broadcasters” of current sentiment in their neighborhood. Variation in degrees can also be viewed as a crude model for the increasing variety of global to local media sources (from “mainstream” publications to influential blogs to small discussion groups).

There is a large literature on the diffusion of opinion in social networks [4,12,10], but the topic is usually studied in the absence of any force towards collective unity. In many contagion-metaphor models, individuals are more or less susceptible to “catching” an opinion or fad from their neighbors, but are

not directly concerned with the global outcome. In contrast, we are specifically interested in scenarios in which individual preferences are present, but are subordinate to reaching a unanimous global consensus.

Our main results are:

- An impossibility result establishing exponential convergence time for the DPP for a broad class of local stochastic updating rules, which includes natural generalizations of the well-studied “voter model” from the diffusion literature (and which is known to converge in polynomial time in the absence of differing individual preferences).
- A new simple and local stochastic updating protocol whose convergence time is provably polynomial on any instance of the DPP. This new protocol allows voters to declare themselves “undecided”, and has a temporal structure reminiscent of periodic polling or primaries.
- An extension of the new protocol that we prove is an approximate Nash equilibrium for a game-theoretic version of the DPP.

2 The Democratic Primary Problem

The Democratic primary problem (DPP) is studied over an undirected graph $G = (V, E)$ with n nodes and m edges, where each node i represents an individual voter. Denote by $\mathcal{N}(i)$ the neighbors of i in G ; we always consider i as a neighbor of himself.

There are two competing choices or opinions, that, without loss of generality, we shall call *blue* and *red* (or b and r for short). A voter i comes with a real-valued weight $w_i \in [0, 1]$ expressing his preference for one of the two opinions; without loss of generality, let $w_i(b) = w_i$ and $w_i(r) = 1 - w_i$ denote his preference for *blue* and *red*, respectively.

Throughout, we make the assumption that $|\sum_{i \in V} w_i(b) - \sum_{i \in V} w_i(r)| > \epsilon$ for some constant ϵ , that is, one opinion is always collectively preferred to the other; and we assume that which opinion is preferred is *not* known a priori to anyone and the goal of the DPP is for the entire population to actually figure this out through a distributed algorithm, or protocol, that is *simple and local*, and converges in time polynomial in n to the collectively preferred consensus. Because of the stochastic nature of the protocol we consider, it is implausible to require that it always converges to the collectively preferred consensus. Instead, we require the protocol does so with *high* probability, by which we mean the probability can differ from 1 by an amount that is at most negligible in n . We summarize the definition of DPP in the following.

DEMOCRATIC PRIMARY PROBLEM (DPP)

Instance: Given an undirected graph $G = (V, E)$ with n nodes, two opinions $\{b, r\}$, and for each $i \in V$, a preference $(w_i(b), w_i(r))$ where $w_i(b), w_i(r) \in [0, 1]$ and $w_i(b) + w_i(r) = 1$. Assume $\exists \alpha \in \{b, r\}$ such that $\sum_i w_i(\alpha) > \frac{n}{2} + \epsilon$ for some constant $\epsilon > 0$ and α is not known a priori.

Objective: Design a *simple and local* distributed protocol that in time polynomial in n lets V converge to α with *high* probability.

We will be considering protocols of the following form:

1. (Initialization) At round 0, each node i in V independently and simultaneously initializes to an opinion in $\alpha \in \{b, r\}$ according to \mathcal{I} , a randomized function that maps i 's local information to an opinion in $\{b, r\}$.
2. (Stochastic Updating) At round $t \geq 1$, a node i is chosen uniformly at random from V ; i then picks a neighbor $j \in \mathcal{N}(i)$ according to a possibly non-uniform distribution over $\mathcal{N}(i)$; this distribution is determined by function \mathcal{F} , which is a randomized function whose arguments are i 's local information. i then converts to j 's opinion.

Therefore a protocol is specified by a pair of functions, $(\mathcal{I}, \mathcal{F})$. This framework by itself does not forbid “unnatural” coding behaviors as discussed in the Introduction; however in the spirit of emphasizing algorithms that are simple and local, we restrict \mathcal{I} and \mathcal{F} to be functions that only depend on simple and local information of a node i . In particular, only the following arguments to either functions are considered: 1) f_i , the distribution of opinions in the neighborhood, where $f_i(b)$ and $f_i(r)$ represent the fractions of neighbors currently holding opinion *blue* and *red*, respectively; 2) i 's intrinsic preferences w_i ; 3) i 's degree $d_i = |\mathcal{N}(i)|$.

3 The Classic and Simplest: The Voter Model

The voter model, which was introduced by Clifford and Sudbury [2] and Holley and Liggett [5], is a well-studied probabilistic stochastic process that models opinion diffusion on social networks in a most basic and natural way. It consists of a class of protocols that satisfy our criterion of being *simple and local*. In fact, this class of protocols is the simplest that we examine in this paper. A voter model protocol is one where in each round, a node i is picked uniformly at random from V , and i in turn picks one of his neighbors uniformly at random and adopts his opinion; it does not specify how the initialization is done. More formally,

Definition 1 (Voter Model). *The voter model is a class of protocols of the form $(\mathcal{I}, \mathcal{F})$ where $\mathcal{F}(f_i) = \alpha$ with probability $f_i(\alpha)$, $\forall \alpha \in \{b, r\}$.*

Importantly, the voter model is a class of protocols in which there are no individual preferences present at all, and the only concern is with reaching unanimity (to either color). This is in sharp contrast to the DPP setting. However, we shall make use of some known results on the voter model, which we turn to now.

Let C_{vm} denote the random variable whose value is the time at which a consensus is reached in a voter model protocol. It can be shown that $E(C_{vm}) = O(\log(n) \max_{i,j} h_{ij})$, where h_{ij} is the expected hitting time of node j of a random walk starting from node i (see [1] for a proof of this). Also it is well-known that for any graph G with self-loops (i.e. $i \in \mathcal{N}(i)$), $h_{ij} = O(n^3)$ for any node i, j [11], so $E(C_{vm}) = O(n^3 \log(n))$. We summarize this in the following theorem.

Theorem 1 ([1]). *For any initialization, it takes $O(n^3 \log(n))$ time in expectation for all the n nodes to converge to a consensus opinion in a voter model protocol.*

Denote by π the stationary distribution of a random walk on G , i.e. $\pi(i) = d_i/2m$ for all $i \in V$ and $\pi(S) = \sum_{i \in S} d_i/2m$ for all $S \subset V$. The next theorem also largely follows from established results in literature. And we omit the proof here.

Theorem 2. *Let $S \subset V$ be the set of nodes initialized to opinion α in a voter model protocol, then after n^5 rounds the probability that an α -consensus is reached differs from $\pi(S)$ by an amount negligible in n .*

Theorem 1 and 2 allow us to conclude that after n^5 rounds into a voter model protocol, with high probability *some* consensus is reached. In particular, let $B, R \in V$ be the set of nodes initialized to *blue* and *red* respectively, the probability of reaching a *b*-consensus (resp., *r*-consensus) differs from $\pi(B)$ (resp., $\pi(R)$) by negligible amount. Recall our goal of solving DPP is to find an efficient protocol that converges to the collectively preferred consensus with high probability. Since the voter model does not even consider w_i , it is clear that it does not solve the DPP. (The voter model does not specify how initialization is done, however it is easy to prove that even if \mathcal{I} is allowed to depend on w_i in an *arbitrary* way, no voter model protocol solves the DPP.)

Therefore, the logical next thing to consider in order to solve the DPP is to allow \mathcal{F} to in addition depend on w_i . And this leads us to the natural extension of the classic voter model: the *biased voter model*.

4 The Biased Voter Model

Discussion from the previous section reveals that in order to solve the DPP, it is necessarily to allow \mathcal{F} to depend on w_i in addition to f_i , so that how an individual changes his opinion is influenced by his neighbors as well as his own intrinsic preferences. A natural class of \mathcal{F} that reflect an individual's preference (or *bias*) are those that let him assume his preferred opinion α with probability higher than $f_i(\alpha)$, which is the probability he assumes opinion α in the voter model. We call the resulting model the *biased voter model* and define it formally as follows.

Definition 2 (Biased Voter Model). *The biased voter model is a class of protocols of the form $(\mathcal{I}, \mathcal{F})$ where for some constant $\epsilon > 0$,*

$$P\{\mathcal{F}(f_i, w_i) = \alpha\} \begin{cases} \geq \min\{f_i(\alpha) + \epsilon, 1\} & \text{if } w_i(\alpha) > 1/2; \\ \leq \max\{f_i(\alpha) - \epsilon, 0\} & \text{otherwise.} \end{cases}$$

and \mathcal{I} is allowed to depend on w_i in an arbitrary way.

Definition 2 is a generic one which only defines biased updating function \mathcal{F} qualitatively without specifying how exactly it is computed. A natural choice is where each agent plays α with probability proportional to the product $f_i(\alpha)w_i(\alpha)$ [8]. In this model an agent balances their preferences with the behavior of their

neighbors in a simple multiplicative fashion and we call this the *multiplicative* biased voter model.

We note the extension to the biased voter model in Definition 2 is fairly general in that \mathcal{F} is allowed to include a broad class of local stochastic updating rules that reflect a node’s preferences; and \mathcal{I} is allowed to be *arbitrary* although it has to be independent of G . These seemingly provide us with a lot of power in the design of protocols; but perhaps surprisingly, in this section we prove that even this broad class of biased voting rules is insufficient to solve the DPP:

Theorem 3. *No biased voter model protocol $(\mathcal{I}, \mathcal{F})$ solves the DPP.*

The rest of this section is organized as follows. In Section 4.1, we prove a technical lemma about a certain Markov chain that can be represented by a line graph. We then use this lemma to prove Theorem 3 in 4.2, by constructing an example where for any voter model protocol $(\mathcal{I}, \mathcal{F})$, it either takes exponential time to converge to the globally preferred color, or convergence is to the globally non-preferred color, both violations of the DPP requirements.

4.1 A Markov Chain Lemma

Consider a Markov Chain on a line graph of n nodes, namely s_1, s_2, \dots, s_n , where transition does not happen beyond adjacent nodes. In this subsection we want to show that if at any state s_i ($1 < i < n$), the Markov chain is more likely to go ‘backward’ to state s_{i+1} than to go ‘forward’ to state s_{i-1} , then starting from state s_i (where $i \geq 2$), it takes exponential time in expectation to hit state s_1 . While this is perhaps intuitive, we will need this result to be in a particular form for the later reduction.

Here are a couple of notations: Let $p_{i,j}$ ($i, j \in [n]$) be the transition probability from node i to j , by construction $p_{i,j} = 0$ if $|i - j| > 1$. Simplify notation by writing $p_i = p_{i,i-1}$ and $q_i = p_{i,i+1}$, which are the ‘forward’ and ‘backward’ transition probability, respectively. Define h_i to be the expected number of rounds for the process to hit state s_1 for the first time, given that it starts from state s_i . Let $\gamma_{max} = \max_{i \in \{2,3,\dots,n-1\}} \frac{q_i}{p_i}$ and $\gamma_{min} = \min_{i \in \{2,3,\dots,n-1\}} \frac{q_i}{p_i}$, we have the following lemma.

Lemma 1. *If $\gamma_{min} \geq 1 + \epsilon$ for some constant $\epsilon > 0$, then h_i ($i \geq 2$) is exponential in n .*

Proof. We first claim that $h_i - h_{i-1} > \frac{\gamma_{min}^{n-i}}{p_n}$. To prove this claim, note h_i satisfies the following linear system

$$h_i = \begin{cases} 0 & (i = 0) \\ 1 + p_i h_{i-1} + q_i h_{i+1} + (1 - p_i - q_i) h_i & (2 \leq i \leq n - 1) \\ 1 + p_n h_{n-1} + (1 - p_n) h_n & (i = n) \end{cases}$$

It is clear $h_j - h_{j-1} > 0$ for all $j > 1$ as a process starting from state s_j has to hit s_{j-1} before hitting s_1 . Let $h_j - h_{j-1} = \lambda_j$, combining it with $h_{j-1} =$

$1 + p_{j-1}h_{j-2} + q_{j-1}h_j + (1 - p_{j-1} - q_{j-1})h_{j-1}$ gives $h_{j-1} - h_{j-2} = \frac{1+q_{j-1}\lambda_j}{p_{j-1}}$, which in turn implies $\lambda_{j-1} > \left(\frac{q_{j-1}}{p_{j-1}}\right)\lambda_j > \gamma_{min}\lambda_j$. Repeating this inductively gives $\lambda_i > \gamma_{min}^{n-i}\lambda_n$. Since $\lambda_n = \frac{1}{p_n}$, this proves the claim.

Immediately following from this claim, we have $h_2 = h_2 - h_1 > \frac{\gamma_{min}^{n-2}}{p_n} \geq (1 + \epsilon)^{n-2}$ if $\gamma_{min} \geq 1 + \epsilon$. Since $h_i > h_2$ whenever $i > 2$, this completes the proof. □

4.2 The Impossibility Result

Our goal in this subsection is to prove Theorem 3. To this end, first consider the biased voter model on the following 3-regular line graph.

A Line Graph. *G is a line graph of 2n nodes, where the left half prefers blue and the right half prefers red. The leftmost and rightmost node each has two self-loops and all the other nodes have one self-loop.*

We prove two lemmas (Lemma 2 and 3) about this particular setting, as a preparation for the proof of the main theorem.

Lemma 2. *For any biased voter model protocol (I, F), given that I results in an initialization where all nodes initialized to blue are to the left of all nodes initialized to red, it takes exponential time in expectation to reach a consensus on the line.*

Proof. We prove this by reducing this stochastic process to the Markov process described in Section 4.1. First observe that since we start with a coloring where all blues are to the left of all reds, this will hold as an invariant throughout the evolution of the whole process and the only way for the coloring to evolve is for the blue node adjacent to a red neighbor to convert to red, or for its red neighbor to convert to blue.

Therefore, we can always describe the state by a pair of integers $(b, 2n - b)$, where b is the number of blue-colored nodes. Now if we lump two states, $(b, 2n - b)$ and $(2n - b, b)$, into one, this model is exactly the Markov process (with $n + 1$ states) described in Section 4.1 with $s_i = \{(i, 2n - i), (2n - i, i)\}$ for $i = \{0, 1, \dots, n\}$.

And by definition of *biased voter model* and the way the Markov chain is constructed in the above, we have $p_i \leq \frac{1/3-\epsilon}{2n}$ and $q_i \geq \frac{1/3+\epsilon}{2n}$ ($i \in \{1, 2, \dots, n-1\}$) for some constant $\epsilon > 0$. Therefore, $\gamma_{min} \geq \frac{1/3+\epsilon}{1/3-\epsilon} = 1 + \delta$, for some constant $\delta > 0$. Invoking Lemma 1 shows that it takes exponential time to hit s_0 starting from state s_i (where $i \geq 1$). Therefore, it takes exponential time to reach a consensus given that one starts with a coloring where all nodes initialized to blue are to the left of all nodes initialized to red. □

Lemma 3. *For any biased voter model protocol (I, F), if I initializes a node i to his preferred opinion with positive probability, then it takes exponential time in expectation to reach a consensus on the line.*

Proof. Note \mathcal{I} is independent of G , therefore whenever it initializes with positive probability, the probability is independent of n . In particular, the probability that \mathcal{I} initializes the leftmost node to *blue* and the rightmost node to *red* is not negligible. Therefore we are through if we can show that *given* the leftmost node is initialized to *blue* and the rightmost to *red*, it takes exponential time to reach a consensus.

Lemma 2 does not differentiate between a b -consensus and a r -consensus. If we concern ourselves only with the outcome of, say a b -consensus, it can be shown that it still takes exponential time to reach given that we start from the same initialization described in Lemma 2 (i.e. all blues are to the right of all reds). We prove this by first observing that, conditioning on that a b -consensus is reached, the time taken is distributed exactly the same as in the modified stochastic process on the same $2n$ -node line graph, with the only difference being making the leftmost node extremely biased towards *blue* so that it always votes for *blue* regardless of his neighbor's opinion. Therefore we only need to prove it takes exponential time for this modified process to reach a consensus (which can only be a *blue* one), and this follows from Lemma 2.

Of course by a similar argument we can show that starting from an initialization where all blues are to the left of all reds, it takes exponential time to reach a r -consensus.

Now consider the initialization where the leftmost node is *blue* and rightmost node is *red* and call this the case of interest. Compare it with the initialization where the leftmost node is *blue* and all the other $n - 1$ nodes are *red*, the r -consensus time of this case is clearly upper bounded by that of the case of interest. By the above discussion, it takes exponential time to reach a r -consensus even when we start with only the leftmost node *blue*; therefore, it takes exponential time to reach a r -consensus for the case of interest. By the same argument, it also takes exponential time to reach a b -consensus for the case of interest. In sum this allows us to conclude that it takes exponential time to reach a consensus given that \mathcal{I} initializes the leftmost node to *blue* and the rightmost to *red*. \square

We are now ready to give a proof for Theorem 3.

Proof. (of Theorem 3) In Lemma 3, we have already shown that any biased voter model protocol $(\mathcal{I}, \mathcal{F})$ fails to solve the DPP (taking exponential time to converge) if we restrict \mathcal{I} to the kind of initialization functions that initializes a node to its preferred opinion with positive probability. It is easy to see that $(\mathcal{I}, \mathcal{F})$ also fails for any \mathcal{I} that does the opposite, in which case \mathcal{I} initializes a node to his not-preferred opinion with probability 1: Simply construct a graph consists solely of nodes that prefer *blue*, and \mathcal{I} initializes it to a r -consensus. Therefore, we conclude that *any* biased voter model protocol $(\mathcal{I}, \mathcal{F})$ fails to solve the DPP. \square

Note since the line graph we construct above is 3-regular, we have actually shown a stronger version of Theorem 3: Even if we allow both \mathcal{I} and \mathcal{F} to depend on d_i , no protocol $(\mathcal{I}, \mathcal{F})$ can solve the DPP. We also note that a similar exponential convergence result can be shown for clique in certain settings.

5 A Protocol for DPP

Previous discussions establish the limitation of the classic voter model protocol and its natural extension to the biased voter model when it comes to solving DPP. We are thus interested in the question: What are the (ideally minimal) extensions to the biased voter model that are needed to obtain a simple, efficient and local protocol for solving the DPP?

In this section, we give one answer to this question, by providing a provable solution to the DPP that employs the following extensions:

1. Introduction of a third choice of opinion, *undecided*, or u for short;
2. Allowing initialization and evolution of a node's opinion be dependent on its degree in G ;
3. Allowing multiple identical copies of the protocol to be run in G and having each node vote for the opinion (between *blue* and *red* only, and ignoring *undecided*) converged to more frequently among the multiple runs. This can be implemented by having a slightly more powerful schedule that after every n^5 steps, re-initializes each node.

It is interesting to note that at least two of these extensions — the ability to temporarily declare oneself undecided, and the notion of an election that is run in multiple phases — have obvious analogues in many actual political processes. In any case, we would argue that our protocol is natural in the sense that it obviously does not engage in any of the kind of coding or signalling schemes mentioned in the Introduction.

We give the protocol in Algorithm 1. This protocol consists of $T = poly(n)$ phases. In each phase, each node simultaneously and independently initializes his opinion to either b , r or u according to some probabilities before launching into the standard unbiased voter model process. These initialization probabilities are properly chosen so that the probability of reaching an α -consensus ($\alpha \in \{b, r\}$) is proportional to $\sum_{i \in V} w_i(\alpha)$. The introduction of the *undecided* opinion u is to allow individuals of high degree deliberately reduce their potentially strong

Algorithm 1. A Simple and Local Voting Protocol

1. Each node i maintains an array R_i of size T
 2. **for** $phase = 1$ to T **do**
 3. Each node i simultaneously and independently initializes its color to b , r and u with probability $\frac{w_i(b)}{d_i}$, $\frac{w_i(r)}{d_i}$ and $\frac{d_i-1}{d_i}$, respectively;
 4. Run the (standard) *voter model* process (on opinion b, r, u) for n^5 rounds
 5. Each node i records his last round opinion of this phase of the voter model process in $R_i[phase]$
 6. **end for**
 7. // Each node now has a record of his local 'outcomes' of all T phases
 8. Each node i ignores all entries in R_i that record u ; among the remaining entries, identify a majority between b and r , breaking ties arbitrarily
 9. Each node i vote for this majority identified as his *final vote*
-

influence on the outcome (We note that in a strategic or game-theoretic setting, high-degree individuals might of course exactly wish to exploit this influence, a topic we examine in Section 6). At the end of each phase, the standard unbiased voter model process is run for n^5 rounds and each node records his opinion in the last round as the ‘outcome’ of this phase. After T phases, each node ignores all phases where his local outcome is *undecided* and identify the majority between *blue* and *red* among the remaining outcomes; he then vote for this majority as his *final vote*.

We now proceed to prove that Algorithm 1 indeed solves the DPP. First we need the following lemma.

Lemma 4. *In each of the T phases of Algorithm 1, with error exponentially small in n , the probability of reaching a b -consensus, r -consensus and u -consensus are $\frac{\sum_{i \in V} w_i(b)}{2m}$, $\frac{\sum_{i \in V} w_i(r)}{2m}$ and $\frac{2m-n}{2m}$, respectively.*

Proof. We give proof for the case of a b -consensus, and the proof for r -consensus and u -consensus follows a similar argument.

Let $B \subseteq V$ be the set of nodes initialized to *blue*, and $P_b = \sum_{B \in 2^V} p(B)p(b | B)$ the probability that a single phase of Algorithm 1 results in a b -consensus. By Theorem 2 we have $|p(b | B) - \pi(B)| = o(c^{-poly(n)})$ for some constant c , or $p(b | B) = \pi(B) \pm o(c^{-poly(n)})$, therefore $P_b = \sum_{B \in 2^V} p(B)p(b | B) = \sum_{B \in 2^V} p(B)(\pi(B) \pm o(c^{-poly(n)})) = \sum_{B \in 2^V} p(B)\pi(B) \pm o(c^{-poly(n)}) = \sum_{k=0}^{2m} p(d(B) = k) (k/2m) \pm o(c^{-poly(n)}) = \frac{E(d_B)}{2m} \pm o(c^{-poly(n)}) = \frac{\sum_{i \in V} w_i(b)}{2m} \pm o(c^{-poly(n)})$. Therefore, we conclude that $|P_b - \frac{\sum_{i \in V} w_i(b)}{2m}|$ is negligible. \square

Recall our goal is to let V converge to the collectively preferred consensus. And by definition of DPP one opinion is significantly preferred than the other, i.e. $|\sum_{i \in V} w_i(b) - \sum_{i \in V} w_i(r)| \geq \epsilon$ for some constant ϵ ; this assumption turns out to be sufficient for Algorithm 1 to achieve this goal if we set $T = poly(n)$ sufficiently large.

Theorem 4. *Setting $T = O(n^3 \log(n))$ in Algorithm 1 solves the DPP.*

Proof. By Lemma 4 we have $P_b \geq \frac{\sum_{i \in V} w_i(b)}{2m} - o(c^{-poly(n)})$ and $P_r \leq \frac{\sum_{i \in V} w_i(r)}{2m} + o(c^{-poly(n)})$. Therefore the gap between P_b and P_r is at least $\frac{\epsilon}{2m} - o(c^{-poly(n)})$, so there exists a positive constant $\delta < \epsilon$ such that the gap between P_b and P_r is at least $\frac{\delta}{2m}$ whenever n is sufficiently large.

Let T_b and T_r be the number of b -consensuses and r -consensuses among the T trials, the bad event happens when $T_b < T_r$. For this bad event to happen, either event $T_b < (P_b - \frac{1}{3} \cdot \frac{\delta}{2m}) T$ or event $T_r > (P_r + \frac{1}{3} \cdot \frac{\delta}{2m}) T$ has to happen. Since $\frac{n}{4m} < P_b < \frac{n}{2m}$ and $P_r < \frac{n}{4m}$, by applying Chernoff bound, it can be shown that $T = O(n^3 \log(n))$ is sufficiently large to guarantee that both of the above two cases happen with negligible probability. \square

Before closing this section, we note that there is an alternative protocol that is a natural variant of Algorithm 1. In this variation, we do not need to introduce

the undecided opinion u , instead we make the degree of G , $d(G) = \max_{i \in V} d_i$, an input to \mathcal{F} . Now that each node is aware of $d(G)$, he can increase his own influence by (conceptually) adding $d(G) - d_i$ self-loops. When each node does so, the graph becomes regular and we can now simply have each node initialize to opinion $\alpha \in \{b, r\}$ with probability $w_i(\alpha)$ and then run the voter model protocol. Using essentially the same analysis it can be shown that this alternative protocol also solves the DPP if repeated sufficiently many times.

6 An ϵ -Nash Protocol for Democratic Primary Game

Our protocol for solving DPP assumes that each individual will actually follow the protocol honestly. However in a strategic setting, an individual may have incentives to deviate from the prescribed protocol. For example, a node i who prefers *blue* may deviate from Algorithm 1 in a way that increases the chance of reaching a *blue*-consensus, even when this consensus is not collectively preferred.

This naturally leads us to consider the Democratic Primary Game (DPG), which is an extension of DPP to the strategic, or game theoretic, setting. In DPG, a node with preference $(w_i(b), w_i(r))$ receives payoff $w_i(b)$ (resp., $w_i(r)$) if the game results in an unanimous global *blue*-consensus (resp., *red*-consensus) and payoff 0 if no consensus is reached. A solution to DPG is a protocol that solves the DPP (which must be *simple and local* and in polynomial time converge to the collectively preferred consensus with high probability) and at the same time is *strategy-proof*, i.e. each node honestly following the protocol constitutes a Nash equilibrium of the game. We note that DPG may also be viewed as a distributed, networked version of the classic “Battle of the Sexes” game, or as a networked coordination game [7].

In the rest of this section, we show the existence of a protocol that is an ϵ -approximate Nash equilibrium, or ϵ -Nash for short, of DPG. This means although a node can deviate unilaterally from this protocol and increases his expected payoff, the amount of this increase is at most ϵ and we show ϵ is negligible in n and can be made arbitrarily small. To this end, we need to make the following mild assumptions.

1. The removal of any node from G leaves the remaining graph connected. Formally, let G_{-i} be the graph induced by $V \setminus \{i\}$, we assume G_{-i} is connected for all $i \in V$.
2. The exclusion of any node does not change the collectively preferred consensus, and moreover, it still leaves a significant (constant) gap between $\sum_{j \in V(G_{-i})} w_j(b)$ and $\sum_{j \in V(G_{-i})} w_j(r)$.
3. Each node i is identified by a unique ID, $ID(i)$, which is an integer in $\{1, 2, \dots, n\}$.

Our ϵ -Nash protocol consists of n runs of the non-Nash protocol Algorithm 1, each on a subgraph G_{-i} . Each run of Algorithm 1 polls the majority opinion of $V \setminus \{i\}$, which by assumption is the same as that of V ; however by excluding i from participating, we prevent him from any manipulation of this particular

Algorithm 2. A Simple and Local Protocol that is ϵ -Nash

-
1. Each node i maintains an array E_i of size $n - 1$
 2. **for** $episode = 1$ to n **do**
 3. Let i be the node such that $ID(i) = episode$
 4. Run Line 1 - 8 of Algorithm 1 on G_{-i}
 5. Each node $j \in V \setminus \{i\}$ records in $E_j[episode]$ the majority between b and r he identifies on Line 8 of (this run of) Algorithm 1
 6. **end for**
 7. // Each node has now participated in $n - 1$ runs of Algorithm 1
 8. **for all** $i \in V$ **do**
 9. **if** both b and r are present in the $n - 1$ entries of E_i **then**
 10. Tossing a private fair coin to decide between b and r , and vote for it as i 's *final vote*
 11. **else**
 12. Vote for the only opinion present as i 's *final vote*
 13. **end if**
 14. **end for**
-

run of the non-Nash protocol. When all the n runs of non-Nash is done, each node ends up with $n - 1$ ‘polls’ and with high probability they should all point to the same collectively preferred consensus. In case it does not, it is strong evidence that some run(s) of the non-Nash protocol had been manipulated and the contingency plan is for each node to ignore all the polling results entirely and toss a (private) fair coin to decide whether to vote for *blue* or *red* — and this turns out to be a sufficient deterrent of unilateral deviation from the non-Nash protocol.

We note conceptually we are making yet another simple extension in the protocol’s expressiveness by allowing it to be run on a subgraph G_{-i} . To implement this, it is important for each node i to be uniquely identified by his neighbors so that they know when to ignore i ; and this is the reason we need assumption 3 listed above. We give this ϵ -Nash protocol in Algorithm 2 and claim the following theorem.

Theorem 5. *Algorithm 2 approximately solves DPG by being an ϵ -Nash equilibrium.*

Proof. Suppose each node follows the protocol faithfully, by our assumption that the exclusion of any node does not change the collectively preferred consensus, say *blue*, the n runs of Algorithm 1 must have all resulted in a b -consensus with high probability. Therefore the final votes result in the collectively preferred b -consensus with high probability.

Now we examine why faithfully executing this protocol is an ϵ -Nash strategy for each node. For a node i that prefers *red* (i.e. the opinion not collectively preferred), assuming everyone else is following Algorithm 2, the expected payoff to i for doing the same is at least his payoff in a b -consensus minus a number negligible in n (because there is a negligible probability that no consensus is reached in the final voting even if every node follows Algorithm 2 faithfully).

Now we consider what happens if he deviates. There are two stages during which i can deviate: the first or the second for-loop in Algorithm 2. i 's effort during the first for-loop is obviously futile if none of the $n - 1$ runs of Algorithm 1 is turned into a r -consensus, and in this case, with high probability all the n runs of Algorithm 1 result in a b -consensus. Therefore, i will have no incentive to deviate during the second for-loop because everyone else is going to vote for b .

Next consider the case where i successfully turns some of the global outcomes of Algorithm 1 into a r -consensus (i.e. all nodes identify r as the majority on Line 8 of Algorithm 1), then with high probability the n runs of Algorithm 1 result in both r -consensus and b -consensus because the single run of it without i participating results in a b -consensus with high probability. In this case, at least $n - 2$ nodes out of $V \setminus \{i\}$ see both *blue* and *red* as outcomes from the $n - 1$ runs of Algorithm 1 they each participated in and will vote for either b or r by tossing a private fair coin, which means the probability of reaching a b -consensus or r -consensus among them, independent of whatever strategy i adopts in the second for-loop, is $(\frac{1}{2})^{(n-2)}$. Therefore, no matter what strategy i adopts in the second for-loop, his expected payoff is negligible and obviously worse than what he would have gotten by not deviating. Therefore, we conclude that executing Algorithm 2 faithfully is actually a Nash strategy for i .

Now consider a node j who prefers a b -consensus. By the same discussion as above, Algorithm 2 results in a b -consensus in the final voting with high probability, therefore the expected payoff to j is at least his payoff in a b -consensus minus a negligible number. Therefore by deviating j can only hope to improve his expected payoff by a negligible amount. And this allows us to conclude that each node following Algorithm 2 faithfully constitutes an ϵ -Nash equilibrium for the game, where ϵ is a negligible number and can be made arbitrarily small. \square

Finally, we note that it is possible for one to construct a distributed protocol that is a Nash equilibrium for DPG, by employing cryptographic techniques developed for distributed computation in a recent work by Kearns et al. [8]. Although the resulting protocol is highly distributed and uses only local information, its use of cryptographic tools, including the broadcast of public keys and secure multiparty function computations, violates our goal of finding simple protocols of the kind we have examined here.

Acknowledgement. We thank Tanmoy Chakraborty for helpful conversations on the work presented here.

References

1. Aldous, D., Fill, J.: Reversible Markov Chains and Random Walks on Graphs. Draft (1994)
2. Cliford, P., Sudbury, A.: A model for spatial conflict. *Biometrika* 60(3), 581–588 (1973)
3. Even-Dar, E., Shapira, A.: A Note on Maximizing the Spread of Influence in Social Networks. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 281–286. Springer, Heidelberg (2007)

4. Granovetter, M.: Threshold Models of Collective Behavior. *American Journal of Sociology* 83(6), 1420–1443 (1978)
5. Holley, R.A., Liggett, T.M.: Ergodic Theorems for Weakly Interacting Infinite Systems and the Voter Model. *Annals of Probability* 3, 643–663 (1975)
6. Lovász, L.: Random Walks on Graphs: A Survey. In: *Combinatorics, Paul Erdős is Eighty*, Budapest. János Bolyai Mathematical Society, vol. 2, pp. 353–398 (1996)
7. Luce, R., Raiffa, H.: *Games and Decisions*. Wiley, Chichester (1957)
8. Kearns, M., Tan, J., Wortman, J.: *Network-Faithful Secure Computation* (preprint, 2007)
9. Kearns, M., Wortman, J.: Learning from Collective Behavior. In: *Proceedings of COLT 2008* (2008)
10. Kleinberg, J.: Behavior in Networks: Algorithmic and Economic Issues. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*, pp. 613–632. Cambridge University Press, Cambridge (2007)
11. Motwani, R., Raghavan, P.: *Randomized Algorithms*. Cambridge University Press, Cambridge (1996)
12. Schelling, T.: *Micromotives and Macrobehavior*. W.W. Norton (1978)
13. Zeleny, J.: Working Together, Obama and Clinton Try to Show Unity. *New York Times* (June 28, 2008),
<http://www.nytimes.com/2008/06/28/us/politics/28unity.html>

Frequent Manipulability of Elections: The Case of Two Voters

Shahar Dobzinski^{1,*} and Ariel D. Procaccia^{2,**}

¹ Hebrew University of Jerusalem, Givat Ram, Jerusalem 91904, Israel
shahard@cs.huji.ac.il

² Microsoft Israel R&D Center, 13 Shenkar Street, Herzeliya 46725, Israel
arielpro@gmail.com

Abstract. The recent result of Friedgut, Kalai and Nisan [9] gives a quantitative version of the Gibbard-Satterthwaite Theorem regarding manipulation in elections, but holds only for neutral social choice functions and three alternatives. We complement their theorem by proving a similar result regarding Pareto-Optimal social choice functions when the number of voters is two. We discuss the implications of our results with respect to the agenda of precluding manipulation in elections by means of computational hardness.

1 Introduction

Can we design a good voting rule that is immune to manipulation? That is, one in which the best strategy for each voter is to report its true preferences, without taking into account complicated strategic issues (“my first-ranked alternative has no chance of winning, so perhaps I should vote for my second best option”)? The classic result of Gibbard and Satterthwaite [10,16] gives us an unfortunate answer: every voting rule that is immune to manipulation must be dictatorial. The question we ask in this paper is: *is there a reasonable voting rule that is mostly immune to manipulation?* That is, can we find a voting rule that cannot be manipulated “most” of the time?

Let us discuss this problem more formally. The basic ingredients of a voting setting are a set of *voters* N , $|N| = n$, and a set of *alternatives* A , $|A| = m$. The preferences of each voter are represented by a ranking of the alternatives, which is the private information of the voter. The collection of the preferences of all the voters is known as a *preference profile*. The setting also consists of a *social choice function* (SCF), which is simply a voting rule: a function that receives the preference profile submitted by the voters, and outputs the winning alternative.

* The author is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities, and by a grant from the Israeli Academy of Sciences.

** This work was done while the author was at the Hebrew University of Jerusalem, and was supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

Motivation and Related Work. Ideally, one would like to design SCFs that are *strategyproof*, i.e., theoretically immune to manipulation. A voter is said to *manipulate* the election if misreporting its preferences improves the outcome (from the voter's point of view). Unfortunately, as mentioned above, the seminal impossibility result of Gibbard [10] and Satterthwaite [16] states that if there are at least three alternatives, an SCF that is strategyproof and onto A must be a dictatorship, in the sense that there is a single voter whose favorite alternative is elected under any preference profile. This devastating theorem (hereinafter, the G-S Theorem) implies that, in theory, it is impossible to design a "reasonable" SCF that is strategyproof.

Nevertheless, several avenues have been suggested for circumventing the G-S Theorem. One approach, introduced by Bartholdi, Tovey and Trick [1], is employing computational complexity. Indeed, Bartholdi et al. suggested that some of the prominent SCFs may be computationally hard to manipulate. The working hypothesis is that, if successfully lying is computationally infeasible, voters would simply report the truth. Since then, and especially in recent years, numerous results about the worst-case hardness of manipulation have been published (see, e.g., [2,4,5,6,8,11,15]).

The foregoing line of work is encouraging, and doubtless being worst-case hard to manipulate is a desirable property in an SCF. However, researchers have pointed out that worst-case hardness may not be a sufficient barrier against manipulation. What one would ideally wish for is an SCF that is *almost always* hard to manipulate, when the instances are drawn according to typical distributions; this notion of hardness of manipulation is closer to the cryptographic notions of hardness.

Recent works have argued that common SCFs are not frequently hard to manipulate with respect to typical distributions. An algorithmic approach to this issue was presented by Procaccia and Rosenschein [14]. This work relies on the arguable (as discussed by Erdélyi et al. [7]) concept of junta distributions, and only deals with manipulation by coalitions in a specific family of SCFs when the number of alternatives is constant. The algorithmic results of Procaccia and Rosenschein were later significantly strengthened by Zuckerman, Procaccia and Rosenschein [20], but this work also deals with specific SCFs and coalitional manipulation. Another algorithmic, general, approach was introduced by Conitzer and Sandholm [3], but in order to apply their results, the SCF has to satisfy a somewhat restrictive property. This property is empirically shown to hold with respect to some SCFs, when the number of alternatives is very small. Yet another approach was proposed by Procaccia and Rosenschein [13], and generalized by Xia and Conitzer [18]. This approach, once again, only deals with manipulation by coalitions and a constant number of alternatives.

An intriguing and ambitious approach to the issue of frequency of hardness in manipulation was presented by Friedgut, Kalai and Nisan [9]. They looked at a setting where each voter votes independently and uniformly at random; this is known in the social choice literature as the *impartial culture assumption*. Friedgut et al. suggested that, under the impartial culture assumption, a potential manipulator

can improve the outcome of the election with nonnegligible probability by simply reporting a random ranking instead of its truthful ranking.

Roughly speaking, define an SCF f to be δ -dictatorial if one must change the outcome of f on at most a δ -fraction of the preference profiles in order to transform f into a dictatorship. We call f ϵ -strategyproof if, given a random profile, a random manipulation succeeds with probability at most ϵ . Friedgut et al. suggested that the following quantitative version of the G-S Theorem is true: If $|A| \geq 3$, then any ϵ -strategyproof SCF that is onto A (and possibly satisfies additional weak properties) is δ -dictatorial, for $\delta = K_1 \cdot \text{poly}(n, m) \cdot \epsilon^{1/K_2}$, where K_1 and K_2 are constants. Such a result would directly imply that if a random manipulation succeeds with only negligible probability, namely ϵ is superpolynomially small, then the SCF must be very close to being dictatorial, that is unreasonable from a social choice point of view. Hence, if this statement is true, it would be of supreme importance to the frequency-of-manipulation agenda.

Friedgut et al. themselves were only able to prove the above result under the assumptions that there are exactly three alternatives, and that the SCF is neutral, i.e. indifferent to the identities of the alternatives. The techniques of Friedgut et al. are beautiful, but it seems to be very difficult to generalize their proof to more than just 3 alternatives. Strictly speaking, neutrality might also be undesirable, since neutrality and anonymity (indifference to the identities of the voters) are sometimes mutually exclusive [12, page 25], and all prominent SCFs are anonymous.

Xia and Conitzer [19] extended the result of Friedgut et al. (via a completely different, involved line of reasoning) to any number of alternatives. However, the quality of their result decreases rapidly with the number of alternatives, so the authors assume that the number of alternatives is constant in order to achieve the ideal of Friedgut et al. In addition, Xia and Conitzer require several very technical and restrictive assumptions with respect to the SCF. Although they show that these assumptions are satisfied by most (but not all) prominent SCFs, the assumptions still severely limit the scope of their result when it comes to the possibility of designing nonstandard SCFs that are usually hard to manipulate.

Our result. We complement the two previous results along the line of work proposed by Friedgut et al. by establishing the desired quantitative version of G-S for an arbitrary number of alternatives m but $n = 2$, namely only two voters. The only assumption we make is that the SCF is Pareto-optimal, i.e., if all voters rank alternative a above b , then b is not elected. Specifically, we prove:

Main Theorem. *Let $\epsilon < \frac{1}{32m^9}$; assume $N = \{1, 2\}$, $m \geq 3$, and let f be an ϵ -strategyproof and Pareto-optimal SCF. Then f is $16m^8\epsilon$ -dictatorial.*

In particular, if the probability of success of a random manipulation is negligible, then f is very close to being dictatorial. As Pareto-Optimality is a very basic requirement, this directly implies that it is impossible to design a reasonable SCF that is frequently hard to manipulate, when each voter votes independently and uniformly at random and $N = \{1, 2\}$.

Discussion. A crucial aspect of our theorem is that it seems to be better than previous results as a first step towards a more general result. Indeed, the proof of Friedgut et al. is fascinating but involved and relies on heavy mathematical machinery: Fourier analysis, isoperimetric inequalities, and so on. The proof of Xia and Conitzer seems to strongly rely on their assumptions, and it is not clear if the same techniques can be used once these assumptions are removed.

On the other hand, our proof is relatively simple and is built “from scratch”. More importantly, Svensson [17] gives an inductive argument that extends the deterministic proof of G-S from two voters to n voters. However, this argument is not “robust”, in the sense that using it directly causes too great a deterioration in the quality of the result with respect to n and m . Certainly, new tricks are needed, but we believe that using clever induction on the number of voters in order to achieve a general result should be possible.

We wish to make some remarks regarding the generality of our result. First, we assume Pareto-optimality, but this assumption can probably be relaxed, since in the deterministic case Pareto-optimality is implied by strategyproofness. Second, our auxiliary monotonicity lemma (Lemma 1) can certainly be generalized to any number of voters n .

Let us briefly examine the significance of our result in its own right (and not as a first step towards a general result). The case of two voters and m alternatives might at first seem less important than the case of n voters and three alternatives that was considered by Friedgut et al. This is true in political elections (where one expects to find more voters than candidates), but not in general (and especially not in computer science). For instance, in settings where multiple agents must decide between joint plans or beliefs the number of alternatives is typically far greater than the number of voters. In addition, when the number of alternatives is constant, a potential manipulator can simply check all the possible rankings, so there is no question of computational complexity. The problem becomes more interesting when the number of alternatives is large, as it is in our case.

As a final remark, we wish to address the impartial culture assumption (voters vote independently and uniformly), also used by Friedgut et al. and Xia and Conitzer. Even if one proves the general quantitative version of G-S (as discussed above), it would not necessarily spell the end of the hardness of manipulation agenda. The rankings of voters are typically not independent nor uniform, but centered around specific strong alternatives. So, the underlying assumption that voters vote independently and uniformly at random may not be realistic. However, this assumption allows for elegant “lower bounds”, as noted by Friedgut et al. Ultimately, the ideal is to obtain results that also hold under a wide range of typical distributions.

Structure of the paper. In Section 2 we formally present the necessary notations and definitions. In Section 3, we formulate and prove our main result.

2 Preliminaries

We deal with a finite set of voters $N = \{1, 2, \dots, n\}$, and a finite set of alternatives A , where $|A| = m$. We denote alternatives by letters such as a, b, c, x, y .

Each voter $i \in N$ holds a strict total order R^i over A , i.e. R^i is a binary relation over A that satisfies irreflexivity, antisymmetry, transitivity and totality. Informally, R^i is a *ranking* of the alternatives. The set $\mathcal{L} = \mathcal{L}(A)$ is the set of all such (linear) orders, so for all $i \in N$, $R^i \in \mathcal{L}$ throughout. A *preference profile* R^N is a vector $\langle R^1, \dots, R^n \rangle \in \mathcal{L}^N$. A *social choice function* (SCF) is a function $f : \mathcal{L}^N \rightarrow A$.

We make the *Impartial Culture Assumption* throughout the paper, that is, we assume that random preference profiles are drawn by independently and uniformly drawing a random ranking for each voter (each possible ranking has a probability of $1/m!$). So, for instance, when we write $\Pr_{R^N} [E]$ we refer to the probability that the event E occurs, when the preferences of each voter R^i are independently and uniformly distributed. Furthermore, when we write, e.g., $\Pr_{R^N, Q^1} [E]$, we mean that the preferences R^1, \dots, R^n and Q^1 are all drawn independently at random.

Definition 1. *Let f be an SCF. f is Pareto-optimal if for all $R^N \in \mathcal{L}^N$, if there exist $x, y \in A$ such that $xR^i y$ for all $i \in N$, then $f(R^N) \neq y$.*

We now define some probabilistic versions of well-known properties of SCFs.

Definition 2. *Let f be an SCF. Voter $i \in N$ is a δ -dictator with respect to $a \in A$ iff*

$$\Pr_{R^N} [f(R^N) \neq a \mid \forall x \in A \setminus \{a\}, aR^i x] \leq \delta.$$

Voter i is a δ -dictator iff it is a δ -dictator with respect to every $a \in A$. f is a δ -dictatorship if there exists a δ -dictator.

The classical definition of a dictatorship corresponds to the definition of a 0-dictatorship under this formulation. Also note that δ -dictatorship under our definition implies δ -far from dictatorship under the definition of Friedgut et al. [9].

Let us turn to a probabilistic definition of strategyproofness. An SCF f is *manipulable* at $R^N \in \mathcal{L}^N$ if there exists a voter $i \in N$ and a ranking Q^i such that $f(Q^i, R^{N \setminus \{i\}})R^i f(R^N)$, where $(Q^i, R^{N \setminus \{i\}})$ is identical to R^N except that R^i is replaced by Q^i . That is, voter i strictly benefits according to its true preferences R^i by reporting false preferences Q^i . An SCF is *strategyproof* if it is not manipulable at any $R^N \in \mathcal{L}^N$.

Definition 3. *An SCF f is ϵ -strategyproof iff for all voters $i \in N$,*

$$\Pr_{R^N, Q^i} [f(Q^i, R^{N \setminus \{i\}})R^i f(R^N)] \leq \epsilon.$$

So, strategyproofness corresponds to 0-strategyproofness according to this probabilistic definition. Our definition of ϵ -strategyproofness is exactly equivalent to all voters having manipulation power at most ϵ according to the definition given by Friedgut et al. [9].

The classic formulation of the Gibbard-Satterthwaite Theorem [10,16] is as follows.

Theorem 1 (Gibbard-Satterthwaite). *Assume $|A| \geq 3$, and let $f : \mathcal{L}^N \rightarrow A$ be a strategyproof SCF that is onto A . Then f is dictatorial.*

Finally, we wish to extend the classic definition of monotonicity. Let $R \in \mathcal{L}$, $a \in A$, and denote

$$I(R, a) = \{Q \in \mathcal{L} : \forall x \in A, aRx \Rightarrow aQx\}.$$

Now, let $R^N \in \mathcal{L}^N$, and denote

$$I(R^N, a) = \{Q^N \in \mathcal{L}^N : \forall i \in N, Q^i \in I(R^i, a)\}.$$

Definition 4. *Let f be an SCF. f is γ -monotonic if*

$$\Pr_{R^N, Q^N} [f(R^N) \neq f(Q^N) \mid Q^N \in I(R^N, f(R^N))] \leq \gamma.$$

In words, f is γ -monotonic if improving a winning alternative harms it with probability at most γ . Monotonicity is equivalent to 0-monotonicity. We wish to point out that monotonicity is a very strong property, as the order of other alternatives can change as long as the winner only improves with respect to other alternatives. In fact, monotonicity is closely related to, and implied by, strategyproofness.

3 Main Theorem

Our aim is to prove a quantitative version of the Gibbard-Satterthwaite Theorem (Theorem 1), under the assumptions that $N = \{1, 2\}$ and that the SCF in question is Pareto-optimal. Note that Pareto-optimality implies surjectivity, as if all the voters rank $x \in A$ first then x must be elected, and this is true for all $x \in A$.

Theorem 2. *Let $\epsilon < \frac{1}{32m^9}$; assume $N = \{1, 2\}$, $m \geq 3$, and let f be an ϵ -strategyproof and Pareto-optimal SCF. Then f is $16m^8\epsilon$ -dictatorial.*

We wish to stress once again that, as in Friedgut et al. [9] and Xia and Conitzer [19], the underlying assumption is the *impartial culture assumption*, that is the voters vote independently and uniformly at random.

Let us now turn to the proof of Theorem 2. The proof follows the lines of the proof of Theorem 1 in Svensson [17]. He gives a very simple and short proof of the G-S Theorem for $N = \{1, 2\}$. Our proof is considerably more involved, but ultimately our main mathematical contribution is to notice that all of Svensson’s arguments are robust, in the sense that they do not greatly restrict the space of preference profiles, and thus survive the transition to the quantitative version. The reader is encouraged to read Svensson’s proof before reading ours.

As noted above, the deterministic notion of strategyproofness implies the deterministic notion of monotonicity. We will require a lemma that gives a quantitative version of this implication. The lemma also presents in detail the type of robustness arguments that we employ throughout the proof of the Theorem.

Lemma 1 (Monotonicity). *Assume $N = \{1, 2\}$, and let f be an ϵ -strategyproof SCF. Then f is $4m^2\epsilon$ -monotonic*

Proof. Since f is ϵ -strategyproof, we have

$$\Pr_{R^N, Q^1} [f(Q^1, R^2)R^1 f(R^N)] \leq \epsilon. \tag{1}$$

We are now about to apply a critical “robustness” argument, which will be central to the proofs of both this lemma and Theorem 2. We first claim that

$$\Pr_{R^N, Q^1} [Q^1 \in I(R^1, f(R^N))] \geq 1/m. \tag{2}$$

Indeed, this is true since any ranking $Q^1 \in \mathcal{L}$ that places $f(R^N)$ on top is a member of $I(R^1, f(R^N))$, and there are $(m - 1)!$ such rankings out of the total $m!$ rankings.

Now, from the basic laws of probability it follows that for two events E_1 and E_2 ,

$$\Pr[E_1] = \Pr[E_1|E_2] \cdot \Pr[E_2] + \Pr[E_1|\neg E_2] \cdot \Pr[\neg E_2] \geq \Pr[E_1|E_2] \cdot \Pr[E_2],$$

and therefore

$$\Pr[E_1|E_2] \leq \frac{\Pr[E_1]}{\Pr[E_2]}. \tag{3}$$

Now, from (1), (2), and (3) we obtain:

$$\begin{aligned} \Pr_{R^N, Q^1} [f(Q^1, R^2)R^1 f(R^N) | Q^1 \in I(R^1, f(R^N))] &\leq \frac{\Pr_{R^N, Q^1} [f(Q^1, R^2)R^1 f(R^N)]}{\Pr_{R^N, Q^1} [Q^1 \in I(R^1, f(R^N))]} \\ &\leq m\epsilon, \end{aligned} \tag{4}$$

where the first inequality follows from (3) and the second inequality follows by using both (1) and (2). By using symmetric arguments and the union bound we have that:

$$\Pr_{R^N, Q^1} [f(Q^1, R^2)R^1 f(R^N) \vee f(R^N)Q^1 f(Q^1, R^2) \mid Q^1 \in I(R^1, f(R^N))] \leq 2m\epsilon. \tag{5}$$

Fix $R^N \in \mathcal{L}^N$ and $Q^1 \in I(R^1, f(R^N))$, and assume that strategyproofness holds “in both directions”, namely the event in (5) does not occur. Let $a = f(R^N)$ and $b = f(Q^1, R^2)$. Assume that $a \neq b$; by strategyproofness aR^1b , and since Q^1 is an improvement of a over R^1 , aQ^1b . Strategyproofness in the other direction implies that bQ^1a , which leads to a contradiction. Hence, $a = b$. To

summarize, we have shown that given that $Q^1 \in I(R^1, f(R^N))$, then $f(R^N) \neq f(Q^1, R^2)$ with probability at most $2m\epsilon$.

Let us extend our arguments to two steps of improvement instead of one. Analogously to (2), we have that:

$$\Pr_{R^N, Q^N} [Q^1 \in I(R^1, f(R^N)) \wedge Q^2 \in I(R^2, f(Q^1, R^2))] \geq \frac{1}{m^2}. \tag{6}$$

Now, similarly to (4) we conclude by ϵ -strategyproofness, (6) and (3) that:

$$\begin{aligned} & \Pr_{R^N, Q^N} [f(Q^N)Q^2 f(Q^1, R^2) | Q^1 \in I(R^1, f(R^N)) \wedge Q^2 \in I(R^2, f(Q^1, R^2))] \\ &= \frac{\Pr_{R^N, Q^N} [f(Q^N)Q^2 f(Q^1, R^2)]}{\Pr_{R^N, Q^N} [Q^1 \in I(R^1, f(R^N)) \wedge Q^2 \in I(R^2, f(Q^1, R^2))]} \\ &= \frac{\Pr_{R^2, Q^N} [f(Q^N)Q^2 f(Q^1, R^2)]}{\Pr_{R^N, Q^N} [Q^1 \in I(R^1, f(R^N)) \wedge Q^2 \in I(R^2, f(Q^1, R^2))]} \\ &\leq m^2\epsilon. \end{aligned}$$

The third equality simply drops R^1 in the probability; this is possible as the event is indifferent to the choice of R^1 . Hence, we can use ϵ -strategyproofness directly on the random preference profile (Q^1, R^2) and the random manipulation Q^2 by voter 2.

By repeating the arguments given above for a single improvement, we get that if we choose R^N and Q^N such that $Q^1 \in I(R^1, f(R^N))$ and $Q^2 \in I(R^2, f(Q^1, R^2))$, then $f(Q^1, R^2) \neq f(Q^N)$ with probability at most $2m^2\epsilon$.

Finally, we apply the union bound one last time to get:

$$\begin{aligned} & \Pr_{R^N, Q^N} [f(R^N) \neq f(Q^N) | Q^N \in I(R^N, f(R^N))] \\ &\leq \Pr_{R^N, Q^N} [(f(R^N) \neq f(Q^1, R^2)) \vee (f(R^N) = f(Q^1, R^2) \wedge f(Q^1, R^2) \neq f(Q^N)) \\ &\quad | Q^N \in I(R^N, f(R^N))] \\ &\leq \Pr_{R^N, Q^1} [(f(R^N) \neq f(Q^1, R^2)) | Q^1 \in I(R^1, f(R^N))] \\ &\quad + \Pr_{R^N, Q^N} [f(R^N) = f(Q^1, R^2) \wedge f(Q^1, R^2) \neq f(Q^N) | Q^N \in I(R^N, f(R^N))] \\ &= \Pr_{R^N, Q^1} [(f(R^N) \neq f(Q^1, R^2)) | Q^1 \in I(R^1, f(R^N))] \\ &\quad + \Pr_{R^N, Q^N} [f(R^N) = f(Q^1, R^2) \wedge f(Q^1, R^2) \neq f(Q^N) \\ &\quad | Q^1 \in I(R^1, f(R^N)) \wedge Q^2 \in I(R^2, f(Q^1, R^2))] \\ &\leq 2m\epsilon + 2m^2\epsilon \leq 4m^2\epsilon. \end{aligned} \tag{7}$$

The third transition follows from the fact that, given that $f(R^N) = f(Q^1, R^2)$ occurred, the events $Q^N \in I(R^N, f(R^N))$ and $Q^1 \in I(R^1, f(R^N)) \wedge Q^2 \in I(R^2, f(Q^1, R^2))$ are one and the same. Formally,

$$\begin{aligned}
 & \Pr_{R^N, Q^N} [f(R^N) = f(Q^1, R^2) \wedge f(Q^1, R^2) \neq f(Q^N) \mid Q^N \in I(R^N, f(R^N))] \\
 &= \Pr_{R^N, Q^N} [f(R^N) = f(Q^1, R^2) \mid Q^N \in I(R^N, f(R^N))] \\
 & \quad \cdot \Pr_{R^N, Q^N} [f(Q^1, R^2) \neq f(Q^N) \mid Q^N \in I(R^N, f(R^N)) \wedge f(R^N) = f(Q^1, R^2)] \\
 &= \Pr_{R^N, Q^N} [f(R^N) = f(Q^1, R^2) \mid Q^1 \in I(R^1, f(R^N)) \wedge Q^2 \in I(R^2, f(Q^1, R^2))] \\
 & \quad \cdot \Pr_{R^N, Q^N} [f(Q^1, R^2) \neq f(Q^N) \mid Q^1 \in I(R^1, f(R^N)) \wedge Q^2 \in I(R^2, f(Q^1, R^2))] \\
 & \quad \quad \quad \wedge f(R^N) = f(Q^1, R^2)] \\
 &= \Pr_{R^N, Q^N} [f(R^N) = f(Q^1, R^2) \wedge f(Q^1, R^2) \neq f(Q^N) \mid Q^1 \in I(R^1, f(R^N)) \\
 & \quad \quad \quad \wedge Q^2 \in I(R^2, f(Q^1, R^2))],
 \end{aligned}$$

where in the second equality above the two left hand side factors are equal since the event $f(R^N) = f(Q^1, R^2)$ is independent of the choice of Q^2 .

The last transition of (7) is true since the probability that

$$f(R^N) = f(Q^1, R^2) \wedge f(Q^1, R^2) \neq f(Q^N)$$

is bounded from above by the probability that $f(Q^1, R^2) \neq f(Q^N)$. □

We are now in a position to prove our main theorem.

Proof (of Theorem 2). Fix two distinct alternatives $a, b \in A$, and define

$$Z(a, b) = \{R^N \in \mathcal{L}^N : \forall x \in A \setminus \{a, b\}, aR^1bR^1x \wedge bR^2aR^2x\}.$$

That is, $Z(a, b)$ is the set of all preference profiles where voter 1 ranks a first and b second, and voter 2 ranks b first and a second. We have that

$$\Pr_{R^N} [R^N \in Z(a, b)] = \left(\frac{(m-2)!}{m!}\right)^2 = 1/m^4. \tag{8}$$

Now, for every $R^N \in Z(a, b)$, we have that $f(R^N) \in \{a, b\}$ from Pareto-optimality. Assume without loss of generality that at least a 1/2-fraction of the profiles in $Z(a, b)$ satisfy $f(R^N) = a$, that is

$$\Pr_{R^N} [f(R^N) = a \mid R^N \in Z(a, b)] \geq \frac{1}{2}. \tag{9}$$

For any $R^N \in Z(a, b)$ such that $f(R^N) = a$, let $Q^2 \in \mathcal{L}$ such that bQ^2xQ^2a for all $x \in A \setminus \{a, b\}$. Let $Y(a, b)$ be the set of all such ordered pairs (R^N, Q^2) , i.e.,

$$Y(a, b) = \{(R^N, Q^2) \in Z(a, b) \times \mathcal{L} : f(R^N) = a \wedge \forall x \in A \setminus \{a, b\}, bQ^2xQ^2a\}.$$

We have that

$$\begin{aligned}
 & \Pr_{R^N, Q^2} [(R^N, Q^2) \in Y(a, b)] \\
 &= \Pr_{R^N} [R^N \in Z(a, b) \wedge f(R^N) = a] \cdot \Pr_{Q^2} [\forall x \in A \setminus \{a, b\}, bQ^2xQ^2a] \\
 &\geq \frac{1}{2m^4} \cdot \frac{1}{m^2} = \frac{1}{2m^6},
 \end{aligned} \tag{10}$$

where the first equality is by the independence of the two events, and the inequality follows from (8) and (9).

At this point we appeal to ϵ -strategyproofness, and apply our recurring robustness argument, namely, in this case, (3) coupled with (10). This gives us:

$$\Pr_{R^N, Q^2} [f(R^1, Q^2)R^2 f(R^N) \mid (R^N, Q^2) \in Y(a, b)] \leq 2m^6\epsilon \leq 1/2,$$

where the last inequality follows from our choice of ϵ . Therefore,

$$\begin{aligned} &\Pr_{R^N, Q^2} [f(R^N)R^2 f(R^1, Q^2) \vee f(R^N) = f(R^1, Q^2) \mid (R^N, Q^2) \in Y(a, b)] \\ &= 1 - \Pr_{R^N, Q^2} [f(R^1, Q^2)R^2 f(R^N) \mid (R^N, Q^2) \in Y(a, b)] \geq 1/2. \end{aligned} \tag{11}$$

From Pareto-optimality we have that for any $(R^N, Q^2) \in Y(a, b)$, $f(R^1, Q^2) \in \{a, b\}$, and, if in addition we have that $f(R^N)R^2 f(R^1, Q^2)$ or $f(R^N) = f(R^1, Q^2)$, then $f(R^1, Q^2) = a$. Indeed, this is true since b is ranked first in R^2 , and by definition $f(R^N) = a$; hence, if $f(R^1, Q^2) = b$ then voter 2 gains by switching from R^2 to Q^2 .

Now, by applying (10) and (11), we obtain:

$$\Pr_{R^N, Q^2} [(R^N, Q^2) \in Y(a, b) \wedge f(R^1, Q^2) = a] \geq \frac{1}{4m^6}.$$

We are now in a position to show that when a preference profile is chosen at random, the probability of obtaining a profile where voter 1 ranks a first, voter 2 ranks a last, and the winner is a is significant. Indeed,

$$\begin{aligned} &\Pr_{R^1, Q^2} [(\forall x \in A \setminus \{a\}, aR^1 x \wedge xQ^2 a) \wedge (f(R^1, Q^2) = a)] \\ &\geq \Pr_{R^1, Q^2} [\exists R^2 \in \mathcal{L} \text{ s.t. } (R^N, Q^2) \in Y(a, b) \wedge f(R^1, Q^2) = a] \\ &\geq \Pr_{R^N, Q^2} [(R^N, Q^2) \in Y(a, b) \wedge f(R^1, Q^2) = a] \geq \frac{1}{4m^6}. \end{aligned} \tag{12}$$

Next, we are finally going to use Lemma 1. We have that

$$\begin{aligned} &\Pr_{R^N, Q^N} [f(Q^N) \neq f(R^N)] \\ &\quad | (\forall x \in A \setminus \{a\}, aR^1 x \wedge xR^2 a) \wedge (f(R^N) = a) \wedge (Q^N \in I(R^N, a))] \\ &\leq \frac{\Pr_{R^N, Q^N} [f(Q^N) \neq f(R^N) \mid Q^N \in I(R^N, a)]}{\Pr_{R^N} [(\forall x \in A \setminus \{a\}, aR^1 x \wedge xR^2 a) \wedge (f(R^N) = a)]} \\ &\leq 4m^2\epsilon \cdot 4m^6 = 16m^8\epsilon. \end{aligned}$$

The first inequality follows from (3), while the second inequality is obtained by applying Lemma 1 and (12). Therefore, there must be some R_0^N that satisfies for all $x \in A \setminus \{a\}$, $aR_0^1 x$ and $xR_0^2 a$, $f(R_0^N) = a$, and

$$\Pr_{Q^N} [f(Q^N) \neq f(R_0^N) = a \mid Q^N \in I(R_0^N, a)] \leq 16m^8\epsilon. \tag{13}$$

Crucially, since in R_0^N voter 1 ranks a first and voter 2 ranks a last, $I(R_0^N, a)$ is exactly the set of preference profiles such that voter 1 ranks a first. In other words, (I3) can be reformulated as:

$$\Pr_{Q^N} [f(Q^N) \neq a \mid \forall x \in A \setminus \{a\}, aQ^1x] \leq 16m^8\epsilon.$$

In words, voter 1 is a $\delta = 16m^8\epsilon$ -dictator with respect to a . If we had assumed that at least half the profiles in $Z(a, b)$ satisfied $f(R^N) = b$, we would have deduced that voter 2 is a $\delta = 16m^8\epsilon$ -dictator with respect to b .

So far, the analysis was for a fixed pair of alternatives $a, b \in A$. By repeating the analysis for every pair of alternatives, we may obtain two sets of alternatives A_1 and A_2 , such that A_i contains all the alternatives for which voter i is a $16m^8\epsilon$ -dictator. First notice that $A_3 = A \setminus (A_1 \cup A_2)$ satisfies $|A_3| \leq 1$, otherwise we could perform the analysis for two alternatives in A_3 and deduce that either the first is in A_1 or the second is in A_2 .

Second, we claim that for two distinct alternatives $a, b \in A$, it can't be the case that $a \in A_1$ and $b \in A_2$. Indeed, otherwise, by the assumption that $\epsilon < 1/(32m^9)$, voter 1 is less than a $1/2m$ -dictator for a , whereas voter 2 is less than a $1/2m$ -dictator for b . This directly implies that:

$$\begin{aligned} & \Pr_{R^N} [f(R^N) \neq a \mid (\forall x \in A \setminus \{a\}, aR^1x) \wedge (\forall x \in A \setminus \{b\}, bR^2x)] \\ & \leq \frac{\Pr_{R^N} [f(R^N) \neq a \mid \forall x \in A \setminus \{a\}, aR^1x]}{\Pr_{R^N} [\forall x \in A \setminus \{b\}, bR^2x]} < \frac{1}{2m} \cdot m = 1/2, \end{aligned}$$

and similarly

$$\Pr_{R^N} [f(R^N) \neq b \mid (\forall x \in A \setminus \{a\}, aR^1x) \wedge (\forall x \in A \setminus \{b\}, bR^2x)] < 1/2.$$

It follows that there exists a profile, where voter 1 ranks a first and voter 2 ranks b first, such that a and b are both winners, which is a contradiction to the definition of f as an SCF.

Now, since $|A_3| \leq 1$ and $m \geq 3$, we must have that one of A_1 or A_2 is empty (it is easily verified that otherwise there must be distinct $x, y \in A$ such that $x \in A_1$ and $y \in A_2$). Our early assumption that at least a $1/2$ -fraction of the profiles in $Z(a, b)$ satisfy $f(R^N) = a$ ultimately led to the conclusion that $a \in A_1$, thus it follows that $A_2 = \emptyset$.

To conclude the proof, we must show that $A_3 = \emptyset$. This is obvious, since if $c \in A_3$, we can repeat the analysis with the pair $\{c, a\}$, and get that either $c \in A_1$ or $a \in A_2$, which implies a contradiction. Hence, it must hold that $A_1 = A$, namely voter 1 is a $16m^8\epsilon$ -dictator. □

References

1. Bartholdi, J., Tovey, C.A., Trick, M.A.: The computational difficulty of manipulating an election. *Social Choice and Welfare* 6, 227–241 (1989)
2. Conitzer, V., Sandholm, T.: Universal voting protocol tweaks to make manipulation hard. In: *Proc. of 18th IJCAI*, pp. 781–788 (2003)

3. Conitzer, V., Sandholm, T.: Nonexistence of voting rules that are usually hard to manipulate. In: Proc. of 21st AAAI, pp. 627–634 (2006)
4. Conitzer, V., Sandholm, T., Lang, J.: When are elections with few candidates hard to manipulate? *Journal of the ACM* 54(3), 1–33 (2007)
5. Elkind, E., Lipmaa, H.: Hybrid voting protocols and hardness of manipulation. In: Deng, X., Du, D.-Z. (eds.) ISAAC 2005. LNCS, vol. 3827, pp. 206–215. Springer, Heidelberg (2005)
6. Elkind, E., Lipmaa, H.: Small coalitions cannot manipulate voting. In: S. Patrick, A., Yung, M. (eds.) FC 2005. LNCS, vol. 3570, pp. 285–297. Springer, Heidelberg (2005)
7. Erdélyi, G., Hemaspaandra, L.A., Rothe, J., Spakowski, H.: On approximating optimal weighted lobbying, and frequency of correctness versus average-case polynomial time. In: Csuhaj-Varjú, E., Ésik, Z. (eds.) FCT 2007. LNCS, vol. 4639, pp. 300–311. Springer, Heidelberg (2007)
8. Faliszewski, P., Hemaspaandra, E., Schnoor, H.: Copeland voting: Ties matter. In: Proc. of 7th AAMAS, pp. 983–990 (2008)
9. Friedgut, E., Kalai, G., Nisan, N.: Elections can be manipulated often. In: Proc. of 49th FOCS (to appear, 2008)
10. Gibbard, A.: Manipulation of voting schemes. *Econometrica* 41, 587–602 (1973)
11. Hemaspaandra, E., Hemaspaandra, L.A.: Dichotomy for voting systems. *Journal of Computer and System Sciences* 73(1), 73–83 (2007)
12. Moulin, H.: *The Strategy of Social Choice*. Advanced Textbooks in Economics, vol. 18. North-Holland, Amsterdam (1983)
13. Procaccia, A.D., Rosenschein, J.S.: Average-case tractability of manipulation in elections via the fraction of manipulators. In: Proc. of 6th AAMAS, pp. 718–720 (2007)
14. Procaccia, A.D., Rosenschein, J.S.: Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research* 28, 157–181 (2007)
15. Procaccia, A.D., Rosenschein, J.S., Zohar, A.: Multi-winner elections: Complexity of manipulation, control and winner-determination. In: Proc. of 20th IJCAI, pp. 1476–1481 (2007)
16. Satterthwaite, M.: Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10, 187–217 (1975)
17. Svensson, L.-G.: The proof of the Gibbard-Satterthwaite theorem revisited. Working Paper No. 1999:1, Department of Economics, Lund University (1999), <http://www.nek.lu.se/NEKlgs/vote09.pdf>
18. Xia, L., Conitzer, V.: Generalized Scoring Rules and the frequency of coalitional manipulability. In: Proc. of 9th ACM-EC, pp. 109–118 (2008)
19. Xia, L., Conitzer, V.: A sufficient condition for voting rules to be frequently manipulable. In: Proc. of 9th ACM-EC, pp. 99–108 (2008)
20. Zuckerman, M., Procaccia, A.D., Rosenschein, J.S.: Algorithms for the coalitional manipulation problem. In: Proc. of 19th SODA, pp. 277–286 (2008)

The Power of Small Coalitions in Cost Sharing^{*}

Florian Schoppmann^{1,2}

¹ Faculty of Computer Science, Electrical Engineering and Mathematics,
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany

fschopp@uni-paderborn.de

² International Graduate School of Dynamic Intelligent Systems

Abstract. In a cost-sharing problem, finitely many players have an unknown preference for some public excludable good (service), and the task is to determine which players to serve and how to distribute the incurred cost. Therefore, incentive-compatible mechanisms are sought that elicit truthful bids, charge prices that recover the cost, and are economically efficient in that they reasonably balance cost and valuations. A commonplace notion of incentive-compatibility in cost sharing is group-strategyproofness (GSP), meaning that not even coordinated deceit is profitable. However, GSP makes strong implications on players' coordination abilities and is known to impose severe limitations on other goals in cost sharing. There is hence good reason to seek for a weaker axiom: In this work, we study the following question: Does relaxing GSP to resilience only against *coalitions of bounded size* yield a richer set of possible mechanisms? Surprisingly, the answer is essentially “no”: We prove that already a mechanism resilient to coalitions of size only two (“2-GSP”) is GSP, once we require that cost shares must only depend on the service allocation (and not directly on the bids). Moreover, we show that even without additional requirements, 2-GSP implies weak group-strategyproofness (WGSP). Consequently, our results give some justification that GSP may, after all, still be desirable in various scenarios. As another benefit, we believe that our characterizations will facilitate devising and understanding new GSP cost-sharing mechanisms. Finally, we relate our findings to other concepts of non-manipulability such as (outcome) non-bossiness [19] and weak utility non-bossiness [13].

1 Introduction

Indivisible units of a *public excludable good* (a service; e.g., connectivity in a network) are to be made available to $n \in \mathbb{N}$ players at non-negative prices. In the *binary-demand* scenario studied in this work, each player has demand for only one unit and is completely characterized by his *valuation* for receiving it. Direct-revelation *cost-sharing mechanisms* are sought that elicit truthful reports of the players' valuations and then determine both an allocation $\mathbf{q} \in \{0, 1\}^n$ of the good and a distribution of the allocation-dependent cost $C(\mathbf{q})$.

^{*} This work was partially supported by the IST Program of the European Union under contract number IST-15964 (AEOLUS).

This work follows the line of examining cost sharing from the viewpoint of incentive compatibility: How can rational selfish players be incentivized to reveal *truthful* information out of self-interest? Cost-sharing problems are fundamental in economics and have a broad area of applications; e.g., distributing volume discounts in electronic commerce, sharing the cost of public infrastructure projects, allocating development costs of low-volume built-to-order products, etc.

There are three essential goals when designing a truthful cost-sharing mechanisms: It should *balance the budget* (i.e., recover the incurred cost with the prices charged), and be *economically efficient* (i.e., trade off the service cost and the excluded players' valuations as good as possible). Practical applications also demand for *polynomial-time computability*.

1.1 Collusion-Resistance

As is common in cost sharing, we assume *quasi-linear* utilities: When served, a player's utility is his valuation minus his payment, otherwise he will not be charged and his utility is zero. The basic notion of truthfulness, *strategyproofness* (SP), requires that no player can improve his utility by false bidding (assuming all other bids are fixed); i.e., truth-telling is always a dominant strategy.

Especially in settings with a large number of players (e.g., in the Internet), not only SP but also resistance against coordinated manipulation is desirable. Several concepts of collusion-resistance are known in the literature: A mechanism is called *group-strategyproof* (GSP) if any defection of a coalition that increases some member's utility inevitably decreases the utility of one of its other members. A weaker notion of collusion resistance is *weak group-strategyproofness* (WGSP) that is fulfilled if any defecting coalition has at least one member whose utility does not strictly improve. Since it is unlikely that players have unlimited means to communicate and make binding agreements with all of their competitors, Serizawa [21] advocated relaxing GSP to *effective pairwise strategy-proofness*, meaning that a mechanism needs only be resilient to pairs of defecting players—and this only if their defection was stable (i.e., none of the two players could betray his partner to further increase his utility). While similar in spirit to this work, Serizawa's findings do not apply to the cost-sharing scenario.

Besides the (coalitional) variants of strategyproofness, there are several other concepts of non-manipulability. Satterthwaite and Sonnenschein [19] suggested a property called (*outcome*) *non-bossiness* (ONB): If a single player changes his bid in a way so that his own outcome does not change, then all other players should also get the same outcome as before. In an unpublished paper, Shenker [22] proved several results on the relationship between various forms of (coalitional) strategyproofness, non-bossiness, and other technical properties. However, his results do not apply in settings with quasi-linear utilities, as in the case of cost sharing. This special domain was later studied by Mutuswami [14]. He introduced a relaxation of ONB called *weak utility non-bossiness* (WUNB), meaning that if a single player changes his bid so that his utility remains the same, then no other player may become better off. Mutuswami [14] showed that SP and ONB together imply WGSP; moreover, SP, ONB, and WUNB together imply

GSP. For scenarios when players are capable of side-payments, Schummer [20] studied *bribe-proof* mechanisms, meaning that no player has an incentive to bribe another player to misreport his type. For the cost-sharing scenario, his results imply that notions of collusion-resistance that include monetary transfers are too strong: They would rule out all but trivial mechanisms where each player's outcome is completely independent of the other players' actions.

1.2 Our Contribution

We concentrate on the question whether reducing the maximum coalition size that a mechanism should withstand allows for a richer set of possible mechanisms. We say a mechanism is k -GSP (or k -WGSP, respectively) if it ensures collusion-resistance up to coalition size k . In detail, our results are:

- While we give (arguably artificial) cost-sharing mechanisms that are k -GSP but not $(k + 1)$ -GSP, we obtain as our main result that already 2-GSP is equivalent to GSP once we require mechanisms to be *separable*, i.e., cost shares must only depend on the service allocation (and not directly on the bids). We remark that no general technique for the design of truthful cost-sharing mechanisms is known that violates separability (cf. Section 1.3). Our result can be seen as a generalization of the main theorem in [14].
- In contrast, WGSP is *not* equivalent to 2-WGSP plus separability.
- Even without separability, 2-GSP always implies WGSP.

We regard the chief asset of our work to be threefold: First, our results indicate that the substantial “jump” in collusion-resistance seems to occur from 1-GSP = SP to 2-GSP. Second, GSP is often felt to be too strong an axiom with unrealistic implications on players' capabilities and behavior (cf. Section 3); now, the fact that GSP is equivalent to merely 2-GSP plus separability gives some a posteriori justification for GSP. Third and last, we firmly believe that our characterizations will facilitate devising and understanding new GSP cost-sharing mechanisms.

1.3 Further Related Work

Arguably the most important result in mechanism design is the family of *Vickrey-Clark-Groves (VCG) mechanisms* (see, e.g., [15]), that are SP and satisfy optimal economic efficiency. Even more, under quite general assumptions, the VCG mechanisms is the only family of mechanisms with these properties [9]. Unfortunately, VCG mechanisms are not resistant against collusion and fail to provide any revenue guarantees. Hence, exact budget balance and optimal efficiency can in general not be simultaneously achieved by *any* SP mechanism. In fact, not even bi-criteria *approximation guarantees* are possible [2], unless economic efficiency is measured in terms of *social cost* (service cost plus excluded valuations) and not the traditional *surplus* (included valuations minus service cost) [18].

Essentially, only one general technique is known for the design of GSP mechanisms, due to Moulin [12]. Its main ingredient are *cross-monotonic* cost shares $\xi_i(S)$ that never decrease when the set of served players S gets larger. Then, a

Moulin mechanism serves the maximal set of players who can afford their corresponding price—due to cross-monotonicity, a unique maximal set always exists. The main advantage of Moulin’s technique is that it reduces the design of GSP mechanism to finding cross-monotonic cost-sharing methods. However, there are several natural cost-sharing problems for which any Moulin mechanisms inevitably suffers poor budget-balance and efficiency [10]. In [4], Bleischwitz et al. give a novel family of GSP mechanisms with good budget balance, yet only for the special case of symmetric costs and at the price of sacrificing efficiency.

A general technique for the design of mechanisms satisfying only the less demanding WGSP is due to Mehta et al. [11]. Their mechanisms are called *acyclic* and are strictly more general than Moulin mechanisms. Bleischwitz et al. [3] devise a special family of acyclic mechanisms that for the broad class of subadditive cost-sharing problems provide exact budget-balance and optimal (cf. [8]) asymptotic efficiency. For some classes of problems, these mechanisms are computable in polynomial time; yet, this does not necessarily hold in general.

To the best of our knowledge, no other general techniques have been proposed in the literature [1]. Besides general design techniques, most other work on cost sharing has focused on finding “good” cross-monotonic cost-sharing methods (see, e.g., [18, 10] an the references therein), obtaining characterization results [12, 6, 7, 17, 4, 10], and finding “good” acyclic mechanisms [11, 3, 5].

2 The Model

Notation. For $n, m \in \mathbb{N}_0$, let $\{n \dots m\} := \{n, n+1, \dots, m\}$ and $[n] := \{1 \dots n\}$. Given vector $\mathbf{v} \in \mathbb{R}^n$, we denote its components by $\mathbf{v} = (v_1, \dots, v_n)$. Two vectors \mathbf{v}, \mathbf{v}' are called *K-variants* if $v_i = v'_i$ for all $i \notin K$. In this case, we write $\mathbf{v}' = (\mathbf{v}_{-K}, \mathbf{v}'_K)$. If $K = \{i\}$, then \mathbf{v} and \mathbf{v}' are *i-variants* and $\mathbf{v}' = (\mathbf{v}_{-i}, v'_i)$.

A *cost-sharing problem* is specified by a *cost function* $C : \{0, 1\}^n \rightarrow \mathbb{R}_{\geq 0}$ that associates all possible *service allocations* to their incurred costs. A service allocation $\mathbf{q} \in \{0, 1\}^n$ together with a distribution of costs $\mathbf{x} \in \mathbb{R}^n$ is called an *outcome*. We denote player i ’s *valuation* for being served by $v_i \in \mathbb{R}$. *Utilities* are *quasi-linear*, i.e., player i ’s utility for outcome (\mathbf{q}, \mathbf{x}) is $v_i \cdot q_i - x_i$.

Definition 1. A cost-sharing mechanism $M = (q, x)$ is a pair of functions $q : \mathbb{R}^n \rightarrow \{0, 1\}^n$ and $x : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that associates any combination of announced bids \mathbf{b} to an outcome $(q(\mathbf{b}), x(\mathbf{b}))$.

Cost-sharing mechanisms are *direct-revelation* mechanisms, since the set of possible bids is equal to the set of possible valuations (types). Given $M = (q, x)$, we write $M_i(\mathbf{b}) := (q_i(\mathbf{b}), x_i(\mathbf{b}))$. We denote player i ’s induced utility by $u_i(\mathbf{b} \mid v_i) := v_i \cdot q_i(\mathbf{b}) - x_i(\mathbf{b})$. When there is no confusion about the true valuation v_i , we simply write $u_i(\mathbf{b})$ instead of $u_i(\mathbf{b} \mid v_i)$. In this work, we only discuss mechanisms that fulfill three standard axiomatic properties: First, *No Positive Transfers (NPT)*

¹ There are other families of GSP mechanisms [10, 16]; however, they are considered implausible here as they force players to participate. See Section 2.

requires that players never get paid. That is, $x_i(\mathbf{b}) \geq 0$. Second, *Voluntary Participation (VP)* means that when served, players never pay more than they bid; otherwise, they are charged nothing. That is, if $q_i(\mathbf{b}) = 1$ then $x_i(\mathbf{b}) \leq b_i$, else $x_i(\mathbf{b}) = 0$. Third and last, *Consumer Sovereignty (CS)* means that each player can bid in a way so that he is served, regardless of the other players' bids. That is, there is a bid $\mathbf{b}^\infty \in \mathbb{R}_{\geq 0}$ such that if $b_i = \mathbf{b}^\infty$ then $q_i(\mathbf{b}) = 1$.

VP and NPT imply that players may opt to not participate (by submitting a negative bid). This property strengthens the collusion-resistance requirements and rules out otherwise implausible and undesirable mechanisms (see [10]).

Definition 2. A mechanism M is strategyproof (SP) if for all true valuations $\mathbf{v} \in \mathbb{R}^n$ and all their i -variants \mathbf{b} it holds that $u_i(\mathbf{b}) \leq u_i(\mathbf{v})$.

Definition 3. Let M be a cost-sharing mechanism. If for all coalitions $K \subseteq [n]$, all true valuations $\mathbf{v} \in \mathbb{R}^n$, and all their K -variants \mathbf{b} it holds that

1. $u_i(\mathbf{b}) \leq u_i(\mathbf{v})$ for at least one $i \in K$, then M is weakly group-strategyproof (WGSP);
2. $u_i(\mathbf{b}) = u_i(\mathbf{v})$ for all $i \in K$ or $u_i(\mathbf{b}) < u_i(\mathbf{v})$ for at least one $i \in K$, then M is group-strategyproof (GSP).

Definition 4 ([19]). A mechanism M is (outcome) non-bossy (ONB) if for all i -variants $\mathbf{b}, \mathbf{b}' \in \mathbb{R}^n$ it holds that $M_i(\mathbf{b}) \neq M_i(\mathbf{b}')$ or $M_{-i}(\mathbf{b}) = M_{-i}(\mathbf{b}')$.

Definition 5 ([13]). A mechanism M is weakly utility non-bossy (WUNB) if for all true valuations $\mathbf{v} \in \mathbb{R}^n$ and all their i -variants \mathbf{b} it holds that $u_i(\mathbf{v}) \neq u_i(\mathbf{b})$ or $u_{-i}(\mathbf{v}) \geq u_{-i}(\mathbf{b})$.

We pay special attention to mechanisms with cost shares that depend only on the service allocation and not directly on the bids (note Proposition 3).

Definition 6. A cost-sharing method is a function $\xi : \{0, 1\}^n \rightarrow \mathbb{R}_{\geq 0}^n$ that associates each service allocation to a vector of cost shares. We say that a cost-sharing mechanism $M = (q, x)$ is separable if there exists a cost-sharing method ξ so that $x = \xi \circ q$, i.e., for all $\mathbf{b} \in \mathbb{R}_{\geq 0}^n : x(\mathbf{b}) = \xi(q(\mathbf{b}))$.

For completeness (despite not the focus of this work), we briefly formalize the two optimization goals *budget balance* and *economic efficiency*. Typically, cost-sharing problems are specified implicitly by combinatorial optimization problems, i.e., $C(\mathbf{q})$ is the minimum-cost value of an optimal solution to the instance that corresponds to \mathbf{q} . Due to the computational complexity, usually only an approximate solution with cost $C'(\mathbf{q})$ can be computed. Still, the revenue of the mechanism should be reasonably bounded, i.e., any computed outcome (\mathbf{q}, \mathbf{x}) should satisfy $C'(\mathbf{q}) \leq \sum_{i=1}^n x_i \leq \beta \cdot C(\mathbf{q})$ for some constant $\beta \geq 1$. Moreover, as a measure for economic efficiency, the incurred cost and the rejected players' valuations should be traded off as good as possible. That is, $C'(\mathbf{q}) + \sum_{i=1}^n (1 - q_i)\bar{b}_i \leq \gamma \cdot \min_{\mathbf{p} \in \{0,1\}^n} \{C(\mathbf{p}) + \sum_{i=1}^n (1 - p_i)\bar{b}_i\}$ for some constant $\gamma \geq 1$, where $\bar{b}_i := \max\{b_i, 0\}$.

3 Resistance against Coalitions of Bounded Size

Demand for some level of collusion-resistance implies assumptions on players’ behavior and their coalition-forming capabilities: For instance, if (a) side-payments are unlikely but (b) players yet have virtually unlimited means to communicate and (c) one expects them to help others even for no personal reward (e.g., by voluntary non-participation in case of indifference), then GSP is an appropriate axiom. Similarly, when players have no means to communicate at all, then simple SP is probably sufficient. One can also think of collusion-resistance at the other end of the spectrum: We use the term “ultimate group-strategyproofness” (UGSP) here if a mechanism even prevents that coalitions can improve their *total utility* by manipulation—thus, UGSP is stronger than bribe-proofness [20]. To conclude, WGSP, GSP, and UGSP imply different levels of transfers that coalitions might accomplish in order to be successful. Figure 1 provides a schematic illustration.




Fig. 1. Dimensions of coalitions’ capabilities

Since it is unlikely that all players can efficiently communicate with each other and make binding agreements on collective deceit, this gives rise to the following natural question: Can we increase the degree of freedom for designing cost-sharing mechanisms by relaxing GSP with respect to coalition sizes? Surprisingly, we show in the rest of this paper that the answer is essentially “no”. (Due to space limitations, we have to omit all proofs. See the full version of this paper.)

Definition 7. A mechanism M is k -GSP (or k -WGSP, respectively) if it satisfies the conditions of Definition 3 for all coalitions K of size up to k .

Note that 2-GSP is equal to pairwise SP from [21] and immediately implies 2-WGSP, SP, and WUNB.

3.1 Some Preliminary Implications by SP and WUNB

We start with some observations needed throughout the paper. Note that the following simple proposition is well-known and a standard fact (see, e.g., [6]).

Proposition 1. A cost-sharing mechanism $M = (q, x)$ is SP if and only if the following holds: For all $i \in [n]$ and all $\mathbf{b}_{-i} \in \mathbb{R}^{[n] \setminus i}$, there is a threshold bid $\tau_i(\mathbf{b}_{-i})$ so that if $b_i > \tau_i(\mathbf{b}_{-i})$ then $q_i(\mathbf{b}) = 1$, if $b_i < \tau_i(\mathbf{b}_{-i})$ then $q_i(\mathbf{b}) = 0$, and if $q_i(\mathbf{b}) = 1$ then $x_i(\mathbf{b}) = \tau_i(\mathbf{b}_{-i})$.

Lemma 1. Let $M = (q, x)$ be a SP cost-sharing mechanism, $\mathbf{v} \in \mathbb{R}^n$ contain the true valuations, and \mathbf{b} be an i -variant. Then:

1. $u_i(\mathbf{b}) < u_i(\mathbf{v})$ and $q_i(\mathbf{v}) = 1 \implies q_i(\mathbf{b}) = 0, u_i(\mathbf{b}) = 0 < u_i(\mathbf{v})$, and $b_i \leq \tau_i(\mathbf{v}_{-i}) = x_i(\mathbf{v}) < v_i$

2. $u_i(\mathbf{b}) < u_i(\mathbf{v})$ and $q_i(\mathbf{v}) = 0 \implies q_i(\mathbf{b}) = 1, u_i(\mathbf{b}) < 0 = u_i(\mathbf{v})$, and $b_i \geq \tau_i(\mathbf{v}_{-i}) = x_i(\mathbf{b}) > v_i$

Lemma 2. *Let M be a WUNB cost-sharing mechanism, $\mathbf{v} \in \mathbb{R}^n$ contain the true valuations, and \mathbf{b} be an i -variant. Then: $M_i(\mathbf{b}) = M_i(\mathbf{v}) \implies u_{-i}(\mathbf{b}) = u_{-i}(\mathbf{v})$.*

Lemma 3. *Let $M = (q, x)$ be a SP and WUNB cost-sharing mechanism, $\mathbf{v} \in \mathbb{R}^n$ contain the true valuations, and \mathbf{b} be an i -variant. Moreover, let $j \in [n] \setminus i$. Then:*

1. $u_j(\mathbf{b}) > u_j(\mathbf{v}) \implies u_i(\mathbf{b}) < u_i(\mathbf{v})$
2. $u_j(\mathbf{b}) < u_j(\mathbf{v})$ and $q_i(\mathbf{v}) = 1 \implies q_i(\mathbf{b}) = 0$ and $b_i < \tau_i(\mathbf{v}_{-i}) = x_i(\mathbf{v}) \leq v_i$
3. $u_j(\mathbf{b}) < u_j(\mathbf{v})$ and $q_i(\mathbf{v}) = 0 \implies q_i(\mathbf{b}) = 1$ and $b_i > \tau_i(\mathbf{v}_{-i}) = x_i(\mathbf{b}) \geq v_i$
4. $u_{-i}(\mathbf{b}) \leq u_{-i}(\mathbf{v})$ or $u_{-i}(\mathbf{b}) \geq u_{-i}(\mathbf{v})$
5. $v_i = \tau_i(\mathbf{v}_{-i}) \implies u_j(\mathbf{b}) \leq u_j(\mathbf{v})$
6. $u_j(\mathbf{b}) > u_j(\mathbf{v}) \implies \tau_j(\mathbf{b}_{-j}) < \tau_j(\mathbf{v}_{-j})$

3.2 k -GSP Is Strictly Weaker Than GSP

Before establishing the link between 2-GSP and GSP in the next sections, we give an example showing that k -GSP is not equivalent to GSP when $k < n$.

Algorithm 1 (3-player mechanism that is 2-GSP but not GSP).

Input: bid vector $\mathbf{b} \in \mathbb{R}^3$ *Output:* allocation $\mathbf{q} \in \{0, 1\}^3$; cost shares $\mathbf{x} \in \mathbb{R}_{\geq 0}^3$

- 1: **if** $\mathbf{b} = (1, 1, 1)$ **then** $\mathbf{q} := (1, 1, 1)$; $\mathbf{x} := (1, 1, 1)$
- 2: **else**
- 3: $\mathbf{q} := (0, 0, 0)$; $\mathbf{x} := (0, 0, 0)$; $\xi := (1, 1, 1)$
- 4: **if** $b_1 > 1$ and $b_2 > 1$ **then** $\xi_3 := 2$
- 5: **for all** $i \in [3]$ with $b_i > \xi_i$ **do** $q_i := 1$; $x_i := \xi_i$

It is easy to see that Algorithm 1 defines a 2-GSP mechanism, as the only player who could ever improve is player 3. In this case, however, $v_1 > 1$ and $v_2 > 1$, so in order to help player 3, both players 1 and 2 have to deviate. Clearly, Algorithm 1 can be generalized for n players so that it is $(n - 1)$ -GSP but not n -GSP.

3.3 Upper Continuity and 2-GSP Together Imply GSP

We first add a continuity condition (see [10]) to our assumptions that determines how to deal with indifferent players (valuation = payment). This condition is fulfilled by almost all general cost-sharing techniques ([4] is an exception).

Definition 8. *A cost-sharing mechanism $M = (q, x)$ is upper continuous if for all players i and all bid vectors \mathbf{b} the following holds: If $q_i(\mathbf{b}_{-i}, x) = 1$ for all $x > b_i$ then also $q_i(\mathbf{b}) = 1$.*

Lemma 4. *Let M be upper continuous and 2-GSP. Then M is also ONB.*

Proposition 2 ([14]). *Let M be a SP, ONB, and WUNB. Then, M is GSP.*

Corollary 1. *Let M be upper continuous and 2-GSP. Then, it is GSP.*

We remark that all results of this section would remain valid if we changed the model to only allow for non-negative bids and valuations.

3.4 Separability and 2-GSP Together Imply GSP

In this section, we generalize [Corollary 1](#) to hold for arbitrary separable mechanisms. Specifically, we will obtain as our main result that a 2-GSP cost-sharing mechanism is GSP if and only if it is separable. We start with an auxiliary lemma, stating that every 2-GSP cost-sharing mechanism is at least resistant against coalitions where deviators either do not participate (submit a negative bid) or bid very much.

Lemma 5. *Let $M = (q, x)$ be a 2-GSP cost-sharing mechanism, $K \subseteq [n]$ be a coalition (of arbitrary size), $\mathbf{v} \in \mathbb{R}^n$ contain the true valuations, and \mathbf{b} be a K -variant so that for all $i \in K : b_i \in \{-1, \mathbf{b}^\infty\}$. Then, either $u_i(\mathbf{b}) = u_i(\mathbf{v})$ for all $i \in K$ or $u_i(\mathbf{b}) < u_i(\mathbf{v})$ for at least one $i \in K$.*

Theorem 1. *Let M be separable and 2-GSP. Then M is GSP.*

Proposition 3 ([\[12\]](#)). *Let M be GSP. Then, M is separable.*

Corollary 2. *A mechanism M is GSP if and only if is 2-GSP and separable.*

3.5 Separability and 2-WGSP Do Not Imply WGSP

Does a statement similar to [Theorem 1](#) also hold for WGSP? The cost-sharing mechanism $M = (q, x)$ defined by [Algorithm 2](#) is a counterexample.

Algorithm 2 (Separable mechanism that is 2-WGSP but not WGSP).

Input: bid vector $\mathbf{b} \in \mathbb{R}^6$ *Output:* allocation $\mathbf{q} \in \{0, 1\}^6$; cost shares $\mathbf{x} \in \mathbb{R}_{\geq 0}^6$

- 1: $\mathbf{q} := (0, \dots, 0), \mathbf{x} := (0, \dots, 0)$
- 2: **for all** $i \in \{4 \dots 6\}$ with $b_i > 1$ or ($b_i = 1$ and $b_{1+(i-3 \bmod 3)} \geq 2$) **do**
- 3: $q_i := 1, x_i := 1$
- 4: **for all** $i \in [3]$ with $b_i \geq 1 + q_{i+3}$ **do** $q_i := 1, x_i := 1 + q_{i+3}$

The unique cost-sharing method ξ of mechanism M is given by

$$\xi_i(\mathbf{q}) := \begin{cases} 0 & \text{if } q_i = 0 \\ 1 & \text{otherwise, if } i \in \{4 \dots 6\} \text{ or } (i \in [3] \text{ and } q_{i+3} = 0) \\ 2 & \text{otherwise, if } i \in [3] \text{ and } q_{i+3} = 1 \end{cases}$$

Since for each player i , there is a threshold bid that does not depend on i 's own bid and that is equal to i 's cost share, [Algorithm 2](#) is SP due to [Proposition 1](#). Moreover, the only players who could ever improve are $[3]$. Now, no subset $S \subset [3]$ of size $|S| = 2$ can jointly improve because there is always a player $i \in S$ whose outcome does not depend on \mathbf{b}_S . Hence, M is 2-WGSP. However, it is not 3-WGSP: Let $\mathbf{v} = (2, 2, 2, 1, 1, 1)$ contain the true valuations and consider $\mathbf{b} = (1, 1, 1, 1, 1, 1)$. Then, $q(\mathbf{v}) = (1, 1, 1, 1, 1, 1)$ and $q(\mathbf{b}) = (1, 1, 1, 0, 0, 0)$, so $[3]$ is a successful coalition.

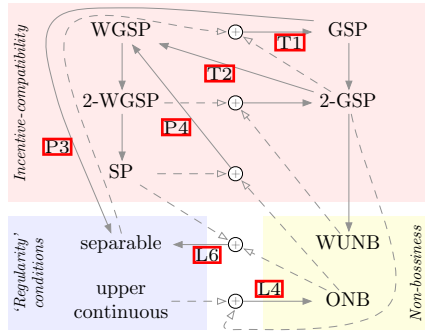


Fig. 2. Overview of the various non-manipulability properties

3.6 2-GSP Implies WGSP

Theorem 2. *Let M be a 2-GSP. Then, M is also WGSP.*

The previous result is a case where a stronger notion of collusion-resistance, yet only for players with limited communication abilities, implies a weaker collusion-resistance against coalitions of arbitrary size.

3.7 Collusion-Resistance and Non-bossiness Properties

Lemma 6. *Let M be a SP and ONB. Then M is separable.*

Consequently, [Theorem 1](#) can be seen as a generalization of [Proposition 2](#) because its requirements (SP, ONB, and WUNB) imply 2-GSP and separability in a relatively straightforward manner. The following example shows that [Theorem 1](#) is *strictly* more general because ONB is not a necessary condition for GSP: Define mechanism $M = (q, x)$ by

$$q(\mathbf{b}) := \begin{cases} (1, 1) & \text{if } (b_1 \geq 1 \text{ and } b_2 > 1) \text{ or } \mathbf{b} = (1, 1) \\ (1, 0) & \text{if } (b_1 \geq 1 \text{ and } b_2 \leq 1) \text{ and } \mathbf{b} \neq (1, 1) \\ (0, 1) & \text{if } b_1 < 1 \text{ and } b_2 > 1 \\ (0, 0) & \text{if } b_1 < 1 \text{ and } b_2 \leq 1 \end{cases} \quad \text{and } x(\mathbf{b}) := q(\mathbf{b}).$$

Obviously, neither of the two players could ever improve. However, the mechanism is not ONB because $M_1(1, 1) = M_1(2, 1)$ but $M_2(1, 1) \neq M_2(2, 1)$.

We conclude by stating another result in [\[14\]](#), which completes our overview of the various notions of non-manipulability and many of their implications (see [Figure 2](#)).

Proposition 4 ([\[14\]](#)). *Let M be SP and ONB. Then M is also WGSP.*

References

1. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: Proc. of the 18th SODA 2007, pp. 189–198. Society for Industrial and Applied Mathematics (2007)

2. Archer, A., Feigenbaum, J., Krishnamurthy, A., Sami, R., Shenker, S.: Approximation and collusion in multicast cost sharing. *Games and Economic Behaviour* 47(1), 36–71 (2004)
3. Bleischwitz, Y., Monien, B., Schoppmann, F.: To be or not to be (served). In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 515–528. Springer, Heidelberg (2007)
4. Bleischwitz, Y., Monien, B., Schoppmann, F., Tiemann, K.: The power of two prices: Beyond cross-monotonicity. In: *MFCS 2007*. LNCS, vol. 4708, pp. 657–668. Springer, Heidelberg (2007)
5. Brenner, J., Schäfer, G.: Singleton acyclic mechanisms and their applications to scheduling problems. In: Monien, B., Schroeder, U.-P. (eds.) *SAGT 2008*. LNCS, vol. 4997, pp. 315–326. Springer, Heidelberg (2008)
6. Deb, R., Razzolini, L.: Voluntary cost sharing for an excludable public project. *Mathematical Social Sciences* 37(2), 123–138 (1999)
7. Deb, R., Razzolini, L.: Auction-like mechanisms for pricing excludable public goods. *Journal of Economic Theory* 88(2), 340–368 (1999)
8. Dobzinski, S., Mehta, A., Roughgarden, T., Sundararajan, M.: Is Shapley cost sharing optimal? In: Monien, B., Schroeder, U.-P. (eds.) *SAGT 2008*. LNCS, vol. 4997, pp. 327–336. Springer, Heidelberg (2008)
9. Green, J., Laffont, J.-J.: Characterizations of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* 45(2), 427–438 (1977)
10. Immorlica, N., Mahdian, M., Mirrokni, V.S.: Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms* 4(2), 1–25 (2008)
11. Mehta, A., Roughgarden, T., Sundararajan, M.: Beyond Moulin mechanisms. In: *Proc. of the 8th EC 2007*, pp. 1–10 (2007)
12. Moulin, H.: Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare* 16(2), 279–320 (1999)
13. Mutuswami, S.: Strategy proof mechanisms for cost sharing. *Economics Discussion Papers* 520, University of Essex, Department of Economics (2000)
14. Mutuswami, S.: Strategyproofness, non-bossiness and group strategyproofness in a cost sharing model. *Economics Letters* 89(1), 83–88 (2005)
15. Nisan, N.: Introduction to mechanisms design (for computer scientists). In: *Algorithmic Game Theory*, ch. 9. Cambridge University Press, Cambridge (2007)
16. Penna, P., Ventrè, C.: More powerful and simpler cost-sharing methods. In: Persiano, G., Solis-Oba, R. (eds.) *WAOA 2004*. LNCS, vol. 3351, pp. 97–110. Springer, Heidelberg (2005)
17. Penna, P., Ventrè, C.: The algorithmic structure of group strategyproof budget-balanced cost-sharing mechanisms. In: Durand, B., Thomas, W. (eds.) *STACS 2006*. LNCS, vol. 3884, pp. 337–348. Springer, Heidelberg (2006)
18. Roughgarden, T., Sundararajan, M.: New trade-offs in cost-sharing mechanisms. In: *STOC 2006*, pp. 79–88 (2006)
19. Satterthwaite, M., Sonnenschein, H.: Strategy-proof allocation mechanisms at differentiable points. *Review of Economic Studies* 48(4), 587–597 (1981)
20. Schummer, J.: Manipulation through bribes. *Journal of Economic Theory* 91(2), 180–198 (2000)
21. Serizawa, S.: Pairwise strategy-proofness and self-enforcing manipulation. *Social Choice and Welfare* 26(2), 305–331 (2006)
22. Shenker, S.: Some technical results on continuity, strategy-proofness, and related strategic concepts (1993),
<ftp://parcftp.parc.xerox.com/pub/net-research/str.ps>

Social Context Games

Itai Ashlagi¹, Piotr Krysta^{2,*}, and Moshe Tennenholtz^{3,4}

¹ Harvard Business School, Harvard University, USA
iashlagi@hbs.edu

² Computer Science Dept., University of Liverpool, UK
p.krysta@liverpool.ac.uk

³ Industrial Engineering & Management, Technion, Israel

⁴ Microsoft Israel R&D Center,
moshet@microsoft.com

Abstract. We introduce *social context games*. A social context game is defined by an underlying game in strategic form, and a social context consisting of an undirected graph of neighborhood among players and aggregation functions. The players and strategies in a social context game are as in the underlying game, while the players' utilities in a social context game are computed from their payoffs in the underlying game based on the graph of neighborhood and the aggregation functions. Examples of social context games are ranking games and coalitional congestion games. In this paper we consider resource selection games as the underlying games, and four basic social contexts. An important property of resource selection games is the existence of pure strategy equilibrium. We study the existence of pure strategy Nash equilibrium in the corresponding social context games. We also show that the social context games possessing pure strategy Nash equilibria are not potential games, and therefore are distinguished from congestion games.

1 Introduction

Game theory has become a standard tool for the analysis of social interactions of self-motivated agents. Naturally, social interactions may be complex and refer to issues such as competition, coordination, and collaboration among agents. The attitude of agents towards other agents is typically captured by their utility functions. However, it may be of interest to separate the payoff the agent receives, by means of e.g. delay, cost, etc., from his social attitude; this will allow to study the possible effects that various social contexts have. Consider for example several service providers who act on behalf of a set of customers. Each service provider suffers a cost, caused by the need for sharing resources with other service providers. This type of situations are typically modeled as a form of congestion game. Indeed, such congestion games have been a central topic of

* The work of the second author is supported by DFG grant Kr 2332/1-3 within Emmy Noether program.

¹ We use the terms agents and players interchangeably.

study in the interplay between computer science and game theory. However, while these games capture the underlying situation, they implicitly assume a very particular social context, where the actual goals of each service provider is to minimize its own cost. In some situations it may be the case that the aim of each agent is to have its payoff ranked as high as possible comparing to other agents' payoffs, as studied in [2]. Another example is when each agent cares about the sum (or average) of payoffs of a set of agents in a coalition it belongs to, as studied in coalitional congestion games [4,5]. In our terminology the above are examples of *social contexts*. Our aim is to study the effects of social contexts on basic properties of fundamental types of underlying games.

Consider a resource selection game. In a resource selection game we have a set of n agents, and a set of m resources. Each agent chooses a resource from among the set of resources, and his cost is a non-decreasing function of the number of agents who have chosen his selected resource. Needless to say, resource selection games are central to work in various communities, such as operations research, computer science, game theory and economics. A resource selection game is a special type of congestion games [8]. A fundamental property of congestion games is that they possess a pure strategy equilibrium. This result is implied by the fact that congestion games possess a potential function. Indeed, the classes of potential games and congestion games coincide [6]. Given the importance of resource selection games, and the desire to consider various social contexts, we will consider the existence of pure strategy Nash equilibrium when the resource selection games are embodied in the following basic social contexts:

1. Rank competition: the agents are partitioned into cliques, where at each clique the agents compete on their relative payoff. This significantly extends upon ranking games, where the whole graph is a single clique.
2. Best-Member Collaboration: there is a given social network, and each agent cares about the highest payoff obtained by him or by one of his neighbors. This is in the spirit of work in congestion games where an agent can choose several resources and cares about the one with the best performance [7]; here the agent cares that either himself or one of his friends will behave (e.g. provide a service) as good as possible.
3. Min-Max Collaboration: there is a given social network, and each agent cares that the worst case payoff obtained by him or by one of his friends will be maximized. This requirement is in the spirit of minmax fairness²; however, it is stated as a social attitude rather than as a system requirement.
4. Surplus Collaboration: there is a social network, and each agent wishes to maximize the average payoff of himself and his friends. This is in the spirit of coalitional congestion games [4,5]; however, here we allow arbitrary graphs rather than a partition of the nodes into cliques.

In the following sections we deal with the social context games generated by resource selection games and the above social contexts. We show:

² Our particular treatment extends upon [3], by allowing arbitrary networks.

- When the resources are identical then any Rank Competition resource selection game possesses a pure strategy equilibrium. This is no longer true when the resources are non-identical.
- Any Best-Member Collaboration game possesses a pure strategy equilibrium whenever the size of the largest independent dominating set of the corresponding graph is smaller than $\frac{n}{2}$; this condition is necessary.
- Any pure strategy Nash equilibrium is an equilibrium also in the corresponding Min-Max collaboration game; the converse is not true.
- Pure strategy equilibrium does not always exist in surplus collaboration resource selection games, even with identical resources and when the graph is a tree. We show a subset of the resource selection games where a pure strategy equilibrium exists.
- We show that all social contexts games above do not possess a potential function, and therefore they are not isomorphic to congestion games.

2 Basic Definitions

A *game in strategic form* (for short, a *game*) is a tuple $H = (N, (A_i)_{i \in N}, (c_i)_{i \in N})$ where $N = \{1, 2, \dots, n\}$ is a finite set of players (agents), for every player $i \in N$, A_i is the action set of player i , and $c_i : \times_{i \in N} A_i \rightarrow \mathbb{R}$ is player i 's cost function. We denote by $\mathbf{A} = \times_{i \in N} A_i$ the set of action profiles. For every non-empty subset of the players $S \subseteq N$ and every vector $\mathbf{b} = (b_1, \dots, b_n)$ we denote by \mathbf{b}_S the vector $(b_i)_{i \in S}$ where $b_{-i} = b_{N \setminus \{i\}}$.³ An action profile $\mathbf{a} \in \mathbf{A}$ is a *Nash equilibrium* (or just an equilibrium)⁴ if for every player i , $c_i(a_i, \mathbf{a}_{-i}) \leq c_i(b_i, \mathbf{a}_{-i})$ for every $b_i \in A_i$.

Given an underlying game H , a social context game is generated by considering a neighborhood graph over the players, and aggregation functions that determine how the game is affected by that graph. Formally, a *social context* is a tuple $F = (G, (f_i)_{i \in N})$, where $G = (N, E)$ is an undirected graph, and for every i , $f_i : G \times \mathbb{R}^N \rightarrow \mathbb{R}$ is an aggregation function. The aggregation function maps a payoff profile of the underlying game into a utility profile, as a function of the graph structure. The aggregation function captures the agent's social attitude. Given an underlying game $H = (N, (A_i)_{i \in N}, (c_i)_{i \in N})$, and a social context $F = (G, (f_i)_{i \in N})$, a social context game $S = S(H, F) = (N, (A_i)_{i \in N}, (t_i)_{i \in N})$ is a game in strategic form, where N is the set of players, A_i is the set of actions available to player i , and $t_i : \mathbf{A} \rightarrow \mathbb{R}$ satisfies that $t_i(\mathbf{a}) = f_i(c_1(\mathbf{a}), \dots, c_n(\mathbf{a}))$ for every $\mathbf{a} \in \mathbf{A}$.

To distinguish between the costs in the games H and S we will refer from now on to the costs in H as *immediate* costs. Notice that the set of players and the set of actions in the social context game are as in H . The following notation will be useful for us. Denote by $v(i)$ the set of neighbors of i in the graph G and let $g(i) = \{i\} \cup v(i)$ be the *group* of player i .

³ Vectors/profiles are denoted in boldface, and their elements in italic.

⁴ We do not consider mixed strategies in this paper.

Throughout this paper we study social context games with the following aggregation functions:

1. Best-Member Collaboration: a player's cost is the minimal immediate cost in her group. Formally, $f_i(G, \mathbf{c} = (c_1, \dots, c_n)) = \min_{j \in g(i)} c_j$.
2. MinMax Collaboration: a player's cost is the maximal immediate cost in her group. Formally, $f_i(G, \mathbf{c} = (c_1, \dots, c_n)) = \max_{j \in g(i)} c_j$.
3. Surplus Collaboration: a player's cost is the average of the immediate costs of her group. Formally, $f_i(G, \mathbf{c} = (c_1, \dots, c_n)) = \frac{1}{|g(i)|} \sum_{j \in g(i)} c_j$.
4. Competitive Ranking: a player cares about his ranking among the players within her group. We assume that each group is a clique of nodes, and the graph is partitioned into cliques. We need a few notations. Let $c = (c_1, c_2, \dots, c_n)$ be a tuple of immediate costs. Let $l_i(\mathbf{c})$ denote the number of players in $g(i)$ that have a lower immediate cost than i , and let $m_i(\mathbf{c})$ denote the number of players in $g(i)$ with identical immediate cost as i in \mathbf{c} . For every player i , $f_i(G, \mathbf{c}) = l_i(\mathbf{c}) + \frac{m_i(\mathbf{c})}{2}$.

Notice that the aggregation function simply counts the number of players who obtain lower immediate costs assuming that ties are broken randomly. This is a standard practice in work on ranking systems [1].

For convenience, we will overload notation, and write $f_i(G, \mathbf{a})$ to refer to $f_i(G, (c_1(\mathbf{a}), \dots, c_n(\mathbf{a})))$.

In this paper we focus on social context games in which the underlying game H is a *resource selection game*. In a resource selection game there is a set of resources $R = \{1, \dots, m\}$ and for every resource $k \in R$, $w_k : \{1, \dots, n\} \rightarrow \mathbb{R}_+$ is resource k 's cost function. Every player i 's action set is $A_i = R$. Finally, the cost of a player in the game who chooses resource k is $w_k(l)$, where l is the number of players that choose resource k . Formally, $c_i(\mathbf{a}) = w_{a_i}(\sigma_{a_i}(\mathbf{a}))$ where $\sigma_k(\mathbf{a}) = |\{j \in N : a_j = k\}|$ denotes the number of players that choose resource k in \mathbf{a} . We assume the resource cost functions are non-decreasing.

Resource selection games belong to a larger class of games called congestion games. It is well-known [8] that every congestion game possesses a (pure strategy) Nash equilibrium.

3 The Competitive Ranking Game

In this section we study social context games in which the aggregation function is the competitive ranking function. We call such a game a competitive ranking game. We will assume that the graph is a partition into cliques. Hence, every player cares about her ranking (with respect to immediate costs) in her clique. We assume there is some order over the cliques, $t = 1, 2, \dots, T$, where $T \geq 1$. Notice that the case $T = 1$ is the special case of ranking games [2].

Theorem 1. *Let $S = (H, F)$ be a competitive ranking game in which H is a resource selection game with identical resources, and G is a partition into cliques. Then, there exists a Nash equilibrium in S .*

Before the proof we need some preparations.

Definition 1. Let \mathbf{a} be an action profile. Let $i \in N$. Resource $j > (<)a_i$ is a right (left) improvement for player i in \mathbf{a} if $f_i(G, \mathbf{a}) > (<)f_i(G, (j, \mathbf{a}_{-i}))$. We say that player i has a right (left) improvement in \mathbf{a} if there exist a resource j which is a right (left) improvement for i in \mathbf{a} .

Resource j is a minimal left improvement for i in \mathbf{a} , if j is a left improvement for i in \mathbf{a} and there is no resource $j' < j$ s.t. j' is left improvement for i in \mathbf{a} .

We say that there is a right (left) improvement in \mathbf{a} if there exist a player i which has a right (left) improvement in \mathbf{a} .

Definition 2. An action profile \mathbf{a} is ordered if for every pair of players $i, i' \in N$ such that $|g(i)| < |g(i')|$ we have that $a_i \leq a_{i'}$.

As we deal with identical resources we let $w() = w_j()$ for every resource j . Let $v(i, \mathbf{a}, j)$ denote the number of neighbors of player i that choose resource j in the action profile \mathbf{a} ; that is, $v(i, \mathbf{a}, j) = |\{i' \in v(i) : a_{i'} = j\}|$.

The following Lemma is the key to the proof of Theorem 1.

Lemma 1. Let S be a social context game as in Theorem 1. Let $k = \lfloor \frac{n}{m} \rfloor$. Let \mathbf{a} be an ordered action profile such that $\sigma_j(\mathbf{a}) \in \{k, k + 1\}$ for every resource j .

1. $f_i(G, (j, \mathbf{a}_{-i})) < f_i(G, \mathbf{a})$ only if $v(i, \mathbf{a}, j) > v(i, \mathbf{a}, a_i)$, $\sigma_{a_i}(\mathbf{a}) = k + 1$ and $\sigma_j(\mathbf{a}) = k$.
2. Suppose there is no right improvement in \mathbf{a} . If j' is a minimal left improvement for some player i , then there is no right improvement in (j', \mathbf{a}_{-i}) .

The proof the above Lemma is omitted due to lack of space.

Proof of Theorem 1. Our proof is by construction. Let $k = \lfloor \frac{n}{m} \rfloor$. Note that there exist an ordered Nash equilibrium in H in which on every resource there are either k or $k + 1$ players, and let \mathbf{a} be such an equilibrium. Since the resources are identical every permutation of the players in \mathbf{a} is an equilibrium in H . Rename the resources such that if $j < j'$ then $\sigma_j(\mathbf{a}) \leq \sigma_{j'}(\mathbf{a})$.

If $\sigma_j(\mathbf{a}) = k$ for every j then by part 1 of Lemma 1 \mathbf{a} is an equilibrium in S and we are done. Observe by part 1 of Lemma 1 that there is no right improvement in \mathbf{a} . If there is no left improvement in \mathbf{a} then \mathbf{a} is an equilibrium in S . Suppose there exist a player i that has left improvement in \mathbf{a} . Let j be a minimal left improvement for i in \mathbf{a} and let $\mathbf{a}^1 = (j, \mathbf{a}_{-i})$. By Lemma 1 part 2 there is no right improvement in \mathbf{a}^1 . If there is no left improvement in \mathbf{a}^1 then \mathbf{a}^1 is an equilibrium in S . Assume otherwise. Note that \mathbf{a}^1 is ordered. Construct \mathbf{a}^2 from \mathbf{a}^1 in a similar fashion as \mathbf{a}^1 was constructed from \mathbf{a} . After each iteration s there is no right improvement in \mathbf{a}^s and \mathbf{a}^s is ordered. Therefore since the number of players and the number of resources are finite this process is finite and we will eventually end with an equilibrium in S . □

Example 1 shows a game where Theorem 1 is not true when H does not have identical resources. The argument showing non-existence of equilibria is omitted.

Example 1. Let there be a single clique with 5 players. The underlying resource selection game has 3 resources with the following resource cost functions: $w_1 \equiv (0, 8, 10, 100, 100)$, $w_2 \equiv (7, 100, 100, 100, 100)$, and $w_3 \equiv (0, 9, 100, 100, 100)$. \square

4 The Best-Member Collaboration Game

In this section we study social context games with the best-member collaboration aggregation functions. We call such a game a best-member collaboration game. Recall that in such games every player wishes to minimize the minimal immediate cost in its group. Before we state our main result we need the following definition:

Definition 3. Let $G = (V, E)$ be an undirected graph. A subset of nodes $Q \subseteq V$ is called a dominating independent set if every node $v \in V \setminus Q$ has an edge connecting to a node in Q and no two nodes in Q are connected by a single edge. The cardinality of the minimum dominating independent set is denoted by $i(G)$.

Theorem 2. Let $S = (H, F)$ be a best-member collaboration game in which the underlying game H is a resource selection game. If $i(G) < \frac{n}{2}$ then there exists a Nash equilibrium in S .

Proof. Let $T \subseteq N$ be a dominating independent set in G such that $|T| = i(G)$. Let $z = |N \setminus T|$. By our assumption $z > \frac{n}{2}$. Let $j' \in \arg \max_j w_j(z - 1)$. Let $H(j', T)$ be the resource selection game with all resources in H excluding resource j' and the set of players is T . Let \mathbf{b} be an equilibrium in $H(j', T)$. Let \mathbf{a} be the action profile in which all players in T choose the same resource as they choose in \mathbf{b} and for every $i \in N \setminus T$, $a_i = j'$. We claim that \mathbf{a} is an equilibrium in S . Let i be a player such that $a_i = j'$. Since $z > \frac{n}{2}$, for every resource $j \neq j'$, $\sigma_j(\mathbf{a}) < z$. Note that since T is a dominating independent set there exist a player $i' \in g(i)$ such that $a_{i'} \neq j'$. Since \mathbf{a}_T is an equilibrium in $H(j', T)$ then $c_i((j, \mathbf{a}_{-i})) > c_{i'}(\mathbf{a})$ for every resource $j \neq j'$. In addition by the definition of j' , $w_{j'}(\sigma_{j'}((j, \mathbf{a}_{-i}))) > c_{i'}(\mathbf{a})$. Therefore i is not better off by deviating. Suppose $a_i = j$ where $j \neq j'$. Since i has no neighbor in T and \mathbf{a}_T is an equilibrium in $H(j', T)$ deviating to a resource $j'' \neq j'$ is not better off for i . Deviating to j' is also not better off for i by the definition of j' and since $z > \sigma_j(\mathbf{a})$. \square

We next provide an example of a best-member collaboration game for which $i(G) = \frac{n}{2}$, which does not possess an equilibrium, even if graph G is connected. This example implies that our theorem is tight.

Example 2. Consider the following best-member collaboration game. The set of players is $\{1, 2, 3, 4, 5, 6\}$, H has two identical resources with strictly increasing cost functions, and G has a 3-clique on players 1, 2, 3, and each of these players $i \in \{1, 2, 3\}$ has an additional edge to player $i + 3$; thus the degree of each vertex $i \in \{1, 2, 3\}$ is 3 and the degree of each vertex 4, 5, 6 is 1. Clearly, $i(G) = 3$. \square

5 The MinMax Collaboration Game

In this section we study social context games in which the aggregation function is the minmax collaboration. We call such game a minmax collaboration game. Hence, every player wishes to minimize the maximal immediate cost in its group. A necessary condition for a deviation by a single player i to be beneficial is that player i has the maximal immediate cost in $g(i)$ before the deviation and strictly reduces its own immediate cost after deviating. Therefore, we get:

Theorem 3. *Let $S = (H, F)$ be a minmax collaboration game in which H is a resource selection game. Then, $NE(H) \subseteq NE(S)$.*

However, there may exist $\mathbf{a} \in NE(S)$ such that $\mathbf{a} \notin NE(H)$. To see this consider three identical resources with strictly increasing cost functions, and a graph which consists of two cliques both of size 4. Let \mathbf{a} be the action profile in which two players from each clique are on resource 1 and all other players are on resource 2. W.l.o.g. let i be a player on resource. Note that deviating to resource 3 is not beneficial to i since there are still two players in $g(i)$ on resource 2.

6 The Surplus Collaboration Game

In this section we study social context games in which the aggregation function is the surplus collaboration. We call such games surplus collaboration games. In such games every player i wishes to minimize the average immediate costs in $g(i)$. We begin with a couple of negative results.

Proposition 1. *There exist a social collaboration game, in which the underlying game is a resource selection game, that does not possess a Nash equilibrium.*

Proof. We define a social context game with 4 players $\{1, 2, 3, 4\}$ and 2 identical resources $\{1, 2\}$. Each resource has the same cost function $(1, 5, 6, 6)$. The graph G has a 3-clique on the vertices 1, 2, 3 and vertex 4 is an isolated vertex (singleton). We will show now that any assignment of the four players to the two resources does not define a Nash equilibrium of this game. W.l.o.g. player 4 is assigned to resource 2. The reader can verify that for each partition of the players over the resources the resulting profile is not in equilibrium. \square

The example in the proof of Proposition 1 is given for identical resources, but uses a disconnected graph. The following example shows that there may not be a NE even when the graph is connected and in particular is a tree. We omit the proof that the corresponding social context game does not possess a NE.

Example 3. Let SCG be a social context game with the following structure. Let G be an undirected tree with one root and 6 children, and let H have 2 identical resources with the cost function $(1, 1, 2.9, 5, 5, 5, 5)$. \square

Notice that while work on coalitional congestion games has shown the non-existence of pure strategy equilibria (when a coalition may be of size greater than

two), in our study this is shown for the case where only unilateral deviations (and not deviations by a whole coalition) are considered. In the next theorem we provide a family of resource selection games and social contexts, which posses a Nash equilibrium in the surplus collaboration social context.

Theorem 4. *Let H be a resource selection game with m identical resources and let G be a tree with maximal degree $m - 2$. Then, there exists a Nash equilibrium in the corresponding surplus collaboration social context game SCG.*

7 (The Lack of) Potential Functions

It is well known that every potential game is a congestion game and vice-versa. Although the underlying game in each of the social context games we have studied is a potential game we show in this section that none of the games in our positive results is a potential game. We will use the following lemma:

Lemma 2 (Monderer-Shapley'96). *Let $H = (N, (A_i)_{i \in N}, (c_i)_{i \in N})$ be a game in strategic form. H is a potential game if and only if for every pair of players $i, j \in N$, for every $\mathbf{a} \in \mathbf{A}_{N \setminus \{i, j\}}$ and every $a_i, b_i \in A_i$ and $a_j, b_j \in A_j$*

$$(c_i(b_i, a_j, \mathbf{a}) - c_i(a_i, a_j, \mathbf{a})) + (c_j(b_i, b_j, \mathbf{a}) - c_j(b_i, a_j, \mathbf{a})) + \\ (c_i(a_i, b_j, \mathbf{a}) - c_i(b_i, b_j, \mathbf{a})) + (c_j(a_i, a_j, \mathbf{a}) - c_j(a_i, b_j, \mathbf{a})) = 0.$$

Theorem 5. *Resource selection games do not possess a potential function in the competitive ranking, best member collaboration, and minmax collaboration social contexts.*

Proof. Due to space limitations we only provide the argument in case of a competitive ranking game. Let $S = (H, F)$ be the following competitive ranking game. The set of players is $\{1, 2, 3\}$. H is resource selection selection game with 2 identical resources with resource cost functions $w(x) = x$. The graph G is partitioned into two cliques; players 1 and 2 form a clique and player 3 is a singleton. The following cycle of action profiles in which only players 2 and 3 change their actions will provides that there is no potential by Lemma 2: $(1, 1, 2), (1, 2, 2), (1, 2, 1)$ and $(1, 1, 1)$. Note that $f_2(G, (1, 1, 2)) - f_2(G, (1, 2, 2)) + f_3(G, (1, 2, 2)) - f_3(G, (1, 2, 1)) + f_2(G, (1, 2, 1)) - f_2(G, (1, 1, 1)) + f_3(G, (1, 1, 1)) - f_3(G, (1, 1, 2)) = 1/2 - 0 + 0 - 0 + 1 - 1/2 + 0 - 0 = 1 \neq 0$. \square

References

- Altman, A., Tennenholtz, M.: Quantifying incentive compatibility of ranking systems. In: Proc. of AAAI 2006 (2006)
- Brandt, F., Fischer, F., Shoham, Y.: On strictly competitive multi-player games. In: AAAI 2006, pp. 605–612 (2006)
- Fotakis, D., Kontogiannis, S.C., Spirakis, P.G.: Atomic congestion games among coalitions. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 572–583. Springer, Heidelberg (2006)

4. Hayrapetyan, A., Tardos, E., Wexler, T.: The effect of collusion in congestion games. In: STOC 2006, pp. 89–98 (2006)
5. Kuniavsky, S., Smorodinsky, R.: Coalitional Congestion Games. Technical report, Technion, Israel (2007)
6. Monderer, D., Shapley, L.S.: Potential games. *Games and Economic Behavior* 14, 124–143 (1996)
7. Penn, M., Polukarov, M., Tennenholtz, M.: Asynchronous congestion games. In: AA-MAS 2008 (2008)
8. Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)

Approximability and Parameterized Complexity of Minmax Values*

Kristoffer Arnsfelt Hansen**, Thomas Dueholm Hansen, Peter Bro Miltersen,
and Troels Bjerre Sørensen

Department of Computer Science, University of Aarhus, Denmark
{arnsfelt,tdh,bromille,trold}@cs.au.dk

Abstract. We consider approximating the minmax value of a multi-player game in strategic form. Tightening recent bounds by Borgs *et al.*, we observe that approximating the value with a precision of $\epsilon \log n$ digits (for any constant $\epsilon > 0$) is **NP**-hard, where n is the size of the game. On the other hand, approximating the value with a precision of $c \log \log n$ digits (for any constant $c \geq 1$) can be done in quasi-polynomial time. We consider the parameterized complexity of the problem, with the parameter being the number of pure strategies k of the player for which the minmax value is computed. We show that if there are three players, $k = 2$ and there are only two possible rational payoffs, the minmax value is a rational number and can be computed *exactly* in linear time. In the general case, we show that the value can be approximated with any polynomial number of digits of accuracy in time $n^{O(k)}$. On the other hand, we show that minmax value approximation is **W[1]**-hard and hence not likely to be fixed parameter tractable. Concretely, we show that if k -CLIQUE requires time $n^{\Omega(k)}$ then so does minmax value computation.

1 Introduction

A game G in *strategic form* between l players is given by a set of players $\{1, \dots, l\}$ and for each player j a finite *strategy space* S_j and a *utility function* $u_j : S_1 \times S_2 \times \dots \times S_l \rightarrow \mathbf{R}$. In this paper, only the utility function for Player 1 is relevant. When the size of S_j is n_j , we shall refer to the game as an $n_1 \times n_2 \times \dots \times n_l$ game. The *minmax* (or *threat*) value of G for Player 1 is given by $\min_{\sigma_{-1} \in \Delta^{(l-1)}} \max_{a \in S_1} E[u_1(a, \sigma_{-1})]$ where $\Delta^{(l-1)}$ is the set of mixed, but uncorrelated, strategy profiles for players $2, \dots, l$. A profile σ_{-1} achieving the minimum in the expression is called an *optimal minmax profile* or an *optimal threat*. The *maxmin* (or *security*) value of G for Player 1 is given by $\max_{\sigma_1 \in \Delta} \min_{a_2, \dots, a_l} E[u_1(\sigma_1, a_2, \dots, a_l)]$ where Δ is the set of mixed strategies for Player 1.

The minmax value of a finite two-player game is a fundamental notion of game theory. Its mathematical and computational properties are extremely well-studied and well-understood, being intimately tied to the theory of linear programming. In particular, the duality theorem of linear programs implies that the

* Work supported by *Center for Algorithmic Game Theory*, funded by the Carlsberg Foundation.

** Supported by a postdoc fellowship from the Carlsberg Foundation.

minmax value equals the maxmin value. Also, the computation of the minmax value of a two-player game in strategic form is essentially equivalent to solving linear programs and can therefore be done in polynomial time (although a strongly polynomial time algorithm remains an open problem).

Minmax values of *multi-player* games are much less well-studied, although these values are arguably also of fundamental interest to game theory. Most importantly, the minmax value plays a pivotal role in the statement and proof of the so-called *folk theorems* that characterize the Nash equilibria of infinitely repeated games. Additionally, the minmax value is the equilibrium payoff of the so-called *team-maxmin equilibria* studied by von Stengel and Koller [15]. For a multi-player game, the maxmin value may be strictly smaller than the minmax value. Computation of the maxmin value easily reduces to the two-player case and can therefore be done efficiently using linear programming. Rather surprisingly, computation of the minmax value of a multi-player game in strategic form was not studied until very recently, where Borgs *et al.* [1] (motivated by computational aspects of the folk theorem) showed that approximating the minmax value of a three-player game within a certain inverse polynomial additive error is **NP**-hard. Our starting point is this important paper.

Given the fundamental nature of the notion of the minmax value, it is important to understand when the **NP**-hardness result can be circumvented by considering special cases or asking for weaker approximations. The purpose of this paper is to provide a number of results along these lines. First, we observe that the inapproximability result of Borgs *et al.* can be tightened and matched with a positive result, using standard techniques.

Theorem 1. *For any constant $\epsilon > 0$, approximating the minmax value of an $n \times n \times n$ game with 0-1 payoffs within additive error $1/n^\epsilon$ is **NP**-hard. On the other hand there is an algorithm that, given a parameter $\epsilon > 0$ and a game in strategic form with l players each having n strategies and all payoffs being between 0 and 1, approximates the minmax value for Player 1 from above with additive error at most ϵ in time $n^{O(l(\log n)/\epsilon^2)}$.*

This suggests the following important problem: *Can the minmax value of a three-player game with payoffs normalized to $[0, 1]$ be approximated within a non-trivial additive constant (say 0.001 or even 0.499) in polynomial (rather than quasi-polynomial) time?* We leave this problem open.

It is of interest to know when the minmax value can be computed exactly. A prerequisite for this is that it is rational. For three-player games, we characterize when the minmax value for Player 1 can be an irrational number, in terms of the number of strategies of Player 1 and the number of distinct (rational) payoffs. For the special case where the value is guaranteed to be rational we present an optimal linear time algorithm for exactly computing the minmax value [4].

¹ As the algorithms of Theorem [1] and Theorem [2] are very simple, we express their complexity in the unit cost random access machine model. E.g., by “linear time” we mean a linear number of atomic operations in the number of real payoffs of the input. On the other hand, the algorithm of Theorem [3] uses sophisticated algorithms from the literature as subroutines and its complexity is better expressed in the Turing machine model, and in terms of bit complexity.

Theorem 2. Consider $k \times n \times n$ three-player games with only l distinct rational payoffs. When either $k \geq 2$ and $l \geq 3$ or $k \geq 3$ and $l \geq 2$ there exists a game such that the minmax value for Player 1 is irrational. Otherwise, when $k = 2$ and $l = 2$ the minmax value for Player 1 is a rational number and we can compute it exactly in time $O(n^2)$ (on a unit cost random access machine).

Thus having observed that the case of few strategies of Player 1 may be easier than the general case, we apply the approach of *parameterized complexity* [8], considering the number of strategies k of Player 1 as the parameter. Combining a classical result of Shapley and Snow [14] with Renegar's decision procedure for the first order theory of the reals [11,12,13] gives rise to a support enumeration algorithm for finding the minmax value and we show the following.

Theorem 3. Given a $k \times n \times \dots \times n$ l -player game G with rational payoffs and a rational number α so that (G, α) has combined bit complexity L , we can decide in time $L^{O(1)} k^{O(kl)} n^{kl}$ (on a Turing machine) if the minmax value of G for Player 1 is at most α . Using the terminology of fixed parameter complexity theory, considering k the parameter, this problem is in $\mathbf{W}[\mathbf{P}]$, and for the case of 0-1 payoffs in $\mathbf{W}[\mathbf{1}]$.

In particular, if l and k are constants, the complexity is polynomial, and we can approximate the minmax value with any polynomial number of bits of accuracy in polynomial time by using the decision procedure in a binary search. As the exponent in the above complexity bound depends linearly on k with impractical bounds for large k as consequence, we next ask if the problem of approximating the minmax value for Player 1 in a three-player game is *fixed parameter tractable*, i.e., if an algorithm solving the problem in time $f(k)n^c$ exists, where f is any function and c is a constant not depending on k . We provide a reduction from k -CLIQUE that gives negative evidence.

Theorem 4. Deciding k -CLIQUE in a graph with n vertices reduces in polynomial time to approximating the minmax value for Player 1 within $1/(4k^2)$ in a three-player $2k \times kn \times kn$ game with payoffs 0 and 1.

Downey and Fellows [7] proved that the k -CLIQUE problem is complete for the class $\mathbf{W}[\mathbf{1}]$, and hence it immediately follows that the problem of approximating the minmax value within $1/k^2$ for Player 1 in a $k \times n \times n$ game with k being the parameter is hard for $\mathbf{W}[\mathbf{1}]$, even when all payoffs are 0 or 1. Combining this with Theorem 3, we in fact have that the 0-1 case is $\mathbf{W}[\mathbf{1}]$ -complete. Readers not well-versed in the theory of parameterized complexity may find the following consequence of the reduction more appealing: The minmax value of a $k \times n \times n$ three-player game with 0-1 payoffs cannot be approximated in time $n^{o(k)}$, unless k -CLIQUE can be solved in time $n^{o(k)}$. If k -CLIQUE could be solved in time $n^{o(k)}$ then as proved by Chen *et al.* [3] it would follow that all problems in the class \mathbf{SNP} (e.g. 3-SAT) could be solved in time $2^{o(n)}$. Thus, under the assumption that all of \mathbf{SNP} cannot be solved in time $2^{o(n)}$, the algorithm of Theorem 3 is essentially optimal for the case of 0-1 payoffs, in the sense that its complexity is $n^{O(k)}$ and $n^{\Omega(k)}$ is a lower bound.

2 Proofs

2.1 Proof of Theorem 1

We first prove the hardness claim. Borgs *et al.* [1, Theorem 1] showed hardness of approximation with additive error $3/n^2$. Now consider, for a positive integer $c \geq 2$ the following “padding” construction: Given an $n \times n \times n$ game G with strategy space S_i of Player i and utility function u_1 for Player 1. Let $n' = n^c$ and define the $n' \times n' \times n'$ game G' with strategy space $S'_i = S_i \times \{1, \dots, n^{c-1}\}$ and utility function for Player 1 being $u'_1((x, a_1), (y, a_2), (z, a_3)) = u_1(x, y, z)$. In words, G' is simply G with each strategy copied n^{c-1} times. Now, G' and G clearly have the same minmax value. Also, for a given $\epsilon > 0$, by picking c to be a large enough constant, we can ensure that $1/(n')^\epsilon < 3/n^2$, so approximating the minmax value of G within $3/n^2$ reduces to approximating the minmax value of G' within $1/(n')^\epsilon$, which concludes the proof of hardness (we remark that this simple padding argument also yields a somewhat simpler proof of Lemma 7.1 of Chen, Teng and Valiant [4]).

We now proceed with the positive approximation result. We only show the result for the case of three players; the general case being very similar. For the proof, we will use the following theorem by Lipton and Young [10, Theorem 2]:

Theorem 5. *For a two-player zero-sum $n \times n$ game with payoffs in $[0, 1]$, there is a simple strategy for each player that guarantees a payoff within ϵ of the value of the game. Here, a simple strategy is one that mixes uniformly on a multiset of $\lceil \ln n / (2\epsilon^2) \rceil$ pure strategies.*

Now consider a given 3-player game G and consider the optimal threat strategy profile (σ_2, σ_3) of Players 2 and 3 against Player 1. Consider σ_3 as fixed and look at the resulting two-player game G' between Player 1 (maximizer) and Player 2 (minimizer). Clearly, this game has value equal to the minmax value for Player 1 in G . Applying Theorem 5 there is a simple strategy σ'_2 for Player 2 that guarantees this value within ϵ . Fix σ'_2 to this strategy and look at the resulting two-player game G'' between Player 1 and Player 3. By construction of σ'_2 , this game has value at most ϵ larger than the value of G' . Applying Theorem 5 again, there is a simple strategy σ'_3 for Player 3 that guarantees this value within ϵ . Thus, if Player 2 and Player 3 play the profile (σ'_2, σ'_3) in the original game, they are guaranteed the minmax value of G plus at most 2ϵ .

Now, given some ϵ' , we let $\epsilon = \epsilon'/2$ and approximate the threat value of Player 1 within ϵ' by exhaustively searching through all pairs of simple strategies for Player 2 and Player 3, finding the best response of Player 1 to each of them and the associated payoff, and returning the lowest such payoff. This completes the proof of the theorem.

It is natural to ask if one can get any non-trivial approximation by considering strategies that mix uniformly over only a constant size multiset, as this would lead to a polynomial time approximation algorithm rather than a quasi-polynomial one. Unfortunately, the answer is negative: For given n and s , let m be maximal such that $\binom{m}{s}^2 \leq n$. Then

$$n < \binom{m+1}{s}^2 \leq \binom{2m}{s}^2 \leq \left(\frac{2me}{s}\right)^{2s}.$$

Consider the $\binom{m}{s}^2 \times m \times m$ game G defined as follows. For every two subsets T_2, T_3 of pure strategies of size s for Player 2 and Player 3 there is a pure strategy, a_{T_2, T_3} for Player 1 so that $u_1(a_{T_2, T_3}, a_2, a_3) = 1$ for $a_2 \in T_2, a_3 \in T_3$ and $u_1(a_{T_2, T_3}, a_2, a_3) = 0$ for $a_2 \notin T_2$ or $a_3 \notin T_3$. If Player 2 and Player 3 both play a uniform mix on their entire strategy spaces, Player 1 can ensure payoff at most $(\frac{s}{m})^2$. On the other hand, if Player 2 and Player 3 play mixed strategies of support size at most s then Player 1 has a reply ensuring payoff 1. We can now, in a similar way as in the construction in the beginning of this section, construct a padded version of the game, obtaining an $n \times n \times n$ game G' such that the minmax value for Player 1 is at most $(\frac{s}{m})^2 < \frac{(2e)^2}{n^{1/s}}$, but for every strategy profile for Player 2 and Player 3 of support size at most s , Player 1 can ensure payoff 1. Thus to approximate the minmax value within some constant $c < \frac{1}{2} - \frac{2e^2}{n^{1/s}}$, we must have $s > \frac{\ln n}{\ln(\frac{4e^2}{1-2c})}$.

2.2 Proof of Theorem 2

First, we give the claimed examples of games for which the minmax value for Player 1 is irrational. We describe each game by a matrix for each action of Player 1, where rows and columns correspond to the actions of Player 2 and Player 3, respectively. That is, we let $u_1(i, j, k) = A_i(j, k)$.

The first game is a $2 \times 2 \times 2$ game where there are 3 distinct payoffs, given by the following matrices.

$$A_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad A_2 = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix}$$

It is easy to see that the minmax strategy profile for Player 2 and Player 3 is the profile where both row 1 and column 1 are played with probability $2 - \sqrt{2}$. This results in the minmax value $6 - 4\sqrt{2}$ for Player 1.

The second game is a $3 \times 2 \times 2$ game where there are 2 distinct payoffs, given by the following matrices.

$$A_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad A_2 = \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix} \quad A_3 = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$$

It is easy to see that the minmax strategy profile for Player 2 and Player 3 is the profile where both row 1 and column 1 are played with probability $\frac{\sqrt{5}-1}{2}$. This results in the minmax value $\frac{3-\sqrt{5}}{2}$ for Player 1.

We now examine the special case where there are only two distinct payoffs, and Player 1 only has two possible actions. For this case, we show that the minmax value is a rational number, and that the optimal threat can be computed in linear time. We assume without loss of generality that the two possible payoffs are 0 and 1. The proof is a case analysis, where each case can be identified and

solved in linear time, assuming that none of the previous cases apply. Case 1 and 2 are the trivial cases where either side has a pure optimal strategy.

Case 1: $\exists i \forall j, k : u_1(i, j, k) = 1$. Player 1 has a “safe” action, i , such that no matter what Players 2 and 3 do, Player 1 achieves value 1. Any strategy profile for Players 2 and 3 is an optimal threat, with minmax value 1.

Case 2: $\exists j, k \forall i : u_1(i, j, k) = 0$. The strategy profile (j, k) is an optimal threat, with minmax value 0.

Case 1 and case 2 can easily be identified and solved in linear time. Notice that when we are not in case 2, we have that $\forall j, k \exists i : u_1(i, j, k) = 1$, and therefore $u_1(i, j, k) = 0 \Rightarrow u_1(i', j, k) = 1$ where $i' \neq i$. This means that Player 1 has a maxmin (security) value of at least $\frac{1}{2}$, which can be achieved by a uniform mix of the two strategies. As the minmax value is at least the maxmin value, any threat with minmax value $\frac{1}{2}$ will be optimal. This is exactly what can be achieved in the next two cases:

Case 3: $\exists j \forall i \exists k : u_1(i, j, k) = 0$. Player 2 has a pure strategy, such that Player 3 can play matching pennies against Player 1. Let k and k' be the strategies of Player 3 achieving payoff 0 against i and i' respectively. Then $(j, (\frac{1}{2}k, \frac{1}{2}k'))$ is an optimal threat, with minmax value of $\frac{1}{2}$.

Case 4: $\exists k \forall i \exists j : u_1(i, j, k) = 0$. Player 3 has a pure strategy, such that Player 2 can play matching pennies against Player 1. Let j and j' be the strategies of Player 2 achieving payoff 0 against i and i' respectively. Then $((\frac{1}{2}j, \frac{1}{2}j'), k)$ is an optimal threat, with minmax value of $\frac{1}{2}$.

Case 3 and case 4 can again easily be identified and solved in linear time.

Case 5: None of the above. The negation of case 1 implies $\forall i \exists j, k : u_1(i, j, k) = 0$. The negation of cases 2, 3 and 4 implies $u_1(i, j, k) = 0 \Rightarrow \forall j', k' : u_1(i', j', k) = u_1(i', j, k') = 1$, where $i' \neq i$. That is, any strategy of Player 2 or 3 can achieve payoff 0 against at most one of Player 1’s strategies, and Players 2 and 3 must agree on which strategy of Player 1 to try to get payoff 0 against. If they disagree, the payoff is 1 no matter what Player 1 does. The best they can hope for is therefore $\min_{p, q \in [0;1]} \max\{1 - pq, 1 - (1 - p)(1 - q)\}$, which gives a lower bound on the minmax value of $\frac{3}{4}$. This value can be achieved in this case in the following way: let $u_1(i, j, k) = u_1(i', j', k') = 0$. Then $((\frac{1}{2}j, \frac{1}{2}j'), (\frac{1}{2}k, \frac{1}{2}k'))$ is an optimal threat, with minmax value of $\frac{3}{4}$.

2.3 Proof of Theorem 3

We prove the theorem for three players, the general case is similar. Shapley and Snow [14] showed that every $k \times n$ zero-sum game has a minmax mixed strategy for Player 2 of support at most k , i.e., using at most k pure strategies. We claim that from this it follows that in every $k \times n \times n$ game there are mixed strategies for Player 2 and Player 3 of support at most k so that the resulting strategy profile σ_{-1} yields the minmax value for Player 1 when Player 1 chooses a best response. Indeed, consider the actual minmax strategy profile $\sigma_{-1}^* = (\sigma_2^*, \sigma_3^*)$. If we consider σ_3^* fixed and consider the resulting two-player game between Player

1 and Player 2, it is clear that σ_2^* is a minmax strategy of this game and that Player 2 will still guarantee the minmax payoff by playing the minmax strategy σ_2^* of support k which is guaranteed to exist by Shapley and Snow’s result. Similarly, we may replace σ_3^* with a strategy of support k without changing the payoff resulting when Player 1 plays a best response.

Our algorithm is a support enumeration algorithm which exhaustively examines each possible support of size k for Player 2 and Player 3. From the above observation it follows that the minmax value of the game is the minimum of the minmax value of each of the resulting $k \times k \times k$ subgames. Therefore, we only have to explain how to compare the minmax value of such a subgame to a given α , and we will be done. For this, we appeal to classical results on the first order theory of the reals.

The decision procedure for the first order theory of the reals due to Renegar [11][12][13] can decide a sentence with $\omega - 1$ quantifier alternations, the k th block of variables being of size n_k , containing m atomic predicates and involving only polynomials of degree at most d with integer coefficients of maximum bit length L using $L(\log L)2^{O(\log^* L)}(md)^{2^{O(\omega)}} \prod_k n_k$ bit operations² and $(md)^{O(\sum_k n_k)}$ evaluations of the Boolean formula of the atomic predicates. We claim that from this it follows that given a $k \times k \times k$ game G with rational payoffs and a rational number α so that (G, α) has combined bit complexity L , we can decide in time $L(\log L)2^{O(\log^* L)}k^{O(k)}$ (on a Turing machine) if the minmax value of G for Player 1 is at most α . We can assume that the payoffs and α are integers at the expense of increasing the bitlength of every number to at most the combined bitlength of the original problem. Define the following polynomials in $2k$ variables.

$$p_l(x_1, \dots, x_{2k}) = \sum_{i=1}^k \sum_{j=k+1}^{2k} u_1(l, i, j - k)x_i x_j \quad , \quad r_i(x_1, \dots, x_{2k}) = x_i$$

$$q_1(x_1, \dots, x_{2k}) = \sum_{i=1}^k x_i \quad , \quad q_2(x_1, \dots, x_{2k}) = \sum_{i=k+1}^{2k} x_i$$

The sentence we must decide is then

$$(\exists x \in \mathbf{R}^{2k}) [p_1(x) \leq \alpha \wedge \dots \wedge p_k(x) \leq \alpha \wedge q_1(x) = 1 \wedge q_2(x) = 1 \wedge r_1(x) \geq 0 \wedge \dots \wedge r_{2k}(x) \geq 0] \ .$$

For this sentence we have $\omega = 1$, $m = 3k + 2$, $d = 2$ and $n_1 = 2k$, and the sentence can thus be decided in the claimed running time using Renegar’s procedure. For the support enumeration algorithm this decision procedure must be invoked for $\binom{n}{k}^2$ different $k \times k \times k$ subgames, and the claimed time bound of the statement of the theorem follows.

Next, we show how to use this algorithm to show $\mathbf{W[P]}$ and $\mathbf{W[1]}$ membership of the two versions of the problem. We use the framework of afpt-programs of

² The bound stated here is the improvement of the bound stated by Renegar due to the recent breakthrough in integer multiplication due to Fürer [9].

Chen, Flum and Grohe [5] and Buss and Islam [2] to do this. To transform the algorithm into an afpt-program showing that the decision problem is in the class $\mathbf{W}[\mathbf{P}]$, we simply replace the enumeration by existential steps guessing the sets of indices of size k giving the support of the strategies of Player 2 and Player 3. In the remainder of this section we will show that for the special case of 0-1 payoffs the decision problem is in the class $\mathbf{W}[\mathbf{1}]$. The idea is to precompute, for every possible $k \times k \times k$ game with 0-1 payoffs, whether the minmax value for Player 1 is at most α . As in the $\mathbf{W}[\mathbf{P}]$ case, indices of the support of the strategies are guessed, but now the $k \times k \times k$ subgame is used as an index in the precomputed table. To see that this can be turned into an appropriate afpt-program, we will formally define the relations used.

Assume that the payoffs of Player 1 are given as a k -tuple of $n \times n$ 0-1 matrices (U^1, \dots, U^k) . Define a unary relation A over k -tuples of $k \times k$ 0-1 matrices as follows: $(M^1, \dots, M^k) \in A$ if and only if the minmax value for Player 1 in the $k \times k \times k$ subgame given by (M^1, \dots, M^k) is at most α .

Define a 6-ary relation B having as first argument a k -tuple of $k \times k$ 0-1 matrices and with the next 3 arguments being indices from 1 to k and the two last arguments being indices from 1 to n as follows.

$$((M^1, \dots, M^k), l, a, b, i, j) \in B \quad \text{if and only if} \quad M_{ab}^l = U_{ij}^l .$$

The algorithm first computes the relations A and B . In the guessing steps the algorithm guesses a k -tuple of matrices (M^1, \dots, M^k) and indices i_1, \dots, i_k and j_1, \dots, j_k . The final checks the algorithm must perform are $(M^1, \dots, M^k) \in A$ and $((M^1, \dots, M^k), l, a, b, i_a, j_b) \in B$ for all $l, a, b \in \{1, \dots, k\}$. The number of steps used for guessing the indices and the final checks is a function depending only on the parameter k , as required.

As discussed by Buss and Islam [2] we can in a generic way transform an algorithm utilizing a constant number of relations into one that only utilizes a single binary relation, thereby obtaining an afpt-algorithm showing that the decision problem is in $\mathbf{W}[\mathbf{1}]$.

2.4 Proof of Theorem 4

Before starting the proof, we remark that the reduction is based on similar ideas as the reduction proving \mathbf{NP} -hardness by Borgs *et al.* However, they reduce from 3-COLORING rather than k -CLIQUE, and in their coloring based games, we don't see how to restrict the strategy space of Player 1 to a small number of strategies, so as to obtain fixed-parameter intractability.

We now describe the reduction. We assume throughout the proof that $k \geq 5$. Given an undirected graph $G = (V, E)$, with $|V| = n$. We construct a $2k \times kn \times kn$ game from G in the following way. Let $S_1 = \{1, \dots, k\} \times \{2, 3\}$ be the strategy space of Player 1 and $S_2 = S_3 = \{1, \dots, k\} \times V$ be the strategy spaces of Player 2 and Player 3. We define the payoff of Player 1 as:

$$u_1((x_1, i), (x_2, v_2), (x_3, v_3)) = \begin{cases} 1 & \text{if } x_1 = x_i \\ 1 & \text{if } x_2 = x_3 \text{ and } v_2 \neq v_3 \\ 1 & \text{if } x_2 \neq x_3 \text{ and } v_2 = v_3 \\ 1 & \text{if } v_2 \neq v_3 \text{ and } (v_2, v_3) \notin E \\ 0 & \text{otherwise} \end{cases}$$

As Player 2 and Player 3 try to minimize the payoff of Player 1 we shall refer to them as *bullies*. One can think of the game as the bullies each choosing a label and a vertex of G . Player 1 then chooses one of the bullies and tries to guess his label. If he guesses correctly he will get a payoff of 1. If not, he will get a payoff of 0, unless the bullies do one of the following:

- (i) Choose the same label, but different vertices.
- (ii) Choose different labels, but the same vertex.
- (iii) Choose a pair of distinct vertices that does not correspond to an edge.

In these cases he will get a payoff of 1. The intuition behind the proof is that the bullies will be able to avoid these three cases if the graph contains a k -clique, thereby better punishing Player 1. Formally, we shall show that the minmax value for Player 1 is exactly $\frac{1}{k}$ if G contains a k -clique and at least $\frac{1}{k} + \frac{1}{4k^2}$ if G does not contain a k -clique.

First, we notice that if G contains a k -clique, the bullies can bring down the payoff of Player 1 to $1/k$ by choosing a vertex from the k -clique uniformly at random and agreeing on a labeling of the vertices. For given strategies of the bullies, let p_{\max} be the highest probability (over j) of any of the bullies choosing a label $j \in \{1, \dots, k\}$. Player 1 will always be able to get a payoff of p_{\max} by choosing j and the corresponding player. It follows that the minmax value of the game is $1/k$ when G contains a k -clique, as desired.

Next, we consider the case when G contains no k -clique. Assume to the contrary that the bullies can force Player 1 to get a payoff less than $1/k + 1/(4k^2)$. For the rest of the proof, consider a fixed strategy profile of the bullies with this property. We have already seen that in this case, $p_{\max} < 1/k + 1/(4k^2)$.

Consider the case where Player 1 always chooses Player 2 and guesses a label uniformly at random. In this case Player 1 will always guess the correct label with probability $1/k$ independently of the actions of the bullies. Let p be the probability of either (i), (ii) or (iii) happening. We have:

$$\frac{1}{k} + \left(1 - \frac{1}{k}\right)p < \frac{1}{k} + \frac{1}{4k^2} \Rightarrow p < \frac{1}{4k(k-1)} . \tag{1}$$

In particular, the probability of (i) happening is less than $1/(4k(k-1))$. Let p_{\min} be the minimum probability assigned to any label by either of the bullies. We have $p_{\min} \geq 1 - (k-1)p_{\max} > 3/(4k) + 1/(4k^2)$. Let $(x, v)_i$ denote the event that Player i chooses the label x and the vertex v . We will use \cdot as a wildcard, such that $(x, \cdot)_i$ denotes the event that Player i chooses the label x and $(\cdot, v)_i$ denotes the event that Player i chooses the vertex v . For $v, w \in V$ we see that:

$$\begin{aligned}
 \frac{1}{4k(k-1)} &> \Pr[x_2 = x_3 \text{ and } v_2 \neq v_3] \\
 &= \sum_{j=1}^k \sum_{v \neq w} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(j, \cdot)_2] \Pr[(\cdot, w)_3 \mid (j, \cdot)_3] \Pr[(j, \cdot)_3] \\
 &\geq p_{\min}^2 \sum_{j=1}^k \sum_{v \neq w} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(\cdot, w)_3 \mid (j, \cdot)_3] \\
 &= p_{\min}^2 \sum_{j=1}^k \left(1 - \sum_{v \in V} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(\cdot, v)_3 \mid (j, \cdot)_3] \right) \\
 &= kp_{\min}^2 - p_{\min}^2 \sum_{j=1}^k \sum_{v \in V} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(\cdot, v)_3 \mid (j, \cdot)_3] .
 \end{aligned}$$

Here, the first inequality follows from (III). The second inequality is by definition of p_{\min} .

We now estimate the probability that Player 2 and Player 3 choose the same vertex, conditioned on the fact that they choose the same label l . We have that for all l in $\{1, \dots, k\}$:

$$\begin{aligned}
 &\sum_{v \in V} \Pr[(\cdot, v)_2 \mid (l, \cdot)_2] \Pr[(\cdot, v)_3 \mid (l, \cdot)_3] \\
 &> k - \frac{1}{4k(k-1)p_{\min}^2} - \sum_{j \neq l} \sum_{v \in V} \Pr[(\cdot, v)_2 \mid (j, \cdot)_2] \Pr[(\cdot, v)_3 \mid (j, \cdot)_3] \\
 &\geq 1 - \frac{1}{4k(k-1)p_{\min}^2} > 1 - \frac{1}{4k(k-1) \left(\frac{3}{4k} + \frac{1}{4k^2}\right)^2} \\
 &= 1 - \frac{4k^3}{(k-1)(3k+1)^2} > \frac{1}{2} .
 \end{aligned}$$

The second inequality follows from observing that all terms in the outer sum are bounded from above by 1. The last inequality follows from $k \geq 5$.

For fixed l , let v_i^l be the vertex chosen with highest probability by player i given that he chooses label l and let $q_i^l = \Pr[(\cdot, v_i^l)_i \mid (l, \cdot)_i]$ be that probability. Since $\sum_{v \in V} \Pr[(\cdot, v)_2 \mid (l, \cdot)_2] \Pr[(\cdot, v)_3 \mid (l, \cdot)_3] > \frac{1}{2}$ and $(\Pr[(\cdot, v)_2 \mid (l, \cdot)_2])_v$ and $(\Pr[(\cdot, v)_3 \mid (l, \cdot)_3])_v$ are probability distributions, we have

$$\forall i \in \{2, 3\}, \forall l \in \{1, \dots, k\} : q_i^l > \frac{1}{2} .$$

Since q_2^l and q_3^l are both strictly bigger than $\frac{1}{2}$, and the probability that Player 2 and Player 3 choose the same vertex is strictly bigger than $\frac{1}{2}$, it is easy to see that we must have $v_2^l = v_3^l$ and we will therefore simply refer to this vertex as v^l . That is, for every label j the bullies agree on some vertex v^j that they both choose with high probability when choosing j . For $j, l \in \{1, \dots, k\}$, it will either be the case that there exists some $v^j = v^l$, with $j \neq l$, or that all the v^j 's are

distinct. In the first case Player 1 will, with high probability, get a payoff of 1 when one of the bullies chooses label j and the other chooses label l (case (ii)). In the second case there will exist a pair of distinct labels j and l , such that there is no edge between v^j and v^l , since the graph doesn't contain a k -clique. Hence, this will cause Player 1 to get a payoff of 1, with high probability, when one of the bullies chooses label j and the other chooses label l (case (iii)). In both cases we get that the probability that (ii) or (iii) holds is at least

$$\sum_{i=0}^1 \Pr[(\cdot, v^j)_{2+i} \mid (j, \cdot)_{2+i}] \Pr[(j, \cdot)_{2+i}] \Pr[(\cdot, v^l)_{3-i} \mid (l, \cdot)_{3-i}] \Pr[(l, \cdot)_{3-i}] > 2p_{\min}^2 \left(\frac{1}{2}\right)^2 > \frac{1}{2} \left(\frac{3}{4k} + \frac{1}{4k^2}\right)^2 = \frac{(3k+1)^2}{32k^4} > \frac{1}{4k(k-1)} .$$

The last inequality follows from $k \geq 5$. This contradicts \square which states that (i), (ii) or (iii) happens with probability less than $1/(4k(k-1))$. Thus, we have completed our proof by contradiction.

3 Conclusions and Open Problems

As mentioned above, an important open problem is achieving a non-trivial approximation of the minmax value of an $n \times n \times n$ game in polynomial time, rather than quasi-polynomial time. Another interesting question comes from the following notions: The *threat point* of a game is defined to be the vector of minmax values for each of its players. We may consider approximating the threat point of a three player game where one of the players has few strategies. For this, we have to consider the problem of approximating the threat value for Player 1 in a three-player $n \times k \times n$ game. That is, it is now one of the two ‘‘bullies’’ rather than the threatened player that has few strategies. We observe that for constant $\epsilon > 0$ such an approximation can be done efficiently by simply discretizing the mixed strategy space of the player with few strategies using a lattice with all simplex points having distance at most ϵ to some lattice point, and then for each lattice point solving the game for the remaining two players using linear programming. Combining this with Theorem 3 gives us the following corollary:

Corollary 1. *There is an algorithm that, given a $k \times n \times n$ game with 0-1 payoffs and an $\epsilon > 0$, computes the threat point within additive error ϵ in time $(n/\epsilon)^{O(k)}$.*

The discretization technique gives algorithms with poor dependence on the desired additive approximation ϵ . We leave as an open problem if the minmax value of an $n \times k \times n$ 0-1 game can be approximated within ϵ in time $(n \log(1/\epsilon))^{O(k)}$.

Acknowledgements

We would like to thank Bernhard von Stengel and Daniel Andersson for valuable help.

References

1. Borgs, C., Chayes, J., Immorlica, N., Kalai, A.T., Mirrokni, V., Papadimitriou, C.: The myth of the folk theorem. In: Proceedings of the 40th annual ACM Symposium on Theory of Computing (STOC 2008), pp. 365–372 (2008)
2. Buss, J.F., Islam, T.: Simplifying the weft hierarchy. *Theoretical Computer Science* 351(3), 303–313 (2006)
3. Chen, J., Chor, B., Fellows, M., Huang, X., Juedes, D.W., Kanj, I.A., Xia, G.: Tight lower bounds for certain parameterized NP-hard problems. *Information and Computation* 201(2), 216–231 (2005)
4. Chen, X., Teng, S.-H., Valiant, P.: The approximation complexity of win-lose games. In: Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007), pp. 159–168 (2007)
5. Chen, Y., Flum, J., Grohe, M.: Machine-based methods in parameterized complexity theory. *Theoretical Computer Science* 339(2–3), 167–199 (2005)
6. Downey, R.G., Fellows, M.R.: Fixed-parameter tractability and completeness I: Basic results. *SIAM Journal on Computing* 24(4), 873–921 (1995)
7. Downey, R.G., Fellows, M.R.: Fixed-parameter tractability and completeness II: On completeness for $W[1]$. *Theoretical Computer Science* 141(1–2), 109–131 (1995)
8. Downey, R.G., Fellows, M.R.: Parameterized complexity. Monographs in computer science. Springer, New York (1999)
9. Fürer, M.: Faster integer multiplication. In: Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC 2007), pp. 57–66 (2007)
10. Lipton, R.J., Young, N.E.: Simple strategies for large zero-sum games with applications to complexity theory. In: Proceedings of the 26th annual ACM Symposium on Theory of Computing (STOC 1994), pp. 734–740 (1994)
11. Renegar, J.: On the computational complexity and geometry of the first-order theory of the reals, part I: Introduction. Preliminaries. The geometry of semi-algebraic sets. The decision problem for the existential theory of the reals. *Journal of Symbolic Computation* 13(3), 255–299 (1992)
12. Renegar, J.: On the computational complexity and geometry of the first-order theory of the reals, part II: The general decision problem. Preliminaries for quantifier elimination. *Journal of Symbolic Computation* 13(3), 301–327 (1992)
13. Renegar, J.: On the computational complexity and geometry of the first-order theory of the reals. part III: Quantifier elimination. *Journal of Symbolic Computation* 13(3), 329–352 (1992)
14. Shapley, L.S., Snow, R.N.: Basic solutions of discrete games. In: Contributions to the Theory of Games. *Annals of Mathematics Studies* 24, pp. 27–35. Princeton University Press, Princeton (1950)
15. von Stengel, B., Koller, D.: Team-maxmin equilibria. *Games and Economic Behavior* 21, 309–321 (1997)

An “Ethical” Game-Theoretic Solution Concept for Two-Player Perfect-Information Games

Joshua Letchford¹, Vincent Conitzer¹, and Kamal Jain²

¹ Department of Computer Science, Duke University, Durham, NC, USA

{jcl, conitzer}@cs.duke.edu

² Microsoft Research, Redmond, WA, USA

kamalj@microsoft.com

Abstract. The standard solution concept for perfect-information extensive form games is subgame perfect Nash equilibrium. However, humans do not always play according to a subgame perfect Nash equilibrium, especially in games where it is possible for all the players to obtain much higher payoffs if they place some trust in each other (and this trust is not violated). In this paper, we introduce a new solution concept for two-player perfect-information games that attempts to model this type of trusting behavior (together with the “ethical” behavior of not violating that trust). The concept takes subgame perfect equilibrium as a starting point, but then repeatedly resolves the game based on the players being able to trust each other. We give two distinct algorithmic definitions of the concept and show that they are equivalent. Finally, we give a fast implementation of one of the algorithms for solving the game, and show that it runs in time $O(n \log n + nh \log(n/h))$.

1 Introduction

Under a typical game-theoretic solution concept, the players pursue nothing other than their own interest at every point in the game. Humans, however, do not always behave this way: depending on what happened earlier in the game, they may feel that they “owe” another player something and act accordingly. We propose a solution concept for two-player extensive-information games that attempts to model this phenomenon.

To illustrate the basic idea, consider the example game in Figure 1. The standard game-theoretic approach to solving this game is to simply use backward induction. If player 2 gets to move, he maximizes his utility by moving left, resulting in the utilities $(0, 2)$. Anticipating this, player 1 will choose to move left in the first move, resulting in the utilities $(1, 0)$. This is the unique subgame perfect equilibrium of the game. We note that both players would prefer the rightmost outcome, which has utilities $(2, 1)$, but the strategic structure of the game prevents this outcome from occurring—at least within the standard game-theoretic approach.

Now, we argue that this is not necessarily the most sensible outcome of the game, assuming that the players have some amount of decency. Suppose player 1 does, in fact, move right. In the standard game-theoretic approach, this would be considered a mistake. However, suppose that it is common knowledge among the players that they

¹ We use “she” for player 1 and “he” for player 2.




Fig. 1. A simple example

understand the game perfectly. Hence, player 2 knows that player 1 did not choose this move by accident, but voluntarily chose to let player 2 pick between the outcome that is better for both of them than the subgame perfect solution, and the outcome that is much better for player 2 but worse for player 1. Player 1 knows very well that she is leaving herself vulnerable to a selfish action by player 2, but chose to move right anyway, with the hope of a better outcome for both. It seems sensible to argue that in this case, it would be unethical for player 2 to move left. Specifically, it seems that player 2 “owes” it to player 1 to give her at least as much utility as she would have received in the subgame perfect equilibrium, especially as player 2 can do so in a way that also gives him at least as much utility as he would have received in the subgame perfect equilibrium. Thus, it seems that the ethical thing to do for player 2 in this situation is to move right; if player 1 believes that player 2 is ethical in this way, then she prefers to move right initially—she “trusts” player 2 to make the “ethical” move. In this paper, we propose a general solution concept corresponding to this ethical type of reasoning.

Incidentally, the simple game above closely resembles a game studied in experimental game theory, called the “trust game.” In the trust game, player 1 has an initial budget. She can choose to give any amount not exceeding this budget to player 2; if she does so, the money will be tripled before player 2 receives it. After receiving the money, player 2 can give any amount back to player 1 (this will not be tripled), and the game ends after this. Again, this game can be solved by backwards induction: it is never in player 2’s interest to give any money back, and hence player 1 should give player 2 no money at all [2]. Experimentally, however, this is not at all what happens [6, 14, 15]. In an experimental study, 85% of subjects in the player 1 role gave at least some money, and 98% of subjects in the player 2 role that received some money gave some back [14]. Also, on average, subjects in the player 1 role gave \$5.52 (out of their initial \$10), and subjects in the player 2 role returned \$6.96 [15]. The full version of our paper discusses what our solution concept prescribes for this game.

A few more remarks are in order. We do not wish to argue that the behavior prescribed by our solution concept is the only behavior that can possibly be described as “ethical.” For example, in a modified version of the trust game where player 2 does not have the option of giving money back at all, our solution concept prescribes that player

² This assumes that a player’s utility is simply the amount of money that the player receives.

I should give no money; but one could perhaps argue that giving money would still be the ethical thing to do, given that the money will be tripled. In fact, under a strict utilitarian framework, one might argue that the ethical thing to do is to transfer all the money. Still, we argue that our solution concept corresponds to a particular, natural (if perhaps limited) type of ethical behavior. For the purposes of this paper, we will avoid discussion of whether our concept is more “rational” than the standard game-theoretic concepts, and hence we will avoid the use of the word “rational.”

Also, while there has been an agenda within game theory of justifying cooperative behavior by showing that cooperation can be sustained as an equilibrium of a repeated game (for instance, in the Prisoner’s Dilemma [13]), philosophically, this paper does not fall under that agenda. (However, because our solutions always Pareto dominate or are equal to a subgame perfect solution, they can in fact be sustained as an equilibrium of the repeated game as well.)

Solution concepts that model this type of ethical behavior potentially have a number of applications. They can be used to predict human behavior. Concepts can be used in artificially intelligent agents, for interacting either with humans or with each other. Indeed, it has been argued that standard game-theoretic solutions do not always perform well in settings where artificially intelligent agents interact with humans [3,4,12]. The design of artificial intelligence that behaves ethically has previously received attention [11,10]. Much of this work relies on humans specifying examples of ethical behavior, which the agent then tries to generalize into more general rules [5,9]. Other work specifies certain *prima facie* duties, and the agent needs to learn from labeled examples how to trade off these duties when they conflict [2]. Our work differs from this prior work in that we define a single concept that is intended to capture a subset of ethical behavior, and all that remains to be done is to find the corresponding solution (no learning is needed).

The rest of this paper is laid out as follows. In Section 2, we study some more complex examples to get some intuition about our solution concept. In Section 3, we give a first definition of our solution concept, which relies on iteratively modifying the agents’ preferences and re-solving for the subgame perfect equilibrium. In Section 4, we give another definition of the concept, which relies on iteratively removing nodes from the game tree and re-solving for the subgame perfect equilibrium; we show that this definition is equivalent to the one from Section 3. Finally, in Section 5, we give a fast algorithm for computing a solution according to our concept.

2 Introductory Examples

In this section, we study two additional example games. The first example shows a seemingly more complex game that can be simplified to be similar to the example in Figure 1. The second example is inherently more complex; however, understanding this example will help significantly to understand the general definition.

Example: A game with moves by Nature. Alice and Bob are sitting next to each other on a plane, and there are not enough pillows on the plane. Alice has a pillow (it was sitting in her seat), and Bob does not. Alice is currently not tired, and Bob is (and, from




Fig. 2. Example: airplane pillows. Key: give pillow, keeppillow, Alice becomes tired, Alice does not become tired, give when Alice becomes tired, keeppillow when Alice becomes tired.

their demeanors, this is common knowledge). Alice could give the pillow to Bob, but she might regret it if she gets tired later on. Of course, Bob could give the pillow back in that case, if he chooses to do so.³ A possible interpretation of this game is shown in Figure 2. On the left side is the full game tree: Alice first decides whether to give the pillow, then Nature decides whether Alice gets tired, and finally Bob decides whether to return the pillow. (We note that if Alice is not tired, she slightly prefers not having the pillow, to have some more space.) We emphasize that this is a perfect-information game. Because of that, we can remove Nature from the game by taking expectations, resulting in the game on the right-hand side. By similar reasoning as that for the example in Figure 1, Alice should give the pillow, and Bob should give it back if Alice is tired. This contrasts with the subgame perfect solution in which Bob would not return the pillow, so that Alice keeps the pillow to herself; the subgame perfect solution is worse for both players.

Example: A more complex game with 6 leaves. We now move on to an example that is fundamentally more complex and that will require some more reflection on what is ethical. Consider the example in Figure 3.




Fig. 3. A more complex example with 6 leaves

Backward induction would tell us that player 2 will move left in each subtree, and hence player 1 should move left, resulting in the unique subgame perfect equilibrium with utilities (2, 0). However, again, we may argue that if player 1 chooses middle or right, then player 2 owes it to player 1 to give her at least 2 (since she could have guaranteed herself this much, and to give her this much player 2 does not need to accept

³ This is ignoring the potential complication that Bob may have fallen asleep on the pillow.

a utility less than the 0 that he would receive in the subgame perfect solution). That is, if player 1 plays middle, player 2 should play right (resulting in utilities (4, 1)); and if player 1 plays right, player 2 should play middle or right—but middle will give him a higher utility, resulting in utilities (3, 4). Hence, at this level of analysis, the best move for player 1 is to move to the middle, resulting in utilities (4, 1). However, we can take this analysis one step further. Now suppose that player 1 moves right anyway. Since (given ethical behavior by player 2) she could have guaranteed herself 4 by choosing middle, it can be argued that player 2 owes her at least 4 (especially because player 2 can do so while still getting at least the 1 that he received at the previous level of analysis). So, at this level, the only ethical thing for player 2 to do is to move right; middle is no longer ethical. Hence, the final solution is for both players to move right.

3 A Definition of Ethical Behavior Based on Iterated Solutions

We now give the general definition of our ethical solution concept. In the example in Figure 3, in a sense, we “solved” the game three times: first, we found the subgame perfect solution; second, we modified the solution based on the notion that player 2 should give player 1 what she could have guaranteed herself in the first (subgame perfect) solution (as long as doing so does not make player 2 worse off than he would have been in the first solution); third, we modified the solution again based on the notion that player 2 should give player 1 what she could have guaranteed herself in the second solution (as long as doing so does not make player 2 worse off than he would have been in the second solution). Furthermore, it is easy to construct examples in which even more levels of analysis are required.

In fact, the second and third solutions can be seen as subgame perfect solutions of a game in which the preferences have been modified based on the payoffs in the previous solution. In particular, let us call the utilities (b_1, b_2) from the previous solution the *base utilities*. Then, player 1’s primary goal is to obtain at least utility b_1 ; player 1’s secondary goal is for player 2 to obtain at least utility b_2 ; her tertiary goal is to maximize her own utility; and her quaternary goal is to maximize player 2’s utility.⁴ That is, given that she achieves her own base utility, player 1 temporarily sets her own interest aside and attempts to ensure that player 2 obtains his base utility; once that has been done, she pursues her own utility again. Player 2’s modified preferences are defined similarly. Formally, we have:

Definition 1. *Given base utilities (b_1, b_2) , we define player 1’s ethical preference relation $\succ_{(b_1, b_2)}^1$ as follows: $(u_1, u_2) \succ_{(b_1, b_2)}^1 (u'_1, u'_2)$ if and only if at least one of the following three conditions applies:*

- $u'_1 < b_1$, and: either $u'_1 < u_1$, or both $u'_1 = u_1$ and $u'_2 < u_2$.
- $u_1 \geq b_1, u'_1 \geq b_1, u'_2 < b_2$, and: $u_2 > u'_2$.
- $u_1 \geq b_1, u_2 \geq b_2$, and: either $u_1 > u'_1$, or $u_1 = u'_1$ and $u_2 > u'_2$.

⁴ The quaternary goal is relevant only for breaking ties and is not essential to our concept; we add it for completeness.

Player 2’s ethical preference relation $\succ_{(b_1, b_2)}^2$ is defined similarly (with the roles of 1 and 2 reversed).

In the special case in which b_1 and b_2 are smaller than any utility in the game, the players simply maximize their own utility (and break ties in favor of the other’s utility).

Now, we obtain a solution as follows: we solve the game, then update the base utilities to be the utilities in that solution, solve the game again with the modified utilities, modify the utilities again, *etc.*, until the solution stops changing.⁵ Formally, we have the following algorithm:

Iterated Backward Induction with Modified Preferences (IBIMP)

1. initialize $b_1 \leftarrow -\infty$
2. initialize $b_2 \leftarrow -\infty$
3. repeat until convergence:
 - (a) solve the game by backward induction with respect to $\succ_{(b_1, b_2)}^1, \succ_{(b_1, b_2)}^2$
 - (b) update b_1, b_2 to be the final utilities in this solution

We have not yet shown that this process will in fact converge, but this will become clear from the alternative characterization in the next section.

For example, for the game in Figure 3, we have the following three solutions:




Fig. 4. IBIMP solves the example game in three iterations. At each nonleaf vertex, an arrow indicates the player’s move in the subgame perfect solution for the modified preferences, and the leaf corresponding to the solution is underlined.

Another way to interpret this process is as follows: a third party repeatedly proposes strategy profiles for both players; the players accept the new proposal if and only if every move is consistent with their ethical preference relation (with respect to the base utilities from the currently accepted proposal). Then, the only sequence of solutions that the third party can successfully propose is the sequence of solutions that results from the algorithm above.

We emphasize again that breaking ties in favor of the other player is not essential to the concept, but it seems natural. (Incidentally, if ties are broken in this way, then the airline pillow example (Figure 2) has the same solution even if player 1 is indifferent between having the pillow or not when she is not tired.)

⁵ It should be noted that in general, a perfect-information game can have multiple subgame perfect Nash equilibria due to ties; finding the optimal one is nontrivial, but can be done in polynomial time [7]. Because we specified a tie-breaking mechanism—breaking ties in the other agent’s favor—we do not need to deal with these issues.

4 An Alternative Characterization Based on Global Pruning

In this section, we present an alternative definition of the solution concept, and show the equivalence between the two definitions. The alternative definition is also algorithmic, and also relies on repeatedly solving games. The difference is in how we modify the game. Instead of modifying the preferences based on the base utilities, we now remove all the leaf nodes for which at least one player’s utility is lower than the base utility.

Iterated Backward Induction with Pruned Leaves (IBIPL)

1. repeat until convergence:
 - (a) solve the game by backward induction (breaking ties in favor of the other player)
 - (b) let b_1, b_2 be the final utilities in this solution
 - (c) remove all the leaves with utilities (u_1, u_2) such that $u_1 < b_1$ or $u_2 < b_2$, and all intermediate nodes that have no children left

For example, for the game in Figure 3, we have the following three solutions:




Fig. 5. IBIPL solves the example game in three iterations. The leaf corresponding to the solution in each iteration is underlined, and removed leaves are crossed out.

We note that for this game, IBIPL’s solution in each stage is the same as in IBIMP. This is true in general, as we will see shortly. First, we note:

Lemma 1. *Under IBIPL, b_1 and b_2 monotonically (weakly) increase.*

Proof. b_1 and b_2 always correspond to a solution, and any leaf l with either $u_1^l < b_1$ or $u_2^l < b_2$ is immediately removed and can hence never be a future solution. ■

Theorem 1. *In each iteration, IBIMP and IBIPL find the same solution. That is, b_1 and b_2 are the same at each stage, and the values at each intermediate node are the same in the solution at each stage (if the intermediate node still exists under IBIPL).*

Proof. The first solutions are the same (both find the backward induction solution in which ties are broken in favor of the other player). We show that if the claim is true for the first k solutions, it is true for the $k + 1$ th solution, proving the claim by induction.

Given an intermediate node v (without loss of generality, one at which player 1 moves) that still occurs in both games, suppose that the utilities for all its children (that

still occur in both games) are the same. We will show that for any child c that still occurs in the IBIMP game but not in the IB IPL tree, that child will not be the most preferred under player 1’s ethical preferences $\succ_{(b_1, b_2)}^1$ (where b_1 and b_2 are the solution values for the k th solution, under both IBIMP and IB IPL by the induction assumption). Because v still occurs in the IB IPL game, it must have at least one child c' left in the IB IPL tree; since it has not been removed, it must have utilities $u_1^{c'} \geq b_1, u_2^{c'} \geq b_2$. (We emphasize again that the utilities are the same under both trees, by assumption.) On the other hand, because c was removed, every leaf l that is a child of c must have been removed; therefore, using the monotonicity property in Lemma 1, l has either $u_1^l < b_1$ or $u_2^l < b_2$. It follows that the utilities at c in the current iteration of IBIMP must have the same property: either $u_1^c < b_1$ or $u_2^c < b_2$. But then, it follows that $(u_1^{c'}, u_2^{c'}) \succ_{(b_1, b_2)}^1 (u_1^c, u_2^c)$. Therefore, under IBIMP, player 1 will not choose c from v .

Hence, both IBIMP and IB IPL choose from the children c' of v for which $u_1^{c'} \geq b_1, u_2^{c'} \geq b_2$. Both of them will choose a child with the highest $u_1^{c'}$, breaking ties to maximize $u_2^{c'}$. It follows that the utilities for v are the same under both IBIMP and IB IPL, and we can repeat this process to show this for all the vertices up to and including the root, thereby establishing that the new b_1 and b_2 will be the same. ■

While the definition of the ethical solution concept corresponding to IBIMP is perhaps more natural and easier to motivate, the equivalent definition corresponding to IB IPL is often easier to work with and prove properties about. The following propositions illustrate this.

Proposition 1. *IB IPL and IBIMP always terminate.*

Proof. IB IPL clearly must terminate, because the tree shrinks in each step (other than the last one). By Theorem 1 it follows that IBIMP must also terminate. ■

Lemma 2. *When IB IPL terminates, all remaining leaves have the same utilities (b_1, b_2) .*

Proof. For each remaining leaf l , we must have $u_1^l \geq b_1$ and $u_2^l \geq b_2$ (otherwise, the leaf would have been eliminated). So, if there is a remaining leaf l with utilities other than (b_1, b_2) , it must Pareto dominate the current solution ($u_1^l > b_1$ and $u_2^l \geq b_2$, or $u_1^l \geq b_1$ and $u_2^l > b_2$). For the sake of contradiction, suppose that such a leaf l exists. Both players break ties in favor of the other, so the utilities (u_1^l, u_2^l) will always be preferred to (b_1, b_2) . Hence, the utilities at the parent of l will be (u_1^l, u_2^l) , or something else that Pareto dominates (b_1, b_2) . The same is true for its parent, etc., up to and including the root. This contradicts (b_1, b_2) being the backward induction solution. ■

Proposition 2. *IB IPL and IBIMP always return a Pareto optimal solution.*

Proof. For the sake of contradiction, suppose that the final solution under IB IPL is not Pareto optimal—that is, there exists a leaf that Pareto dominates the solution. This leaf cannot have been eliminated, based on Lemma 2. But then, we have two remaining

leaves with different utilities, which contradicts Lemma 2. By Theorem 1 it follows that IBIMP also returns a Pareto optimal solution. ■

Proposition 3. *IBIPL and IBIMP always return a solution in which both players' utilities are at least as high as in any subgame perfect equilibrium where players break ties in each other's favor.*

Proof. Under IBIPL, after the first iteration, b_1 and b_2 are equal to the utilities from such a subgame perfect equilibrium; by Lemma 1 the utilities in later iterations can only be higher. By Theorem 1 the property also holds for IBIMP. ■

5 A Fast Implementation of IBIPL

The pseudocodes for IBIMP and IBIPL give us some basic (albeit not fully specified) algorithms for finding the ethical solution. In this section, we present a fast implementation of IBIPL with a runtime of $O(n \log n + nh \log(n/h))$, where n is the number of leaf nodes, and h is the height of the tree.

The algorithm and analysis assume, without loss of generality, a game tree in which all nonleaf nodes have at least two children. The basic idea is to maintain a data structure corresponding to the game tree, which maintains the optimal action at each nonleaf node. When in successive iterations, leaves are deleted, we only need to update the ancestors of those leaves (in fact, we may not need to update all of them).

A fast implementation of IBIPL

1. Initialize two arrays A_1, A_2 of pointers to the leaf nodes
2. Sort A_1 by the first player's utility, and A_2 by the second player's utility (ties can be broken arbitrarily)
3. Initialize index pointers i_1, i_2 to the first element of A_1 and A_2 , respectively
4. Using A_1 , compute the smallest value by which any two distinct values of u_1 differ, divide this number by twice the largest value of u_2 , and call the result ϵ_1 ; compute ϵ_2 similarly
5. Solve the game by backward induction (breaking ties in favor of the other player); in the process, at each node v , create a Fibonacci max heap whose elements are v 's children, ordered by their values for $u_1 + \epsilon_1 u_2$ if player 1 controls v , and by $u_2 + \epsilon_2 u_1$ if player 2 controls v (the ϵ terms are used to break ties in the other player's favor); the top child's u_1, u_2 values become v 's values
6. Let b_1, b_2 be the values at the root
7. Repeat until convergence:
 - (a) In array A_1 use binary search to find the first element for which $u_1 \geq b_1$; let its location be i_1'
 - (b) For every element in A_1 in a location $\{i_1, i_1 + 1, \dots, i_1' - 1\}$ do:
 - i. If the corresponding leaf v has not been marked deleted, mark it deleted and call **Update** $_1(P(v), u_1(v), u_2(v))$, where $P(v)$ is v 's parent
 - (c) Repeat the previous two steps with array A_2
 - (d) Let $i_1 \leftarrow i_1'$ and $i_2 \leftarrow i_2'$
 - (e) Update b_1, b_2 to the new values of u_1, u_2 at the root

Update₁(w, u_1, u_2)

1. From w 's Fibonacci max heap, remove (u_1, u_2)
2. If the max heap has become empty, call **Update₁**($P(w), u_1, u_2$)
3. Otherwise, if the values at the top of the heap have changed, update w 's values $u_1(w), u_2(w)$; if w is not the root, then call **Update₂**($P(w), u_1, u_2, u_1(w), u_2(w)$)

Update₂(w, u_1, u_2, u_1', u_2')

1. Let u_1'', u_2'' be the current values of w
2. From w 's Fibonacci max heap, remove (u_1, u_2) , and insert (u_1', u_2')
3. If the values at the top of the heap have changed, update w 's values $u_1(w), u_2(w)$; if w is not the root, then call **Update₂**($P(w), u_1'', u_2'', u_1(w), u_2(w)$)

Theorem 2. *The fast implementation of IB IPL runs in $O(n \log n + nh \log(n/h))$ time.*

Proof. Creating the sorted arrays will take $O(n \log n)$ time.

The first subgame perfect solution takes $O(n)$ time to generate; this includes creating and populating all of the Fibonacci max heaps (for which the amortized insertion time is $O(1)$). Calculating ϵ_1 and ϵ_2 also takes $O(n)$ time, because we have sorted arrays and thus only need to compare n adjacent pairs to find the smallest difference.

There are at most n iterations of the loop: each iteration other than the last must delete at least one leaf node. Within each iteration, we must find i_1' and i_2' , which takes $O(\log n)$ time using binary search. Thus, this takes a total of $O(n \log n)$ time.

We still need to consider the time needed for the deletions and updates. Each leaf node can be deleted at most once, so we have $O(n)$ deletions in total. Finding the leaves that need to be deleted only requires us to advance through the arrays from i_1 to i_1' and from i_2 to i_2' . Hence, finding the leaf nodes to delete requires $O(n)$ time in total.

Each individual leaf deletion can result in a number of updates (including both **Update₁**s and **Update₂**s); however, it can result in at most h updates, because a node can only call **Update₁** or **Update₂** on its parent. Letting b_v be the branching factor (number of children) of v , updating node v requires $O(\log(b_v))$ time for (at most) an insertion and a deletion into a Fibonacci max heap. At worst, we have h nodes v_1, \dots, v_h that require updating as a result of a single leaf deletion, resulting in a total update time of $\log(b_{v_1}) + \dots + \log(b_{v_h})$. We know, however, that $b_{v_1} + \dots + b_{v_h} \leq 2n$: this is because there are at most $2n$ nodes in the tree in total (because we assume that each node has a branching factor of at least 2), and the children of different nodes do not overlap. The optimization problem **maximize** $\log(b_{v_1}) + \dots + \log(b_{v_h})$ **subject to** $b_{v_1} + \dots + b_{v_h} \leq 2n$ is solved by setting $b_{v_i} = 2n/h$ for every i , because the log function is concave. It follows that the total time required for updates as a result of a single leaf deletion is $O(h \log(n/h))$, resulting in a bound of $O(nh \log(n/h))$ for the total time for updates.

Adding everything together, our total runtime bound is $O(n \log n + n + n \log n + n + nh \log(n/h)) = O(n \log n + nh \log(n/h))$. ■

For the purpose of reducing the runtime bound (and in its own right), it is interesting to consider how many iterations a particular type of tree can require. In the proof of the runtime bound above, we only used the fact that there are at most n iterations. We already know that the centipede game requires $\Omega(n)$ iterations, but of course the centipede game tree is extremely unbalanced. In the full version of our paper, we give lower bounds of $\Omega(\sqrt{n})$ iterations for both the case where the tree is a balanced binary tree, and the case where the tree has depth two. We also construct a game in which the solution in the first iteration is to move left at the root, in the second iteration it is to move right, and in the third it is once again to move left.

6 Conclusions

In this paper, we introduced a new solution concept for two-player perfect-information games that attempts to model a type of trusting behavior (together with the “ethical” behavior of not violating that trust). The concept takes subgame perfect Nash equilibrium as a starting point, but then repeatedly resolves the game based on the players being able to trust each other. We gave two distinct algorithmic definitions of the concept and showed that they are equivalent. Finally, we gave a fast implementation of one of the algorithms for solving the game, and showed that it runs in time $O(n \log n + nh \log(n/h))$.

There exist a large number of directions for future research. First, the validity of the concept should be evaluated. While we believe that the two equivalent definitions provide a strong normative justification of our concept, there may be other axiomatizations of the concept that make it even more convincing. However, as we have said previously, our concept only considers one particular type of ethical behavior, and other types of ethical behavior may lead to other natural solution concepts. It would also be interesting to investigate in more detail to what extent our solution concept models human behavior, taking a more descriptive approach rather than the normative approach discussed so far.

Another interesting direction is to try to generalize the concept to 3+-player games and/or games with imperfect information. Neither of these generalizations seem trivial. For example, if there is third player that barely affects the outcome of the game at all, then are the first two players still obliged to maintain player 3’s utility at at least the same level across iterations? And, if (due to imperfect information) it is not clear to player 1 whether player 2 took a “trusting” move, is player 1 obliged to assume that player 2 took such a move or not? Hence, it is not clear whether the 3+-player and/or imperfect-information cases admit as clean of a concept as the 2-player perfect-information case. Another issue is that our concept in some sense assumes that it is common knowledge that both players will behave ethically, and it is not clear what should be done if this is not the case. (One might model this as a game of imperfect information in which Nature first decides which players are ethical.)

Finally, how can we use this concept to approximately solve games that are so large that it is not possible to write down the entire tree? AI techniques for such games are usually based on limited-depth lookahead and heuristics to evaluate the nodes at this limited depth. For our concept, it is not clear whether the correct approach is to use

this type of limited-depth search on the full remaining tree within each iteration of the algorithm; or, to run the algorithm (all the iterations) on a limited-depth tree; or to do something entirely different. It also seems that if the two players do not use the same heuristics or depths, this can cause significant difficulties, because from one player’s perspective the other may not be acting ethically.

Acknowledgments

We thank the National Science Foundation and the Alfred P. Sloan Foundation for support (through award number IIS-0812113 and a Research Fellowship, respectively).

References

1. Anderson, M., Anderson, S.L.: The status of machine ethics: a report for the AAAI symposium. *Minds and Machines* 17, 1–10 (2007)
2. Anderson, M., Anderson, S.L., Armen, C.: Medethex: A prototype medical ethics advisor. In: AAAI, pp. 1759–1765 (2006)
3. Gal, Y., Pfeffer, A.: Modeling reciprocity in human bilateral negotiation. In: AAAI (2007)
4. Grosz, B.J., Kraus, S., Talman, S., Stossel, B., Havlin, M.: The influence of social dependencies on decision-making: Initial investigations with a new game. In: AAMAS, pp. 782–789 (2004)
5. Guarini, M.: Particularism and the classification and reclassification of moral cases. *IEEE Intelligent Systems* 21(4), 22–28 (2006)
6. King-Casas, B., Tomlin, D., Anen, C., Camerer, C.F., Quartz, S.R., Montague, P.R.: Getting to know you: Reputation and trust in a two-person economic exchange. *Science* 308(5718), 78–83 (2005)
7. Littman, M.L., Ravi, N., Talwar, A., Zinkevich, M.: An efficient optimal-equilibrium algorithm for two-player game trees. In: UAI (2006)
8. McKelvey, R.D., Palfey, T.R.: An experimental study of the centipede game. *Econometrica* 60(4), 803–836 (1992)
9. McLaren, B.M.: Computational models of ethical reasoning: Challenges, initial steps, and future directions. *IEEE Intelligent Systems* 21(4), 29–37 (2006)
10. Moor, J.H.: The nature, importance, and difficulty of machine ethics. *IEEE Intelligent Systems* 21(4), 18–21 (2006)
11. Nagel, R., Tang, F.F.: An experimental study on the centipede game in normal form—an investigation on learning. *Journal of Mathematical Psychology* 42, 356–382 (1998)
12. Talman, S., Gal, Y., Hadad, M., Kraus, S.: Adapting to agents personalities in negotiation. In: AAMAS (2005)
13. Tucker, A.: A two-person dilemma. In: Rasmusen, E. (ed.) *Readings in Games and Information*, pp. 7–8. Blackwell Publishing, Malden (2001); Originally written in 1950 (unpublished)
14. Zak, P.J.: The neurobiology of trust. *Scientific American*, 88–95 (June 2008)
15. Zak, P.J., Kurzban, R., Matzner, W.T.: Oxytocin is associated with human trustworthiness. *Hormones and Behavior* 48, 522–527 (2005)

The Secretary Problem with a Hazard Rate Condition

Mohammad Mahdian, Randolph Preston McAfee, and David Pennock

Yahoo! Research

{mahdian,mcafee,pennockd}@yahoo-inc.com

Abstract. In the classical secretary problem, the objective is to select the candidate of maximum value among a set of n candidates arriving one by one. The value of the candidates come from an unknown distribution and is revealed at the time the candidate arrives, at which point an irrevocable decision on whether to select the candidate must be made. The well-known solution to this problem, due to Dynkin, waits for n/e steps to set an “aspiration level” equal to the maximum value of the candidates seen, and then accepts the first candidate whose value exceeds this level. This guarantees a probability of at least $1/e$ of selecting the maximum value candidate, and there are distributions for which this is essentially the best possible. One feature of this algorithm that seems at odds with reality is that it prescribes a long waiting period before selecting a candidate. In this paper, we show that if a standard hazard rate condition is imposed on the distribution of values, the waiting period falls from n/e to $n/\log(n)$, meaning that it is enough to observe a diminishingly small sample to set the optimal aspiration level. This result is tight, as both the hazard condition and the optimal sampling period bind exactly for the exponential distribution.

1 Introduction

The classical secretary problem, dating back to late 1950’s and early 1960’s, asks whether one can hire the best secretary among a set of n candidates with unknown values arriving for interviews in a random order. Only one offer can be made, and the decision about hiring a candidate has to be made immediately after the interview and before seeing the next candidate. The surprising result (often attributed to Dynkin [2]) is that it is indeed possible to select the best secretary with a constant probability using the following strategy: do not make any offer to the first n/e candidates that arrive, and after that, make an offer to the first candidate whose value exceeds the value of all the candidates seen so far. In other words, the algorithm starts with an *exploration phase* that sees the first n/e candidates and sets an *aspiration level* equal to the highest value seen in the exploration phase. After that, the algorithm hires the first candidates that exceeds the aspiration level. We call such a strategy an “*aspiration strategy*”. Many variants of the secretary problem and this solution, and their applications to problems such as online mechanism design for auctions are studied in the literature; see, for example, the survey papers [14].

An alternative formulation of the secretary problem (and according to Ferguson [3], the formulation closer to the original “game of googol” puzzle due to Martin Gardner in 1960 [5]) is as follows: An adversary selects the distribution from which the n values will be drawn independently.¹ The algorithm learns about this distribution (but not the actual values), and then sees the values one by one. The algorithm must decide when to stop, at which point it receives the last value seen. The objective is to maximize the probability that this value is the maximum value, or, equivalently, the expectation of the ratio of this value and the ex post maximum value in the sequence. As observed by Ferguson [3], this formulation of the problem is equivalent to the original one, and the same aspiration strategy achieves a factor of $1/e$. Furthermore, this result is tight, that is, there is a distribution of values for which, even knowing the distribution, the optimal strategy is the aspiration strategy with a factor of $1/e$.

One unsatisfactory feature of the aspiration strategy is that it requires a long exploration phase (observing a constant fraction of the candidates) before setting the aspiration level. One might suspect that this is due to the fact that there is no penalty (except the possibility of missing the maximum value) in the secretary problem for waiting, but surprisingly, similar results hold for variants of the secretary problem where waiting is punished by discounting the values or by setting an explicit cost for each period that no one is hired [4,7]. The main message of this paper is to show that a long exploration phase is indeed an artifact of the assumption that the values are coming from a worst-case distribution. We show that if we impose a hazard rate condition, frequently used in engineering analyses and the analysis of auctions, the picture changes drastically: even the worst value distribution satisfying this condition requires an exploration phase of length $n/\log(n)$, significantly smaller than n/e .

2 The Model

Consider a setting where one needs to hire a secretary among a set of n potential candidates. These potential candidates arrive one per time period. Candidate i has a value v_i . v_i 's are independently and identically distributed according to a distribution F with cdf f . The value of v_i is revealed when the i 'th candidate arrives. The objective is to give a *termination rule* (i.e., an algorithm that after observing the first i values decides whether to stop or to continue) that stops at the candidate with the maximum value among the n candidates. The algorithm does not have any knowledge of the distribution F .

A standard approach to solve the above problem is to collect data by observing v_i 's for k periods, set an *aspiration level* equal to the maximum value observed

¹ To be precise, Gardner's formulation does not assume the values are selected independently; instead, it assumes that the adversary selects the values from any joint distribution that is permutation invariant (exchangeable). However, the assumption that the values are chosen independently is nearly without loss of generality since an infinite sequence of exchangeable random variables is conditionally independent by Kingman's theorem [6].

during those k periods, and then accept the first candidate observed with value exceeding the aspiration level, or the last candidate. We call such a strategy an *aspiration strategy*. A classical result due to Dynkin [2] shows that an aspiration strategy with $k = n/e$ picks the candidate with the maximum value among all n candidates with probability at least $1/e$. Furthermore, this is essentially the best possible, that is, there are distributions for which no algorithm can guarantee a probability more than $1/e$ of picking the candidate with the maximum value.

In this paper, we focus on a special class of value distributions F . First, we assume that values are all positive, that is, F has support on $[0, \infty)$ and $F(0) = 0$. This is essentially without loss of generality, as any candidate with negative or zero value can be discarded. More importantly, we assume that the distribution F satisfies the *monotone hazard rate condition*, which means that the hazard function $\frac{f(y)}{1-F(y)}$ is a monotone non-decreasing function of y . This assumption is standard in engineering applications and auction theory [8], and intuitively means that given that the value of a candidate is not less than y , the likelihood that it is equal to y increases as y increases. As an example, Gaussian, uniform, and exponential distributions satisfy this property.

In addition to the above, we make the simplifying assumption that the values are bounded. Therefore, the inverse function $F^{-1}(z) := \inf\{y \in [0, \infty) : F(y) \geq z\}$ is well-defined on its domain $[0, 1]$ and satisfies $F^{-1}(0) = 0$ and $F^{-1}(1) < \infty$ [2].

3 Main Result

Our main result is to analyze the optimal sample size in an aspiration strategy, assuming that the value distribution satisfies the hazard rate condition. This is stated in the following theorem.

Theorem 1. *Suppose the value distribution F satisfies the above conditions. Then the optimal sample size in an aspiration strategy does not exceed*

$$\frac{n-1}{\sum_{i=1}^{n-1} \frac{1}{i}} - 1.$$

Proof. Consider the optimal aspiration strategy, and let y denote the realized maximum value of the first k observations. The probability density function associated with this event can be written as $kF(y)^{k-1}f(y)$. The strategy then performs a search from candidates $k+1$ to n . Period j ($j > k$) is reached if none of the earlier values are acceptable. This happens with probability $F(y)^{j-k-1}$, and in this case if the value x of the j 'th candidate exceeds y (or if $j = n$), it is accepted. Thus the payoff in an aspiration strategy associated with setting the aspiration level with k observations is

² It is not hard to see that our proof works for a more general class of distributions, satisfying the condition that $\lim_{z \rightarrow 1} (1-z)F^{-1}(z) = 0$. This includes distributions such as the exponential distribution. Details of the proof is left to the final version of the paper.

$$\begin{aligned}
 \pi_k &= \int_0^\infty kF(y)^{k-1}f(y)\left(\sum_{j=k+1}^{n-1} F(y)^{j-k-1} \int_y^\infty xf(x)dx + F(y)^{n-k-1} \int_0^\infty xf(x)dx\right)dy \\
 &= \int_0^\infty kf(y)\left(\sum_{j=k+1}^{n-1} F(y)^{j-2} \int_y^\infty xf(x)dx + F(y)^{n-2} \int_0^\infty xf(x)dx\right)dy \\
 &= k \sum_{j=k+1}^{n-1} \frac{F(y)^{j-1}}{j-1} \int_y^\infty xf(x)dx \Big|_{y=0}^\infty + k \int_0^\infty \left(\sum_{j=k+1}^{n-1} \frac{F(y)^{j-1}}{j-1} yf(y)\right) dy \\
 &\quad + \frac{k}{n-1} F(y)^{n-1} \Big|_{y=0}^\infty \int_0^\infty xf(x)dx \\
 &= 0 + k \sum_{j=k+1}^{n-1} \int_0^\infty y \frac{F(y)^{j-1}}{j-1} f(y)dy + \frac{k}{n-1} \int_0^\infty xf(x)dx \\
 &= k \sum_{j=k}^{n-2} \int_0^1 F^{-1}(z) \frac{z^j}{j} dz + \frac{k}{n-1} \int_0^1 F^{-1}(z) dz,
 \end{aligned}$$

where the third equation follows from integration by parts, and the last equation is using the change of variables $z = F(y)$. Thus,

$$\begin{aligned}
 \pi_{k+1} - \pi_k &= (k+1) \sum_{j=k+1}^{n-2} \int_0^1 F^{-1}(z) \frac{z^j}{j} dz - k \sum_{j=k}^{n-2} \int_0^1 F^{-1}(z) \frac{z^j}{j} dz + \frac{1}{n-1} \int_0^1 F^{-1}(z) dz \\
 &= \sum_{j=k+1}^{n-2} \int_0^1 F^{-1}(z) \frac{z^j}{j} dz - \int_0^1 F^{-1}(z) z^k dz + \frac{1}{n-1} \int_0^1 F^{-1}(z) dz \\
 &= \int_0^1 F^{-1}(z) \left(\sum_{j=k+1}^{n-2} \frac{z^j}{j} - z^k + \frac{1}{n-1}\right) dz \\
 &= F^{-1}(z) \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1}}{j(j+1)} - \frac{z^{k+1}}{k+1} + \frac{z}{n-1}\right) \Big|_{z=0}^1 \\
 &\quad - \int_0^1 F^{-1}'(z) \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1}}{j(j+1)} - \frac{z^{k+1}}{k+1} + \frac{z}{n-1}\right) dz \\
 &= - \int_0^1 F^{-1}'(z) \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1}}{j(j+1)} - \frac{z^{k+1}}{k+1} + \frac{z}{n-1}\right) dz, \tag{1}
 \end{aligned}$$

where the last equation follows from the fact that $F^{-1}(1) < \infty$ and the identity

$$\frac{1}{n-1} + \sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} - \frac{1}{k+1} = 0. \tag{2}$$

Using the same identity, we can write

$$\begin{aligned} \pi_{k+1} - \pi_k &= - \int_0^1 F^{-1'}(z) \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1} - z}{j(j+1)} - \frac{z^{k+1} - z}{k+1} \right) dz \\ &= \int_0^1 F^{-1'}(z)(1-z) \left(\sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} \sum_{i=1}^j z^i - \frac{1}{k+1} \sum_{i=1}^k z^i \right) dz. \end{aligned}$$

This can be written as

$$\pi_{k+1} - \pi_k = \int_0^1 F^{-1'}(z)(1-z)\beta(z)dz, \tag{3}$$

where

$$\beta(z) = \sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} \sum_{i=1}^j z^i - \frac{1}{k+1} \sum_{i=1}^k z^i.$$

Since $F^{-1'}(z) = \frac{1}{f(F^{-1}(z))}$, the function $(1-z)F^{-1'}(z) = \frac{1-z}{f(F^{-1}(z))}$ is the composition of the function $\frac{1-F(y)}{f(y)}$, which is non-increasing by the hazard rate condition, and the function F^{-1} , which is non-decreasing. Therefore, $(1-z)F^{-1'}(z)$ is a non-increasing function of z . The function $\beta(z)$ can be simplified as follows:

$$\begin{aligned} \beta(z) &= \sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} \sum_{i=k+1}^j z^i + \left(\sum_{j=k+1}^{n-2} \frac{1}{j(j+1)} - \frac{1}{k+1} \right) \sum_{i=1}^k z^i \\ &= \sum_{i=k+1}^{n-2} \sum_{j=i}^{n-2} \frac{1}{j(j+1)} z^i - \frac{1}{n-1} \sum_{i=1}^k z^i \\ &= \sum_{i=k+1}^{n-2} \left(\frac{1}{i} - \frac{1}{n-1} \right) z^i - \frac{1}{n-1} \sum_{i=1}^k z^i \\ &= \sum_{i=k+1}^{n-2} \frac{z^i}{i} - \frac{1}{n-1} \sum_{i=1}^{n-2} z^i, \end{aligned} \tag{4}$$

where the second and the third equations follow from the identity (2). This function has the following properties:

1. $\beta(0) = 0$.
2. For $z > 0$ sufficiently small, $\beta(z) < 0$. This follows from $\beta'(0) = \frac{-1}{n-1} < 0$.
3. There is a $z^* \in (0, 1]$ with $\beta(z) > 0$ if and only if $z > z^*$.

The last property follows from expressing β as

$$\beta(z) = z^k \left(\sum_{i=k+1}^{n-2} z^{i-k} \left(\frac{1}{i} - \frac{1}{n-1} \right) - \frac{1}{n-1} \sum_{i=1}^k z^{i-k} \right)$$

and noting that the term in parentheses is increasing in z , so that once the term is positive it remains positive.

Claim. Assume β is a function satisfying the above properties. If $\int_0^1 \beta(z)dz \leq 0$, then for every non-negative non-increasing function $x(\cdot)$ we have $\int_0^1 x(z)\beta(z) dz \leq 0$.

Proof. Consider the value z^* guaranteed to exist by property (3). We have

$$\begin{aligned} \int_0^1 x(z)\beta(z)dz &= \int_0^{z^*} x(z)\beta(z)dz + \int_{z^*}^1 x(z)\beta(z)dz \\ &\leq \int_0^{z^*} x(z^*)\beta(z)dz + \int_{z^*}^1 x(z^*)\beta(z)dz \\ &= x(z^*) \int_0^1 \beta(z)dz \\ &\leq 0. \end{aligned}$$

The first inequality replaced x in the first integral with the lowest possible value given the other values, increasing the objective function since β was negative, and x in the second integral with the highest possible, again increasing the objective function since β was positive.

Now apply the claim with $x = F^{-1'}(z)(1 - z)$. By (4), we have

$$\begin{aligned} \int_0^1 \beta(z)dz &= \sum_{i=k+1}^{n-2} \int_0^1 \frac{z^i}{i} dz - \frac{1}{n-1} \sum_{i=1}^{n-2} \int_0^1 z^i dz \\ &= \sum_{i=k+1}^{n-2} \frac{1}{i(i+1)} - \frac{1}{n-1} \sum_{i=1}^{n-2} \frac{1}{i+1} \\ &= \frac{1}{k+1} - \frac{1}{n-1} \sum_{i=1}^{n-1} \frac{1}{i} \end{aligned}$$

Therefore, when $k \geq \frac{n-1}{\sum_{i=1}^{n-1} \frac{1}{i}} - 1$, by the above equation $\int_0^1 \beta(z)dz \leq 0$. Therefore, since $x = F^{-1'}(z)(1 - z)$ is a non-negative and non-increasing function, by Claim 3 and Equation (3) we obtain

$$\pi_{k+1} - \pi_k = \int_0^1 x(z)\beta(z)dz \leq 0.$$

Thus, the sample size k that maximizes π_k cannot be larger than $\frac{n-1}{\sum_{i=1}^{n-1} \frac{1}{i}} - 1$.

Note that $\sum_{i=1}^{n-1} \frac{1}{i} \geq \ln n$. Thus, an upper bound for the optimal k is

$$\frac{n-1}{\ln n} - 1 < \frac{n}{\ln n}.$$

Table 1. Comparison between three aspiration strategies with exponential distribution

n	n/e		$n/\ln n$		Optimal Search
	Value	Higher	Value	Higher	
50	64%	22%	69%	12%	88%
100	64%	27%	71%	15%	89%
500	64%	36%	76%	20%	92%
1000	63%	38%	78%	22%	92%
5,000	63%	43%	81%	25%	94%
10,000	63%	44%	82%	26%	94%

The upper bound is strict in the sense that it is taken on exactly when the distribution is exponential and satisfies the hazard rate constraint with equality.

Our result shows that, for example, when n is a million, it is enough to explore 7.2% of the candidates before setting an aspiration level, rather than 36.8% as in the standard result. Moreover the share of observations devoted to learning falls as the sample size rises, which is a natural conclusion, as the standard mechanism does a lot of search.

How well does the improved calculation do? Table 1 reports estimated values using aspiration strategies with exploration phases of length n/e and $n/\ln n$, for various values of n and an exponential distribution. All values are expressed as a percentage of the expected maximum realized value

$$\int_0^\infty nx F(x)^{n-1} f(x) dx.$$

The column “Higher” gives the proportion of times a strategy beats the other strategy. Thus the n/e strategy wins over $n/\ln n$ frequently, but the $n/\ln n$ produces a higher value.

The last column of the table sets out the average values associated with optimal search, conditional on the exponential distribution. Optimal search involves setting a critical value c_t in each time period t and accepting an offer if it exceeds that critical value. The critical values c_t can be computed using a dynamic programming approach as follows: clearly $c_n = 0$ as any offer better than nothing is acceptable in the last step. This produces a last period expected value v_n of

$$\int_0^\infty x f(x) dx.$$

Let v_t be the value of searching starting from period t . This value is

$$v_t = \int_{c_t}^\infty x f(x) dx + F(c_t) v_{t+1}.$$

Clearly v_t is maximized by $c_t = v_{t+1}$, which gives

$$v_t = \int_{v_{t+1}}^\infty x f(x) dx + F(c_t) v_{t+1} = v_{t+1} + \int_{v_{t+1}}^\infty (1 - F(x)) dx.$$

It is not hard to see that the upper bound on the length of the exploration phase is tight for the exponential distribution, but can be too large for other distributions. For example, for the uniform distribution, by Equation (II) we have

$$\begin{aligned}\pi_{k+1} - \pi_k &= - \int_0^1 \left(\sum_{j=k+1}^{n-2} \frac{z^{j+1}}{j(j+1)} - \frac{z^{k+1}}{k+1} + \frac{z}{n-1} \right) dz \\ &= - \sum_{j=k+1}^{n-2} \frac{1}{j(j+1)(j+2)} + \frac{1}{(k+1)(k+2)} - \frac{1}{2(n-1)} \\ &= \frac{1}{2} \left(\frac{1}{(k+1)(k+2)} - \frac{1}{n} \right).\end{aligned}$$

Thus, the optimal k for the uniform distribution is less than \sqrt{n} . These results give insight into optimal experimentation. The standard approach utilizes 37% of the data for experimental purposes. In contrast, when tails are not too fat, one should use at most a $1/\ln n$ fraction for learning about the distribution. With the uniform distribution, a $1/\sqrt{n}$ fraction is the appropriate length of the experimentation phase.

References

1. Babaioff, M., Immorlica, N., Kempe, D., Kleinberg, R.: Online auctions and generalized secretary problems. In: ACM SIGecom Exchanges (to appear)
2. Dynkin, E.B.: The optimal choice of the stopping moment for a markov process. Dokl. Akad. Nauk. SSSR 150, 238–240 (1963)
3. Ferguson, T.S.: Who solved the secretary problem? *Statistical Science* 4(3), 282–289 (1989)
4. Freeman, P.R.: The secretary problem and its extensions: A review. *International Statistical Review* 51, 189–206 (1983)
5. Gardner, M.: *Scientific American* (February/March 1960)
6. Kingman, J.F.C.: Uses of exchangeability. *Annals of Probability* 6, 183–197 (1978)
7. Mahdian, M., McAfee, R.P., Pennock, D.: Secretary problem with durable employment. Working paper (2008)
8. McAfee, R.P., McMillan, J.: Auctions and bidding. *Journal of Economic Literature* 25(2), 699–738 (1987)

Impact of QoS on Internet User Welfare

Galina Schwartz, Nikhil Shetty, and Jean Walrand

Department of Electrical Engineering and Computer Sciences (EECS),
University of California Berkeley,
Cory Hall, Berkeley, CA 94720-1770, USA
{schwartz,nikhils,wlr}@eecs.berkeley.edu

Abstract. In this paper, we investigate the welfare effects of transition from a single-service class to two-service classes in the Internet. We consider an ISP who offers network access to a fixed user base, consisting of users who differ in their quality requirements and willingness to pay for the access. We model user-ISP interactions as a game in which the ISP makes capacity and pricing decisions to maximize his profits and the users only decide which service to buy, if any. Our model provides robust pricing for networks with single- and two-service classes. Our results indicate that transition to multiple service classes is socially desirable, but could be blocked due to the unfavorable distributional consequences that it inflicts on the existing network users. To facilitate the transition, we propose a simple regulatory tool that alleviates the political economic constraints and thus makes the transition feasible.

1 Introduction

In today's Internet, despite the technological possibility of providing differentiated services¹, no such services are actually offered by the ISPs. We outline five main reasons for the persistence of this situation. First, ISPs claim that the demand is uncertain and hence the risks involved render such a QoS provision unprofitable². Therefore, ISPs lack incentives to explore offering QoS.

Second, there is a lack of coordination between the ISPs in QoS provision. It is possible that each ISP would individually like to offer QoS. But, the provision of QoS depends on all ISPs on a path between two users, not just on a single ISP. Lack of coordination between the ISPs and conflicts about dividing the QoS revenue between the involved ISPs preclude them from offering QoS⁴.

Third, QoS provision appears to be an inferior investment relative to plain capacity expansion. For e.g. [5] asserts that over-engineering the current network by investing in capacity is more profitable than investing in provision of multiple service classes while [6] concludes that the upfront costs might be too high and call for very simple QoS mechanisms.

¹ For the sake of brevity, we use the term QoS to refer to such differentiated services.

² Indeed, this profound uncertainty posits considerable difficulties for estimating demand [1,2,3].

Fourth, the current threat of network neutrality regulations hampers ISPs' incentives for QoS. Indeed, in the current regime, the ISPs are "at their best behavior," i.e., they suffer from self-imposed constraints [7,8,9]. These constraints preclude them from investing into developing QoS. Indeed, in the existing political climate, such investments are subject to regulatory uncertainty. For example, if an ISP evaluates the NPV (net present value) of its project to provide a specific QoS product, a positive probability of imposition of a neutral regime clearly reduces the project's NPV relative to the situation where the threat of regulation is not present. It is likely that these self-imposed constraints also create disincentives to invest in capacity expansion.

Lastly, fifth, the existing QoS pricing research [10,11] indicates the difficulties of robust pricing of QoS. The idea of utilizing DiffServ to provide multiple service classes by pricing them differently was around for a while [12]. Still, recent research attempts [10] suggest that the only socially optimal capacity division allocates zero capacity share for non-premium services. This paper mostly focuses on the fifth reason, but also suggests a connection to the third and fourth.

We make the following contributions to the literature. First, we develop robust pricing for a single ISP providing two service classes. Second, we investigate the political economic considerations that may constrain the feasibility of transition from the current network (where QoS is not provided) to the network with QoS provision. Third, we propose a simple regulatory tool that permits to alleviate the political economic constraints and make the transition feasible.

In this paper, we develop a model that permits robust pricing of differentiated services, based on the network architecture similar to the Paris Metro proposal (PMP) [12]. Other closely related papers modeling PMP are [13,14]. Although the authors in [13] include capacity choice in the description of the game, they assume zero capacity costs, and mostly focus on a subgame in which capacities are fixed. While [13] assumes zero costs of capacity, in [14], capacity is costly, with capacity costs increasing and convex. On one hand, [13] demonstrates that monopolistic ISP will indeed provide two service classes, and suggest that the lack of QoS provision on the Internet could be a consequence of competition among the ISPs. On the other hand, [14] finds that, in equilibrium, the two competitors have different prices and congestion levels, with the most expensive one being the least congested.

Both, [13,14] focus on ISP competition, with network access provided by duopolists. They do not study the effects of ISP choices on user welfare. We assume that capacity is costly, and provide full analysis of capacity choice and its division and its consequences for user welfare, but we do it for a monopolistic ISP. Our model is extendable to multiple ISPs, and our preliminary results [15] indicate that ISP competition does not necessarily preclude QoS provision.

In this paper, we also connect the ongoing network neutrality debate with ISPs' incentives to invest into QoS provision. We argue that the threat of neutrality regulations hampers ISPs' incentives for QoS and hence denies society the higher welfare effects associated with multiple service classes. In our related work [16], we explore an inexpensive regulatory tool that simplifies the division

of capacity and alleviates investment disincentives of ISPs by establishing property rights over a small fraction of their capacity. This tool facilitates the social planner to reduce the harmful effects of transition to multiple service classes and enables the deployment of QoS and development of novel applications.

2 Model

We start with a simple model, in which a single ISP (a monopolist) offers connectivity to a user base of fixed size, and we let N (which we assume to be large) be the total number of users. First, the ISP chooses his capacity $C > 0$ that he builds at a constant unit cost $\tau > 0$. Investment in capacity is irreversible. Second, once the capacity is sunk, the ISP makes a pricing decision after which each user decides whether to adopt the service. The ISP's objective is to maximize his profit Π_{total} which equals his revenue net of his expense on capacity:

$$\Pi_{total} = \max_{C,p} \{pZ - \tau C\},$$

where Z is the number of users who adopt the service, and p is the ISP access price for users. If the ISP offers multiple service classes and allocates a capacity $C_i > 0$ for the provision of service i at a price $p_i > 0$, his objective becomes:

$$\Pi_{total} = \max_{C_i, p_i} \left\{ \sum_i p_i Z_i - \tau C \right\},$$

where $C = \sum_i C_i$ and Z_i is the number of users who adopt service i .

We assume that, on average, each user sends an identical unit amount of traffic. We define the *quality of service* q observed by users as $q = 1 - Z/C$, if Z users are multiplexed in capacity C . This definition of quality reflects the common perception about service quality. As Z decreases and capacity remains the same, the quality of service improves, i.e., as the capacity per user increases, the quality increases as well. Further, we assume that each user in the user base is characterized by his type θ , which is a random variable uniformly distributed in $[0, 1]$. To characterize the user demand, we make the following assumptions. For a user with type θ , the lowest *acceptable* service quality is $q = \theta$; and his highest *affordable* price is $p = \theta$. Thus, a user buys a service only if this service is acceptable and affordable, i.e., $p < \theta \leq q$. (See [17] for a related discussion of this model.) For a user of type θ , if the quality of service $q \geq \theta$, his surplus is the difference between the price and his willingness to pay (which is θ itself). Thus, for the user with type θ , the surplus U_θ is given by

$$U_\theta = (\theta - p)I(q - \theta), \text{ where } I(y) = \begin{cases} 1 & \text{if } y \geq 0 \\ 0 & \text{if } y < 0 \end{cases}. \quad (1)$$

θ represents the quality of the application that a user is interested in. Thus, in our model, user adoption is determined by the availability of the most quality intensive application that his type θ utilizes. Indeed, if a user adopts the network

service for e-mail only, he gains no extra surplus from the fact that the actual network quality permits him to use streaming video (which he does not utilize). Compare this to the user demand in [18]. In general, for a distribution $g(\theta)$ of user types $\theta \in [0, 1]$, the aggregate user surplus becomes $U_{total} = \int_0^1 U_\theta g(\theta) N d\theta$, and with our assumption of uniformly distributed user types, we have $U_{total} = \int_0^1 U_\theta N d\theta$. Let c denote capacity per user in the user base ($c = C/N$) and let z denote the fraction of users adopting the service of quality q ($z = Z/N$). Then, $Z/C = z/c$ and from our definition of service quality: $q = 1 - z/c$. Let also $\Pi = \frac{\Pi_{total}}{N}$ and $U = \frac{U_{total}}{N}$, where Π is the profit and U is the surplus per user in the base. Thus, the ISP objective and user surplus become:

$$\Pi = \max_{c,p} \{pz - \tau c\} \text{ and } U = \int_0^1 U_\theta d\theta. \tag{2}$$

Per user in the base, the social surplus S is the sum of user surplus and provider profit, i.e., $S = U + \Pi$. Below, we analyze the ISP’s optimal choices for scenarios where the ISP provides (i) a single service, and (ii) two different services, possibly in the presence of a regulator who constrains the ISP.

3 Analysis

3.1 Single Service Class

In this section, we assume that the ISP provides only a single service class. This means that his entire capacity is offered at a single price p . To start, we let ISP capacity be fixed at C and solve the problem of optimal pricing. Once capacity is sunk, the ISP’s objective is to maximize his revenue R (per user in the base): $R(c) = \max_p pz$. From (II), a user with type θ will adopt the access (service) if and only if $p < \theta \leq q$, where $q = 1 - z/c$, with z being the fraction of users who adopt the service. Clearly, the service is affordable to all users with type $\theta > p$. As more users adopt the service, z increases and q decreases until it becomes equal to the user type at some critical value of θ . Let users with types $\theta \in (\underline{\theta}, \bar{\theta}]$ adopt the service. Then, we obtain (see [19])

$$\underline{\theta} = p \text{ and } \bar{\theta} = \frac{p+c}{1+c}, \text{ and } z = \frac{c}{1+c}(1-p). \tag{3}$$

Using (3), we express the revenue as

$$R(c,p) = pz = \frac{c}{1+c}p(1-p). \tag{4}$$

To find the optimal price (which corresponds to maximum ISP revenue), we differentiate (4) w.r.t. p and obtain $p = 1/2$. Thus, for any c , the revenue is maximized at $p^\dagger = 1/2$, and $R(c) = \frac{c}{4(1+c)}$. Henceforth, we will use the superscript \dagger to designate the values that correspond to the ISP’s optimal choice in the single service class case. Next, we use the expression for $R(c)$ to simplify the ISP objective to $\Pi = \max_c \frac{c}{4(1+c)} - \tau c$, which leads to (see [19]):

$c^\dagger = \frac{1}{2\sqrt{\tau}} - 1$ and $p^\dagger = 1/2$; $\theta_l^\dagger = \frac{1}{2}$ and $\bar{\theta}^\dagger = \frac{\frac{1}{2}+c}{1+c}$. Notice that, c^\dagger is positive for $\tau \in [0, 0.25]$ only. The optimal values Π^\dagger , U^\dagger and S^\dagger are derived in the [19].

3.2 Two Service Classes (Divided Capacity)

Next, we assume that the ISP divides his capacity $C = cN$ into two parts and provides two services. Let $C_l = c_l N$ and $C_h = c_h N$ be the capacities utilized for each service (l and h respectively), and p_l and p_h be the prices charged for the respective services; we assume $p_l < p_h$. Further, let z_l and z_h be the fraction of users who adopt the respective services l and h . From our definition of service quality (Section 2), $q_i = 1 - z_i/c_i$, where $c_i = C_i/N$, $i = l, h$ and $\sum_i c_i = c$. Let x denote the capacity fraction allocated to the provision of service h . Then,

$$c_h = xc \text{ and } c_l = (1 - x)c, \tag{5}$$

and the ISP objective can be written as $\Pi = \Pi_{total}/N = \max_{c,x,p_h,p_l} (\sum_{i=l,h} p_i z_i - \tau c)$. From (5), one can easily switch between the use of (c_l, c_h) and (c, x) as choice variables. In fact, the expressions in terms of c_l and c_h easily transform to the ones in terms of c and x . Although we will express the ISP's objective in terms of c and x , to simplify the presentation we express the equations in c_l and c_h .

Theorem 1. *For any fixed c and x , in the ISP optimum,*

$$p_l(c, x) = \frac{1}{2} - \frac{c_h c_l}{2[(1+c_l)(1+c_h)c_l+c_h]}, p_h = \frac{p_l+c_l}{1+c_l} \tag{6}$$

$$\theta_l = p_l, \bar{\theta}_l = \theta_h = p_h, \bar{\theta}_h = \frac{p_h+c_h}{1+c_h}, \tag{7}$$

where users with $\theta \in (\theta_l, \bar{\theta}_l]$ and $\theta \in (\theta_h, \bar{\theta}_h]$ adopt service l and h respectively.

For proof, see [19]. From Theorem 1, we have $\bar{\theta}_h > \bar{\theta}_l$, which implies that service h , that has a higher price ($p_l < p_h$) has higher quality ($q_l < q_h$) too. Henceforth, we will denote the ISP's optimal choices in two service classes by \ddagger .

Corollary 1. *For any fixed c and x in the ISP optimum, we have*

$$p_l(c, x) < \frac{1}{2} \text{ and } p_h(c, x) > \frac{1}{2}. \tag{8}$$

For proof, see [19]. From (1) and Corollary 1, we deduce that in the case of a transition from a single service class to two service classes, all existing users who adopt service l gain surplus while those who adopt service h lose surplus.

From Theorem 1 (using (6) and (7)), we derive the maximum revenue $R(c, x)$

$$R(c, x) = \frac{[(1 + c_h)c_l + c_h]^2}{4[(1 + c_l)(1 + c_h)c_l + c_h](1 + c_h)}. \tag{9}$$

Expression (9) is too cumbersome to carry out investigation analytically. Hence, we use MATLAB[®] to obtain the solution numerically, for results see [19].

3.3 Network Regulations

In the absence of regulation, from Section 3.2, the ISP objective could be written as a function of c and x as follows: $\Pi = \max_{c,x} \Pi(c, x) = \max_{c,x} \{R(c, x) - \tau c\}$. We will say that the network is regulated when x is chosen by an outside party (regulator). We assume that the regulator’s choice variable is x , i.e., the regulator only affects the ISP by constraining him from dedicating more than a fraction α of the entire capacity to service h . Then, the regulated ISP’s profit maximization can be expressed as

$$\Pi = \max_{c, p_l, p_h} \{p_l z_l + p_h z_h - \tau c\} \text{ where } c_l \geq (1 - \alpha)c \text{ and } c_h \leq \alpha c.$$

We consider three regulatory scenarios. Regulator 1 (a social planner) maximizes social surplus (sum of aggregate user surplus and ISP profit), regulator 2 maximizes user surplus and regulator 3 maximizes the surplus of the users served under a single service class (i.e., users with type $\theta \in (\theta^\dagger, \bar{\theta}^\dagger]$). For the regulators 1 - 3, the respective objectives S_1, S_2 and S_3 are:

$$S_1 = \max_x \{U + \Pi\} \text{ and } S_2 = \max_x U \text{ and } S_3 = \max_x U|_{\theta \in (\theta^\dagger, \bar{\theta}^\dagger]}, \quad (10)$$


where U is defined in (2). Let $\alpha_1, \alpha_2, \alpha_3$ be the values of x chosen by the regulators 1-3 respectively. From (10), it is intuitive that: $\alpha_3 < \alpha_2 < \alpha_1 < x^\ddagger$.

The case of a single service class is identical to the imposition of $x = 0$. We do not consider explicit regulations in the case of single service class, but we believe that the lack of QoS provision by ISPs in the current Internet reflects the tacit presence of such a regulatory threat. The ongoing network neutrality debate reflects that the threat is indeed real. We argue that this regulatory threat makes the ISPs to act as if the constraint $x \equiv 0$ is imposed [7]. This threat could be one of the reasons why QoS is not provided currently.

4 Results

Fig. 1(a) depicts the surplus gain of the single class users (users with type $\theta \in (\theta^\dagger, \bar{\theta}^\dagger]$). From Fig. 1(a), in the absence of regulation or with regulator 1, the transition to two service classes results in aggregate surplus loss for these users. Although, with regulator 2, these users do gain surplus, this might be insufficient to make the transition politically feasible, especially if the percentage of losing single service class users is high. We suggest that users who lose surplus from transition are likely to resist the change. Hence, when this fraction of users is high, these users may block such a transition. To assure feasibility, the regulator may need to reduce the fraction of such losing users. We use regulator 3 to show that it is possible to keep the fraction of users with surplus loss under 10%.

Fig. 1(b) depicts the percentage of existing single service users whose surplus decreases after the transition to two service classes. This fraction should be taken into account and kept sufficiently low by the regulator to assure that these users do not block the transition. This is exactly the reason why we consider regulator 3. Indeed, with the unregulated ISP and regulator 1, more than 50%



(a) Percent Surplus Change for Existing Single Class users ($\tau = 0.02$) (b) Percentage of Single Class users with loss of surplus

Fig. 1. Comparison of Regulatory Regimes

of single service class users lose from the transition to two service classes. But, for regulator 3, the percentage of losing users is at most 10%. For regulator 2, the percentage of users with surplus loss from transition to two service classes is 40% at $\tau = 0.1$. But, as τ decreases, this fraction decreases and reaches 10% at $\tau = 0.01$. To sum up, regulator 3 makes the transition to two service classes politically viable and this transition is socially desirable (despite the fact that only a limited fraction of capacity is allocated to service h).

5 Discussion and Conclusion

We make the following contributions to the literature. First, we develop robust pricing for the network with two service classes. Second, we investigate the political economic considerations that may constrain the feasibility of transition from the current network (where QoS is not provided) to the network with QoS provision. Third, we propose a simple regulatory tool that permits to alleviate the political economic constraints and make the transition feasible.

Specifically, from our analysis, the transition to two service classes is socially desirable, but it could be blocked due to unfavorable distributional consequences that the transition inflicts on some fraction of current network users. We introduce a regulator (regulator 3), whose objective is to maximize the surplus of the existing users. We show that with this regulator, at any unit cost of capacity at most 10% of these users will experience surplus loss by the transition to two service classes. Regulator 3 reaches this outcome via the imposition of a ceiling on the fraction of installed capacity that the ISPs are allowed to allocate to the provision of QoS services.

References

- [1] Gupta, A., Jukic, B., Li, M., Stahl, D.O., Whinston, A.B.: Estimating Internet Users' Demand Characteristics. *Computational Economics* 17(2-3), 203–218 (2001)
- [2] Varian, H.R.: The Demand for Bandwidth: Evidence from the INDEX Project (2001), <http://www.sims.berkeley.edu/~hal/Papers/brookings.pdf>
- [3] Altmann, J., Rupp, B., Varaiya, P.: Effects of Pricing on Internet User Behavior. *Netnomics* 3(1), 67–84 (2001)
- [4] He, L., Walrand, J.: Pricing and Revenue Sharing Strategies for Internet Service Providers. In: *Proc. IEEE INFOCOM 2005*, pp. 205–216 (2005)
- [5] Odlyzko, A.: The Economics of the Internet: Utility, Utilization, Pricing, and Quality of Service. Technical report (1999)
- [6] Crowcroft, J., Hand, S., Mortier, R., Roscoe, T., Warfield, A.: QoS's Downfall: At the Bottom, or Not at All!. In: *RIPQoS 2003: Proc. ACM SIGCOMM workshop on Revisiting IP QoS*, pp. 109–114. ACM, New York (2003)
- [7] Felten, E.: Nuts and Bolts of Network Neutrality. Working Paper (2006), <http://itpolicy.princeton.edu/pub/neutrality.pdf>
- [8] Yoo, C.S.: Network Neutrality and the Economics of Congestion. *Georgetown Law Journal* 94 (June 2006)
- [9] Sidak, J.G.: A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet. SSRN eLibrary
- [10] Stahl, D.O., Dai, R., Whinston, A.B.: An Economic Analysis of Multiple Internet QoS Channels (2003)
- [11] He, L., Walrand, J.: Pricing Differentiated Internet Services. In: *Proc. IEEE INFOCOM 2005*, vol. 1, pp. 195–204 (March 2005)
- [12] Odlyzko, A.: Paris Metro Pricing for the Internet. In: *EC 1999: Proc. of the 1st ACM conference on Electronic commerce*, pp. 140–147. ACM, New York (1999)
- [13] Gibbens, R., Mason, R., Steinberg, R.: Internet Service Classes under Competition 18(12), 2490–2498 (2000)
- [14] de Montmarin, M.D.M.: A Result Similar to the Odlyzko's Paris Metro Pricing. *Applied Economics* 38, 1821–1824 (2006)
- [15] Schwartz, G., Shetty, N., Walrand, J.: Impact of QoS on Internet User Welfare: Effect of ISP Competition (in preparation, 2008), www.eecs.berkeley.edu/~nikhils/SSW-MultiQOS.pdf
- [16] Schwartz, G., Shetty, N., Walrand, J.: Network Neutrality: Avoiding the Extremes. In: *Forty-sixth Annual Allerton Conference* (2008)
- [17] Walrand, J.: *Economic Models of Communication Networks*, pp. 57–87. Springer, Heidelberg (2008)
- [18] Hermalin, B.E., Katz, M.L.: The Economics of Product-Line Restrictions With an Application to the Network Neutrality Debate (July 2006)
- [19] Schwartz, G., Shetty, N., Walrand, J.: Modeling the Impact of QoS on Internet User Welfare and ISP Incentives. Working Paper (2008), www.eecs.berkeley.edu/~nikhils/SSW-QOS.pdf

Nonlinear Pricing with Network Externalities

Dawen Meng^{1,3,*}, Guoqiang Tian^{2,3}, and Lei Sun³

¹ International Trade Research Center, Shanghai Institute of Foreign Trade,
Shanghai, 201600, P.R. China

devinmeng@gmail.com

² Department of Economics, Texas A&M University, College Station,
Texas 77843, USA

³ Shanghai University of Finance and Economics, Shanghai 200433, P.R. China

Abstract. This paper considers the screening problem faced by a monopolist of a network good in a general setting. We fully characterize the optimal contracts in the joint presence of network externalities and asymmetric information about agents' types. We find that the pattern of consumption distortion crucially depends on the degree of network congestability. It is shown that an optimal consumption scheme exhibits a two-way distortion, no distortion on the top, or one-way distortion if and only if network is congestible, neutral-congestible or discongestible.

Keywords: Nonlinear Pricing, Network Externalities, Distortion.

1 Introduction

Two prevailing rules in the classical adverse selection model are “*no distortion on the top*” and “*one-way distortion for less efficient types*”, which means under asymmetric information, the most efficient agent has the first-best allocation, all the other types have allocations lower than the first-best one. These rules were supported by many seminal studies such as those in Maskin and Riley (1984), Baron and Myerson (1982), Mussa and Rosen (1978), etc. However, these rules may not be true in the presence of *externalities*. In this paper, we will discuss this issue in a nonlinear pricing environment.

An externality is present whenever the well-being of a consumer or the production possibilities of a firm are directly affected by the actions of another in the economy. In consumption sectors, externalities arise whenever a person's utility function includes arguments controlled by the other consumers. Among all forms of consumption externalities, network externality is the most typical one. It is defined as a change in the benefit, or surplus, that an agent derives from a good when the total consumption of the same kind of goods changes. Such a phenomenon might arise for any of a number of reasons: because the usefulness of the product depends directly on the size of the network (e.g., telephones, fax machines); or because of bandwagon effect, which means the desire to be in style: to have a good because almost everyone else has it; or indirectly

* Corresponding author.

through the availability of complementary goods and services (often known as the “hardware-software paradigm”).

Most studies on consumption externalities in principal-agent setup find that “no distortion on the top” rule does not hold any more, but the consumptions still exhibit “one-way distortion”. Hahn (2003) built a model of telecommunication to examine the role of call and network externalities in nonlinear pricing. He concludes that in equilibrium all types end up with suboptimal quantities. He attributes this result to the existence of call externalities which is the benefit of incoming calls to a subscriber who does not have to pay for the calls. Segal (1999, 2003) developed a general model of contracting with externalities and characterize the nature of the arising inefficiencies. When externalities are positive, he shows that each agent’s consumption level is smaller in the resulting equilibrium allocation than in the socially efficient one. Strategic complementarity is identified as the factor accounting for this result. However, his analysis makes two additional assumptions in identifying the direction of distortion: first, the consumers are identical, and second, total welfare depends only on aggregate trade and not on its allocation across consumers. Csorba (2008) showed that underconsumption result holds even without these two assumptions if externalities are positive. Applying monotone comparative static tools, he demonstrates that the joint presence of asymmetric information and positive network effects leads to a strict downward distortion for all consumers in the quantities provided.

It will be seen in this paper that the one-way distortion results in Segal (1999, 2003), and Csorba (2008) only allow the moderate complementarity of consumers. If the consumers are strongly complementary (i.e., network is strongly discongestible) to one another, then the concavity of objective function guaranteeing the existence of global maximizer may be violated, and thus there may exist no optimal contract for the principal. Furthermore, if the consumers are rivals or substitutes to each other (i.e., network is congestible), the one-way distortion results are no longer true. We give a full characterization of the pattern of distortions in the joint presence of network externalities and asymmetric information about agents’ types. It is shown that the pattern of consumption distortion crucially depends on the degree of network congestability. Specifically, we show that an optimal contract exhibits a “two-way distortion”, “no distortion on the top”, or “one-way distortion” if and only if network is congestible, neutral-congestible or mildly discongestible.

The reminder of this paper is organized as follows: Section 2 sets up economic environments. Section 3 considers the design of optimal nonlinear pricing contract in the presence of network externalities. Finally, concluding remarks are offered in Section 4.

2 Economic Environments

Consider a principal-agent model in which the principal is a monopolist of a network good with marginal production cost c and total output q . The principal’s

payoff function is given by $V = t - cq$, where t is the payment received from consumers.

Consumers have heterogenous preferences for the good, suppose there are n different types of consumers, so that $\theta_1 < \theta_2 < \dots < \theta_n$ and let Θ denote the set of types. Let $f(\theta_i)$ represents the proportion of θ_i type consumers in the whole population. Denote by $\Delta\theta \equiv \theta_i - \theta_{i-1} > 0$ the difference of adjacent types and $F(\theta_i) = \sum_{j \leq i} f(\theta_j)$ the cumulative frequency. Then by Law of Large Numbers, it is mathematically equivalent to the framework with a single agent whose type θ is distributed with c.d.f $F(\cdot)$ and p.d.f $f(\cdot)$ on $\Theta = \{\theta_1, \dots, \theta_n\}$.

A consumer of θ_i type is assumed to have an utility function of $U_i = \theta_i V(q_i) + \Psi(Q) - t_i$, where q_i is the amount of the network good he consumes, $Q = \sum_i f(\theta_i)q_i$ is the total amount of network good in the economy (network magnitude) and t_i is the tariff charged for q_i by the principal. The utility function can be divided into two terms additively: $\theta_i V(q_i)$ is the *intrinsic value* of consuming, while $\Psi(Q)$ is the *network value*. Note that, we assume the network effect is homogeneous among all the consumers, namely, the network value is independent of individual preference θ_i and individual consumption q_i .

It is assumed that $V'(q) > 0$ and $V''(q) < 0$. The degree of network congestability is defined by the sign of $\Psi''(Q)$.

Definition 1. *The network is congestible, neutral-congestible, or discongestible if and only if $\Psi''(Q) < 0$, $\Psi''(Q) = 0$, or $\Psi''(Q) > 0$, respectively.*

Remark 1. When the network capacity is large and the maintaining technology is advanced enough, an increase in one consumer’s consumption will increase the marginal utilities of others, and so $\Psi''(Q) > 0$. When network capacity and maintaining technology are limited, consumers are rivals to one another in the sense that an increase in one consumer’s consumption will decrease the marginal utilities of others, and thus $\Psi''(Q) < 0$. Finally, when the expansion of network benefits all the consumers with constant margin, the network value term is a linear function with $\Psi''(Q) = 0$.

The objective of the monopolist is to design a menu of incentive-compatible and self-selecting quantity-price pairs $\{q(\hat{\theta}), t(\hat{\theta})\}$ to maximize her own expected revenue, where $\hat{\theta} \in \Theta$ is the consumer’s announcement. The timing of contracting is given by the following procedures:

- Stage 1. The consumer observe his own “type” θ .
- Stage 2. The monopolist offers a contract $\{q(\hat{\theta}), t(\hat{\theta})\}$.
- Stage 3. The consumer accepts or refuses the contract.
- Stage 4. The contract is executed, both parties get their respective payoffs.

If the monopolist should offer the same menu of contracts for all consumers, all the consumers except of the lowest type θ_1 have incentive to mimic the lower types, as in the standard adverse selection model, an incentive-compatible menu $\{q_i, t_i\}_{i=1}^n$ should satisfy participation constraints (IR_i) and incentives constraints (IC_{ij}):

$$\begin{aligned}
 IR_i &: \theta_i V(q_i) + \Psi(Q) - t_i \geq 0 \\
 IC_{ij} &: \theta_i V(q_i) + \Psi(Q) - t_i \geq \theta_i V(q_j) + \Psi(Q) - t_j^1
 \end{aligned}$$

or equivalently, written in the form of information rents as:

$$\begin{aligned}
 IR_i &: U_i \geq 0, \\
 IC_{ij} &: U_i \geq U_j + (\theta_i - \theta_j)V(q_j), \forall i, j.
 \end{aligned}$$

Here the reservation utilities of the consumers are normalized to be zero.

As in standard incentive theory literature, we first analyze the set of constraints to find the binding ones².

Lemma 1. *In the second-best optimal contract there are n binding constraints: IR_1 , the individual rationality constraint of the lowest-type consumer; and $IC_{i(i-1)}$, for $i = 2, 3, \dots, n$, the downward local incentive constraints.*

Proof. The proof is omitted since it is standard.

3 Economies with Network Externalities

In this section, we considers the principal’s contract-designing problem in the presence of network externalities. To get a reference system for comparison, let us first suppose that there is no asymmetry of information between the monopolist and consumer. The first-best³ consumption levels are obtained by equating the consumer’s marginal utility and the monopolist’s marginal cost. Hence, we have the following first-order conditions:

$$\theta_i V'(q_i) + \Psi'(Q) = c, \forall i \in \{1, 2, \dots, n\}. \tag{1}$$

Under asymmetric information, with the incentive compatible constraints and participation constraints, the monopolist’s optimization program can be represented as:

$$(P) \left\{ \begin{array}{l} \max_{\{U_i, q_i\}} \left\{ \sum_{i=1}^n f(\theta_i) [\theta_i V(q_i) - cq_i] + \Psi \left(\sum_{i=1}^n f(\theta_i) q_i \right) - \sum_{i=1}^n f(\theta_i) U_i \right\} \\ \text{s.t. } IR_i : U_i \geq 0, \\ IC_{ij} : U_i \geq U_j + (\theta_i - \theta_j)V(q_j), \forall i, j \end{array} \right.$$

By adding incentive constraints ($IC_{i+1,i}$) and ($IC_{i,i+1}$), we see that $q_{i+1} \geq q_i$, which is necessary for a implementable mechanism. From Lemma 1 we have $U_i = U_{i-1} + \Delta\theta V(q_{i-1}) = \dots = \Delta\theta \sum_{j=1}^{i-1} V(q_j)$. By substituting into these

¹ Note that, there is a continuum of consumers in each type, so the individual misreport has no effect on the total consumption Q .

² Here binding means holding with equality.

³ In this paper we denote by “first-best” the case with complete information, “second-best” the case with asymmetric information.

information rent functions, the principal’s objective function in (P) simplifies to the following form:

$$\sum_{i=1}^n \left[\theta_i - \frac{1 - F(\theta_i)}{f(\theta_i)} \Delta\theta \right] f(\theta_i) V(q_i) + \Psi \left(\sum_{i=1}^n f(\theta_i) q_i \right) - c \sum_{i=1}^n f(\theta_i) q_i . \quad (2)$$

The first-best and second-best consumptions can be thus regarded as the solutions to the following unconstrained program parameterized on ϵ :

$$\max_{\mathbf{q} \in \mathbb{R}_+^n} \Pi(\mathbf{q}, \epsilon) , \quad (3)$$

where

$$\Pi(\mathbf{q}, \epsilon) = \sum_{i=1}^n \left[\theta_i + \epsilon H(\theta_i) \right] f(\theta_i) V(q_i) + \Psi \left(\sum_{i=1}^n f(\theta_i) q_i \right) - c \sum_{i=1}^n f(\theta_i) q_i , \quad (4)$$

$H(\theta) \equiv \frac{1 - F(\theta)}{f(\theta)} \Delta\theta$, $\mathbf{q} = (q_1, q_2, \dots, q_n) \in \mathbb{R}_+^n$ and $\epsilon \in [-1, 0]$. If $\epsilon = 0$, the solution to (3) is the first-best consumption given in (11); if $\epsilon = -1$, it is the second-best consumption characterized by the following proposition. We ignore the implementability conditions $q_n \geq \dots \geq q_1$ for the moment, and check at the end whether they are satisfied in equilibrium.

Proposition 1. *If the weak monotone hazard rate condition $\frac{d}{d\theta} \left[\frac{1 - F(\theta)}{f(\theta)} \right] \leq 0$ is satisfied, and the Hessian matrix Π_{qq} is negative definite at $(q_1^{SB}, q_2^{SB}, \dots, q_n^{SB})$ for $\epsilon = -1$, then the second-best consumption q_i^{SB} satisfies*

$$\left[\theta_i - H(\theta_i) \right] V'(q_i) + \Psi' \left(\sum_{i=1}^n f(\theta_i) q_i \right) = c, \forall i \in \{1, 2, \dots, n\} . \quad (5)$$

Proof. Expression (5) can be attained directly from the first order condition of (3). The weakly hazard rate property $H'(\theta) \leq 0$ ensures the implementability conditions $q_i^{SB} \leq q_{i+1}^{SB}, \forall i \in \{1, 2, \dots, n\}$, and the negative definiteness of Hessian matrix Π_{qq} ensures the second order sufficient condition of optimization. Thus (5) gives the second-best consumptions.

We now compare the first-best and the second-best consumptions in the following proposition.

Proposition 2. *Suppose that $V'(\cdot) > 0, V''(\cdot) < 0$, and the weakly monotone hazard rate property $\frac{d}{d\theta} \left[\frac{1 - F(\theta)}{f(\theta)} \right] \leq 0$ holds. Then distortion way of consumptions depends on the degree of network congestibility.*

1. *If the network is mildly discongestible, i.e., the value of $\Psi''(Q) > 0$ is not too big such that the Hessian matrix Π_{qq} is negative definite for all $\mathbf{q} \in \mathbb{R}_+^n$ and $\epsilon \in [-1, 0]$,⁴ then the consumption exhibits one-way distortion: $q_i^{SB} < q_i^{FB}, \forall i$.*

⁴ If the Hessian matrix is not negative definite, there might exist no global maximizer for the principal’s program (3).

2. If the network is congestible with $\Psi''(Q) < 0$, then it exhibits two-way distortion. That means there exists a threshold value $i^* \in \{1, 2, \dots, n\}$, for $i > i^*$, $q_i^{SB} > q_i^{FB}$; for $i < i^*$, $q_i^{SB} < q_i^{FB}$.
3. If the network is neutral-congestible with $\Psi''(Q) = 0$, then it exhibits “one-way distortion” and “no distortion on the top”: $q_i^{SB} < q_i^{FB}, \forall i < n$ and $q_n^{SB} = q_n^{FB}$.

In all these cases the network magnitude is downsized: $Q^{SB} < Q^{FB}$.

Proof. Let

$$\Gamma \equiv \text{diag}\{\lambda_1, \lambda_2, \dots, \lambda_n\}, \tag{6}$$

$$\gamma \equiv (f(\theta_1), f(\theta_2), \dots, f(\theta_n))^T, \tag{7}$$

where $\lambda_i = [\theta_i + \epsilon H(\theta_i)] f(\theta_i) V''(q_i)$, then

$$\Pi_{qq} \equiv \Gamma + \Psi''(Q) \gamma \gamma^T \tag{8}$$

$$\Pi_{q\epsilon} \equiv \left(H(\theta_1) f(\theta_1) V'(q_1), H(\theta_2) f(\theta_2) V'(q_2), \dots, H(\theta_n) f(\theta_n) V'(q_n) \right)^T. \tag{9}$$

The first-order condition to (3) is $\Pi_{q\epsilon} = 0$. The negative definiteness of Hessian matrix Π_{qq} acts as the second-order sufficient condition of the principal’s optimization program, and the monotone hazard rate property ensures the implementability conditions $q_{i+1} \geq q_i, \forall i$.

Differentiating the above first-order condition with respect to parameter ϵ , we attain

$$\Pi_{qq} \frac{dq}{d\epsilon} + \Pi_{q\epsilon} = 0, \tag{10}$$

which implies

$$\frac{dq}{d\epsilon} = -(\Pi_{qq})^{-1} \Pi_{q\epsilon}. \tag{11}$$

Substituting expressions (8) and (9) into the above expression, we have

$$\frac{dq}{d\epsilon} = -\left[\Gamma + \Psi''(Q) \gamma \gamma^T \right]^{-1} \Pi_{q\epsilon} = -\left[\Gamma^{-1} - \Psi''(Q) \frac{\Gamma^{-1} \gamma \cdot \gamma^T \Gamma^{-1}}{1 + \Psi''(Q) \gamma^T \Gamma^{-1} \gamma} \right] \Pi_{q\epsilon}. \tag{12}$$

The i th elements of the LHS is then given by

$$\frac{dq_i}{d\epsilon} = \frac{\rho - H(\theta_i) V'(q_i)}{[\theta_i + \epsilon H(\theta_i)] V''(q_i)}. \tag{13}$$

where

$$\rho \equiv \frac{\Psi''(Q) \sum_{j=1}^n \frac{f(\theta_j) H(\theta_j) V'(q_j)}{[\theta_j + \epsilon H(\theta_j)] V''(q_j)}}{1 + \Psi''(Q) \sum_{j=1}^n \frac{f(\theta_j)}{[\theta_j + \epsilon H(\theta_j)] V''(q_j)}}.$$

Note that the Hessian matrix Π_{qq} and hence its inverse Π_{qq}^{-1} are negative definite. So for non-zero vector γ we have

$$\gamma^T \Pi_{qq}^{-1} \gamma = \gamma^T \left[\Gamma^{-1} - \Psi''(Q) \frac{\Gamma^{-1} \gamma \cdot \gamma^T \Gamma^{-1}}{1 + \Psi''(Q) \gamma^T \Gamma^{-1} \gamma} \right] \gamma = \frac{\gamma^T \Gamma^{-1} \gamma}{1 + \Psi''(Q) \gamma^T \Gamma^{-1} \gamma} < 0.$$

$V''(\cdot) < 0$ implies every element of the diagonal matrix Γ^{-1} is negative, so $\gamma^T \Gamma^{-1} \gamma < 0$, and thus the denominator of ρ is positive, namely,

$$1 + \Psi''(Q) \gamma^T \Gamma^{-1} \gamma = 1 + \Psi''(Q) \sum_{j=1}^n \frac{f(\theta_j)}{[\theta_j + \alpha H(\theta_j)] V''(q_j)} > 0.$$

The sign of ρ and hence the sign of $\frac{dq_i}{d\epsilon}$ are determined by the sign of $\Psi''(Q)$.

1. If $\Psi''(Q) > 0$, then $\rho < 0$, $\frac{dq_i}{d\epsilon} > 0$, which implies $q_i^{SB} < q_i^{FB}$.
2. If $\Psi''(Q) < 0$, then $\rho > 0$. From conditions $H'(\theta) \leq 0$, $q_i \leq q_{i+1}$, and $V''(\cdot) < 0$, we get $H(\theta_1)V'(q_1) > H(\theta_2)V'(q_2) > \dots > H(\theta_n)V'(q_n)$, then we have the following inequalities:

$$0 = H(\theta_n)V'(q_n) < \rho < \frac{\sum_{j=1}^n \frac{f(\theta_j)H(\theta_j)V'(q_j)}{[\theta_j + \alpha H(\theta_j)]V''(q_j)}}{\sum_{j=1}^n \frac{f(\theta_j)}{[\theta_j + \alpha H(\theta_j)]V''(q_j)}} < H(\theta_1)V'(q_1).$$

Then there exists a unique $i^* \in \{1, 2, \dots, n\}$ such that, when $i > i^*$, we have $\rho > H(\theta_i)V'(q_i)$, $\frac{dq_i}{d\epsilon} < 0$, and consequently $q_i^{SB} > q_i^{FB}$; when $i < i^*$, we have $\rho < H(\theta_i)V'(q_i)$, $\frac{dq_i}{d\epsilon} > 0$, and consequently $q_i^{SB} < q_i^{FB}$.

3. If $\Psi''(Q) = 0$, then $\rho = 0$, consequently $\frac{dq_n}{d\epsilon} = 0$, $\frac{dq_i}{d\epsilon} > 0, \forall i < n$, which implies $q_n^{SB} = q_n^{FB}$ and $q_i^{SB} < q_i^{FB}, \forall i < n$.

The derivative of Q with respect to ϵ is

$$\frac{dQ}{d\epsilon} = \sum_{i=1}^n \left\{ \frac{[\rho - H(\theta_i)V'(q_i)] f(\theta_i)}{[\theta_i + \epsilon H(\theta_i)] V''(q_i)} \right\} = \frac{-\sum_{i=1}^n \frac{f(\theta_i)H(\theta_i)V'(q_i)}{[\theta_i + \epsilon H(\theta_i)]V''(q_i)}}{1 + \Psi''(Q) \sum_{i=1}^n \frac{f(\theta_i)}{[\theta_i + \epsilon H(\theta_i)]V''(q_i)}} > 0. \tag{14}$$

Thus, we have $Q^{SB} < Q^{FB}$ for all cases.

Q.E.D.

The result can be interpreted as follows. In order to minimize the information rent captured by the higher-type consumers, the allocation of lower-type agents should be reduced relative to the first-best level. This is the basic tradeoff between allocation efficiency and rent extraction in the standard adverse selection settings. In contrast to the canonical settings, the consumptions of different consumers in our model interact through network effects. In a mildly discongestible network, they are complementary to one another in the sense that the consumptions of higher type are also distorted downward. On the contrary, in a congestible network, the consumers are rivals to one another, and then the consumptions of the higher-type consumers are distorted upward to “counteract” the decrease of that of the lower type. If the network is neutral-congestible, this “feedback” effect of the lower-type users consumption on that of the higher-type users disappears so that the classical results are still obtained.

4 Conclusion

This paper presents a model of nonlinear pricing in the joint presence of network externalities and asymmetric information. We give a full characterization of the

nonlinear pricing contract. We showed that in mildly discongestible network, network externalities and asymmetric information together lead to a downward distortion for all consumers' consumption levels, which is in line with Hahn(2003), Segal(1999,2003) and Csobra (2008) ; while in a congestible network, the consumptions of different types of consumers will be distorted in opposite directions; in a neutral-congestible network, the results obtained in canonical settings are still available.

Acknowledgments. We wish to thank two anonymous referees for helpful comments and suggestions. The second author thanks the National Natural Science Foundation of China and Private Enterprise Research Center at Texas A&M University for Financial support.

References

1. Baron, D., Myerson, R.: Regulating a Monopolist with Unknown Costs. *Econometrica* 50, 911–930 (1982)
2. Csorba, G.: Screening Contracts in The Presence of Positive Network Effects. *International Journal of Industrial Organization* 26, 213–226 (2008)
3. Hahn, J.: Nonlinear Pricing of Telecommunication with Call and Network Externalities. *International Journal of Industrial Organization* 21, 949–967 (2003)
4. Lockwood, B.: Production Externalities and Two-way Distortion in Principal-multi-agent Problems. *Journal of Economic Theory* 92, 144–166 (2000)
5. Maskin, E., Riley, J.: Monopoly with Incomplete Information. *Rand Journal of Economics* 15, 171–196 (1984)
6. Mussa, M., Rosen, S.: Monopoly and Product Quality. *Journal of Economic Theory* 18, 301–317 (1978)
7. Myerson, R.: Optimal Auction Design. *Mathematics of Operations Research* 6, 58–73 (1981)
8. Segal, I.: Contracting with Externalities. *Quarterly Journal of Economics* 114, 337–388 (1999)
9. Segal, I.: Coordination and Discrimination in Contracting with Externalities: Divide and Conquer. *Journal of Economic Theory* 113, 147–181 (2003)

Author Index

- Aggarwal, Gagan 621
Agrawal, Shipra 126
Anshevich, Elliot 386
Apt, Krzysztof 426
Arcaute, Esteban 263
Asadpour, Arash 477
Ashlagi, Itai 675
Auletta, Vincenzo 278
- Bilò, Davide 251
Bilò, Vittorio 70
Bonifaci, Vincenzo 239
Brandes, Ulrik 394
Brandt, Felix 198
Bu, Tian-Ming 138
- Cai, Mao-cheng 454
Caragiannis, Ioannis 374
Chakrabarty, Deeparnab 498, 566
Chakraborty, Tanmoy 548
Chalkiadakis, Georgios 307
Chang, Kevin L. 41
Chaudhuri, Kamalika 522
Chen, Yiling 556
Chung Graham, Fan 15, 522
Conitzer, Vincent 295, 426, 696
- Dang, Chuangyin 31
Daskalakis, Constantinos 186
Deng, Xiaotie 138
Dobzinski, Shahar 653
Dumrauf, Dominic 94
- Elkind, Edith 307
Ergun, Özlem 61
Even-Dar, Eyal 577
- Fanelli, Angelo 70
Fang, Qizhi 454
Feldman, Jon 577, 621
Feldman, Michal 109
Fischer, Felix 198
Flammini, Michele 70
Fleischer, Lisa 19
- Gallo, Edoardo 322
Garcia-Molina, Hector 53
Garg, Rahul 19
Gassner, Elisabeth 82
Ghosh, Arpita 556
Giotis, Ioannis 629
Goel, Gagan 498
Gonen, Rica 597
Gourvès, Laurent 465
Gradwohl, Ronen 539
Grigoriev, Alexander 362
Gualà, Luciano 251
Gui, Luyi 61
Gujar, Sujit 438
Guo, Mingyu 426
- Hansen, Kristoffer Arnsfelt 684
Hansen, Thomas Dueholm 490, 684
Harks, Tobias 239
Hatzl, Johannes 82
Heydenreich, Birgit 414
Hofer, Martin 286, 394
Holzer, Markus 198
Huang, Wan 506
Huberman, Bernardo A. 334
- Jackson, Matthew O. 4
Jain, Kamal 696
Jain, Shaili 342
Jamall, Mohammad Shoaib 522
Jennings, Nicholas R. 307
Johari, Ramesh 263
Johnson, Aaron 41
- Kaklamanis, Christos 374
Kanellopoulos, Panagiotis 374
Kanoulas, Dimitrios 222
Kapoor, Sanjiv 19
Karlin, Anna R. 629
Kearns, Michael 548, 639
Kempe, David 585, 609
Khandekar, Rohit 19
Kontogiannis, Spyros C. 210
Krumke, Sven O. 82
Krysta, Piotr 675

- Lambert, Nicolas 154
 Lau, Lawrence J. 2
 Lelarge, Marc 178
 Letchford, Joshua 696
 Lin, Qianya 138
 Lu, Pinyan 231, 402
 Lukose, Rajan 566
 Luo, Zhi-Quan (Tom) 16

 Mahdian, Mohammad 585, 708
 Mannor, Shie 263
 Mansour, Yishay 577
 Markakis, Evangelos 307, 426
 Maskin, Eric 1
 McAfee, Randolph Preston 556, 708
 Meng, Dawen 724
 Milgrom, Paul 12
 Miltersen, Peter Bro 684
 Mishra, Debasis 414
 Monien, Burkhard 94
 Monnot, Jérôme 465
 Moscardelli, Luca 70, 278
 Müller, Rudolf 414
 Mungamuru, Bobji 53
 Muthukrishnan, S. 577, 621

 Naor, Joseph (Seffi) 514
 Nazerzadeh, Hamid 477
 Nick, Bobo 394
 Nisan, Noam 531

 Olbrich, Lars 286

 Pál, Martin 621
 Papadimitriou, Christos 17
 Parkes, David C. 342
 Penna, Paolo 278
 Pennock, David 556, 708
 Persiano, Giuseppe 278
 Procaccia, Ariel D. 653
 Proietti, Guido 251

 Qi, Qi 138

 Raz, Danny 514
 Roth, Aaron 118

 Saberi, Amin 19, 477
 Salek, Mahyar 609

 Scalosub, Gabriel 514
 Scarf, Herbert E. 18
 Schäfer, Guido 239
 Schapira, Michael 351, 531
 Schoppmann, Florian 665
 Schwartz, Galina 716
 Shetty, Nikhil 716
 Shoham, Yoav 154
 Singer, Yaron 351
 Skopalik, Alexander 286
 Sørensen, Troels Bjerre 684
 Sperber, Heike 82
 Spirakis, Paul G. 210, 222
 Sun, Lei 724

 Tamir, Tami 109
 Tan, Jinsong 639
 Telelis, Orestis A. 490
 Tennenholtz, Moshe 675
 Tian, Guoqiang 724
 Tsaknakis, Haralampos 222

 Uetz, Marc 362, 414

 van Loon, Joyce 362
 Varian, Hal R. 13
 Vassilvitskii, Sergei 597
 Vazirani, Vijay V. 498
 von Stengel, Bernhard 506

 Walrand, Jean 716
 Wang, Lei 498
 Wang, Zizhuo 126
 Widmayer, Peter 251
 Wilfong, Gordon 386
 Woeginger, Gerhard J. 82
 Wu, Fang 334
 Wu, John 166

 Yadati, Narahari 438
 Ye, Yinyu 14, 31, 126
 Yu, Changyuan 231, 402, 498
 Yu, Guosong 446

 Zhang, Guochuan 446
 Zhou, Yunhong 566
 Zhu, Zhisu 31
 Zohar, Aviv 531