

Stefano Leonardi (Ed.)

ARCoSS

LNCS 5929

# Internet and Network Economics

Advanced Research in Computing and Software Science

5th International Workshop, WINE 2009  
Rome, Italy, December 2009  
Proceedings



Springer

*Commenced Publication in 1973*

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

## Editorial Board

David Hutchison, UK

Josef Kittler, UK

Alfred Kobsa, USA

John C. Mitchell, USA

Oscar Nierstrasz, Switzerland

Bernhard Steffen, Germany

Demetri Terzopoulos, USA

Gerhard Weikum, Germany

Takeo Kanade, USA

Jon M. Kleinberg, USA

Friedemann Mattern, Switzerland

Moni Naor, Israel

C. Pandu Rangan, India

Madhu Sudan, USA

Doug Tygar, USA

## Advanced Research in Computing and Software Science

Subline of Lectures Notes in Computer Science

### Subline Series Editors

Giorgio Ausiello, *University of Rome 'La Sapienza', Italy*

Vladimiro Sassone, *University of Southampton, UK*

### Subline Advisory Board

Susanne Albers, *University of Freiburg, Germany*

Benjamin C. Pierce, *University of Pennsylvania, USA*

Bernhard Steffen, *University of Dortmund, Germany*

Madhu Sudan, *Massachusetts Institute of Technology, MA, USA*

Deng Xiaotie, *City University of Hong Kong*

Jeannette M. Wing, *Carnegie Mellon University, Pittsburgh, PA, USA*

Stefano Leonardi (Ed.)

# Internet and Network Economics

5th International Workshop, WINE 2009  
Rome, Italy, December 14-18, 2009  
Proceedings

Volume Editor

Stefano Leonardi  
Sapienza University of Rome  
Department of Information and Systems  
Via Ariosto 25, 00185 Roma, Italy  
E-mail: Stefano.Leonardi@dis.uniroma1.it

Library of Congress Control Number: 2009939834

CR Subject Classification (1998): H.4, K.4.4, J.1, J.3, 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 3-642-10840-7 Springer Berlin Heidelberg New York  
ISBN-13 978-3-642-10840-2 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.com

© Springer-Verlag Berlin Heidelberg 2009  
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India  
Printed on acid-free paper SPIN: 12809240 06/3180 5 4 3 2 1 0



# Preface

This volume contains the papers presented at WINE 2009: the 5th International Workshop on Internet and Network Economics held December 14–18, 2009, in Rome, at the Department of Computer and System Sciences, Sapienza University of Rome.

Over the past decade there has been growing interaction between researchers in theoretical computer science, networking and security, economics, mathematics, sociology, and management sciences devoted to the analysis of problems arising in the Internet and the worldwide web. The Workshop on Internet and Network Economics (WINE) is an interdisciplinary forum for the exchange of ideas and results arising in these varied fields.

There were 142 submissions to the workshop including regular and short papers. All submissions were rigorously peer reviewed and evaluated on the basis of the quality of their contribution, originality, soundness, and significance. Almost all submissions were reviewed by at least three Program Committee members. The committee decided to accept 34 regular papers and 29 short papers. The Best Student Paper award sponsored by Google Inc. was given to Saeed Alaei and Azarakhsh Malekian for the paper “An Analysis of Troubled Assets Reverse Auction.”

The program also included three invited talks by S. Muthukrishnan (Google Inc. and Rutgers University), H. Peyton Young (Oxford and Johns Hopkins University) and Eva Tardos (Cornell University). Three tutorials were also offered on the days before the workshop, from Andrei Broder (Yahoo! Research) on Computational Advertising, Nikhil Devanur and Kamal Jain (Microsoft Research) on Computational Issues in Market Equilibria, and Tim Roughgarden (Stanford University) on Bayesian and Worst-Case Revenue Maximization.

We would like to thank Google Inc., Microsoft Research, Yahoo! Research, Fondazione Ugo Bordoni and Sapienza University of Rome for the generous financial support to WINE 2009. We would also like to thank the Department of Computer and System Sciences, Sapienza University of Rome, for hosting the event. Vincenzo Bonifaci and Piotr Sankowski offered their precious help for the review process, the conference website and the workshop proceedings. We also acknowledge EasyChair, a fantastic, robust, easy to use, freely available system for managing the work of the Program Committee and the production of workshop proceedings.

# Conference Organization

## Program Committee Chair

Stefano Leonardi                      Sapienza University of Rome, Italy

## Program Committee

Moshe Babaioff	Microsoft Research
Patrick Briest	Cornell University, USA
Andrei Broder	Yahoo! Research
Shuchi Chawla	Wisconsin University, USA
Ning Chen	Nanyang Technological University, Singapore
Bruno Codenotti	National Research Council, Italy
Artur Czumaj	Warwick University, UK
Xiaotie Deng	City University of Hong Kong
Amos Fiat	Tel-Aviv University, Israel
Michele Flammini	L'Aquila University, Italy
Martin Gairing	University of California at Berkeley, USA
Arpita Ghosh	Yahoo! Research
Nicole Immorlica	Northwestern University, USA
Matthew Jackson	Stanford University, USA
Kamal Jain	Microsoft Research
Ehud Kalai	Northwestern University, USA
Piotr Krysta	Liverpool University, UK
Ron Lavi	The Technion, Israel
Aranyak Mehta	Google Inc.
Paul Milgrom	Stanford University, USA
S. Muthukrishnan	Google Inc. and Rutgers University, USA
Rudolf Müller	Maastricht University, The Netherlands
Tim Roughgarden	Stanford University, USA
Guido Schaefer	CWI, Amsterdam, The Netherlands
Paul Spirakis	CTI, Patras, Greece
Sergei Vassilvitskii	Yahoo! Research
Berthold Voecking	RWTH Aachen University, The Netherlands
Yinyu Ye	Stanford University, USA

## Local Organization

Giorgio Ausiello  
Vincenzo Bonifaci  
Stefano Leonardi (Chair)  
Alberto Marchetti-Spaccamela  
Piotr Sankowski

## External Reviewers

Abraham, Ittai  
 Augustine, John  
 Aziz, Haris  
 Berenbrink, Petra  
 Bernstein, Shai  
 Bilo', Vittorio  
 Bonifaci, Vincenzo  
 Bouveret, Sylvain  
 Brenner, Janina  
 Bu, Tian-Ming  
 Buchbinder, Niv  
 Buchbinder, Niv  
 Budinich, Michele  
 Carlsson, John  
 Cavallo, Ruggiero  
 Cesa-Bianchi, Nicolo  
 Chen, Ho-Lin  
 Chen, Xi  
 Chen, Yiling  
 Christodoulou, Giorgos  
 Codenotti, Paolo  
 Dai, Decheng  
 Daskalakis, Costis  
 Delage, Erick  
 Dhangwatnotai, Peerapong  
 Du, Ye  
 Dumrauf, Dominic  
 Eberhart, Robert  
 Englert, Matthias  
 Fanelli, Angelo  
 Feldman, Michal  
 Feldmann, Rainer  
 Fortnow, Lance  
 Fu, Hu  
 Garg, Dinesh  
 Goldberg, Paul  
 Gravin, Nikolai  
 Guo, Mingyu  
 Harks, Tobias  
 Hoefer, Martin  
 Hua, Xia  
 Huo, Yumei  
 Kanellopoulos, Panagiotis  
 Kintali, Shiva  
 Lachish, Oded  
 Lahaie, Sebastien  
 Lambert, Nicolas  
 Leyton-Brown, Kevin  
 Ligett, Katrina  
 Lin, Genie  
 Lu, Pinyan  
 Lucier, Brendan  
 Mahini, Hamid  
 Melideo, Giovanna  
 Mobius, Markus  
 Monaco, Gianpiero  
 Morgan, John  
 Moscardelli, Luca  
 Nadav, Uri  
 Ohseto, Shinji  
 Pap, Gyula  
 Piliouras, George  
 Procaccia, Ariel  
 Qi, Qi  
 Raecke, Harry  
 Rothe, Joerg  
 Röglin, Heiko  
 Saia, Jared  
 Salek, Mahyar  
 Sankowski, Piotr  
 Sauerwald, Thomas  
 Schapira, Michael  
 Schoppmann, Florian  
 Sharma, Aneesh  
 Singer, Yaron  
 Slivkins, Aleksandrs  
 So, Anthony Man-Cho  
 Stiller, Sebastian  
 Sundararajan, Mukund  
 Telelis, Orestis  
 Tscheuschner, Tobias  
 Vredeveld, Tjark  
 Wang, Zizhuo  
 Xia, Lirong  
 Yu, Changyuan  
 Zhang, Jie  
 Zhang, Qin  
 Zhao, Hairong

# Table of Contents

## Invited Talks

Ad Exchanges: Research Issues . . . . .	1
<i>S. Muthukrishnan</i>	
Adaptive Learning in Systems of Interacting Agents . . . . .	13
<i>H. Peyton Young</i>	
Quantifying Outcomes in Games (Abstract) . . . . .	17
<i>Éva Tardos</i>	

## Regular Papers

Competitive Routing over Time . . . . .	18
<i>Martin Hoefer, Vahab S. Mirrokni, Heiko Röglin, and Shang-Hua Teng</i>	
On 2-Player Randomized Mechanisms for Scheduling . . . . .	30
<i>Pinyan Lu</i>	
On Stackelberg Pricing with Computationally Bounded Consumers . . . . .	42
<i>Patrick Briest, Martin Hoefer, Luciano Gualà, and Carmine Ventre</i>	
On Best Response Dynamics in Weighted Congestion Games with Polynomial Delays . . . . .	55
<i>Angelo Fanelli and Luca Moscardelli</i>	
Parametric Packing of Selfish Items and the Subset Sum Algorithm . . . . .	67
<i>Leah Epstein, Elena Kleiman, and Julián Mestre</i>	
An Online Multi-unit Auction with Improved Competitive Ratio . . . . .	79
<i>Sourav Chakraborty and Nikhil Devanur</i>	
Prediction Mechanisms That Do Not Incentivize Undesirable Actions . . . . .	89
<i>Peng Shi, Vincent Conitzer, and Mingyu Guo</i>	
Pricing Strategies for Viral Marketing on Social Networks . . . . .	101
<i>David Arthur, Rajeev Motwani, Aneesh Sharma, and Ying Xu</i>	
Consistent Continuous Trust-Based Recommendation Systems . . . . .	113
<i>Ola Rozenfeld and Moshe Tennenholtz</i>	

The Stackelberg Minimum Spanning Tree Game on Planar and Bounded-Treewidth Graphs . . . . .	125
<i>Jean Cardinal, Erik D. Demaine, Samuel Fiorini, Gwenaël Joret, Ilan Newman, and Oren Weimann</i>	
Tighter Bounds for Facility Games . . . . .	137
<i>Pinyan Lu, Yajun Wang, and Yuan Zhou</i>	
Degrees of Guaranteed Envy-Freeness in Finite Bounded Cake-Cutting Protocols . . . . .	149
<i>Claudia Lindner and Jörg Rothe</i>	
Approximate Pure Nash Equilibria via Lovász Local Lemma . . . . .	160
<i>Thành Nguyen and Éva Tardos</i>	
Externalities in Keyword Auctions: An Empirical and Theoretical Assessment . . . . .	172
<i>Renato Gomes, Nicole Immorlica, and Evangelos Markakis</i>	
Covering Games: Approximation through Non-cooperation . . . . .	184
<i>Martin Gairing</i>	
Contract Auctions for Sponsored Search . . . . .	196
<i>Sharad Goel, Sébastien Lahaie, and Sergei Vassilvitskii</i>	
Bidding for Representative Allocations for Display Advertising . . . . .	208
<i>Arpita Ghosh, Preston McAfee, Kishore Papineni, and Sergei Vassilvitskii</i>	
Social Networks and Stable Matchings in the Job Market . . . . .	220
<i>Esteban Arcaute and Sergei Vassilvitskii</i>	
Maximizing the Minimum Load: The Cost of Selfishness . . . . .	232
<i>Leah Epstein, Elena Kleiman, and Rob van Stee</i>	
Competitive Repeated Allocation without Payments . . . . .	244
<i>Mingyu Guo, Vincent Conitzer, and Daniel M. Reeves</i>	
Pseudonyms in Cost-Sharing Games . . . . .	256
<i>Paolo Penna, Florian Schoppmann, Riccardo Silvestri, and Peter Widmayer</i>	
Computing Optimal Contracts in Series-Parallel Heterogeneous Combinatorial Agencies . . . . .	268
<i>Yuval Emek and Michal Feldman</i>	
Nash Equilibria for Voronoi Games on Transitive Graphs . . . . .	280
<i>Rainer Feldmann, Marios Mavronicolas, and Burkhard Monien</i>	

Selfish Scheduling with Setup Times . . . . .	292
<i>Laurent Gourvès, Jérôme Monnot, and Orestis A. Telelis</i>	
Computational Aspects of Multimarket Price Wars . . . . .	304
<i>Nithum Thain and Adrian Vetta</i>	
The Impact of Social Ignorance on Weighted Congestion Games . . . . .	316
<i>Dimítris Fotakis, Vasilis Gkatzelis, Alexis C. Kaporis, and Paul G. Spirakis</i>	
The Complexity of Models of International Trade . . . . .	328
<i>Christopher A. Wilkens</i>	
The Geometry of Truthfulness . . . . .	340
<i>Angelina Vidali</i>	
Optimal Incentives for Participation with Type-Dependent Externalities . . . . .	351
<i>Michal Feldman, Ran Tessler, and Yoav Wilf</i>	
Nash Dynamics in Congestion Games with Similar Resources . . . . .	362
<i>Anand Bhargat, Tanmoy Chakraborty, and Sanjeev Khanna</i>	
Online Ad Assignment with Free Disposal . . . . .	374
<i>Jon Feldman, Nitish Korula, Vahab Mirrokni, S. Muthukrishnan, and Martin Pál</i>	
Truthful and Quality Conscious Query Incentive Networks . . . . .	386
<i>Devansh Dikshit and Narahari Yadati</i>	
An Analysis of Troubled Assets Reverse Auction . . . . .	398
<i>Saeed Alaei and Azarakhsh Malekian</i>	
Direction Preserving Zero Point Computing and Applications: (Extended Abstract) . . . . .	410
<i>Xiaotie Deng, Qi Qi, and Jie Zhang</i>	

## Short Papers

Continuity Properties of Equilibria in Some Fisher and Arrow-Debreu Market Models . . . . .	422
<i>Vijay V. Vazirani and Lei Wang</i>	
Route Distribution Incentives . . . . .	430
<i>Joud Houry, Chaouki T. Abdallah, Kate Krause, and Jorge Crichigno</i>	
Wiretapping a Hidden Network . . . . .	438
<i>Haris Aziz, Oded Lachish, Mike Paterson, and Rahul Savani</i>	

Refining the Cost of Cheap Labor in Set System Auctions . . . . .	447
<i>Ning Chen, Edith Elkind, and Nick Gravin</i>	
On the Power of Mediators . . . . .	455
<i>Josep Díaz, Dieter Mitsche, Navin Rustagi, and Jared Saia</i>	
Strong Nash Equilibria in Games with the Lexicographical Improvement Property . . . . .	463
<i>Tobias Harks, Max Klimm, and Rolf H. Möhring</i>	
A Note on Strictly Competitive Games . . . . .	471
<i>Ilan Adler, Constantinos Daskalakis, and Christos H. Papadimitriou</i>	
The Efficiency of Fair Division . . . . .	475
<i>Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou</i>	
Sequential Bidding in the Bailey-Cavallo Mechanism . . . . .	483
<i>Krzysztof R. Apt and Evangelos Markakis</i>	
Subsidized Prediction Markets for Risk Averse Traders . . . . .	491
<i>Stanko Dimitrov, Rahul Sami, and Marina Epeľman</i>	
Envy, Multi Envy, and Revenue Maximization . . . . .	498
<i>Amos Fiat and Amiram Wingarten</i>	
Nudging Mechanisms for Technology Adoption . . . . .	505
<i>Jacomo Corbo and Yevgeniy Vorobeychik</i>	
Mechanism Design for Complexity-Constrained Bidders . . . . .	513
<i>Ravi Kumar, Mohammad Mahdian, and Amin Sayedi</i>	
Priority Right Auction for Komi Setting (Extended Abstract) . . . . .	521
<i>Xiaotie Deng and Qi Qi</i>	
Randomized Online Algorithms for the Buyback Problem . . . . .	529
<i>B.V. Ashwinkumar and Robert Kleinberg</i>	
Envy-Free Allocations for Budgeted Bidders . . . . .	537
<i>David Kempe, Ahuva Mu'alem, and Mahyar Salek</i>	
On the Inefficiency Ratio of Stable Equilibria in Congestion Games . . . . .	545
<i>Arash Asadpour and Amin Saberi</i>	
Betting on the Real Line . . . . .	553
<i>Xi Gao, Yiling Chen, and David M. Pennock</i>	
Characterization of Strategy-Proof, Revenue Monotone Combinatorial Auction Mechanisms and Connection with False-Name-Proofness . . . . .	561
<i>Taiki Todo, Atsushi Iwasaki, and Makoto Yokoo</i>	

2D-TUCKER Is PPAD-Complete . . . . .	569
<i>Dömötör Pálvölgyi</i>	
Bidder Optimal Assignments for General Utilities . . . . .	575
<i>Paul Dütting, Monika Henzinger, and Ingmar Weber</i>	
A New Ranking Scheme of the GSP Mechanism with Markovian Users . . . . .	583
<i>Xiaotie Deng and Jiajin Yu</i>	
Mediated Equilibria in Load-Balancing Games . . . . .	591
<i>Joshua R. Davis, David Liben-Nowell, Alexa Sharp, and Tom Wexler</i>	
On the Impact of Strategy and Utility Structures on Congestion-Averse Games . . . . .	600
<i>Thomas Voice, Maria Polukarov, Andrew Byde, and Nicholas R. Jennings</i>	
On Strong Equilibria in the Max Cut Game . . . . .	608
<i>Laurent Gourvès and Jérôme Monnot</i>	
Stability and Convergence in Selfish Scheduling with Altruistic Agents . . . . .	616
<i>Martin Hoefer and Alexander Skopalik</i>	
Gaming Dynamic Parimutuel Markets . . . . .	623
<i>Qianya Lin and Yiling Chen</i>	
News Posting by Strategic Users in a Social Network . . . . .	632
<i>Mangesh Gupte, MohammadTaghi Hajiaghayi, Lu Han, Liviu Iftode, Pravin Shankar, and Raluca M. Ursu</i>	

## Erratum

Prediction Mechanisms That Do Not Incentivize Undesirable Actions . . .	E1
<i>Peng Shi, Vincent Conitzer, and Mingyu Guo</i>	
<b>Author Index</b> . . . . .	641



# Ad Exchanges: Research Issues

S. Muthukrishnan

Google Inc.

**Abstract.** An emerging way to sell and buy display ads on the Internet is via ad exchanges. RightMedia [1], AdECN [2] and DoubleClick Ad Exchange [3] are examples of such real-time two-sided markets. We describe an abstraction of this market. Based on that abstraction, we present several research directions and discuss some insights.


## 1 Introduction

Activities on the Internet can be abstracted as interactions due to three parties. There are

- the *users* who navigate to various web pages,
- the *publishers* who control the web pages and generate the content in them, and
- are *advertisers* who wish to get the attention of the users using the publishers as the channel for placing ads on the pages.

Precisely what ads show when a user accesses a page is a detailed process. A central issue is matching advertisers to publishers, where the number of advertisers and publishers is very large. Direct negotiation between advertisers and publishers might work for large companies, but currently there are *intermediaries* such as ad agencies, ad networks and publisher networks that aggregate several parties. An advertisers might use an ad agency to develop the marketing campaign and use an ad network that pools many advertisers to negotiate with a publisher network that places ads with many publishers; alternatively, for a large advertiser, the ad agency might directly negotiate contracts with a few large publishers. There are several intermediaries and their services overlap with networks comprising both advertisers and publishers, or agencies that are also ad networks, etc. Thus there are several ad paths between an advertiser and a user; some ad paths are shown in Figure 1.

An emerging way of selling and buying ads on the Internet is via an *exchange* that brings sellers (publishers) and buyers (advertisers) together to a common marketplace. There are exchanges in the world for trading financial securities to currency, physical goods, virtual credits, and much more. Exchanges serve many purposes from bringing efficiency, to eliciting prices, generating capital, aggregating information etc. Market Microstructure is the area that studies all aspects of such exchanges [4,5]. Ad exchanges are recent. Ad exchanges offer ad networks and publishers to transact centrally for ads. RightMedia [1], AdECN [2] and DoubleClick [3] are examples.



**Fig. 1.** Ad paths

- Publishers expect to get the best price from the exchange, better than from any specific ad network; in addition, publishers get liquidity.
- Advertisers get access to a large inventory at the exchange, and in addition, the ability to target more precisely across web pages.
- Finally, the exchange is a clearing house ensuring the flow of money.

In many ways, these ad exchanges are modeled after financial stock exchanges. Since 2005 when RightMedia appeared, ad exchanges have become popular. In Sept 2009, RightMedia averaged 9 billion transactions a day with 100's of thousands of buyers and sellers. Recently, DoubleClick announced their new ad exchange. It seems ad exchanges are likely to become a major platform for trading ads.

We abstract a model for ad exchanges. Based on the model, we present research problems in auction theory, optimization and game theory. The goal is to present a blueprint for research in design, analyses and use of ad exchanges.

## 2 Ad Exchange

We present an abstract *AdX* model to describe ad exchanges. It is defined as a sequence of events.

1. User  $u$  visits the webpage  $w$  of publisher  $p(w)$ .  
For now, we assume page  $w$  has a single slot for ads.
2. Publisher  $p(w)$  contacts the exchange  $E$  with  $(w, u, \rho)$  where  $\rho$  is the *minimum price*  $p(w)$  is willing to take for the slot in  $w$ .

We assume that  $E$  knows all the contents of  $w$  as well as the various specifics of the ad slot in it, including its dimensions and inappropriate ads for that slot as agreed on with  $p(w)$ . We also assume that the exchange manages information about user  $u$  in a manner agreed upon with  $p(w)$ . This is reasonable because  $p(w)$  has moral and contractual relationship with its viewers  $u$  while also having incentive to help advertisers target users suitably. Also,  $E$  can independently crawl contents of  $w$  if needed, in many cases, we might as well assume that in the model.

3. The exchange  $E$  contacts ad networks  $a_1, \dots, a_m$  with  $(E(w), E(u), \rho)$ , where  $E(w)$  is information about  $w$  provided by  $E$ , and likewise,  $E(u)$  is the information about  $u$  provided by  $E$ .

$E(u)$  is the information  $E$  is able to provide about  $u$  to the ad networks as agreed upon with  $p(w)$ . When  $p(w)$  entrusts  $E$  to reveal  $w$ , we assume that ad networks know contents of  $w$  and also derived information such as topics

of contents in  $w$ . The ad networks may have the resources to obtain this on their own or resourceful exchanges can do it for them. This is modeled as  $E(w)$ . There are instances when  $p(w)$  does not wish ad networks to know  $w$ , in which case,  $E(w)$  will be only derived information about  $w$  and  $w$ 's identity will remain unknown to  $a_i$ 's.

4. Each ad network  $a_i$  returns  $(b_i, d_i)$  on behalf of its customers which are the advertisers;  $b_i$  is its *bid*, that is, the maximum it is willing to pay for the slot in page  $w$  and  $d_i$  is the ad it wishes to be shown. The ad networks may also choose not to return a bid.

It is assumed that  $b_i \geq \rho$ , else no bid is returned. It is also assumed that  $d_i$  is suitable for the ad slot. Further, it is assumed that  $a_i$  targets ads based on its contracts with its customers and negotiates prices for the service with them. Ad  $d_i$  may be passed by reference as supported by common http protocol.

5. Exchange  $E$  determines a winner  $i^*$  for the ad slot among all  $(b_i, d_i)$ 's and its price  $c_{i^*} \leq b_{i^*}$  via an *auction* and returns  $(c_{i^*}, d_{i^*})$  to  $p(w)$ .

It is assumed that the winning network  $i^*$  becomes aware of the outcome including the price, and in some but not necessarily all instances, the losing networks can determine that too.  $E$  is responsible that  $d_i$  is suitable for  $w, u$  according to its contract with  $p(w)$ . This may be accomplished in a variety of ways from pre-verification to outsourcing the task to the networks.  $E$  negotiates pricing for its service with various  $p(w)$ 's and  $a_i$ 's. Also,  $E$  generates bills and reports, collects payments from  $a_i$ 's, and makes payments to  $p(w)$ 's.


6. The publisher  $p(w)$  serves webpage  $w$  with ad  $d_{i^*}$  to user  $u$ . This is known as an *impression* of ad  $d_{i^*}$ .

The ad  $d_{i^*}$  is rendered by user's browser and the user interacts with the ad in a variety of ways.

We have left out system level details in the model above. For example, internet protocols are used for forwarding, referencing, accounting, etc. Further, for efficiency reasons, steps may be optimized out and or made offline. For example, ads and bids may be loaded onto the Exchange *a priori* so networks are accessed only infrequently. Also, due to time constraints, crawling and processing  $w$  may be overlooked and  $E(w)$  will be minimal or there may be timeouts, etc. The flow of the model is shown in Figure 2.

The AdX model captures the essence of ad exchanges that the instrument traded is *ad impressions*. So, AdX model has real time auction for each impression and the bids are CPM (cost per mille or thousand impressions). The entire execution above begins when a user arrives at a webpage and must be completed before the page is rendered on their screen. This sets an upper bound for the execution in 10's or 100's of milliseconds.<sup>1</sup> Thus AdX model provides a spot market for ad impressions. The assumption is that higher order campaign goals such as maximize number of impressions subject to budget criteria, or reach as many target users as possible, and other goals will be executed by optimized bidding

<sup>1</sup> AdECN claims a 12 milliseconds execution 2.



**Fig. 2.** AdX Model

by the networks, and indeed one can design more sophisticated instruments on top of the spot market.

**AdX vs Financial Exchanges.** As in financial exchanges, the buyers and sellers come to AdX only when they choose. Further, individual advertisers access AdX through intermediaries like in financial exchanges. However, a significant difference is in the nature of what is traded. Unlike financial securities, ad impressions are heterogeneous, significantly differing from one instance to another in their value depending on their impact on the user which varies widely across millions of users. Further, impressions are supremely perishable: if the trade does not happen within the browser performance anticipated by the user, the opportunity to place an ad is lost. As a result, AdX has two important tasks. First is informational. AdX embodies the effort by various parties to help advertisers discover users to target, eg via  $E(w), E(u)$ . Second is economics. With vastly heterogeneous goods, many of the pricing methods face challenges and AdX enables the market to discover prices by making trading automatic and via auctions.



**Fig. 3.** Sponsored Search

**AdX vs Sponsored Search.** The AdX model is distinct from *sponsored search*. In sponsored search, a user poses a query at a search engine and gets search results together with ads arranged top to bottom. The assignment of ads to positions is by an auction among all advertisers who placed a cost-per-click (CPC) bid on a keyword that matches the query. If the user clicks on an ad, that advertiser pays the search engine the auction price. The ad path is simple and shown in Figure 3. The dynamics are simpler since there is a single publisher and a one-sided marketplace of buyers. On the other hand, sponsored search

aligns the incentives of advertisers and search engines with the quality of ads for the users, and hence, the publisher faces the challenge of monitoring and maintaining quality. The underlying auction for sponsored search is analyzed in [8,9]. There are many outstanding research issues in sponsored search [6], but they are not the focus here. We will henceforth focus on the AdX model.

### 3 Research Issues

We describe several research issues using examples. The issues when formalized more precisely will lead to different research problems.

#### 3.1 Basic Auction at the Exchange

A central question concerns the auction at the exchange. There is a single slot of ad being auctioned, and the bidders know  $\rho$  and participate only if their bid is above  $\rho$ . Further it is a sealed bid auction since each network bids directly with the exchange. The setting appears to be a standard auction in one-sided market. Revenue equivalence theorem [7] would then indicate that first price, second price and several other auctions will all yield identical expected revenue in Bayesian equilibrium under certain natural assumptions. Still, second price or Vickrey auctions are preferable since they encourage more stable dynamics. So, we will assume second price auction as a benchmark when needed. We now discuss a nuance.

**Example.** Say network  $a_1$  has 2 advertisers  $a_{11}$  and  $a_{12}$  with bids 10 and 8 resp., while network  $a_2$  has a single advertiser  $a_{21}$  with bid 5.  $a_1$  forwards 10 and  $a_2$  forwards 5 to the exchange. Second price auction at the exchange will declare  $a_1$  the winner with price 5.  $a_{11}$  wins the ad slot. Now  $a_1$  may charge  $a_{11}$  either 5 (the exchange price) or 8 (the second price in  $a_1$ ). Instead, if  $a_1$  revealed both bids 10 and 8 to the exchange, second price auction at the exchange would charge 8 to the winner  $a_{11}$ . Thus since the exchange does not know the entire book of the ad networks, a network has an excess to disburse and the publisher gets less than the best price from the book. ■

*Book value* is the second largest value of all the bidders.

**Problem 1.** *Assuming  $\rho$  is exogenous and assuming the advertisers reveal their bids truthfully to the networks, is there a possibly truthful auction at the exchange that will extract a large fraction of the book value?*

The tricky case is when the bid that sets the price of the winner is in the winner's network. A straightforward approach is to generalize the protocol and ask each network  $i$  to forward *two* bids  $(b_1^i, b_2^i)$  where  $b_1^i$  is the largest of the bids in  $i$  and  $b_2^i$  the second largest. A naive strategy is to declare  $i^* = \operatorname{argmax}_i b_1^i$  the winner with price  $\max\{b_2^{i^*}, \max_{j \neq i^*} b_1^j\}$ ; then, ad networks have no incentive to declare nonzero  $b_2^i$ . A simple allocation strategy would be to declare largest  $b_1^i + b_2^i$  as the winner. This incentivizes networks to bid  $((b_1^i + b_2^i)/2, (b_1^i + b_2^i)/2)$  and is not truthful.

### 3.2 Auction and Bidding by Ad Networks

Consider mechanisms for the ad networks. In general, the exchange can not assume ad networks will follow any particular mechanism. Still, exchange's choice of auction will impact mechanisms ad networks will use, and choice of mechanisms of ad networks will influence auction at the exchange.

**Problem 2.** *Assuming  $\rho$  is exogenous, the exchange runs a second price auction with reserve price  $r \geq \rho$ , and advertisers are captive, that is, remain with their choice of ad network throughout, what is a revenue optimal mechanism for an ad network?*

Seen from an ad network  $i$ 's point of view, there are several advertisers bidding to be chosen for the impression. The questions are, which ad to choose and what price to charge? This appears to be identical to the standard auction framework where optimal mechanisms are known [11]. Under suitable assumptions, the solution is to run a second price auction with an appropriately chosen reserve price  $r_i$ . The key difference here is that the ad network  $i$  does not have a guaranteed good to auction, rather it has a *contingent* good. More formally, there is some probability  $\alpha(b)$  that it will win the impression if it bid  $b$  at the exchange. This probability function  $\alpha$  is determined by the bidding strategy of other networks and their advertisers' valuations.

Under standard assumptions that advertisers have values drawn from a known distribution, one can extend the theory of optimal auctions [11] to derive an optimal mechanism for the ad networks [12]. For example, one can show that the revenue-optimal mechanism for a network to randomize its bids in some range  $[l_r, u_r]$ , dependent on  $r$ . A more detailed understanding of the equilibrium will be of great interest. Even without the complication of what auction to run in the network, the problem of bidding for higher campaign goals is challenging in presence of  $\alpha(\cdot)$ . Recently, this was studied in [13].

### 3.3 Auctioning with Heterogenous Valuations

While the exchange auctions impressions, there are instances when an impression is valued very differently by the bidders. For example, an impression to a particular user may be far more valuable to one bidder than others. This targeting is enabled by the information  $E(u)$  shared by the publisher  $p(w)$  with the ad networks [2].

**Example.** We have two bidders  $A$  and  $B$ ,  $A$  has value  $v_A = 100$  and  $B$  has value  $v_B = 1$  for the ad slot. In equilibrium of the first price auction,  $A$  wins by bidding  $1 + \epsilon$  for some  $\epsilon > 0$  and exchange's revenue is  $1 + \epsilon$ . This outcome holds for second price auction as well. On the other hand, maximum value to be extracted is  $\max v_A, v_B = 100$ . What is the maximum revenue a mechanism can extract in equilibrium? ■

<sup>2</sup> [http://www.adecn.com/faq\\_4.html](http://www.adecn.com/faq_4.html)

Say bidder  $i$  has value  $v_i$ , bids  $b_i$ , and there is a single item to be sold. Then, the maximum revenue is  $R^* = \max_i v_i$ . Classical results [11] would provide most revenue assuming the distribution from which  $v_i$ 's are drawn is known. Here, instead, we focus on *prior-free* case where distributions are not known *a priori*, not even the maximum of individual distributions.<sup>3</sup> We know from [14] that no truthful mechanism exists with (roughly) expected revenue  $\Omega(\frac{R^*}{\log R^*})$ ; on the other hand, they show an auction with expected revenue  $\Omega(\frac{R^*}{(\log R^*)^{1+\epsilon}})$  for fixed  $\epsilon > 0$ . Their auction [14] is as follows. With probability  $1 - \delta$ , use the second price auction. With probability  $\delta$ , choose  $r$  according to the distribution below and if  $b_1 \geq r$ , highest bidder wins at price  $r$ :

$$f(x) = \frac{\epsilon}{x(\log(x/b_2) + 1)^{1+\epsilon}}, \quad x \in [b_2, \infty).$$

**Problem 3.** *Design a non-truthful mechanism for prior-free auction of a single slot with near-optimal revenue, but with good equilibrium properties.*

An approach is to use quasi-proportional allocation in which the good is allocated to one bidder using a distribution where  $i$ th bidder is picked with probability  $\frac{f(b_i)}{\sum_i f(b_i)}$ , for some suitable function  $f$ . If  $f(x) = x$ , we have the well-known proportional allocation. Revenue properties of quasi-proportional mechanisms are not well-studied. In [17], authors show that with  $f(x) = \sqrt{x}$ , one can extract  $(R^*)^\gamma$  revenue for some  $\gamma > 0$  under certain conditions. Further, using [18], the equilibrium can be characterized for quasi-proportional auctions. A deeper study of quasi-proportional auctions will be of great interest. In particular, are there functions  $f$  for which quasi-proportional allocation generates more revenue in its equilibrium than the lower bound in [14]?

For AdX, we need a generalization of the standard prior-free setting to ad networks as well as the exchange (much as [12] extends [11]). For an economics perspective on market clearing and role of intermediaries in prior-free auctions, see [15].

**Problem 4.** *Design (even non-truthful) mechanisms for prior-free bidding of ad networks in AdX, with good equilibrium properties and (near-)optimal revenue.*

### 3.4 Callout Optimization

In Step 3,  $E$  seeks bids from ad network  $a_i$ 's. This may be accomplished at a system level by (a) having  $a_i$ 's preload their ad campaigns into  $E$  so suitable ad campaigns and their bids that apply to  $(E(w), E(u), \rho)$  can be retrieved locally, (b) hosting  $a_i$ 's bidding software in  $E$ 's machines so the software can manage

<sup>3</sup> In practice, one would argue that in repeated auctions as in the exchange, one can learn the distributions. It is an interesting exercise what can be learned from the data – not necessarily the entire distribution, just enough to extract revenue – and how, as well as its impact in engineering the system, in particular, when learning is likely to be only approximate.

the network's campaigns and generate suitable  $(b_i, d_i)$ , or (c) making http calls to  $a_i$ 's servers, and awaiting  $(b_i, d_i)$  to be determined by the network. (a) and (b) are essentially not real-time because networks cannot update bids, bidding logic, key parameters and relevant data impression to impression. We focus on (c) and discuss an optimization problem.

It is resource-intensive for  $E$  to call out to each network for each impression. In practice, ad networks can describe  $(E(w), E(u), \rho)$ 's of cumulative interest to their customers, and therefore for each impression, only a subset  $S_{(E(w), E(u), \rho)}$  of ad networks need to be called. Still, we may assume that it will be difficult if not impossible for each network to take all the http calls from  $E$ . So, here is an optimization problem  $E$  faces.

**Problem 5.** *Each ad network  $i$  has bandwidth budget  $B_i$ . Say  $E$  has bandwidth budget of  $B$ . Design an online algorithm for  $E$  that for each incoming call  $(w_j, u_j, \rho_j)$ , chooses a subset  $S_j \subseteq S_{(E(w_j), E(u_j), \rho_j)}$  of networks to call such that no ad network  $i$  gets more than  $B_i$  calls per second,  $E$  make fewer than  $B$  calls per second, and optimizes the expected*

- number of bids, ie, number of nonempty  $(b_i(j), d_i(j))$ 's received at  $E$ , or
- efficiency  $\sum_j \max_i b_i(j)$ , or
- sales revenue  $\sum_j \max_i |b_i(j) \neq \max_i b_i(j)| b_i(j)$ , or
- profit for  $E$ .

In the problem above, we need stochastics. The algorithm has to know about likelihood of ad network  $i$  making a bid for  $(E(w_j), E(u_j), \rho_j)$ . In particular,  $\Pr(b_i(j) \geq \rho_j | E(w_j), E(u_j))$  and  $Exp(b_i(j) | b_i(j) \geq \rho_j)$  are useful. Hence,  $E$  needs to estimate these terms for each  $j$ . Given these parameters, initial approximation results for the problem are in [19].

### 3.5 Publisher Optimization and Strategies

Publisher  $p(w)$  has many decisions to make when user  $u$  visits  $w$ .

*Accessing AdX.*  $p(w)$  has several sources for ads including contracts with advertisers signed by their own sales teams, contracts with specific ad networks and the ability to reach AdX on any impression. Filling many impressions with highest price ads from AdX will have long term impact on other sources, and failure to deliver on contracts; filling contract commitments might lead to loss of revenue from AdX spot market. The task is to design an online algorithm to commit impressions to different ad sources in order to honor contracts as well as maximize revenue. Some initial results are in [21] where solutions are based on the publisher virtually bidding on behalf of ad sources. More detailed understanding of strategies will be of great use to individual publishers.

*Form of inventory.* If publisher  $p(w)$  wishes,  $E$  can limit information via  $E(w)$ , not reveal identity of  $w$  or  $p(w)$ , and merely provide information about the



nature of the site eg., sports/baseball, etc. We call this *undisclosed* inventory<sup>4</sup>; in contrast, when identity of  $w$  is shared with the ad networks, we call it *disclosed* inventory. Publishers may choose to keep some of their inventory undisclosed, in order to avoid conflicts with their other sales channels. It is assumed that undisclosed inventory fetches less than disclosed inventory. Then,  $p(w)$  has an online optimization problem: for each impression, how to choose disclosed or undisclosed option in  $E$  to trade off short term vs long term revenue. A suitable model to address this problem will be of interest.

*Price.* Publisher  $p(w)$  has to choose  $\rho$  while accessing  $E$ : large  $\rho$  may not generate a bid, and a small  $\rho$  may undervalue the inventory. One way to choose  $\rho$  is as the maximum over other sales channels for an ad at the slot, but in general, one needs an online  $\rho$  setting algorithm that endogenizes the demand and supply in the system adaptively. Again, a clean model and implementable algorithm is of interest.

Here is the combined problem of inventory management and accessing AdX.

**Problem 6.** *Given models for impressions inventory  $(w, u)$ , models for bids  $(b_{i^*}, d_{i^*})$  from  $E$ , models of ad sales and prices through other channels, design an algorithm that on each impression (a) decides whether to go to AdX, (b) chooses disclosed or undisclosed inventory at AdX, and (c) selects min price  $\rho$ , in order to optimize the expected overall (long term) revenue.*

### 3.6 Arbitrage Bidding and Risk Analysis

AdX model trades impressions and uses CPM prices. Ad networks can sell other pricing methods, such as pay-per-click with cost-per-click (CPC) prices to their advertisers. This comes with an arbitrage opportunity and associated risk as described below.

**Example.** Define click-through-rate (CTR) as average fractional number of clicks per impression. Consider an ad with CTR 0.1 and CPC of \$1. Ideally, the network should bid \$100 CPM; then spend is \$100 for 1000 impressions, revenue is \$100, and they are even. However, in any empirical run of auctions, CTR estimates are not precise. (1) Say CTR was overestimated to be 0.2. Then the network bids \$200 CPM, spend is \$200 for 1000 impressions, and revenue is \$100. This assumes the network will get at least 1000 impressions if they overbid, which is reasonable given large inventory in ad slots. (2) Say CTR was underestimated to be 0.05 and the network bids \$50 CPM. Then spend is \$50 $x$  for  $x$  fraction of 1000 impressions, and revenue is \$100 $x$ . The assumption is the network will only get  $x$  fraction of 1000 impressions. To summarize, the outcome is  $(200 - 100)$  loss if network overbid vs  $(100 - 50)x$  profit if network underbid and loss of  $(1 - x)$  fraction of impressions. ■

<sup>4</sup> See blind vs disclosed inventory in [http://www.adecn.com/faq\\_5.html](http://www.adecn.com/faq_5.html), or branded vs anonymous in [http://www.doubleclick.com/products/advertisingexchange/benefits\\_for\\_sellers.aspx](http://www.doubleclick.com/products/advertisingexchange/benefits_for_sellers.aspx)

**Problem 7.** *Consider advertisers who contract with the network for CPC bids, certain reach (number of distinct users reached) and frequency counts (the number of times a user sees an ad). Design an algorithm for ad networks to place CPM bids into AdX that for a given risk level and volatility, maximizes expected revenue and guarantees contract counts. Take into account bidding into the exchange for contingent good as in Section [3.2](#).*

### 3.7 AdX Integrity

Exchanges in general strive for transparency. The ability of participants to understand the inventory they buy or be convinced of the integrity of underlying process will, in general, induce more to access the exchange. In reality, systems need a balance of overall trust and key parts with legal or technological checks if needed. While this balance will be worked out by the marketplace, here is a theoretical problem.

**Problem 8.** *Design a cryptographically sound real-time auction protocol so that any participating party in AdX can verify that (a) all communication, accounting and computations were performed correctly, and (b) auction was closed envelope, that is, no bidder sees others' bids prior to the auction. This has to work for repeated auction of impressions in AdX where some information is revealed between impressions.*

Secure and collusion free auction design has been explored by the cryptography community, eg using homomorphic encryption scheme [\[20\]](#). However, one still needs realtime methods for AdX application, methods which will run with one http round trip time. Also, one needs a clear model to prove strong cryptographic properties of the various protocols. Some progress is in [\[22\]](#).

### 3.8 Configuration Auctions

Even when there is a single slot in  $w$ , different configurations of ads — eg 1 video ad or 2 image ads or 4 text ads — might fit there. While the exchange will return the most efficient configuration in total CPM bids, an interesting question is how to spread the price among the ads in the winning configuration (Problem 6 in [\[6\]](#)). Further, inherently, the bids from the networks are for entire configurations which makes it more like configuration auctions and associated externalities [\[23\]](#).

We go beyond to look at cases where  $w$  has multiple ad slots. How should they be sold? One approach is to sell the slots independently. This is easy on the system but advertisers may wish to appear in only one slot on a page which can not be guaranteed. Another approach is to sell all the slots in one block, but such exclusive packages may not be popular except on premium sites. A realistic approach is to make AdX more sophisticated and consider all slots together in some auction. In reality, advertisers' have complex preferences. For example, an advertiser may value a slot on the right highly if they did not win the top slot. That is, advertisers have *conditional values*.

**Problem 9.** *Devise a suitable bidding language for advertisers and ad networks to express (conditional) preferences for slots, and design suitable auction mechanisms for the exchange as well as for ad networks to allocate substitutable winners if any.*

We propose a simple approach. Publisher  $p(w)$  presents a list of slots in  $w$  in the order in which will be auctioned. Then ad networks can return one bid for each slot, and specify the maximum number of impressions they wish. The auction proceeds slot by slot; winners are removed from consideration for successive slots if they are fulfilled. We call this *linearized auction* (LA). LA trades off efficiency for expressiveness and lets bidders represent some conditional preferences but not others. Studying suitable LA mechanisms including the impact of publisher's choice of the list will be of interest. Also, comparison of LA mechanisms with richer tree bidding [24] will be of interest.

## 4 Concluding Remarks

This writeup provides some insight into the research issues in AdX that models exchanges like RightMedia, AdGCN and DoubleClick. These exchanges are emerging as major new platforms to sell display ads on the Internet, but are still nascent. Progress on research issues listed here will likely impact the design and growth of not only these ad exchanges but also the “ecosystem” of bidders, optimizers and quantifiers around them. We list some more issues:

1. *Game theory of advertisers.* Advertisers may go to multiple networks, or choose networks strategically. What are the resulting dynamics? How does their strategic behavior affect competition within their campaigns along multiple ad paths, across advertisers, across ad networks and ultimately, across exchanges?
2. *Ad Quality.* In sponsored search, quality of an ad is correlated with click-through, and so is the pricing and incentives of the advertiser. We need a similar quality metric for impressions and endogenize that in the auction to align advertiser incentives. A proposal is to generate a suitable Markov model for users that will capture even the long term impact of ad impressions.

*Acknowledgements.* I sincerely thank my colleagues Rahul Bafna, Eyal Manor, Yishay Mansour, Noam Nisan and Scott Spencer at Google for weekly AdX lessons. Thanks to Mallesh Pai, David Parkes, Al Roth, Hal Varian and Ricky Vohra for pointers, and Sudipto Guha, Aranyak Mehta and Mukund Sundarajan for comments. Thanks to Michael Rabin and Moti Yung for an enjoyable collaboration.

## References

1. RightMedia, <http://www.rightmedia.com/> More info at, <http://www.rightmedia.com/right-media-101/>
2. AdECN. <http://www.adecn.com/>. Whitepaper at, [http://www.adecn.com/resources/ A True Exchange for Online Advertising.pdf](http://www.adecn.com/resources/A%20True%20Exchange%20for%20Online%20Advertising.pdf). Presentation at, <http://www.adecn.com/resources/AdECNPresentation.pps>

3. DoubleClick Ad Exchange,  
<http://www.doubleclick.com/products/advertisingexchange/index.aspx>
4. O'Hara, M.: *Market Microstructure Theory*. Blackwell, Oxford (1995)
5. Madhavan, A.: *Market Microstructure: A Survey*. *Journal of Financial Markets* 3(3), 205–258 (2000)
6. Muthukrishnan, S.: *Internet Ad Auctions: Insights and Directions*. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) *ICALP 2008, Part I. LNCS*, vol. 5125, pp. 14–23. Springer, Heidelberg (2008)
7. Klemperer, P.: *Auction Theory: A Guide to the Literature*. *J. of Economic Surveys*, 227–286 (1999)
8. Varian, H.: *Position Auctions*. *Intl. J. of Industrial Organization* (2006)
9. 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)
10. Vickrey, W.: *Counterspeculation, Auctions, and Competitive Sealed Tenders*. *Journal of Finance* 16, 8–37 (1961)
11. Myerson, R.: *Optimal auction design*. *Math of Operations Research* 6, 58–73 (1981)
12. Feldman, J., Mirrokni, V., Muthukrishnan, S., Pai, M.: *Charaterizing Equilibria of Auctions with Mediators* (manuscript, 2009)
13. Ghosh, A., Rubinstein, B., Vassilvitskii, S., Zinkevich, M.: *Adaptive bidding for display advertising*. In: *Proc. WWW* (2009)
14. Lu, P., Teng, S., Yu, C.: *Truthful Auctions with Optimal Profit*. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) *WINE 2006. LNCS*, vol. 4286, pp. 27–36. Springer, Heidelberg (2006)
15. Baliga, S., Vohra, R.: *Market Research and Market Design* (2003),  
<http://www.kellogg.northwestern.edu/faculty/baliga/htm/mrandmd.pdf>
16. Aggarwal, G., Goel, G., Mehta, A.: *Efficiency of (revenue-)optimal mechanisms*. In: *ACM Conference on Electronic Commerce (EC)*, pp. 235–242 (2009)
17. Mirrokni, V., Muthukrishnan, S., Nadav, U.: *Quasi-Proportional Mechanisms: Prior-free Revenue Maximization*. *arXivs* (2009),  
<http://arxiv1.library.cornell.edu/pdf/0909.5365v1>
18. Even Dar, E., Mansour, Y., Nadav, U.: *Convergence in Proportional Games*. In: *Proc. STOC* (2009)
19. Chakraborty, T., Even-Dar, E., Guha, S., Mansour, Y., Muthukrishnan, S.: *Callout optimization* (manuscript, 2009)
20. Parkes, D., Rabin, M., Shieber, S., Thorpe, C.: *Practical secrecy-preserving, verifiably correct and trustworthy auctions*. In: *Proc. EC*, pp. 70–81 (2006)
21. Ghosh, A., McAfee, P., Papineni, K., Vassilvitskii, S.: *Bidding for Representative Allocations for Display Advertising*. In: Leonardi, S. (ed.) *WINE 2009. LNCS*, vol. 5929, pp. 208–219. Springer, Heidelberg (2009)
22. Rabin, M., Muthukrishnan, S., Yung, M.: *Fast, communication efficient, provably secure protocols for real-time Vickrey auctions* (manuscript, 2009)
23. Muthukrishnan, S.: *Bidding on Configurations in Internet Ad Auctions*. In: Ngo, H.Q. (ed.) *COCOON 2009. LNCS*, vol. 5609, pp. 1–6. Springer, Heidelberg (2009)
24. Lahaie, S., Parkes, D., Pennock, D.: *An Expressive Auction Design for Online Display Advertising*. In: *AAAI*, pp. 108–113 (2008)

# Adaptive Learning in Systems of Interacting Agents

H. Peyton Young

University of Oxford

A learning rule is *adaptive* if it is simple to compute, requires little information about the actions of others, and is plausible as a model of behavior [1, 2]. In this paper I survey a family of adaptive learning rules in which experimentation plays a key role. These rules have the property that, in large classes of games, agents' individual behavior results in Nash equilibrium behavior by the group a high proportion of the time. Agents need not *know* that Nash equilibrium is being played -- indeed they need not know anything about the structure of the game in which they are embedded. Instead, equilibrium evolves as an unintended consequence of individual adaptation. The theory is particularly relevant to modeling systems of interacting agents that are very large and complex, so that one cannot reasonably expect that players would try to optimize based on their beliefs about the state of the system. Concrete examples include drivers adjusting to urban traffic patterns, or people sending and receiving information in large networks. While such rules can be viewed as a descriptive model of how humans adapt in such situations, they can also be taken as design elements in engineered systems, such as distributed sensors or robots, where the 'agents' are programmed to behave in a way that leads to desirable system-wide outcomes.

In its simplest form, *experimentation* means that an agent occasionally tries out new strategies (drawn according to a probability distribution over the set of feasible strategies), and keeps the new strategy if and only if it results in a higher payoff than his current strategy. It can be shown that in potential games, and more generally in weakly acyclic games, this leads to system-wide Nash equilibrium behavior in the following approximate sense: a Nash equilibrium will be played a large proportion of the time when the experimentation rate is sufficiently small [3].

Note that this is weaker than saying that the system will arrive at a Nash equilibrium after some finite time and stay there; it is also weaker than saying that it is *probable* that the system will arrive at a Nash equilibrium and stay there. Rather, it says that if  $p(t)$  is the proportion of all time periods  $t' \leq t$  at which a Nash equilibrium occurs, then for some suitably small  $\varepsilon > 0$ ,  $\liminf_{t \rightarrow \infty} p(t) \geq 1 - \varepsilon$  almost surely.

In more general games, the simple "hill-climbing" form of experimentation described above does not generally lead to Nash equilibrium. However there is a natural variant, called *interactive trial and error learning*, that works even when there is no potential function [4]. In this procedure, an agent's state at any given time  $t$  is described by three variables: a benchmark action  $a_t$ , a benchmark utility or aspiration level  $u_t$ , and an internal state or "mood"  $m_t$ . There are two principal moods: content and discontent. A *content* agent searches occasionally and deliberately, adopting a new action only if it is strictly better than his current strategy just as in the simple

experimentation model described earlier. If the experiment is successful he adjusts his benchmark action and utility to the new values. A *discontent* agent flails around, trying out new actions at random. However, after each search there is a positive probability that he spontaneously stops searching and adopts the most recent action and payoff as his new benchmarks. We assume that the stopping probability increases monotonically with the current payoff, and that it is bounded away from zero and one in all states.

The transition between moods is triggered by the pattern of recent payoffs. If a content agent experiences an increase in payoff *without experimenting*, he enters a ‘hopeful’ state for one period; if his payoff continues to be higher than his current benchmark for another period he becomes content again with the higher payoff as his new benchmark. If, however, a content agent experiences a *decrease* in payoff through no fault of his own (i.e., without experimenting), he enters a transitional ‘watchful’ state for one period; if his realized payoff continues to be lower than his benchmark for another period he becomes discontent, and flailing around ensues.

Conceptually this approach is similar to adaptive procedures that have been proposed in other branches of the learning literature. In computer science, for example, the algorithm WoLF (Win or Lose Fast) has a similar flavor: stay with a strategy that is doing well and change quickly if it does poorly [5]. In biology the foraging behavior of bees has been shown to have a broadly similar character [6].

It can be shown that the version of the rule described above leads to Nash equilibrium in a wide variety of games. Specifically, consider an  $n$ -person, finite, normal-form game  $G$  with generic payoffs and at least one pure Nash equilibrium. *Given any small  $\varepsilon > 0$ , if all players use interactive trial and error learning with sufficiently small experimentation probability  $\varepsilon' < \varepsilon$ , then almost surely a Nash equilibrium will be played at least  $1 - \varepsilon$  of the time* [4].

Since agents respond only to their *own* payoffs, it follows that they do not even need to know that they are involved in a game for Nash equilibrium to be a reasonable prediction of their long-run behavior.

There remains the question of how long it takes for adaptive learning to reach equilibrium, or something close to it, starting from out-of-equilibrium conditions. In other words, can one design learning algorithms that are *efficient*? The answer depends on the structure of the game and on its ‘size’ as measured by the number of players and/or the number of strategies. It can be shown that some games take an exponentially long time to learn by any adaptive learning rule that is ‘uncoupled,’ i.e., the rule does not need other agents’ utility functions as inputs [7]. But there are other types of games – such as coordination games – that can be learned relatively quickly. Here I shall briefly outline one result of this nature.

Consider a symmetric  $2 \times 2$  coordination game  $G$  that is played on a network. The network has one agent located at each node and its undirected edges describe the pairwise interactions that occur between agents. In each period  $t$ , each agent  $i$  chooses one action  $a_i(t)$  from his feasible action set. His payoff,  $u_i(a_i(t), a_{-i}(t))$ , is the sum of the payoffs he gets from playing this action once against each of his neighbors (given their choices  $a_{-i}(t)$ ).

Fix a symmetric two-person game  $G$  such that each player has two actions,  $a, b$ , coordinating on action  $a$  has a strictly higher payoff than coordinating on  $b$ , and miscoordinating has zero payoff. Given a network  $N$  and a small number  $\delta > 0$ , let  $T_{\delta, N}$  be the maximum expected waiting time (over all initial states) such that in every period  $t \geq T_{\delta, N}$ , at least  $1 - \delta$  of the agents are playing the Pareto superior action with probability at least  $1 - \delta$ . For the given game  $G$  we shall say that a learning rule is *efficient at level  $\delta$  on a family of networks  $\mathcal{N}$* , if  $T_{\delta, N}$  is bounded above for all networks  $N \in \mathcal{N}$ .

As a particular example, consider the following adaptive procedure known as log-linear learning [8]. At the start of period  $t$ , each agent gets an independent update signal from a Poisson random variable, where the signals are i.i.d. among agents. When an agent  $i$  receives a signal, he chooses an action according to a logit function of the action's payoff conditional on the current actions of  $i$ 's neighbors, that is,

$$P(a_i(t+1) = a) = e^{\beta u_i(a, a_{-i}(t))} / (e^{\beta u_i(a, a_{-i}(t))} + e^{\beta u_i(b, a_{-i}(t))}). \quad (1)$$

Here  $\beta > 0$  measures the *sensitivity of response*: if  $\beta$  is high the agent chooses a best response with high probability; if  $\beta$  is low the response is close to being random.

This process can be quite inefficient. Suppose, for example, that the network is large and complete: everyone is a neighbor of everyone else. If the process starts in the inferior equilibrium in which everyone plays  $b$ , the expected waiting time until the process comes close to the  $a$ -equilibrium is exponential in the number of agents (for a given choice of the response parameter  $\beta$ ). The reason is that the large number of players, combined with their high degree of interconnectedness, acts as a drag on the choice of each individual, hence change in the aggregate takes a very long time.

Suppose, by contrast, that the network consists of many small disjoint groups (e.g., families or tribes) of size  $s$ . Assume that each member of a group interacts only with other members of the group and has no connections with outsiders. In any given group it suffices for a few people to choose action  $a$  "by mistake" for the others to want to choose  $a$  also. If the response parameter  $\beta$  is sufficiently large, then once a given group has reached the all- $a$  equilibrium, they will be at this equilibrium with high probability in every subsequent period. Since the process is operating independently across groups, it follows that the expected waiting time until a *high proportion* of the agents is playing  $a$  with high probability is bounded above irrespective of the total number of agents.

The most interesting situation arises when agents are spatially distributed and interact only with their near neighbors. For example, suppose they are located at the vertices of a two-dimensional grid, and each agent interacts with his four neighbors. (Assume the grid is embedded on the torus so that it is regular of degree four.) It turns out that in this case too the waiting time for logit learning is bounded irrespective of the size of the grid. This follows from a topological condition that ensures efficiency of log-linear learning on general networks. The basic idea is that every agent should be in *some* group of bounded size whose members interact mainly with each other as opposed to outsiders (for details see [9, 10]). Such a family of networks is said to be

*close-knit*. A great variety of spatial distributions have this property. It can be shown that *log-linear learning of 2 x 2 coordination games is efficient on any close-knit family of networks*.

In summary, there are simple adaptive learning rules that come close to Nash equilibrium behavior at the system-wide level, and that require no special computational powers on the part of the agents or knowledge of the game they are playing. The efficiency of these processes depends importantly on both the payoff structure of the game and on the topology of the agents' interactions. While some games take exponentially long to learn, there are others of practical importance that can be learned quite quickly. It is a challenging open problem to extend these types of results to other forms of learning and more general classes of games.

## References

1. Young, H.P.: Learning and Evolution in Games: Adaptive Heuristics. In: Durlauf, S.N., Blume, L.E. (eds.) *The New Palgrave Dictionary of Economics*, 2nd edn. MacMillan, London (2008)
2. Hart, S.: Adaptive Heuristics. *Econometrica* 73, 1401–1430 (2005)
3. Marden, J.R., Young, H.P., Arslan, G., Shamma, J.S.: Payoff-Based Dynamics for Multiplayer Weakly Acyclic Games. *SIAM Journal on Control and Optimization* 48, 373–396 (2009)
4. Foster, D.P., Young, H.P.: Regret Testing: Learning to Play Nash Equilibrium Without Knowing You Have an Opponent. *Theoretical Economics* 1, 341–367 (2006)
5. Young, H.P.: Learning by Trial and Error. *Games and Economic Behavior* 65, 626–643 (2009)
6. Bowling, M., Veloso, M.: Multiagent Learning Using a Variable Learning Rate. *Artificial Intelligence* 136, 215–250 (2002)
7. Thuijssman, F., Peleg, B., Amitai, M., Shmida, A.: Automata, Matching, and Foraging Behavior in Bees. *Journal of Theoretical Biology* 175, 305–316 (1995)
8. Hart, S., Mansour, Y.: The Communication Complexity of Uncoupled Equilibrium Procedures. In: *Proceedings of the 39<sup>th</sup> Annual ACM Symposium on Theory of Computing*, pp. 345–353 (2007)
9. Blume, L.E.: The Statistical Mechanics of Best-Response Strategy Revision. *Games and Economic Behavior* 11, 111–145 (1995)
10. Young, H.P.: *Individual Strategy and Social Structure*, ch. 6. Princeton University Press, Princeton (1998)
11. Young, H.P.: The Diffusion of Innovations in Social Networks. In: Blume, L.E., Durlauf, S.N. (eds.) *The Economy as an Evolving Complex System*, vol. III, pp. 267–281. Oxford University Press, Oxford (2003)



# Quantifying Outcomes in Games

Eva Tardos

Department of Computer Science,  
Cornell University

**Abstract.** Network games play a fundamental role in understanding behavior in many domains, ranging from communication networks through markets to social networks. In this talk we'll study the degradation of quality of solution caused by the selfish behavior of users in a number of different games including congestion games that model routing or cost-sharing, and games that model Ad-Auctions. In each setting our goal is to quantify the degradation of quality of solution caused by the selfish behavior of user. We compare the selfish outcome to a centrally designed optimum both in terms of the quality of Nash equilibria and also the quality of outcomes of learning behavior by the users.

# Competitive Routing over Time

Martin Hoefer<sup>1,\*</sup>, Vahab S. Mirrokni<sup>2</sup>, Heiko Röglin<sup>3,\*\*</sup>, and Shang-Hua Teng<sup>4</sup>

<sup>1</sup> Department of Computer Science, RWTH Aachen University  
mhoefer@cs.rwth-aachen.de

<sup>2</sup> Google Research New York  
mirrokni@google.com

<sup>3</sup> Department of Quantitative Economics, Maastricht University  
heiko@roeglin.org

<sup>4</sup> Computer Science Department, University of Southern California  
shanghua.teng@gmail.com

**Abstract.** Congestion games are a fundamental and widely studied model for selfish allocation problems like routing and load balancing. An intrinsic property of these games is that players allocate resources simultaneously and instantly. This is particularly unrealistic for many network routing scenarios, which are one of the prominent application scenarios of congestion games. In many networks, load travels along routes over time and allocation of edges happens sequentially. In this paper we consider two frameworks that enhance network congestion games with a notion of time. We propose *temporal network congestion games* that use coordination mechanisms — local policies that allow to sequentialize traffic on the edges. In addition, we consider *congestion games with time-dependent costs*, in which travel times are fixed but quality of service of transmission varies with load over time. We study existence and complexity properties of pure Nash equilibria and best-response strategies in both frameworks. In some cases our results can be used to characterize convergence for various distributed dynamics.

## 1 Introduction

As an intuitive game-theoretic model for competitive resource usage, *network congestion games* have recently attracted a great deal of attention [1,2,3]. These games are central in modeling routing and scheduling tasks with distributed control [4]. Such games can be described by a routing network and a set of players who each have a source and a target node in the network and choose a path connecting these two nodes. The quality of a player's choice is evaluated in terms of the total delay or latency of the chosen path. For this, every edge  $e$  has a latency function that increases with the number of players whose paths include edge  $e$ . Ignoring the *inherent delay* in transmitting packets in networks

---

\* Supported by DFG through UMIC Research Centre at RWTH Aachen University.

\*\* Supported by a fellowship within the Postdoc-Program of the German Academic Exchange Service (DAAD).

or routing cars in road networks, this model implicitly assumes that players use all edges on their paths instantaneously and simultaneously.

Depending on the application, it might not be reasonable to assume that a player instantaneously allocates all edges on his chosen path. Consider for instance a road traffic network, in which players route cars to their destinations. Clearly, a traffic jam that delays people at rush hour might be harmless to a long distance traveler who reaches the same street hours later. In this case, it is more natural to assume that edges are allocated consecutively, and players take some time to pass an edge before they reach the next edge on their path. In particular, each edge may have a *local queueing policy* to schedule the players traversing this edge.

In this paper, we study two different models that extend the standard model of network congestion games by a temporal component. In our first model, we incorporate the assumption that on each edge, the traffic over the edge must be sequentialized which in turn results in a *local scheduling problem with release times* on each edge, and requires a formal description of the local scheduling or queueing policy on each edge. To model these local scheduling policies, we use the idea of *coordination mechanisms* [5,6,7,8] that have been introduced in the context of machine scheduling and selfish load balancing [9]. In selfish load balancing, each player has a task and has to assign it to one of several machines in order to minimize his completion time. A coordination mechanism is a set of local scheduling policies that run locally on machines. Given an assignment of tasks to machines, the coordination mechanism run on a machine  $e$  gets as input the set of tasks assigned to  $e$  and their processing times on  $e$ . Based on this information, it decides on a preemptive or non-preemptive schedule of the tasks on  $e$ . The local scheduling policies of the coordination mechanism do not have access to any global information, like, e.g., the set of all tasks and their current allocation.

Applying the idea of coordination mechanisms to network congestion games results in the definition of *temporal congestion games*, which are studied in Section 3. We assume that each edge in a network congestion game is a machine equipped with a local scheduling policy, and each player has a task and chooses a path. Starting from their source, tasks travel along their path from one edge to another until they reach the target. They become available on the next edge of their path only after they have been processed completely on the previous edges. The player incurs as latency the total travel time that his task needs to reach the target. Each player then strives to pick a path that minimizes his travel time.

In our second model, which we term *congestion games with time-dependent costs* and study in Section 4, we assume that the travel time along each edge is a constant independent of the number of players using that edge. This model captures the property that increased traffic yields decreased quality of service for transmitting packets. We model this via a time-dependent cost function. We assume time is discretized into units (e.g., seconds), and the cost of an edge during a second depends on the number of players currently traveling on the

edge. Each player now strives to pick a path that minimizes the total time-dependent costs during the travel time along the edges.

Our games extend atomic congestion games, which were initially considered by Rosenthal [3]. They are a vivid research area in (algorithmic) game theory and have attracted much research interest, especially over the last decade. A variety of issues have been addressed, most prominently complexity of computing equilibria [1,2,3] and bounding their inefficiency [10,11,12]. For an overview and introduction to the topic we refer to the recent expositions by Roughgarden [4] and Vöcking [9]. Addressing the notion of time in congestion games has only been started very recently in a number of papers [13,14,15]. Koch and Skutella [15] present a general model for flows over time using queueing models. Similarly, Anshelevich and Ukkusuri [13] derive a number of related results for a similar model of flows over time. In contrast to our work both papers address non-atomic congestion games, in which players are infinitesimally small flow particles. Farzad et al. [14] consider a priority-based scheme for both, non-atomic and atomic games. In their model players have priorities, and a resource yields different latencies depending on the priority of players allocating it. This includes an approach of Harks et al. [16] as a special case. While there can be different latencies for different players, this model does not include a more realistic “dynamic” effect that players delay other players only for a certain period of time. This is the case in our paper, as well as in [13,15] for the non-atomic case.

## 1.1 Our Contribution

For temporal congestion games, we study four different (classes of) coordination mechanisms: (1) *FIFO*, in which tasks are processed non-preemptively in order of arrival. (2) *Non-preemptive global ranking*, in which there is a global ranking among the tasks that determines in which order tasks are processed non-preemptively (e.g., Shortest-First or Longest-First). (3) *Preemptive global ranking*, in which there is a global ranking that determines in which order tasks are processed and higher ranked tasks can preempt lower ranked tasks. (4) *Fair Time-Sharing*, in which all tasks currently located at an edge get processed simultaneously and each of them gets the same share of processing time.

For the FIFO policy (in unweighted symmetric games) and the Shortest-First policy (in weighted symmetric games) we show an interesting contrast of positive and negative results: even though computing a best response is NP-hard, there always exists an equilibrium, which can be computed in polynomial time. Moreover, the equilibrium is not only efficiently computable, but we present natural dynamics in which uncoordinated agents are able to find an equilibrium quickly even without solving computationally hard problems. We then show that Shortest-First is the only global ranking that guarantees the existence of Nash equilibria in the non-preemptive setting. That is, for any other global ranking (e.g., Longest-First) there exist temporal congestion games without equilibria. In contrast to this, we show that preemptive games are potential games for every global ranking and that uncoordinated agents reach an equilibrium quickly. Finally, we show that even though Fair Time-Sharing sounds like an appealing

coordination mechanism it does not guarantee the existence of equilibria, not even for unweighted symmetric games.

For the second model, congestion games with time-dependent costs, we prove that these games can be reduced to standard congestion games. Hence, they are potential games, and in addition the known results on the price of anarchy carry over. We prove that computing a best response in these games is NP-hard in general. Even for a very restricted class of games with polynomially bounded delays and acyclic networks computing an equilibrium is PLS-complete. Due to space limitations, some proofs are deferred to the full version of this paper.

## 2 Notation

A *network congestion game* is described by a directed graph  $G = (V, E)$ , a set  $\mathcal{N} = \{1, \dots, n\}$  of *players* with *source nodes*  $s_1, \dots, s_n \in V$  and *target nodes*  $t_1, \dots, t_n \in V$ , and a non-decreasing *latency function*  $\ell_e: [n] \rightarrow \mathbb{R}_{\geq 0}$  for each edge  $e$ . We will only consider linear latency functions of the form  $\ell_e(x) = a_e x$  in this paper. For such functions, we call  $a_e$  the *speed* of edge  $e$ . The *strategy space*  $\Sigma_i$  of a player  $i \in \mathcal{N}$  is the set of all simple paths in  $G$  from  $s_i$  to  $t_i$ . We call a network congestion game *weighted* if additionally every player  $i$  has a weight  $w_i \geq 1$ , and *unweighted* if  $w_1 = \dots = w_n = 1$ . Given a *state*  $P = (P_1, \dots, P_n) \in \Sigma = \Sigma_1 \times \dots \times \Sigma_n$  of a network congestion game, we denote by  $n_e(P) = \sum_{i: e \in P_i} w_i$  the *congestion* of edge  $e \in E$ . The *individual latency* that a player  $i$  incurs is  $\ell_i(P) = \sum_{e \in P_i} \ell_e(n_e(P))$ , and every player is interested in choosing a path of minimum individual latency. We call a congestion game *symmetric* if every player has the same source node and every player has the same target node. If not explicitly mentioned otherwise, we consider unweighted asymmetric congestion games.

We incorporate time into the standard model in two different ways. Formally, this alters the individual latency functions  $\ell_i$ . The specific definitions will be given in the sections below. For our altered games we are interested in stable states, which are pure strategy Nash equilibria of the games. Such an equilibrium is given by the condition that each player plays a best response and has no unilateral incentive to deviate, i.e.,  $P$  is a pure *Nash equilibrium* if for every player  $i$  and every state  $Q$  that is obtained from  $P$  by replacing  $i$ 's path by some other path, it holds  $\ell_i(P) \leq \ell_i(Q)$ , where  $\ell_i$  denotes the (altered) latency function of player  $i$ . We will not consider mixed Nash equilibria in this paper, and the term Nash equilibrium will refer to the pure version throughout.

## 3 Coordination Mechanisms

In this section we consider *temporal network congestion games*. These games are described by the same parameters as standard weighted network congestion games with linear latency functions. However, instead of assuming that a player allocates all edges on his chosen path instantaneously, we consider a scenario in which players consecutively allocate the edges on their paths. We assume that

each player has a weighted task that needs to be processed by the edges on his chosen path.

Formally, at each point in time  $\tau \in \mathbb{R}_{\geq 0}$ , every task  $i$  is located at one edge  $e_i(\tau)$  of its chosen path, and a certain fraction  $f_i(\tau) \in [0, 1]$  of it is yet unprocessed on that edge. The coordination mechanism run on edge  $e$  has to decide in each moment of time which task to process. If it decides to work on transmitting task  $i$  for  $\Delta\tau$  time units starting at time  $\tau$ , then the unprocessed fraction  $f_i(\tau + \Delta t)$  of task  $i$  at time  $\tau + \Delta t$  is  $\max(0, f_i(\tau) - \Delta\tau/(a_e w_i))$ . In total, task  $i$  needs  $a_e w_i$  time units to finish on edge  $e$ . Once  $f_i(\tau) = 0$ , task  $i$  arrives at the next edge on its path and becomes available for processing. The coordination mechanism can base the decision on which task to process next for how long only on local information available at the edge — such as the weights and arrival times of those tasks that have already arrived at the edge. The individual latency  $\ell_i(P)$  of player  $i$  in state  $P$  is the time at which task  $i$  is completely finished on the last edge of  $P_i$ .

### 3.1 The FIFO Policy

One of the most natural coordination mechanisms is the FIFO policy. If several tasks are currently located at the same edge, then the one that has arrived first is executed non-preemptively until it finishes. In the case of ties, there may be an arbitrary tie-breaking that is consistent among the edges.

**Unweighted and Symmetric Games.** In this section we treat unweighted symmetric temporal network congestion games. For these games we obtain an interesting contrast of positive and negative results: even though computing a best response is NP-hard, there always exists a Nash equilibrium, which can be computed in polynomial time. Moreover, the equilibrium is not only efficiently computable, but uncoordinated agents are able to find it quickly even without solving computationally hard problems.

**Theorem 1.** *For unweighted symmetric temporal network congestion games with the FIFO policy a Nash equilibrium always exists. Moreover, a Nash equilibrium can be computed efficiently.*

*Proof.* Let us assume without loss of generality that players are numbered according to their rank in tie-breaking, i.e. 1 is the highest ranked player, and  $n$  is the lowest ranked player. Assume that we start in an arbitrary state of the game in which the players have chosen arbitrary paths. Below we define a subclass of best responses, which we call *greedy best responses*. We claim that we obtain an equilibrium if we let the players  $1, 2, \dots, n$  play each one greedy best response in this order. To prove this, assume that the players  $1, \dots, i$  are already playing greedy best responses, and now let player  $i + 1$  also change his strategy to a greedy best response. We show that after this strategy change the players  $1, \dots, i$  are still playing greedy best responses, which proves by induction that a Nash equilibrium is reached once every player has played a greedy best response. We

prove the following invariant: if the players  $1, \dots, i$  play greedy best responses, then none of them can be delayed at any node by a lower ranked player  $j > i$ . Furthermore, the current paths of the players  $1, \dots, i$  are (greedy) best responses no matter which paths the other players  $j > i$  choose. For the first player, every best response is defined to be a greedy best response. Given this definition, we argue that the aforementioned claim is true for  $i = 1$ : We consider the network  $G = (V, E)$  as a weighted graph in which every edge  $e \in E$  has weight  $a_e$ . Let  $P_1$  denote a shortest path in this weighted graph from the source  $s$  to the target  $t$  and let  $a^*$  denote its length. If the highest ranked player chooses path  $P_1$ , then he cannot be delayed at any node  $v$  by any other player  $j$ , as otherwise,  $j$  would have found a shorter path from  $s$  to  $v$ , contradicting the choice of  $P_1$  as shortest path from  $s$  to  $t$ . Hence, when player 1 chooses path  $P_1$  his total latency is  $a^*$  no matter which paths the other players choose. Clearly, the length  $a^*$  is also a lower bound on the time it takes any player to reach the target, and hence, choosing  $P_1$  is a (greedy) best response for player 1. Moreover, any (greedy) best response of player 1 corresponds to a shortest path  $P_1$  in the aforementioned weighted graph. Now let us recursively define what a greedy best response is for player  $i + 1 > 1$ . For this, assume that the players  $1, \dots, i$  play already greedy best responses. Based on the paths chosen by these players, we construct a distance function  $d: V \rightarrow \mathbb{R}_{\geq 0}$  for the network  $G = (V, E)$ , which eventually tells us for every node how long it takes player  $i + 1$  to get there. The construction of this distance function follows roughly Dijkstra's algorithm: Let  $I \subseteq V$  denote the set of nodes that have already an assigned distance. We start with  $I = \{s\}$  and  $d(s) = 0$ . For extending the set  $I$ , we crucially use the fact that the players  $1, \dots, i$  cannot be delayed by other players, which means that every edge  $e \in E$  has a fixed schedule saying when it is used by the players  $1, \dots, i$  and when it is available for player  $i + 1$ . These fixed schedules imply in particular that for every node  $v \in V$  there exists a shortest path  $s, v_1, \dots, v_k = v$  for player  $i + 1$  from  $s$  to  $v$  such that every subpath  $s, v_1, \dots, v_{k'}$  is a shortest path from  $s$  to  $v_{k'}$ . Hence, taking into account the fixed schedules and the possible delays that they induce on player  $i + 1$ , we can extend the set  $I$  as in Dijkstra's algorithm, that is, we insert the node  $v \in V \setminus I$  into  $I$  that minimizes  $\min_{u \in I} d(u) + \ell(u, v)$ , where  $\ell(u, v)$  denotes the time it takes player  $i + 1$  to get from  $u$  to  $v$  if he arrives at node  $u$  at time  $d(u)$ . The distance  $d(v)$  assigned to node  $v$  is  $\min_{u \in I} d(u) + \ell(u, v)$ . This algorithm constructs implicitly a path from  $s$  to any other node. Any path from  $s$  to  $t$  that can be constructed by this algorithm (the degree of freedom is the tie-breaking) is called a greedy best response of player  $i + 1$ .

It is easy to see that any such greedy best response is really a best response for player  $i + 1$  if only the players  $1, \dots, i + 1$  are present, because there is no quicker way to reach the target  $t$  from the source  $s$  given the current paths of the players  $1, \dots, i$ . To complete the proof we need to argue that player  $i + 1$  cannot be delayed by players  $j > i + 1$ . Assume there is a node  $v$  and a player  $j > i + 1$  such that  $j$  arrives earlier at node  $v$  than  $i + 1$ . This contradicts the construction of the path as it implies that there is a faster way to get from the

source  $s$  to the node  $v$ . Again this argument crucially uses the property that the players  $1, \dots, i$  cannot be delayed by lower ranked players.  $\square$

Let us remark that greedy best responses defined in the previous proof are the discrete analogon of subpath-optimal flows introduced by Cole et al. [17]. Basically, a path  $s, v_1, \dots, v_k, t$  is a greedy best response for player  $i$  if any subpath  $s, v_1, \dots, v_{k'}$  is a shortest path from  $s$  to  $v_{k'}$ . Note that this is not the case in arbitrary best responses: it could, for example, be the case that player  $i$  has to wait at some node  $v_{k'}$  because he is blocked by a player with a higher rank. Then, the subpath from  $s$  to  $v_{k'}$  is not necessarily the shortest possible path in every best response. However, we believe that the restriction to greedy best responses is a natural assumption on the players' behavior.

The previous result shows not only that a Nash equilibrium always exists, but it also shows that players reach it in a distributed fashion using different forms of dynamics. Consider the following Nash dynamics among the players. At each point in time, one player is picked and allowed to change his strategy. We show below that in general it is NP-hard for this player to decide whether it can decrease his latency by changing his path. In that case, the player might stick to his current path or make an arbitrary strategy change, following some heuristic. However, at each point in time there is one player who can easily find a (greedy) best response, namely the highest ranked player  $i + 1$  that does not play a greedy best response, but the players  $1, \dots, i$  do. We assume that this player changes to a greedy best response when he becomes activated. We also assume that a player who is already playing a greedy best response does not change his strategy when he becomes activated. A *round* is a sequence of activations in which every player gets at least once the chance to change his strategy. From the proof of Theorem 1 it follows easily that a Nash equilibrium is reached after at most  $n$  rounds. We are interested in particular in the *random greedy best response dynamics*, in which in each iteration the activated player is picked uniformly at random, and the *concurrent best response dynamics*, in which in each iteration all players are simultaneously allowed to change their strategy, each one with some constant probability  $0 < p_i \leq 1$ . In both these dynamics, rounds are polynomially long with high probability. Summarizing, we obtain the following corollary.


**Corollary 2.** *In every unweighted symmetric temporal network congestion game with the FIFO policy it takes at most  $n$  rounds to reach a Nash equilibrium. In particular, the random and concurrent greedy best response dynamics reach a Nash equilibrium in expected polynomial time.*

Finally, we turn to the hardness result.

**Theorem 3.** *Computing best responses is NP-hard in unweighted symmetric temporal network congestion games with the FIFO policy.*

**Weights and Asymmetric Players.** Now we show that any relaxation of the restrictions in the previous sections leads to games without equilibria.





**Fig. 1.** (a) Asymmetric temporal network congestion game without Nash equilibrium for FIFO. Edge labels indicate the speeds  $a_e$ . For all unlabeled edges  $e$ , we have  $a_e = 1$ . (b) Unweighted symmetric game without Nash equilibrium for Time-Sharing.

**Theorem 4.** *There exist temporal congestion games with the FIFO policy that do not possess Nash equilibria and (1) are weighted and symmetric, or (2) are unweighted and asymmetric.*

*Proof.* The example for the first case is simple; it consists of three edges: there are three nodes  $s$ ,  $v$ , and  $t$  and two parallel edges from  $s$  to  $v$  (if multi edges are not allowed, they can be split up into two edges each by inserting intermediate nodes) and one edge from  $v$  to  $t$ . All edges have speed 1. Assume that there are two players with weights 2 and 3, and assume that the player with weight 3 has higher priority. If both players use the same edge from  $s$  to  $v$ , then the player with weight 2 has an incentive to switch to the free edge. If they use different edges, the player with weight 3 has an incentive to use the same edge as the other player.

Now let us turn to the second case. We consider the instance shown in Figure 1 (a). In this game there are three unweighted players, and each player  $i$  has two possible strategies: the vertical three edges (denoted by  $A_i$ ) and another path (denoted by  $B_i$ ). The following sequence of moves constitutes a cycle in the best response dynamics:  $(A_1, A_2, A_3) \rightarrow (B_1, A_2, A_3) \rightarrow (B_1, B_2, A_3) \rightarrow (B_1, B_2, B_3) \rightarrow (A_1, B_2, B_3) \rightarrow (A_1, A_2, B_3) \rightarrow (A_1, A_2, A_3)$ . It is easy to verify that the remaining configurations  $(A_1, B_2, A_3)$  and  $(B_1, A_2, B_3)$  are no Nash equilibria either.  $\square$

### 3.2 Non-preemptive Global Ranking

Another natural approach is to assume that there is a global ranking  $\pi: [n] \rightarrow [n]$  on the set of tasks with  $\pi(1)$  being the task with the highest priority and so on. In this case, tasks are scheduled non-preemptively according to this ranking. When an edge  $e$  becomes available, the highest ranked task  $i$  that is currently located at the edge is processed non-preemptively. It exclusively uses  $e$  for  $a_e w_i$  time units. After that, task  $i$  moves to the next edge on its path, and  $e$  selects

the next task if possible. In this section, we consider mainly weighted games and assume without loss of generality that  $w_1 \leq w_2 \leq \dots \leq w_n$ .

**Shortest-First Policy.** In this section we consider the identity ranking  $\pi(i) = i$ , which corresponds to the *Shortest-First policy*. It is easy to see that Theorem 1 and Corollary 2 carry over to this case. The proof for FIFO was essentially based on the observation that once all players  $1, \dots, i$  play a (greedy) best response, they cannot be affected by the lower ranked players. This is even more true for the Shortest-First policy as the lower ranked players now face the additional disadvantage of having a longer processing time.

**Theorem 5.** *In every weighted symmetric temporal network congestion game with the Shortest-First policy a Nash equilibrium exists. Moreover, a Nash equilibrium can be computed efficiently, and it takes at most  $n$  rounds to reach a Nash equilibrium. In particular, the random and concurrent greedy best response dynamics reach a Nash equilibrium in expected polynomial time.*

Also the hardness result in Theorem 3 carries over easily.

**Theorem 6.** *In (unweighted) symmetric temporal network congestion games with the Shortest-First policy computing a best response is NP-hard.*

Although the previous arguments guarantee existence and convergence to a Nash equilibrium for the Shortest-First policy, such games are not necessarily potential games.

**Proposition 7.** *There is a symmetric temporal network congestion game with the Shortest-First policy that is no potential game.*

**Other Global Rankings or Asymmetric Players.** Now we consider the case of more general rankings.

**Theorem 8.** *For any given set of player task weights  $w_1 \leq \dots \leq w_n$  and any ranking  $\pi$  other than the identity, there exist a graph and latency functions such that the resulting symmetric temporal congestion game does not possess a Nash equilibrium.*

The proof is given in the full version of the paper. It relies on the fact that for rankings other than the identity a larger task can delay smaller tasks near the source due to the ranking, but smaller tasks can delay larger tasks near the sink due to faster travel time. The same result holds for asymmetric games with the Shortest-First policy. We can simply add a separate source for each player and connect it via a single edge to the original source. By appropriately adjusting the delays  $a_e$  on these edges, we can ensure that smaller tasks are suitably delayed before arriving at the original source. This results in the same incentives.

**Corollary 9.** *For any given set of task weights  $w_1, \dots, w_n$  and the Shortest-First policy, there exist a graph and latency functions such that the resulting asymmetric temporal congestion game does not have a Nash equilibrium.*

### 3.3 Preemptive Global Ranking

When we assume a global ranking and allow preemptive execution, it is possible to adapt the arguments of Theorem 1 to weighted asymmetric games. Indeed, all arguments in this section work for a very general class of preemptive games with unrelated edges. That is, every player  $i$  has its own processing time  $p_{ie}$  for every edge  $e$ . These processing times may even depend on the time at which player  $i$  reaches edge  $e$ . The only assumption we need to make is that the processing times are monotone in the sense that if task  $i$  reaches edge  $e$  at time  $t$ , then it does not finish later than when it reaches edge  $e$  at time  $t' > t$ .

**Theorem 10.** *Every asymmetric temporal congestion game with preemptive policy  $\pi$  is a potential game. A Nash equilibrium exists and can be computed in polynomial time. For any state and any player, a best response can be computed in polynomial time.*

Similarly we can adapt the previous observations in Corollary 2 and show that various improvement dynamics converge in polynomial time.

**Corollary 11.** *In every asymmetric temporal network congestion game with any preemptive policy  $\pi$ , it takes at most  $n$  rounds to reach a Nash equilibrium. The expected number of iterations to reach a Nash equilibrium for random and concurrent best response dynamics is bounded by a polynomial.*

### 3.4 Fair Time-Sharing

In this section we consider fair time-sharing, a natural coordination mechanism based on the classical idea of uniform processor sharing [18]. When multiple tasks are present at an edge  $e$ , they are all processed simultaneously, and each one of them gets the same share of bandwidth or processing time. As in generalized processor sharing [19] we assume round-robin processing with infinitesimal time slots. Even though such a fairness property is desirable, the following theorem shows that Nash equilibria are not even guaranteed to exist for symmetric unweighted games.

**Theorem 12.** *There is an unweighted symmetric temporal network congestion game with the Time-Sharing policy that does not have a Nash equilibrium.*

*Proof.* The instance shown in Figure 1(b) has three players. As the three edges leaving the source  $s$  are very slow, in any Nash equilibrium all three players will use different edges leaving the source. We assume without loss of generality that the first player chooses the upper edge, the second player chooses the middle edge, and the third player chooses the lower edge. Then players 1 and 3 have still two alternatives how to continue, whereas the path of player 2 is already determined. The speeds of the edges are chosen such that player 1 wants to use the edge with speed  $5 + \varepsilon$  if and only if player 3 does not use the edge with speed  $4 - \varepsilon$ . On the other hand, player 3 wants to use the edge with speed  $4 - \varepsilon$  if and only if player 1 uses the edge with speed  $5 + \varepsilon$ , which completes the proof.  $\square$

Dürr and Nguyen [20] show that Time-Sharing on parallel links always yields a potential game, even for unrelated machines (edges). That is, for parallel links Nash equilibria always exist. Their potential function can be rewritten as the sum of the completion times (individual latencies) of the player. It is known [21] that a schedule minimizing this sum can be computed in polynomial time. Such a global minimum of the potential function must obviously be a pure Nash equilibrium for the Time-Sharing policy, yielding the following corollary.

**Corollary 13.** *For games on parallel links with unrelated tasks and the Time-Sharing policy a Nash equilibrium can be computed efficiently.*

## 4 Constant Travel Times and Quality of Service

Now let us consider *network congestion games with time-dependent costs*. Again, players consecutively allocate the edges on their paths. However, the travel time along an edge  $e$  in the network is fixed to a constant delay  $d_e$ . If a player chooses a path along the edges  $e_1, e_2, \dots$ , then he arrives at  $e_2$  at time  $d_1$  and at  $e_3$  at time  $d_1 + d_2$ , and so on. This travel time through the network is independent of how many other players allocate any of the edges. In this section, we only consider asymmetric games. For the strategic part we assume that each edge generates a separate usage cost  $c_e$  per time unit. This could, for instance, measure the quality of service that is enjoyed by the players during transmission. The cost depends on the number of players allocating the edge at a given point in time. In particular, edge  $e$  has a cost function  $c_e: [n] \rightarrow \mathbb{N}$  that describes the cost for allocating it for one second in terms of the current number of players. If for a state  $P$  an edge  $e$  is shared at time  $\tau$  by  $n_e(\tau, P)$  players, all these players get charged cost  $c_e(n_e(\tau, P))$ . The cost incurred by player  $i$  on a path  $P_i = (e_1, \dots, e_l)$  is then  $\ell_i(P) = \sum_{j=1}^l \sum_{\tau=\tau_j}^{\tau_j+d_{e_j}-1} c_{e_j}(n_{e_j}(\tau, P))$ , where  $\tau_1 = 0$  and  $\tau_j = \sum_{k=1}^{j-1} d_{e_k}$ .

It turns out that this model is equivalent to a regular congestion game. For each edge and each time unit we introduce a resource  $r_{e,\tau}$  and modify the strategy spaces as follows: For a path  $P = (e_1, \dots, e_l)$  the new strategy includes all resources  $r_{e_j,\tau}$  for  $\tau = \tau_j, \dots, \tau_j + d_{e_j} - 1$  and  $j = 1, \dots, l$ . This is a regular congestion game with latencies given by the time costs. Hence, results on the existence of Nash equilibria and the price of anarchy carry over.

**Corollary 14.** *Network congestion games with time-dependent costs are equivalent to a class of regular congestion games. In particular, there is a pure Nash equilibrium in every game, and any better-response dynamics converges.*

However, as the standard congestion game obtained by this reduction might have a large number of resources and as it is not necessarily a network congestion game, complexity results do not carry over.

**Theorem 15.** *Computing a best response in network congestion games with time-dependent costs is NP-hard. For games with polynomially bounded delays and acyclic networks, best responses can be computed efficiently, but computing a Nash equilibrium is PLS-complete.*

## References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. *Journal of the ACM* 55(6) (2008)
2. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The complexity of pure Nash equilibria. In: *Proc. 36th Symp. Theory of Computing (STOC)*, pp. 604–612 (2004)
3. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
4. Roughgarden, T.: Routing Games. In: *Algorithmic Game Theory*, pp. 461–486. Cambridge University Press, Cambridge (2007)
5. Azar, Y., Jain, K., Mirrokni, V.S.: (Almost) optimal coordination mechanisms for unrelated machine scheduling. In: *Proc. 19th Symp. Discrete Algorithms (SODA)*, pp. 323–332 (2008)
6. Caragiannis, I.: Efficient coordination mechanisms for unrelated machine scheduling. In: *Proc. 20th Symp. Discrete Algorithms (SODA)*, pp. 815–824 (2009)
7. 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)
8. Immorlica, N., Li, L., Mirrokni, V.S., Schulz, A.S.: Coordination mechanisms for selfish scheduling. In: Deng, X., Ye, Y. (eds.) *WINE 2005*. LNCS, vol. 3828, pp. 55–69. Springer, Heidelberg (2005)
9. Vöcking, B.: Selfish Load Balancing. In: *Algorithmic Game Theory*, pp. 517–542. Cambridge University Press, Cambridge (2007)
10. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: *Proc. 37th Symp. Theory of Computing (STOC)*, pp. 67–73 (2005)
11. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: *Proc. 37th Symp. Theory of Computing (STOC)*, pp. 57–66 (2005)
12. Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: *Proc. 41st Symp. Theory of Computing, STOC* (2009)
13. Anshelevich, E., Ukkusuri, S.: Equilibria in dynamic selfish routing. In: *Proc. 2nd Intl. Symp. Algorithmic Game Theory, SAGT* (2009)
14. Farzad, B., Olver, N., Vetta, A.: A priority-based model of routing. *Chicago Journal of Theoretical Computer Science*, Article no. 1 (2008)
15. Koch, R., Skutella, M.: Nash equilibria and the price of anarchy for flows over time. In: *Proc. 2nd Intl. Symp. Algorithmic Game Theory, SAGT* (2009)
16. Harks, T., Heinz, S., Pfetsch, M.: Competitive online multicommodity routing. In: Erlebach, T., Kaklamanis, C. (eds.) *WAOA 2006*. LNCS, vol. 4368, pp. 240–252. Springer, Heidelberg (2007)
17. Cole, R., Dodis, Y., Roughgarden, T.: Bottleneck links, variable demand, and the tragedy of the commons. In: *Proc. 17th Symp. Discrete Algorithms (SODA)*, pp. 668–677 (2006)
18. Kleinrock, L.: *Queueing Systems. Computer Applications*, vol. 2. Wiley, Chichester (1976)
19. Parekh, A., Gallager, R.: A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Transactions on Networking* 1(3), 344–357 (1993)
20. Dürr, C., Thang, N.K.: Non-clairvoyant scheduling games. In: *Proc. 2nd Intl. Symp. Algorithmic Game Theory, SAGT* (2009)
21. Bruno, J., Coffman Jr., E.G., Sethi, R.: Scheduling independent tasks to reduce mean finishing-time (extended abstract). *SIGOPS Oper. Syst. Rev.* 7(4), 102–103 (1973)

# On 2-Player Randomized Mechanisms for Scheduling

Pinyan Lu

Microsoft Research Asia  
pinyanl@microsoft.com

**Abstract.** In this paper, we study randomized truthful mechanisms for scheduling unrelated machines. We focus on the case of scheduling two machines, which is also the focus of many previous works [12,13,6,4]. For this problem, [13] gave the current best mechanism with an approximation ratio of 1.5963 and [14] proved a lower bound of 1.5. In this work, we introduce a natural technical assumption called *scale-free*, which says that the allocation will not change if the instance is scaled by a global factor. Under this assumption, we prove a better lower bound of  $\frac{25}{16}$  ( $= 1.5625$ ). We then study a further special case, namely scheduling two tasks on two machines. For this setting, we provide a correlation mechanism which has an approximation ratio of 1.5089. We also prove a lower bound of 1.506 for all the randomized scale-free truthful mechanisms in this setting.

## 1 Introduction

Mechanism design has become an active area of research both in Computer Science and Game Theory. In the mechanism design setting, players are selfish and wish to maximize their own utilities. To deal with the selfishness of the players, a mechanism should both satisfy some game-theoretical requirements such as *truthfulness* and some computational properties such as *good approximation ratios*. The study of their algorithmic aspect was initiated by Nisan and Ronen in their seminal paper “Algorithmic Mechanism Design” [15]. The focus of this paper was on the scheduling problem on unrelated machines, for which the standard mechanism design tools (e.t. VCG mechanisms [5,7,16]) do not suffice. They proved that no deterministic mechanism can have an approximation ratio better than 2 for this problem. This bound is tight for the case of two machines. However if we allow randomized mechanisms, this bound can be beaten. In particular they gave a 1.75-approximation randomized truthful mechanism for the case of two machines. This bound has since been improved to 1.6737 [12] and then to 1.5963 [13] by Lu and Yu. In [14], Mu’alem and Schapira proved a lower bound of 1.5 for this setting. The focus of this paper is to explore the exact bound between 1.5 and 1.5963.

In [13], Lu and Yu also proved a lower bound of  $\frac{11}{7}$  ( $\approx 1.5714$ ) for all the *task independent* truthful mechanisms. A recent work [6] by Dobzinski and Sundararajan showed that any truthful mechanism for two machines with a finite

approximation ratio is *task-independent*. However the definitions of these two “task-independence” are not identical. The lower bound of [13] requires a strong version of “task-independence” and the characterization theorem in [6] only works for a weak version of “task-independence”. Formal definitions of these two versions of “task independence” are given in the next section. This gives an interesting open problem: is there any weak task-independent mechanism which can beat all the strong task independent ones (In particular has an approximation ratio which is better than  $\frac{11}{7}$ )? We note that all the previous known mechanisms in this setting are strong task-independent [15, 12, 13]. Roughly speaking, in a strong task-independent mechanism, the random bits used by the allocation algorithm for different tasks are independent, while in weak task-independent mechanisms, they may have some correlation. In section 4, we provide such a correlation mechanism. This is the first truthful mechanism for this problem which is not strong task-independent. This mechanism provides an approximation ratio of 1.5089 for the case of two task, which is strictly better than all the strong task independent mechanisms in the same setting. We note that the lower bound of  $\frac{11}{7}$  already holds even for the special case of scheduling two tasks in two machines.

The main focus of this paper is on the lower bound side. We introduce a natural assumption called *scale-free*, which says that the allocation will not change if a instance of the problem is scaled by a global factor. The property of *scale-free* is very natural for an allocation algorithm since a global factor only reflects the unit used for the running times. For example, if we change the unit from “hour” to “min”, we will scale the instance by a factor of 60, a reasonable allocation algorithm should be identical on these two instances (since they are in fact the same instance). We provide a refined characterization for all the scale-free truthful mechanisms with finite approximation ratio. Based on this characterization, we prove a lower bound of 1.5625 using Yao’s min-max principle [17]. We design a distribution of instances and argue that any scale-free deterministic truthful mechanisms cannot get an expected approximation ratio which is better than 1.5625. In order to get a better lower bound, we use a limitation argument and this value of 1.5625 holds when the number of tasks approaches infinity. So this lower bound only works for instances with a sufficiently large number of tasks. As we have a better mechanism for scheduling 2 tasks, we also study the lower bound of this special case under the assumption of scale-free. The instances used in the general lower bound cannot give a bound which is better than 1.5 when each of them only contains 2 tasks. So we choose a more carefully designed instances distribution to get a lower bound of 1.506. All these lower bound suggests that the lower bound of 1.5 may not be tight. However, it remains open to prove a better lower bound without any assumption.

A lot of technical effort in this work is given to parameter optimization both for the mechanism design part and lower bound proof part. Such optimization is also critical in this problem since the gap between the known upper bound and lower bound is already quite tiny. For example, for the 2 task case, the approximation ratio of the correlation mechanism we provided is 1.5089, while



the lower bound is 1.5. There is only a gap of 0.0089. A very carefully designed instance distribution gives a better lower bound of 1.506.

Despite the fact that our lower bounds rely on a technical assumption, we feel it is interesting for several reasons. First we think the assumption of scale-free is very natural, it is hard to imagine that some scale dependant mechanism can beat all the scale free mechanisms. So we conjecture that these lower bounds are also true for all the randomized truthful mechanisms. On the other hand, if one believes that a better mechanism exists, one has to look for really new mechanisms which are not scale-free. In both cases, we believe that our work in this paper is an important step toward the exact bound.

## 1.1 Related Work

Scheduling unrelated machines is one variant of the most fundamental scheduling Problems. For this NP-hard optimization problem, there is a polynomial-time 2-approximation algorithm, and unless  $P = NP$ , it is impossible to approximate the optimum within a factor less than  $3/2$  in polynomial time [11]. However there is no corresponding payment strategy to make the above allocation algorithm truthful.

In the mechanism design setting, Lavi and Swamy considered a restricted variant, where each task  $j$  only has two values of running time (small time  $L_j$  or big time  $H_j$ ), and gave a 3-approximation randomized truthful mechanism [10]. They first use the cycle monotonicity in designing mechanisms and applied the LP rounding idea based on [9].

For the lower bounds side, Christodoulou, Koutsoupias and Vidali improved the lower bound from 2 to  $1 + \sqrt{2}$  when the number of machines is at least 3 [3], and then to 2.618 when the number of machines is sufficiently large [8]. In a recent beautiful work by Ashlagi, Dobzinski and Lavi, an optimal lower bound ( $m$ ) was proved for all anonymous truthful mechanisms [1].

Christodoulou, Koutsoupias and Vidali gave a characterization for all truthful mechanisms in the same setting as this paper, including those with unbounded approximation ratio [4].

In [2], Christodoulou, Koutsoupias and Kovács considered the fractional version of this problem, in which each task can be split among the machines. For this version, they gave a lower bound of  $2 - 1/m$  and an upper bound of  $(m + 1)/2$ , where  $m$  is the number of machines. We remark that these two bounds are closed for the case of two machines as in the integral deterministic version. So to explore the exact bound for the randomized version seems very interesting and desirable.

## 2 Notations and Preliminaries

In this section we review some definitions and results on mechanism design and the scheduling problem. In the following, for a generic matrix  $\mathbf{a} = (a_{ij})$ , we use  $a_i$  to denote the  $i$ -th row of the matrix, and  $\mathbf{a}_{-i}$  to denote the matrix obtained



from  $\mathbf{a}$  deleting  $\mathbf{a}_i$ . We also use  $(\mathbf{v}, \mathbf{a}_{-i})$  to denote the matrix obtained from  $\mathbf{a}$  by replacing  $\mathbf{a}_i$  with vector  $\mathbf{v}$ . We use  $\mathbb{R}_+$  to denote the set of non-negative real numbers.

In a scheduling problem, there are  $n$  tasks and  $m$  machines (in this paper, we mainly consider the case where  $m = 2$ ), where each machine  $i \in [m]$  needs  $t_{ij}$  units of time to perform task  $j \in [n]$ . We usually use the matrix  $\mathbf{t} = (t_{ij})$  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  $\mathbf{t}_i$ , and we call the vector  $\mathbf{t}_i$  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 the least amount of tasks as possible, so players may misreport their types. We use  $\mathbf{b}_i \in \mathbb{R}_+^n$  to denote player  $i$ 's reported type, and call it player  $i$ 's bid. Obviously  $\mathbf{b}_i$  may not equal to  $\mathbf{t}_i$  if that helps 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  $\mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_m)$ , outputs an allocation denoted by a matrix  $\mathbf{x} = (x_{ij})$ .  $x_{ij}$  is 1 if task  $j$  is assigned to machine  $i$ , and 0 otherwise. Every task must be completely assigned, hence  $\sum_{i \in [m]} x_{ij} = 1, \forall j \in [n]$ .
- **A payment algorithm:** the payment algorithm  $P$ , given the input of players bid matrix  $\mathbf{b}$ , outputs a vector  $\mathbf{p} = (p_1, \dots, p_m)$ , where  $p_i$  denotes the money that player  $i$  receives from the mechanism.

A mechanism is deterministic if both its allocation and payment algorithms are deterministic. If at least one of the algorithms uses random bits, the mechanism is called randomized.

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  $\mathbf{t}_i$  over an allocation  $\mathbf{x}$  and money  $p_i$  is defined as:

$$u_i(\mathbf{x}, p_i | \mathbf{t}_i) = p_i - \sum_{j \in [n]} x_{ij} t_{ij}.$$

In deterministic mechanisms, both  $\mathbf{x}$  and  $p_i$  are functions of bid matrix  $\mathbf{b}$ , we can also write the utility as

$$u_i(\mathbf{b} | \mathbf{t}_i) = p_i(\mathbf{b}) - \sum_{j \in [n]} x_{ij}(\mathbf{b}) t_{ij}.$$

Recalling 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  $\mathbf{b}_{-i}$  of all other players, we have

$$u_i((\mathbf{t}_i, \mathbf{b}_{-i}) | \mathbf{t}_i) \geq u_i((\mathbf{b}_i, \mathbf{b}_{-i}) | \mathbf{t}_i), \quad \forall \mathbf{b}_i \in \mathbb{R}_+^n$$

In randomized mechanisms, both  $x_{ij}$  and  $p_i$  are random variables. There are two versions of truthfulness for randomized mechanisms. 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. In this paper, we focus on the stronger version, universally truthful.

For a truthful mechanism  $M = (X, P)$ , we may assume that all the players will report their true types, hence  $\mathbf{b} = \mathbf{t}$ . Now, how can we evaluate the performance of the mechanism's allocation algorithm  $X$ ? We consider the makespan, which is the maximum load of all the machines. Given input  $\mathbf{t}$ , the makespan of mechanism  $M$  is denoted by  $l_M(\mathbf{t})$ , and  $l_M(\mathbf{t}) = \max_{i \in [m]} \sum_{j \in [n]} x_{ij} t_{ij}$ . We use  $l_{opt}(\mathbf{t})$  to denote the optimum, and  $l_{opt}(\mathbf{t}) = \min_x \max_{i \in [m]} \sum_{j \in [n]} x_{ij} t_{ij}$ . A mechanism  $M$  is called a  $c$ -approximation mechanism if for any instance  $\mathbf{t}$ , we have  $l_M(\mathbf{t}) \leq c \cdot l_{opt}(\mathbf{t})$ . For randomized mechanism  $M$ , we require  $E[l_M(\mathbf{t})] \leq c \cdot l_{opt}(\mathbf{t})$ , where the expectation is over the random bits used in the mechanism.

**Definition 1** (Task-Independent Mechanisms). *A deterministic mechanism  $M$  is task-independent, if for any two bid matrices  $\mathbf{b}, \mathbf{b}'$  such that  $b_{ij} = b'_{ij}$  for all  $i \in [m]$ , the allocation of task  $j$  is identical, i.e.  $x_{ij}(\mathbf{b}) = x_{ij}(\mathbf{b}'), \forall i \in [m]$ .*

For randomized mechanisms, there are also two versions of task-independence. One is a weak task-independent mechanism, which is a distribution of several task-independent deterministic mechanisms. The other is a strong task-independent mechanism, which satisfies that not only does the allocation of task  $j$  not change as long as  $j$ 's column of  $\mathbf{b}$  does not change, but also all the random variables  $x_{ij}$  are independent between different tasks.

We quote a theorem from [6] (Theorem 4.5 in [6]), which gives a characterization for truthful mechanisms for scheduling two machines.

**Theorem 1** ([6]). *Let  $M$  be a mechanism for minimizing the makespan for 2 machines that provides a finite approximation ratio. Then  $M$  is task independent.*

This theorem implies that any randomized truthful mechanism with a finite approximation ratio is weak task-independent. In [13], Lu and Yu proved a lower bound of  $\frac{11}{7}$  ( $\approx 1.5714$ ) for all the *strong task-independent* truthful mechanisms. Given these two facts, we have the following interesting open question:

*Question 1.* Does there exist a weak task-independent randomized truthful mechanism which provides a better approximation ratio ( $< \frac{11}{7}$ )?

**Definition 2** (Scale-Free Mechanisms). *We call an allocation algorithm scale-free if for any instance  $\mathbf{b}$  and any non-zero constant  $\lambda$ , the outputs of the algorithm on the input  $\mathbf{b}$  and  $\lambda\mathbf{b}$  are identical. A mechanism is called scale-free if its allocation algorithm is scale-free. A randomized mechanism is called scale-free if it is a distribution of deterministic scale-free mechanisms.*

Together with the properties of scale-free and task-independent, the allocation of a task  $j$  only depends on the ratio of the two bids  $\frac{b_{1j}}{b_{2j}}$  for this task. Then using the monotone theorem of truthful mechanism [15], we have the following characterization of scale-free task-independent truthful mechanisms.

**Theorem 2.** *A deterministic scale-free task-independent truthful mechanism for scheduling two unrelated machines is of the following form: there are  $n$  thresholds  $\alpha_j$  ( $j \in [n]$ ) for  $n$  tasks. For every task  $j \in [n]$ , the mechanism allocates it to the first machine iff  $b_{1j} < \alpha_j b_{2j}$  (or  $b_{1j} \leq \alpha_j b_{2j}$ ).*

For randomized mechanisms, these thresholds are random variables that do not depend on player's bid. In strong task-independent mechanisms, these random variables are further required to be independent. While in weak task-independent mechanisms, they may have some correlation.

Our lower bounds in this paper are proved by Yao's min-max principle [17], which is a typical tool used to prove lower bounds of randomized mechanisms (algorithms, protocols, etc). Based on the characterization Theorem 1. We state the principle in our setting as following.

**Lemma 1.** *Given a distribution of instances, if any (scale-free) task independent deterministic mechanism cannot have an expected approximation ratio better than  $\alpha$ , then  $\alpha$  is a lower bound for all the (scale-free) universal truthful randomized mechanisms.*

### 3 A Lower Bound of 1.5625

In this section, we prove a lower bound of 1.5625 for all the randomized scale-free truthful mechanisms. Let  $k$  be an integer and  $a > 1$  be a parameter specified later. We consider the following instances distribution. There are  $k + 1$  instances, each containing  $k + 1$  tasks. All the instances have equal probability, i.e. a probability of  $\frac{1}{k+1}$ . The  $i$ -th ( $1 \leq i \leq k + 1$ ) instance is as following: the running times of the  $i$ -th task are  $ka$  and  $ka^2$  for the first and second machines respectively; and the running times of the other  $k$  tasks are 1 and  $a$  for the first and second machines respectively.

For the  $i$ -th instance, the optimal solution is to allocate the  $i$ -th task to the first machine and the remaining  $k$  tasks to the second machine; the optimal makespan is  $ka$  for every instance. Now we consider the performance of a deterministic scale-free task independent mechanism on these instances. Since in every instance and for every task, the ratio of two running times is the same (i.e. equal to  $a$ ), every deterministic scale-free task-independent mechanism will allocate the same task (tasks with the same number) in different instances in the same way. This means that if the mechanism assigns task 1 to the first machine in the first instance, then it must assigns task 1 to the first machine in all the instances. Now we assume that the mechanism assign  $t$  tasks in the first instance to the first machine, then the behavior of the mechanism on all these instances is completely fixed. By the symmetry of the tasks, w.o.l.g, we can assume that the mechanism assigns the first  $t$  tasks to the first machine. Now we can calculate the expected approximation ratio of the mechanism on this distribution of instances.

For the first  $t$  instances, the makespan is the load of the first machine, which is  $ka + (t - 1) \times 1 = ka + t - 1$ . For the other  $k + 1 - t$  instances, the makespan is the load of the second machine, which is  $ka^2 + (k + 1 - t - 1)a = ka^2 + (k - t)a$ .

Therefore the expected approximation ratio  $R$  of the mechanism on these instances is

$$\frac{t(ka+t-1)}{ak(k+1)} + \frac{(k+1-t)(ka^2+(k-t)a)}{ak(k+1)} = \frac{a+1}{ak(k+1)}(t^2 - (ak+1)t + a(k^2+k)).$$

For any fixed  $k$  and  $a > 1$ , this value  $R$  is a quadratic polynomial of  $t$ . So we have

$$R \geq \frac{a+1}{ak(k+1)}(a(k^2+k) - \frac{(ak+1)^2}{4}).$$

For a sufficiently large  $k$ , the ratio in the RHS approaches the ratio of the  $k^2$  terms which is  $a+1 - \frac{a(a+1)}{4}$ . When  $a = \frac{3}{2}$ , this expression reaches its maximum value  $\frac{25}{16} = 1.5625$ .

By Yao's min-max principle, this instance's distribution gives a lower bound of 1.5625. We remark that this lower bound only occurs for a sufficiently large number of tasks.

**Theorem 3.** *Any randomized scale-free truthful mechanism for scheduling two unrelated machines can not have an approximation ratio which is better than 1.5625.*

## 4 Correlation Gives Better Mechanisms

In this and the next sections, we study a further restricted case, namely scheduling two tasks on two unrelated machines. This seems a very special setting, but we believe it is still very interesting for several seasons. First, we will prove that previous lower bounds (1.5 in general and 1.57 for strong task-independent mechanisms) both hold even for this special case. Second, from a pure mathematical point of view, this is the simplest non-trivial setting, however the exact bound for this simplest case is still unclear. Third, the techniques and ideas developed here for studying this special setting may extend to more general settings. For example, the characterization in [4] is first proved for the 2 task case and then extends to many tasks.

The proof for the lower bound of 1.5 in [14] requires at least 3 tasks. Here we prove that this is also true for two tasks.

**Lemma 2.** *Any randomized truthful mechanisms for scheduling two tasks on two machines cannot have an approximation ratio that is better than 1.5.*

*Proof.* We consider a distribution of the following two instances, each with probability of  $\frac{1}{2}$ .

	task 1	task 2
machine 1	1	1
machine 2	1	2

	task 1	task 2
machine 1	1	2
machine 2	1	1

Any deterministic task-independent mechanism will assign task 1 in these two instances in the same way. By symmetry, we can assume that the mechanism assigns task 1 to machine 1. Then the makespan of the mechanism for the first instance is at least 2 no matter how it allocates task 2. However, the optimal makespan is 1; therefore, the expected approximation ratio of this mechanism on these two instances is at least  $\frac{2+1}{2} = 1.5$ . By Yao's min-max principle, 1.5 is a lower bound for all the randomized truthful mechanisms.

The proof for a lower bound of  $\frac{11}{7}$  in [13] only uses instances with 2 tasks, so this bound also holds for the this special setting.

**Lemma 3.** *Any strong task-independent randomized truthful mechanism scheduling two tasks on two machines cannot have an approximation ratio that is better than  $\frac{11}{7}$  ( $\approx 1.5714$ ).*

Given these two lower bounds. It is interesting to see if this bound of  $\frac{11}{7}$  can be beaten. The answer is affirmed by the following correlation mechanism, which also partially answers Question 1.

Let  $f : \mathbb{R}^+ \rightarrow [0, 1]$  be a non-decreasing monotone function, where  $\mathbb{R}^+ = \{x \in \mathbb{R} | x \geq 0\}$ ,  $f(0) = 0$ ,  $\lim_{x \rightarrow \infty} f(x) = 1$  and  $f(x) + f(1/x) = 1$ . The correlation mechanism for scheduling two tasks on two machines is described in Figure 1.

<p><b>Input:</b> The reported bid matrix <math>b</math>.</p> <p><b>Output:</b> A randomized allocation <math>x</math> and a payment <math>p = (p_1, p_2)</math>.</p> <p><b>Allocation and Payment Algorithm:</b></p> <p><math>x_{1j} \leftarrow 0, x_{2j} \leftarrow 0, j = 1, 2</math>.</p> <p><math>p_1 \leftarrow 0; p_2 \leftarrow 0</math>.</p> <p>Choose <math>\alpha \in \mathbb{R}^+</math> randomly according to function <math>f</math>.</p> <p>if <math>b_{11} &lt; \alpha b_{21}</math>,</p> <p style="padding-left: 2em;"><math>x_{11} \leftarrow 1, p_1 \leftarrow p_1 + \alpha b_{21}</math>;</p> <p>else</p> <p style="padding-left: 2em;"><math>x_{21} \leftarrow 1, p_2 \leftarrow p_2 + \alpha^{-1} b_{11}</math>.</p> <p>if <math>b_{22} &lt; \alpha b_{12}</math>,</p> <p style="padding-left: 2em;"><math>x_{22} \leftarrow 1, p_2 \leftarrow p_2 + \alpha b_{12}</math>;</p> <p>else</p> <p style="padding-left: 2em;"><math>x_{12} \leftarrow 1, p_1 \leftarrow p_1 + \alpha^{-1} b_{22}</math>.</p>
--

**Fig. 1.** The Correlation Mechanism

It is easy to show that this mechanism is universally truthful for any function  $f$  with the properties listed above. When the random variable  $\alpha$  is fixed, it is a task-independent mechanism and for each task it is simply a weighted VCG mechanism. The main new idea in this mechanism is that there are some correlation of randomness for different tasks. Here the random variable  $\alpha$  is used both in the mechanisms for the first task and the second task. The intuitive

argument is like this. If  $\alpha > 1$ , then the mechanism is biased to the first machine for the first task and to the second machine for the second task. If  $\alpha < 1$ , it is the other way round. The different bias for different tasks makes the allocation more balanced. Here the requirement of  $f(x) + f(1/x) = 1$  makes the mechanism symmetrical for the two machines and for the two tasks.

Now we analyze the performance formally. We consider the following generic instance.

	task 1	task 2
machine 1	$b_{11}$	$b_{12}$
machine 2	$b_{21}$	$b_{22}$

The expected makespan  $t$  of the correlation mechanism on this instance is

$$t = (b_{11} + b_{12})Pr(\alpha > \frac{b_{11}}{b_{21}}, \alpha \leq \frac{b_{22}}{b_{12}}) + (b_{21} + b_{22})Pr(\alpha \leq \frac{b_{11}}{b_{21}}, \alpha > \frac{b_{22}}{b_{12}}) + \max(b_{11}, b_{22})Pr(\alpha > \frac{b_{11}}{b_{21}}, \alpha > \frac{b_{22}}{b_{12}}) + \max(b_{12}, b_{21})Pr(\alpha \leq \frac{b_{11}}{b_{21}}, \alpha \leq \frac{b_{22}}{b_{12}}).$$

Since all the probabilities in the above expression can be expressed by function values of  $f$ , its performance can be estimated at least numerically (and by computer) when the function is given. Here we specify the following simple function  $f$  so that the analysis can be done analytically (and by hand). It is a case-by-case analysis and is omitted here due to space limitation.

$$f(x) = \begin{cases} 1, & x \geq A, \\ \frac{1}{2} + \frac{x-1}{2(A-1)}, & 1 \leq x < A, \\ \frac{1}{2} - \frac{\frac{1}{x}-1}{2(A-1)}, & \frac{1}{A} \leq x < 1, \\ 0, & 0 \leq x < \frac{1}{A}. \end{cases} \tag{1}$$

Despite the complicated appearance in the above expression, this function is a simple and natural one.  $A$  is a threshold, when input is beyond that, the function value is always 1.  $f(x)+f(1/x) = 1$  requires that  $f(1) = \frac{1}{2}$ . The function between 1 and  $A$  is simply the the line segment connecting these two end points  $(1, \frac{1}{2})$  and  $(A, 1)$ . The function below 1 is determined by the function above 1 and the requirement  $f(x) + f(1/x) = 1$ .

**Theorem 4.** *By using the function as (1), where  $A = -\frac{1}{2} + \sqrt{3} + \frac{1}{2}\sqrt{25 - 12\sqrt{3}}$  ( $\approx 2.26$ ), the approximation ratio of the Correlation Mechanism is  $\frac{1}{6}(\sqrt{25 - 12\sqrt{3}} + 7)$  ( $\approx 1.5089$ ).*

We remark that the function of (1) is only used to illustrate the idea of correlation mechanisms. It is by no means the best choice. However its bound (1.5089) is already very close to the lower bound (1.506) we will prove in the next section.

## 5 The Ratio of 1.5 Is Not Achievable

Given the success of using correlation in the previous section, and also noticing that the instances used to prove the lower bound in Section 3 cannot get anything beyond 1.5 for the case of two tasks, one may make a point that we can choose some suitable function  $f$  in the correlation mechanism to achieve an approximation ratio of exactly 1.5. In this section, we prove that this is impossible.

**Theorem 5.** *For any non-decreasing monotone function  $f : \mathbb{R}^+ \rightarrow [0, 1]$  which satisfies  $\forall x \in \mathbb{R}^+, f(x) + f(1/x) = 1$ , there exists  $c, d \in \mathbb{R}^+$  such that  $c \geq 1$ ,  $c \geq d$  and*

$$c + f(c) + df(c) - cf(c) - df(d) > 1.5.$$

*Proof.* We assume for contradiction that there exists a function  $f$  such that for all  $c, d \in \mathbb{R}^+$  satisfying  $c \geq 1$ ,  $c \geq d$ , we have

$$c + f(c) + df(c) - cf(c) - df(d) \leq 1.5. \tag{2}$$

For any fixed  $d$ , let  $c \rightarrow \infty$ , we have  $c + f(c) + df(c) - cf(c) \geq 1 + d$ . So for any  $x \in \mathbb{R}^+$ , we have

$$f(x) \geq 1 - \frac{1}{2x}.$$

Using this and the fact that  $\forall x \in \mathbb{R}^+, f(x) + f(1/x) = 1$ , we have

$$f(x) \leq \frac{x}{2}.$$

Let  $d = 1/c$  in (2), we have

$$c + f(c) + \frac{f(c)}{c} - cf(c) - \frac{1}{c}(1 - f(c)) \leq 1.5.$$

This implies

$$(c - 2)(c + 1)f(c) \geq (c - 2)(c + 1/2).$$

So for  $c > 2$  we have  $f(c) \geq 1 - \frac{1}{2(c+1)}$  and for  $1 \leq c < 2$  we have  $f(c) \leq 1 - \frac{1}{2(c+1)}$ . Together with the fact that  $f$  is a non-decreasing monotone, these two inequalities enforce that  $f(2) = 1 - \frac{1}{2 \times (2+1)} = \frac{5}{6}$ .

Now let  $c = 2$  in (2), we have

$$2 + \frac{5}{6} + \frac{5}{6}d - 2 \times \frac{5}{6} - df(d) \leq 1.5.$$

This implies that for any  $x \leq c = 2$ ,  $f(x) \geq \frac{5}{6} - \frac{1}{3x}$ . But  $f$  cannot simultaneously satisfy this and  $f(x) \leq \frac{x}{2}$ . For example choosing  $x = \frac{4}{5}$ ,  $f(\frac{4}{5}) \geq \frac{5}{6} - \frac{1}{3x} = \frac{5}{12}$ , but on the other hand  $f(\frac{4}{5}) \leq \frac{x}{2} = \frac{2}{5} < \frac{5}{12}$ , a contradiction.

We can improve this theorem by proving that any scale-free mechanisms cannot achieve 1.5 even for this very special cases.

**Theorem 6.** *Any randomized scale-free truthful mechanism for scheduling two tasks on two unrelated machines cannot have an approximation ratio that is better than 1.506.*

*Proof.* We use a matrix  $\begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix}$  to denote the following instance with two machines and two tasks

	task 1	task 2
machine 1	$b_{11}$	$b_{12}$
machine 2	$b_{21}$	$b_{22}$

Now we consider the following distribution of 12 instances:

- $\begin{bmatrix} 1 & 1 \\ a & 1 \end{bmatrix}$ ,  $\begin{bmatrix} 1 & 1 \\ 1 & a \end{bmatrix}$ ,  $\begin{bmatrix} a & 1 \\ 1 & 1 \end{bmatrix}$ ,  $\begin{bmatrix} 1 & a \\ 1 & 1 \end{bmatrix}$ , each with probability  $p_1$ ;
- $\begin{bmatrix} 1 & b \\ \frac{1}{b} & 1 \end{bmatrix}$ ,  $\begin{bmatrix} 1 & \frac{1}{b} \\ b & 1 \end{bmatrix}$ ,  $\begin{bmatrix} b & 1 \\ 1 & \frac{1}{b} \end{bmatrix}$ ,  $\begin{bmatrix} \frac{1}{b} & 1 \\ 1 & b \end{bmatrix}$ , each with probability  $p_2$ ;
- $\begin{bmatrix} 1 & b \\ \frac{1}{a} & 1 \end{bmatrix}$ ,  $\begin{bmatrix} 1 & \frac{1}{a} \\ b & 1 \end{bmatrix}$ ,  $\begin{bmatrix} b & 1 \\ 1 & \frac{1}{a} \end{bmatrix}$ ,  $\begin{bmatrix} \frac{1}{a} & 1 \\ 1 & b \end{bmatrix}$ , each with probability  $p_3$ .

$a, b, p_1, p_2, p_3$  are parameters to be specified later and satisfy  $1 \leq b \leq 2 \leq a$  and  $p_1 + p_2 + p_3 = 1/4$ .

For these instances, the possible running-time-ratios of tasks are only  $1/a, 1/b, 1, a, b$ . For each task, the mechanism can choose a threshold (only 6 possible different thresholds). But by symmetry, we can always assume that the thresholds for the first task are above 1 (so there are 3 possible different thresholds). Overall of there are 18 possible different mechanisms. We can choose the parameters such that the expected approximation ratios of them are all larger than a given value, then this is our lower bound.

By choosing  $a = 2.125, b = 1.88, p_1 = 0.1346, p_2 = 0.0796, p_3 = 0.0358$ , we can get a lower bound of 1.506. We can prove this formally and also argue that these are the best parameters we can choose. The details are omitted.

## 6 Conclusion and Discussion

The main results of this paper are two new lower bounds and one new upper bound. Two direct interesting open questions are to get rid of the technical assumption for these lower bounds and to generalize the correlation mechanism to general cases. It is quite surprising that the exact bound for this simple 2-player mechanism has not been settled after a couple of work. We recall that quite simple mechanisms and relatively easy lower bound proofs already match both in the corresponding deterministic and fraction version. We believe that our work in this paper is an important step toward the final answer.

In the general case ( $m$  machines), the gap between the best lower bounds (constants) and the best upper bounds ( $\Theta(m)$ ) is huge both in deterministic and randomized versions. Any improvement in either direction is highly desirable. We hope that the technique and ideas we and others developed for this special case can extend to the general case.



## Acknowledgement

I am grateful to Wei Chen, Yajun Wang, and Changyuan Yu for helpful discussion on this topic.

## References

1. Ashlagi, I., Dobzinski, S., Lavi, R.: An Optimal Lower Bound for Anonymous Scheduling. In: Proceedings of EC 2009, pp. 169–176 (2009)
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. Christodoulou, G., Koutsoupias, E., Vidali, A.: A lower bound for scheduling mechanisms. In: The Proceedings of SODA 2007, pp. 1163–1170 (2007)
4. Christodoulou, G., Koutsoupias, E., Vidali, A.: A Characterization of 2-Player Mechanisms for Scheduling. In: Halperin, D., Mehlhorn, K. (eds.) Esa 2008. LNCS, vol. 5193, pp. 297–307. Springer, Heidelberg (2008)
5. Clarke, E.H.: Multipart Pricing of Public Goods. *Public Choice* 11, 17–33 (1971)
6. Dobzinski, S., Sundararajan, M.: On Characterizations of Truthful Mechanisms for Combinatorial Auctions and Scheduling. In: Proceedings of EC 2008, pp. 38–47 (2008)
7. Groves, T.: Incentives in Teams. *Econometrica* 41, 617–631 (1973)
8. Koutsoupias, E., Vidali, A.: A Lower Bound of  $1 + \phi$  for Truthful Scheduling Mechanisms. In: Kučera, L., Kučera, A. (eds.) MFCS 2007. LNCS, vol. 4708, pp. 454–464. Springer, Heidelberg (2007)
9. Lavi, R., Swamy, C.: Truthful and Near-Optimal Mechanism Design via Linear Programming. In: Proceedings of FOCS 2005, pp. 595–604 (2005)
10. Lavi, R., Swamy, C.: Truthful mechanism design for multi-dimensional scheduling via cycle-monotonicity. In: Proceedings EC 2007, pp. 252–261 (2007)
11. Lenstra, J.K., Shmoys, D.B., Tardos, É.: Approximation algorithms for scheduling unrelated parallel machines. *Math. Program.* 46(3), 259–271 (1990)
12. Lu, P., Yu, C.: An Improved Randomized Truthful Mechanism for Scheduling Unrelated Machines. In: Proceedings of STACS 2008, pp. 527–538 (2008)
13. Lu, P., Yu, C.: Randomized Truthful Mechanisms for Scheduling Unrelated Machines. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 402–413. Springer, Heidelberg (2008)
14. Mu’alem, A., Schapira, M.: Setting lower bounds on truthfulness. In: Proceedings of SODA 2007, pp. 1143–1152 (2007)
15. Nisan, N., Ronen, A.: Algorithmic mechanism design (extended abstract). In: Proceedings of STOC 1999, pp. 129–140 (1999)
16. Vickrey, W.: Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 8–37 (1961)
17. Yao, A.: Probabilistic computations: Toward a unified measure of complexity. In: Proceedings of FOCS 1977, pp. 222–227 (1977)

# On Stackelberg Pricing with Computationally Bounded Consumers<sup>★</sup>

Patrick Briest<sup>1,★★</sup>, Martin Hoefer<sup>2,★★★</sup>, Luciano Guala<sup>3</sup>, and Carmine Ventre<sup>4,†</sup>

<sup>1</sup> Department of Computer Science, University of Paderborn, Germany  
patrick.briest@upb.de

<sup>2</sup> Lehrstuhl Informatik I, RWTH Aachen University, Germany  
mhoefer@cs.rwth-aachen.de

<sup>3</sup> Dipartimento di Matematica, Università di Roma “Tor Vergata”, Italy  
guala@mat.uniroma2.it

<sup>4</sup> Department of Computer Science, University of Liverpool, UK  
Carmine.Ventre@liverpool.ac.uk

**Abstract.** In a Stackelberg pricing game a *leader* aims to set prices on a subset of a given collection of items, such as to maximize her revenue from a *follower* purchasing a feasible subset of the items. We focus on the case of computationally bounded followers who cannot optimize exactly over the range of all feasible subsets, but apply some publicly known algorithm to determine the set of items to purchase. This corresponds to general multi-dimensional pricing assuming that consumers cannot optimize over the full domain of their valuation functions but still aim to act rationally to the best of their ability.

We consider two versions of this novel type of Stackelberg pricing games. Assuming that items are weighted objects and the follower seeks to purchase a min-cost selection of objects of some minimum weight (the MIN-KNAPSACK problem) and uses a simple greedy 2-approximate algorithm, we show how an extension of the known single-price algorithm can be used to derive a polynomial-time  $(2 + \varepsilon)$ -approximation algorithm for the leader’s revenue maximization problem based on so-called *near-uniform* price assignments. We also prove the problem to be strongly NP-hard.

Considering the case that items are subsets of some ground set which the follower seeks to cover (the SET-COVER problem) via a standard primal-dual approach, we prove that near-uniform price assignments fail to yield a good approximation guarantee. However, in the special case of elements with frequency 2 (the VERTEX-COVER problem) it turns out that exact revenue maximization can be done in polynomial-time. This stands in sharp contrast to the fact that revenue maximization becomes APX-hard already for elements with frequency 3.

## 1 Introduction

The problem of *multi-dimensional pricing* consists of assigning revenue maximizing prices to a set of distinct items given information about the preferences of potential

---

\* Work supported by DFG grant Kr 2332/1-2 within Emmy Noether Program.

\*\* Work done in part while the author was at Cornell University, supported by a scholarship of the German Academic Exchange Service (DAAD).

\*\*\* The author is supported by UMIC Research Center at RWTH Aachen University.

† Work also supported by EPSRC grant EP/F069502/1 and EU through IP AEOLUS.

customers. A natural way to describe consumer preferences is via *valuation functions* that map subsets of items to non-negative real numbers describing how much a set is valued by a certain customer. Given fixed item prices a rational customer acting according to *quasi-linear utilities* chooses to purchase the subset of items maximizing her *utility*, i.e., the difference between her value for the set and the sum of prices of items contained in it. It is known that general multi-dimensional pricing with unlimited supply of each item allows polynomial-time approximation algorithms achieving ratios that are logarithmic in the number of customers or linear in the number of distinct items via considering only *single-price solutions* [4,8]. Lower bounds of the same order of magnitude have also been proven for approximation guarantees in both parameterizations [7,14].

For the known algorithmic results it is sufficient that valuation functions can be accessed via *demand oracles*, i.e., consumers are treated as black boxes that can answer the question: “Given some vector of prices, which set of items would you choose to buy?” But how much can it help us to have more detailed information about the structure of customer preferences? If the number of customers is large, unfortunately, the answer is “not at all” as follows readily from the known lower bounds, which hold for special cases of the problem in which the exact preferences can be elicited via demand queries. However, the situation is different when the number of distinct customers is small.

In a 2-player *Stackelberg Pricing Game*, named after the underlying market model due to Heinrich Freiherr von Stackelberg [27], we are given a collection of items, some of which have fixed-costs. A so-called *leader* may assign prices to the remaining items. A *follower* then purchases a *feasible* set of items of minimum cost and pays the leader for the priceable items in the set. This problem is equivalent to multi-dimensional pricing with a single follower if the follower’s feasible sets are unrestricted. However, a standard assumption when considering Stackelberg games is that the follower has to be able to optimize in polynomial time over her feasible sets. As an example, think of items as edges in a graph and the follower’s feasible sets as all possible paths connecting some vertices  $s$  and  $t$ . For some kinds of followers, e.g., buying a min-cost vertex cover in a bipartite graph, it has been shown that improved approximation guarantees are possible [8].

In this paper we initiate the study of a closely related question: What if the follower is unable to exactly optimize over her feasible sets, because the problem is computationally hard, but is still guaranteed to act rationally to the best of her ability? More formally, we will assume that when prices have been fixed the follower applies a publicly known approximation algorithm to find a near-optimal feasible set of items to purchase. This assumption is quite reasonable when followers are actually software agents with known implementation details. To the best of our knowledge, this is the first analysis of multi-dimensional pricing with follower preferences that are neither *single-minded* or *unit-demand*, nor expressible as exact optimization over the full domain of the valuation function. Before describing our results in detail, we review some important related work and introduce the notation used throughout the paper.

**Related Work.** Algorithmic aspects of multi-dimensional pricing problems, which are important in the context of pure optimization as well as the design of revenue-maximizing auction mechanisms [3], were first studied by Aggarwal et al. [1] and Guruswami et al.

[20]. Subsequently, quite a number of improved algorithmic results for special cases of the problem [29][2][3][5][6][21] and complexity theoretic lower bounds [7][4][10] have been derived.

Our introductory example of shortest-path Stackelberg pricing was first introduced by Labbé et al. [23], who derived a bi-level LP formulation of the problem and proved NP-hardness. Subsequently, Roch et al. [24] presented a first polynomial time approximation algorithm with a provable (logarithmic) approximation guarantee. More recently, Cardinal et al. [11] investigated the corresponding minimum spanning tree game, proving that this version of the problem is APX-hard and that the very simple *single-price algorithm* achieves a logarithmic approximation guarantee. Finally, Briest et al. [8] extended the analysis of the single-price algorithm to Stackelberg pricing in general. Stackelberg pricing in which the follower purchases a single-source shortest path tree has been considered in [6].

Stackelberg pricing can also be considered with objectives other than revenue maximization. When prices are tolls on network arcs the problem of congestion minimization has received considerable attention. Karakostas and Kolliopoulos [22], Fleischer, Jain and Mahdian [17], Fleischer [18] and Swamy [25] show how tolls can be used to enforce low-congestion Nash equilibria in selfish network routing games.

**Preliminaries.** In this paper we consider games falling in the following general class. There are two players in the game, one *leader* and one *follower*. There is also a set of items  $\mathcal{I}$  that is partitioned into fixed-cost items  $\mathcal{F}$  and priceable items  $\mathcal{P}$ . Each fixed-cost item  $i \in \mathcal{F}$  has a fixed-cost  $c(i) \geq 0$ . For each priceable item  $i \in \mathcal{P}$  the leader can specify a price  $p(i) \geq 0$ . The follower has a set  $\mathcal{S} \subset 2^{\mathcal{I}}$  of *feasible subsets* and is interested in buying some subset in  $\mathcal{S}$ . The cost of a subset  $S \in \mathcal{S}$  is given by the cost of fixed-cost items and the price of priceable items:  $cost(S) = \sum_{i \in S \cap \mathcal{F}} c(i) + \sum_{i \in S \cap \mathcal{P}} p(i)$ . The *revenue* of the leader from subset  $S$  is given by the prices of the priceable items that are included in  $S$ , that is,  $r(S) = \sum_{i \in S \cap \mathcal{P}} p(i)$ . We let  $S_A(p)$  be the feasible subset in  $\mathcal{S}$  chosen by the follower when she uses polynomial-time algorithm  $A$  given prices  $p$ . Naturally, the follower would like  $A$  to return the minimum-cost subset in  $\mathcal{S}$ , but this could be a hard task to solve in polynomial-time. We capture this intuition by making no assumption on optimality of the algorithm:  $A$  can return a suboptimal subset in  $\mathcal{S}$ . Our interest is to find the pricing function  $p^*$  for the leader that generates maximum revenue when the follower uses algorithm  $A$ , i.e.,  $p^* \in \arg \max_p r(S_A(p))$ . We denote this maximum revenue by  $r^*$ . To guarantee that the revenue is bounded and the optimization problem is non-trivial, we assume that there is at least one feasible subset that is composed only of fixed-cost items and that the follower algorithm outputs it under certain circumstances. Towards this aim, we further assume that for each priceable item there is a threshold price above which no subset including it will be output by the follower algorithm. This last assumption holds for every follower algorithm with bounded approximation ratio.

In the above class of games, we will consider the MIN-KNAPSACK pricing problem and the SET-COVER pricing problem. In the knapsack pricing problem, the set of items is a set of weighted objects  $\mathcal{O}$ . A subset of  $\mathcal{O}$  is feasible if the total weight of the object comprising it is at least a given bound  $W$ . We will refer to the revenue optimization problem for the knapsack pricing problem by STACKKP. In the set cover pricing

problem, following our terminology, given some ground set to be covered, every item corresponds to a given subset of the ground set. A subset of items is feasible for the follower whenever it covers all the elements of the ground set. We will refer to the revenue optimization problem for the set cover pricing problem by `STACKSC` and denote the special case of vertex cover pricing by `STACKVC`.

**Contributions.** The focus of this paper are followers applying approximation algorithms that are (i) structurally simple and (ii) sufficiently suboptimal to ensure that revenue maximization has to take into account the algorithm’s exact structure. We first consider `STACKKP` and assume that the follower uses the well known greedy algorithm (see Section 2) to compute a 2-approximate solution to the minimization version of the knapsack problem she needs to solve. Even though structurally quite simple, this problem seems to capture many of the fundamental problems of Stackelberg pricing with computationally limited followers.

We show that in this case a careful adaptation of the known single-price strategy termed *near-uniform pricing*, which essentially assigns a single price to a subset of items and removes the remaining ones from the market by assigning a sufficiently high price (see Section 2 for a formal definition), can be used to approximate optimal revenue in most of the solution space. Adding some fairly standard enumeration techniques we are able to derive a polynomial-time  $(2 + \varepsilon)$ -approximation for the revenue maximization problem. The main technical difficulty lies in the fact that our analysis needs to argue about the exact computation done by the follower for a given price vector rather than using some global optimality condition. We point out that our algorithm is best possible among all algorithms based on near-uniform price assignments. Finally, we show that the revenue maximization problem in this setting is strongly NP-hard.

We then turn our attention to `STACKSC` and assume that the follower is using the primal-dual schema based approximation algorithm (see Section 3) to find a selection of sets to purchase. We view the problem in its equivalent formulation of `VERTEX-COVER` in hypergraphs and start by investigating the special case of regular `VERTEX-COVER` in standard graphs. We prove that while near-uniform price assignments cannot achieve better than logarithmic approximation guarantee in this case, exploiting the algorithm’s structure nevertheless allows for exact revenue maximization in polynomial time. To the best of our knowledge, this is the second example of polynomial-time revenue maximization being possible for a class of Stackelberg pricing games. Previously, it was shown that games with a follower purchasing a min-cost vertex cover in a bipartite graph in the special case that all priceable vertices are located on one side of the bipartition allows for polynomial-time revenue maximization [8]. It would be very interesting to see whether there is a deeper connection between these two problems. Turning to general hypergraphs it turns out that revenue maximization (`STACKSC`) becomes hard already with edges of cardinality 3 (or elements of frequency 3 in the `SET-COVER` view) and is APX-hard in general. This is quite surprising given that the follower’s primal-dual algorithm achieves approximation guarantee  $f$  for any frequency  $f$ , i.e., the approximation complexity of the underlying problem scales quite smoothly. We also argue that in this general case neither near-uniform price vectors nor our algorithm from the `VERTEX-COVER` case can guarantee any sub-exponential approximation ratio.

**Knowing the Follower’s Algorithm.** A central assumption in our analysis is that the leader has full knowledge of the follower’s algorithm. This assumption might appear quite strong. It turns out, however, that this non-black-box attitude on the leader’s side is necessary to achieve any reasonable approximation guarantees. Suppose the leader only knows the approximation guarantee of the follower’s algorithm  $A$  and is given black-box-access to it, but no specific details are revealed. In this case it is easy to argue that for both `STACKKP` and `STACKVC` no algorithm can achieve a finite approximation guarantee. We omit the simple proof of the following fact due to space limitations.

**Proposition 1.** *For any constant  $\rho > 1$ , there are instances of `STACKKP` and `STACKVC` in which no leader’s algorithm can achieve a finite approximation guarantee given only information about  $\rho$  and black-box-access to the follower’s  $\rho$ -approximation algorithm.*

Similarly, it is necessary to assume that the follower decides on the algorithm to be used in advance. If the follower is allowed to choose the algorithm (from a known set of alternatives) once the leader has set the prices, then an impossibility result similar to the one above applies.

## 2 Knapsack Pricing

In the `MIN-KNAPSACK` problem we are given a set  $\mathcal{O}$  of  $n$  objects, some of them with fixed cost and some priceable. Each object  $o \in \mathcal{O}$  has weight  $w(o) \in \mathbb{N}$  and we are given an integer weight bound  $W$ . Following the general framework given above, each subset  $X$  of  $\mathcal{O}$  has a cost which is defined as the sum of the cost of the fixed-cost objects in  $X$  and the prices of the priceable objects in  $X$ . The follower wants to purchase a set of objects of minimum cost whose weight is at least  $W$ . We assume that the follower uses the standard greedy algorithm outlined below to find an approximation of such a minimum-cost set.

**The Follower’s Algorithm.** An object’s cost-efficiency (below referred to as efficiency for brevity) is defined as  $\phi(o) = c(o)/w(o)$  or  $\phi(o) = p(o)/w(o)$ , depending on whether it is fixed-cost or priceable. Algorithm [1](#) below proceeds as follows. First, order all objects by non-decreasing efficiency (breaking ties by decreasing weight). Then add objects to the knapsack in this order. If an object makes the weight of the solution it completes at least  $W$ , remember this (feasible) solution and discard the object. Finally, return the best solution found. Note that we assume that ties are broken according to decreasing weight, i.e., larger objects are considered first given identical efficiency. This is a natural tie breaking rule, as it aims at minimizing the *overlap* of objects that exceed the knapsack’s remaining capacity when they are considered.

**Transforming the Optimal Solution.** Let  $p^*$  be the optimal price assignment and  $\mathcal{P}^*$  be the set of priceable objects that are selected by Algorithm [1](#) given these prices.

The key ingredient for our approximation algorithm for knapsack pricing is the observation that price assignments that result in a large number of priceable objects being bought by the follower can be approximated by almost uniform price assignments at the expense of reducing overall revenue by no more than a constant factor.

**Algorithm 1.** The greedy approximation algorithm for MIN-KNAPSACK

---

Let  $o_1, o_2, \dots, o_n$  be the objects ordered by non-decreasing efficiency, i.e.,

$$\phi(o_1) \leq \dots \leq \phi(o_n).$$

$$X \leftarrow Y \leftarrow \emptyset.$$

$$c_Y \leftarrow +\infty.$$

**for**  $i = 1, \dots, n$  **do**

$$X \leftarrow X \cup \{o_i\}.$$

**if**  $w(X) \geq W$  **then**

**if**  $\text{cost}(X) < c_Y$  **then**

$$Y \leftarrow X.$$

$$c_Y \leftarrow \text{cost}(X).$$

$$X \leftarrow X \setminus \{o_i\}.$$

**Return**  $Y$ .

---

**Definition 1.** We call a price assignment  $p$  near-uniform, if there exists a single efficiency  $\phi > 0$ , such that  $p(o) = w(o) \cdot \phi$  or  $p(o) = +\infty$  for every  $o \in \mathcal{P}$ .

We call an object  $b$  blocking, if the weight of the current solution exceeds  $W$  when it is added by the greedy algorithm. Let  $\mathcal{B} = \{b_1, \dots, b_l\}$  denote the blocking objects with  $\phi(b_1) \leq \dots \leq \phi(b_l)$ . Since every blocking object corresponds to a unique solution checked by the greedy algorithm, our approach to relating the algorithm's behavior on two different price vectors is to relate the sets of blocking objects in both cases.

**Theorem 1.** Let  $p^*$  be the optimal price assignment for a given STACKKP instance and  $r^*$  the resulting revenue. Furthermore, assume that given prices  $p^*$  the follower purchases at least  $k \in \mathbb{N}$  priceable objects, i.e.,  $|\mathcal{P}^*| \geq k$ . Then there exists a near-uniform price assignment  $\tilde{p}$  with revenue at least  $r^*(k-1)/(2k)$ .

*Proof.* Define  $w^* = \sum_{o \in \mathcal{P}^*} w(o)$ ,  $r^* = \sum_{o \in \mathcal{P}^*} p^*(o)$  and let  $\phi_{ave}^* = r^*/w^*$ . Let  $c^*$  denote the total cost of the solution bought by the follower given prices  $p^*$ . We define a near-uniform price assignment  $\tilde{p}$  by  $\tilde{p}(o) = (1/2)w(o)\phi_{ave}^*$  for all  $o \in \mathcal{P}^*$ , and  $\tilde{p}(o) = +\infty$  else.

There is a one-to-one correspondence between the blocking objects  $\mathcal{B} = \{b_1, \dots, b_l\}$  given prices  $p^*$  and solutions checked by the greedy algorithm. Because of the fact that blocking objects do not influence the set of solutions checked by the greedy algorithm (beyond the one they belong to themselves), it is w.l.o.g. to assume that there is at most a single priceable blocking object given prices  $p^*$ , and if it exists it belongs to  $\mathcal{P}^*$ .

Proving the claimed revenue guarantee for  $\tilde{p}$  consists of two parts. First, we show that with prices  $\tilde{p}$  the blocking objects of efficiency less than  $\phi_{ave}^*/2$  are still blocking with with prices  $\tilde{p}$  and their corresponding solutions have cost greater than  $c^* - (1/2)r^*$ . We then show that among the solutions with blocking objects of efficiency greater or equal than  $\phi_{ave}^*/2$  there exist some with cost at most  $c^* - (1/2)r^*$  and the cheapest among these contains priceable objects of value at least  $r^*(k-1)/(2k)$ .

We first deal with blocking objects with efficiency  $\phi(b_j) < \phi_{ave}^*/2$  and argue that they remain blocking. Consider  $b_j$  with  $\phi(b_j) = c(b_j)/w(b_j) < \phi_{ave}^*/2$ . Since we have at most one blocking priceable object (of efficiency at least  $\phi_{ave}^*$ ),  $b_j$  must be



fixed-cost. Given prices  $p^*$ , let  $s_{<j}$  be the remaining unfilled capacity of the knapsack when  $b_j$  is considered and let  $w_{<j}$  and  $r_{<j}$  denote the weight of priceable objects in the knapsack at this point and their total price, respectively. Similarly, let  $w_{>j}$  and  $r_{>j}$  denote the weight and price of priceable objects from  $\mathcal{P}^*$  that are considered after  $b_j$  and let  $\phi_{>j} = r_{>j}/w_{>j}$  be their average efficiency. Finally, define  $f_{>j}$  to be the total cost of fixed-cost objects in the optimal solution with higher efficiency than  $b_j$ .

We are going to argue that moving all priceable objects to a position behind  $b_j$  (as happens with near-uniform pricing  $\tilde{p}$ ) will not cause  $b_j$  to become non-blocking. First observe that  $r_{<j} < w_{<j}\phi_{ave}^*/2 \leq w^*\phi_{ave}^*/2 = (1/2)r^*$  and, thus, we have that  $w_{>j}\phi_{>j} > (1/2)r^*$ . By the fact that  $b_j$  is not part of the cheapest solution, we know that  $f_{>j} + w_{>j}\phi_{>j} \leq c(b_j)$ . It follows that

$$\begin{aligned} w_{>j}\phi_{>j} &\leq c(b_j) - f_{>j} \leq c(b_j) - (s_{<j} - w_{>j})\phi(b_j), \text{ since } f_{>j} \geq (s_{<j} - w_{>j})\phi(b_j) \\ &= (w(b_j) - s_{<j} + w_{>j})\phi(b_j). \end{aligned}$$

Assume now that  $b_j$  becomes non-blocking if we remove total weight  $w_{<j}$  from the knapsack, i.e.,  $w(b_j) - s_{<j} \leq w_{<j}$ . We may then write that

$$\phi_{>j} \leq \left(1 + \frac{w(b_j) - s_{<j}}{w_{>j}}\right) \phi(b_j) \leq \left(1 + \frac{w_{<j}}{w_{>j}}\right) \phi(b_j).$$

Using that  $w_{>j}\phi_{>j} > (1/2)r^*$  and  $\phi(b_j) < (1/2)\phi_{ave}^*$  we obtain  $\frac{1}{2}r^* < w_{>j}\phi_{>j} \leq (w_{>j} + w_{<j})\phi(b_j) < \frac{1}{2}w^*\phi_{ave}^* = \frac{1}{2}r^*$ , a contradiction. We conclude that  $b_j$  remains a blocking element.

Note, that the fact that blocking objects of efficiency at most  $\phi_{ave}^*/2$  remain blocking also implies that no non-blocking objects with efficiency below  $\phi_{ave}^*/2$  can become blocking, since under prices  $\tilde{p}$  the knapsack contains less total weight when such an object is considered. Also observe that given prices  $\tilde{p}$ , the fact that  $r_{<j} < (1/2)r^*$  implies that every solution found by the greedy algorithm with a blocking object  $b_j$  of efficiency less than  $\phi_{ave}^*/2$  has total cost greater than  $c^* - (1/2)r^*$ .

We continue by considering blocking objects with efficiency  $\phi(b_j) \geq \phi_{ave}^*/2$ . If  $\mathcal{P}^*$  did not contain a blocking object given prices  $p^*$ , then changing prices to  $\tilde{p}$  does not change the set of blocking objects of efficiency at least  $\phi_{ave}^*/2$ . To see this, note that with prices  $\tilde{p}$  the knapsack's remaining capacity is smaller than before at any point, so no previously blocking object can become non-blocking. On the other hand, all the non-blocking objects left enough room to pack the objects from  $\mathcal{P}^*$ , so changing the order in which they are considered cannot create new blocking objects either. It follows that all solutions found by the greedy algorithm with blocking objects of efficiency at least  $\phi_{ave}^*/2$  contain priceable objects with a total price of  $r^*/2$ . The previously cheapest solution is still found and has cost  $c^* - (1/2)r^*$ .

If  $\mathcal{P}^*$  did contain a blocking element with prices  $p^*$ , but does not given prices  $\tilde{p}$ , we consider two cases. If the previously optimal solution is still found by the greedy algorithm, the same argumentation as above guarantees revenue  $r^*/2$ . If it is not, it must be the case that some previously non-blocking element has now become blocking. Consider the first such element  $b_j$ . Since all objects in the knapsack at the time  $b_j$  is considered and  $b_j$  itself were part of the cheapest solution given prices  $p^*$ , we have



again found a solution of total cost at most  $c^* - (1/2)r^*$ . Since all priceable objects are in the knapsack at this point and remain in it, a lower bound of  $r^*/2$  on the revenue follows immediately.

Finally, assume that  $\mathcal{P}^*$  contains a blocking element with both prices  $p^*$  and  $\tilde{p}$ . In this case, the solution with priceable blocking object is guaranteed to have cost at most  $c^* - (1/2)r^*$ . By the algorithm's tie-breaking rule every solution with a blocking object of efficiency larger than  $\phi_{ave}^*/2$  found at a later time contains all but the smallest object from  $\mathcal{P}^*$  and, by the fact that  $|\mathcal{P}^*| \geq k$ , contains priceable objects of total value at least  $(1 - 1/k)w^*(1/2)\phi_{ave}^* = r^*(k - 1)/(2k)$ .  $\square$

---

**Algorithm 2.** A  $(2 + \varepsilon)$ -approximation algorithm for STACKKP
 

---

Let 1 be the minimum,  $\Phi$  the maximum efficiency of fixed-cost objects,  $W$  the knapsack capacity,  $\sqrt{1 + \varepsilon/2} - 1 \geq \delta > 0$ ,  $r_{max} \leftarrow 0$ .

Choose  $k \geq (2(2 + \varepsilon))/(2 + \varepsilon - 2(1 + \delta)^2)$  and let

$\Lambda = \{\phi(o) \mid o \in \mathcal{F}\} \cup \{(1 + \delta)^j \mid j = 0, \dots, \lceil \log_{1+\delta} \Phi \rceil\}$ .

**foreach** set  $S \subseteq \mathcal{P}$  of priceable objects with  $|S| \leq k - 1$  **do**

**foreach** price assignment  $p$  with  $p(o)/w(o) \in \Lambda$  for all  $o \in S$  and  $p(o) = +\infty$  else

**do**

        Let  $r$  be the resulting revenue.

$r_{max} \leftarrow \max\{r_{max}, r\}$ .

**foreach**  $0 \leq i \leq \lceil \log_{1+\delta} \Phi \rceil$  **do**

    Let  $\phi_i = (1 + \delta)^i$ .

**foreach**  $0 \leq j \leq \lceil \log_{1+\delta} W \rceil$  **do**

        Let  $S$  be a set of at least  $k$  priceable objects with total weight between  $(1 + \delta)^j$  and  $(1 + \delta)^{j+1}$ , if such set exists, else  $S = \emptyset$ .

        Set  $p(o) = w(o)\phi_i$  for all  $o \in S$  and  $p(o) = +\infty$  for all other priceable objects.

        Let  $r$  be the resulting revenue.

$r_{max} \leftarrow \max\{r_{max}, r\}$ .

Return  $r_{max}$ .

---

**Approximation Algorithm for Revenue Maximization and Hardness Result.** We are now ready to present our  $(2 + \varepsilon)$ -approximate algorithm for STACKKP. Algorithm 2 proceeds in two stages. First it checks for some given constant  $k \in \mathbb{N}$  all possible price assignments to sets of at most  $k - 1$  priceable objects. Then for each possible weight  $w$ , it finds a set with  $k$  or more priceable objects with total weight (roughly)  $w$  (if such a set exists), and considers all near-uniform price assignments to that set.

Note that in both stages, checking all possible efficiencies cannot be done in polynomial time. Instead we restrict our attention to the efficiencies of the fixed-cost objects plus all powers of  $(1 + \delta)$  for some  $\delta > 0$  to guarantee that we efficiently find a near-optimal price assignment in which all objects are considered by the greedy algorithm in the right order. Similarly, in the second stage we cannot enumerate all possible weights for the objects in our near-uniform price assignments, but again have to settle for powers of  $(1 + \delta)$ . The proof of Theorem 2 is omitted due to space limitations.

**Theorem 2.** Algorithm 2 computes in polynomial time a  $(2 + \varepsilon)$ -approximation for STACKKP.

We mention that we can provide an instance showing that the above analysis is essentially tight. Finally, it turns out that revenue maximization against a follower using Algorithm 1 is strongly NP-hard. Details are left for the full version of this paper.

**Theorem 3.** *STACKKP is strongly NP-hard.*

### 3 Set Cover Pricing

In this section we consider the vertex and set cover pricing problems. In the simplest case the follower wants to purchase a minimum vertex cover of a graph  $G = (V, E)$ . Let  $V = \mathcal{P} \cup \mathcal{F}$ , where as above  $\mathcal{F}$  and  $\mathcal{P}$  denote the set of fixed-cost and priceable vertices respectively. More generally, for set cover we consider an equivalent formulation of vertex cover in hypergraphs. Namely, given the universe and its subsets of the set cover instance, we define an hypergraph where the vertices are the universe subsets and hyperedges are the elements of the universe. In this case there can be hyperedges in  $E$  connecting more than two vertices. We assume that the follower uses the standard primal-dual algorithm to find an approximation to the minimum vertex cover.

**The Follower's Algorithm.** The algorithm iteratively considers uncovered edges and raises budgets at the incident vertices until (at least) one of the vertices becomes tight. The algorithm is given for the case of regular vertex cover as Algorithm 3 for more details see [26, Ch. 15].

---

**Algorithm 3.** Primal-dual algorithm of the follower

---

$\gamma_e \leftarrow 0$  for all edges  $e \in E$   
 For all  $v \in V$  let  $c'(v) \leftarrow c(v)$  if  $v \in \mathcal{F}$  and  $c'(v) \leftarrow p(v)$  otherwise  
 Fix an order of edges  $e_1, \dots, e_m$   
**for**  $i = 1, \dots, m$  **do**  
     Let  $e_i = (u, v)$  and  $\sigma_u \leftarrow c'(u) - \sum_{e: e=(u,v')} \gamma_e$   
      $\sigma_v \leftarrow c'(v) - \sum_{e: e=(u',v)} \gamma_e$   
      $\gamma_e \leftarrow \min(\sigma_u, \sigma_v)$   
 Add every vertex  $v$  with  $c'(v) = \sum_{e: e=(u,v)} \gamma_e$  to the cover.

---

**Revenue Maximization for Vertex Cover.** At first, let us consider the approximation ratio of single-price and near-uniform price assignments. For the latter, as from Definition 1 we have a  $\phi \in \mathbb{R}_0^+$  and  $p(v) \in \{\phi, \infty\}$  for all  $v \in \mathcal{P}$ . It turns out that for these pricings there is a simple logarithmic lower bound (proof omitted).

**Theorem 4.** *Single-price and near-uniform price assignments yield an approximation factor of  $\Omega(\log n)$  for STACKVC, where  $n$  is the number of priceable vertices.*

Instead, we present a natural greedy algorithm to compute optimal prices for the seller. It simulates a run of the primal-dual algorithm and raises prices of the priceable vertices in the same manner as the dual budgets  $\gamma$  are raised by the follower. In this way the algorithm greedily tries to sell a vertex to the follower as soon as she is willing to pay for it.

**Algorithm 4.** Greedy pricing for STACKVC

---

```

 $\gamma_e \leftarrow 0$  for all edges  $e \in E$ 
for each edge  $e = e_1, e_2, \dots, e_m$  in the order of the follower do
  if  $e = (u, v)$  is incident to a priceable vertex  $v$  then
     $\gamma_e \leftarrow c(u) - \sum_{e': e'=(u, v')} \gamma_{e'}$ 
     $p(v) \leftarrow \sum_{e': e'=(u, v)} \gamma_{e'}$ 
  else
     $\sigma_u \leftarrow c(u) - \sum_{e': e'=(u, v')} \gamma_{e'}$ 
     $\sigma_v \leftarrow c(v) - \sum_{e': e'=(u', v)} \gamma_{e'}$ 
     $\gamma_e \leftarrow \min(\sigma_u, \sigma_v)$ 

```

---

**Theorem 5.** Algorithm 4 solves STACKVC in polynomial time.

*Proof.* Consider an optimum pricing  $p^*$ , which yields a strictly larger revenue than the greedy pricing  $p_g$  computed with our algorithm. For such a given set of prices  $p^*$ , we denote by  $\gamma_e^*$  the dual contribution of edge  $e$ , which we call the *budget* of edge  $e$ . This contribution is the result of applying the follower's algorithm to the instance using  $p^*$ . We restrict our attention to an optimal pricing  $p^*$  for which the follower purchases all priceable vertices. The existence of such a pricing follows from the next lemma (proof omitted).

**Lemma 1.** Given any pricing  $p$ , there is a pricing  $p_l$  with  $p_l(v) \leq p(v)$  for all  $v \in \mathcal{P}$ , for which the leader obtains at least as much revenue as for  $p$  and the follower purchases every vertex  $v \in \mathcal{P}$ .

Now consider the smallest  $i'$ , for which edge  $e' = e_{i'} = (u, v')$  has  $\gamma_{e'}^* \neq \gamma_{e'}$ . It is easy to note that this edge must be incident to a priceable vertex  $v'$ , and the difference in budgets is a result from setting different prices. As in both  $p^*$  and  $p_g$  all vertices are bought by the follower, it must be the case that  $p^*(v) < p_g(v)$ , and hence  $\gamma_{e'}^* < \gamma_{e'}$ . We now compare the revenue of pricing  $p^*$  to a pricing  $p'$  with  $p'(v) = p^*(v)$  for every vertex  $v \neq v'$ , and for which  $p'(v') = p^*(v') + \gamma_{e'} - \gamma_{e'}^*$ . In  $p'$  the budgets  $\gamma_e$  are equivalent to  $\gamma_e$  (i.e., the budgets generated by the greedy pricing  $p_g$ ) for every edge  $e_1, \dots, e_{i'-1}$  and also  $e_{i'} = e'$ .

We call  $\delta^j(v) = \sum_{e_i: e_i=(v, u), i \leq j} \gamma_{e_i} - \gamma_{e_i}^*$  the *reservation* that is created by  $p^*$  at vertex  $v$  at the end of processing edge  $e_j$ . The budget of  $e'$  is raised to a smaller amount in  $p^*$  than in  $p'$ , so after processing  $e'$  there is positive reservation at the other endpoint  $u$  of  $e'$ , i.e.  $\delta^{i'}(u) = \delta^{i'}(v')$  at vertex  $u$ . No other vertex except  $u$  and  $v'$  has reservation at this point, so it holds that  $\sum_{v \neq v'} |\delta^{i'}(v)| \leq \delta^{i'}(v')$ . This will be our invariant, and in the following we prove it for the remaining edges  $j > i'$  and the remaining iterations of the algorithm with pricing  $p'$  (proof omitted).

**Lemma 2.** For any iteration  $j \geq i'$  we have that  $\sum_{v \neq v'} |\delta^j(v)| \leq \delta^j(v')$ .

The lemma above shows that the sum of absolute values of reservation at all vertices except  $v'$  at any point during the remaining runs of the followers algorithm is at most

the initial reservation  $\delta^{i'}(v')$ . Note that  $v'$  is bought in both cases. In  $p^*$  all priceable vertices are bought by the follower, but this might not be true for  $p'$  and vertices  $v \neq v'$ .  $p'$  might lose revenue there. By Lemma 11 this can be fixed by reducing the price in  $p'$  of every priceable vertex to the sum of the budgets of incident edges. Note that in  $p^*$  every vertex was bought, which implies every price  $p^*(v) = \sum_{e:e=(u,v)} \gamma_e^*$  is also the sum of budgets of incident edges. As the total absolute reservation in the end of the algorithm is at most  $\delta^{i'}(v')$ , the total decrease in revenue that is lost on vertices  $v \neq v'$  in this step is at most  $\delta^{i'}(v')$ . This is exactly the revenue surplus that  $p'$  generates over  $p^*$  at vertex  $v'$ . Thus, pricing  $p'$  yields at least as much revenue as  $p^*$ . This implies that we can transform any pricing by iteratively adjusting the prices without decreasing revenue, such that all budgets of the edges are equal to those generated by the greedy pricing  $p_g$ . In particular, this implies that  $p_g$  is an optimal pricing.  $\square$

The next proposition, whose proof is left for the full version of this paper, shows that knowing the order of the edges when pricing for the primal-dual algorithm in STACKVC is essential.

**Proposition 2.** *For every constant  $\varepsilon > 0$ , there exists an instance of STACKVC such that if the order of the edges in the follower's primal-dual 2-approximation algorithm is unknown to the leader, then every pricing  $p$  yields an approximation ratio of  $\Omega(1/\varepsilon)$ .*

**Hardness Results.** The main argument in the previous section works only for the case of regular vertex cover. Let us turn to the case of set cover with elements contained in at least three sets, i.e. elements with *frequency* at least three. We understand them as hyperedges incident to more than two vertices. In this case it might be profitable to reduce the price for a vertex from the value in the greedy pricing. Indeed, we show that set cover pricing problem is much harder to solve. The proof is omitted due to space limitations.

**Theorem 6.** *STACKSC is APX-hard even if all elements have maximum frequency 3 and all fixed-cost sets have cost 1.*

Finally, we derive a devastating lower bound on greedy, single-price, and near-uniform price assignments in set cover pricing. Once again, we leave the proof for the full version of this paper.

**Theorem 7.** *The greedy, single-price, and near-uniform price assignments yield an approximation factor of  $2^{\Omega(|I|)}$  for STACKSC, where  $|I|$  is the size of the representation.*

**Acknowledgments.** We wish to thank Guido Proietti for the helpful discussions in the starting phase of this study. We are very grateful to Piotr Krysta for several insightful discussions relating to a number of results presented in this paper during a productive week of research in Liverpool.

## References

1. Aggarwal, G., Feder, T., Motwani, R., Zhu, A.: Algorithms for Multi-Product Pricing. In: Proc. of 31st ICALP (2004)
2. Balcan, M., Blum, A.: Approximation Algorithms and Online Mechanisms for Item Pricing. In: Proc. of 7th EC (2006)
3. Balcan, N., Blum, A., Hartline, J., Mansour, Y.: Mechanism Design via Machine Learning. In: Proc. of 46th FOCS (2005)
4. Balcan, M., Blum, A., Mansour, Y.: Item Pricing for Revenue Maximization. In: Proc. of 9th EC (2008)
5. Berman, P., Karpinski, M.: Improved Approximation Lower Bounds on Small Occurrence Optimization. ECCC-Report TR03-008 (2003)
6. Bilò, D., Gualà, L., Proietti, G., Widmayer, P.: Computational Aspects of a 2-Player Stackelberg Shortest Paths Tree Game. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 251–262. Springer, Heidelberg (2008)
7. Briest, P.: Uniform Budgets and the Envy-Free Pricing Problem. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part I. LNCS, vol. 5125, pp. 808–819. Springer, Heidelberg (2008)
8. Briest, P., Hoefer, M., Krysta, P.: Stackelberg Network Pricing Games. In: Proc. of 25th STACS (2008)
9. Briest, P., Krysta, P.: Single-Minded Unlimited-Supply Pricing on Sparse Instances. In: Proc. of 17th SODA (2006)
10. Briest, P., Krysta, P.: Buying Cheap is Expensive: Hardness of Non-Parametric Multi-Product Pricing. In: Proc. of 18th SODA (2007)
11. Cardinal, J., Demaine, E., 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)
12. Chawla, S., Hartline, J., Kleinberg, R.: Algorithmic Pricing via Virtual Valuations. In: Proc. of 8th EC (2007)
13. Chen, N., Ghosh, A., Vassilvitskii, S.: Optimal Envy-free Pricing with Metric Substitutability. In: Proc. of 9th EC (2008)
14. Demaine, E.D., Feige, U., Hajiaghayi, M.T., Salavatipour, M.R.: Combination Can Be Hard: Approximability of the Unique Coverage Problem. In: Proc. of 17th SODA (2006)
15. Elbassioni, K., Raman, R., Ray, S.: On Profit-Maximizing Pricing for the Highway and Toll-booth Problems. Technical report arXiv:0901.1140v1 (2009)
16. Elbassioni, K., Sitters, R., Zhang, Y.: A Quasi-PTAS for Envy-Free Pricing on Line Graphs. In: Proc. of 15th ESA (2007)
17. Fleischer, L., Jain, K., Mahdian, M.: Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games. In: Proc. of 45th FOCS (2004)
18. Fleischer, L.: Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Computer Science* 348(2-3), 217–225 (2005)
19. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-completeness*. Freeman, New York (1979)
20. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On Profit-Maximizing Envy-Free Pricing. In: Proc. of 16th SODA (2005)
21. Hartline, J., 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)

22. Karakostas, G., Kolliopoulos, S.: Edge Pricing of Multicommodity Networks for Heterogeneous Users. In: Proc. of 45th FOCS (2004)
23. Labbé, M., Marcotte, P., Savard, G.: A Bilevel Model of Taxation and its Application to Optimal Highway Pricing. *Management Science* 44(12), 1608–1622 (1998)
24. Roch, S., Savard, G., Marcotte, P.: An Approximation Algorithm for Stackelberg Network Pricing. *Networks* 46(1), 57–67 (2005)
25. Swamy, C.: The Effectiveness of Stackelberg Strategies and Tolls for Network Congestion Games. In: Proc. of 18th SODA (2007)
26. Vazirani, V.: *Approximation Algorithms*. Springer, Heidelberg (2001)
27. von Stackelberg, H.: *Marktform und Gleichgewicht (Market and Equilibrium)*. Verlag von Julius Springer, Vienna (1934)

# On Best Response Dynamics in Weighted Congestion Games with Polynomial Delays<sup>\*</sup>

Angelo Fanelli<sup>1,2</sup> and Luca Moscardelli<sup>3</sup>

<sup>1</sup> Department of Computer Science, RWTH Aachen University, Germany  
fanelli@cs.rwth-aachen.de

<sup>2</sup> Division of Mathematical Sciences, Nanyang Technological University, Singapore

<sup>3</sup> Department of Science, University of Chieti-Pescara, Italy  
moscardelli@sci.unich.it

**Abstract.** We investigate the speed of convergence of best response dynamics to approximately optimal solutions in weighted congestion games with polynomial delay functions. In [1] it has been shown that the convergence time of such dynamics to Nash equilibrium may be exponential in the number of players  $n$  even for unweighted players and linear delay functions. Nevertheless, extending the work of [1], we show that  $\Theta(n \log \log W)$  (where  $W$  is the sum of all the players' weights) best responses are necessary and sufficient to achieve states that approximate the optimal solution by a constant factor, under the assumption that every  $O(n)$  steps each player performs a constant (and non-null) number of best responses.

## 1 Introduction

Congestion games have attracted a good deal of attention because they are a well established approach to model scenarios in which selfish agents individually strive to allocate shared resources as effectively as possible. In such games a set of  $m$  resources is available to a set of  $n$  players. Each player comes along with a set of strategies. A strategy of a player corresponds to the selection of a subset of the resources. In the unweighted setting, the delay of a particular resource depends on the number of players choosing that resource; in the more general setting of weighted players, it depends on the sum of the weights of the players choosing the considered resource. The cost of each player is the sum of the delays associated with the selected resources. A state of the game is any combination of strategies for the players and its social cost denotes its quality from a global perspective, which is typically defined as the sum of the players' costs.

On the one hand, Rosenthal [18] has shown, by a potential function argument, that for unweighted congestion games the natural decentralized mechanism known as *Nash dynamics*, in which at each step some player performs an improvement step switching her strategy to a better alternative, is guaranteed to converge to a pure Nash equilibrium [17], i.e. a fixed point in which no

---

<sup>\*</sup> This research was partially supported by the grant NRF-RF2009-08 “Algorithmic aspects of coalitional games”.

player can perform an improvement step (note that in a Nash dynamic players play their improvement steps sequentially, and not in parallel). On the other hand, weighted congestion games do not necessarily admit a Nash equilibrium [15] unless specific settings are considered (linear delay functions [13], singleton congestion games [12], matroid congestion games [2]).

Even in the unweighted setting in which Nash equilibria are guaranteed to exist, a recent result of Fabrikant et al. [10] shows that such dynamics may require a number of steps exponential in the number of players  $n$  in order to reach such an equilibrium. Their analysis relates congestion games to local search problems by showing that it is PLS-complete [9] to compute a Nash equilibrium for general unweighted congestion games. Moreover from their completeness proof and from previous results about local search problems, it follows that there exist congestion games with initial states such that any improvement sequence starting from these states needs an exponential number of steps in order to reach a Nash equilibrium. More recently, Ackermann et al. [1] show that the previous negative result holds even in the restricted case of linear unweighted congestion games. Furthermore, as already remarked, in the more general setting of weighted congestion games Nash equilibria may not exist. Therefore, both for the unweighted and the weighted setting, the natural arising question is that of determining how many improvement step are required in order to reach a solution (not necessarily being an equilibrium) with a social cost not far from the optimal one.

**Our Contribution.** In this paper we address the question of bounding the social cost of the state reached after a sequence of *best responses*, i.e. the case in which the players select the best available strategy. In particular, we extend the results of Fanelli et al. [11] relative to linear unweighted congestion games to the more general setting of weighted congestion games with polynomial delay functions having maximum degree  $d$ . Note that if a player performs a best response, she does not necessarily switch to another strategy, as the best strategy may correspond with the one currently selected by the player. In such a case the best response is not an improvement step. To this aim, we must consider sequences in which each player performs at least a best response. Therefore, we can generally assume any sequence of best responses structured in terms of subsequences which we call *covering walks*, corresponding to dynamics in which each player performs at least a best response. In order to state our results let us define the approximation ratio of a state as the ratio between the social cost of such state and the optimal one. Moreover, let us call a covering walk  $\beta$ -bounded if  $\beta$  is the maximum number of best responses each player can perform during the walk. We prove fast convergence of best response Nash dynamics to constant factor approximate solution in weighted congestion games under the assumption that such dynamics are structured in terms of  $O(1)$ -bounded covering walks, thus considering a more general game evolution dynamic with respect to the one analyzed by Fanelli et al. [11] that deals only with the more restricted case of 1-bounded covering walks. On the one hand, we point out that the constant value of  $\beta$  has an important role in the fast convergence. In fact, as proved by Awerbuch et al. [5], if  $\beta$  is polynomial in  $n$ , there exists a (unweighted and with



linear delay functions) congestion game with an initial state such that, after an exponential number of polynomially bounded covering walks starting from such state, it is possible to reach solutions whose social cost is very far from the optimal solution. On the other hand,  $\beta = O(1)$  is a reasonable assumption since it ensures a sort of fairness between the players, as the ratio between the number of best responses of any two players in a  $O(1)$ -bounded covering walk is constant. In particular, we show that the approximation ratio achieved after a sequence of  $k$   $O(1)$ -bounded covering walks is  $O\left(W^{d\left(\frac{d}{d+1}\right)^{k-1}}\right)$  and  $\Omega\left(\frac{W^{d\left(\frac{d}{d+1}\right)^{k-1}}}{k}\right)$  (which is asymptotically matching for constant values of  $k$ ), where  $W$  is the sum of the players' weights. As a consequence, we prove that, for any given  $d$ ,  $\Theta(\log \log W)$   $O(1)$ -bounded covering walks are necessary and sufficient to achieve a constant factor approximate solution.

Our results raise the important open question of the minimum degree of coordination among the players, that is the exact order of  $\beta$ , needed in a covering walk in order to achieve a polynomial time convergence to constant approximated solutions.

**Related Works.** The performances of Nash equilibria in unweighted and weighted congestion games with polynomial delay functions have been investigated by Aland et al. [3], who have provided a tight price of anarchy both for the unweighted case and the weighted one. In particular they have improved the previous results due to Awerbuch et al. [4] (for the weighted case) and Christodoulou and Koutsoupias [7] (for the unweighted case).

Convergence issues to Nash equilibria or approximately optimal solutions in congestion games have been recently investigated in different contexts. Fabrikant et al. [10] have studied the complexity of computing Nash equilibria in general congestion games proving that such a problem is PLS-complete. Ackermann et al. [1] have shown that such result still holds if we restrict to unweighted congestion games with linear delay functions. From these results it follows that there exist linear congestion games with initial states such that any improvement sequence starting from these states needs an exponential number of steps to reach a Nash equilibrium.

Chien and Sinclair [6] have introduced the notion of  $\epsilon$ -Nash dynamics (and the corresponding concept of  $\epsilon$ -Nash equilibrium). In  $\epsilon$ -Nash dynamics only  $\epsilon$ -improvement steps are permitted, i.e. steps that improve the cost of a player by a multiplicative factor of more than  $\epsilon$ . For symmetric unweighted congestion games in which the delay functions satisfy the so-called “bounded-jump” condition, they show that such dynamics converge to a stable configuration, i.e. an  $\epsilon$ -Nash equilibrium, after a number of steps that is polynomial in the number of players and  $\epsilon^{-1}$ . The bounded-jump condition is a weak condition which states that when a new player is added to a resource, the cost of all the players using that resource cannot increase by more than a factor polynomially bounded by  $n$ . Unfortunately, as shown by Skopalik and Vöcking [19],  $\epsilon$ -Nash dynamics do not guarantee polynomial convergence in the asymmetric unweighted setting, even if the delay functions satisfy a bounded-jump condition. However, Awerbuch et al. [5] have proved

that for asymmetric unweighted congestion games with delay functions satisfying a bounded-jump condition,  $\epsilon$ -Nash dynamics rapidly converge to approximately optimal solutions.

Concerning the best response Nash dynamics, Christodoulou et al. [8] have initiated the study of the approximation ratio of the solution achieved in unweighted congestion games with linear delay functions after a sequence of 1-bounded covering walks starting from an arbitrary state. Such a study was completed by Fanelli et al. [11], who have shown that  $\Theta(\log \log n)$  1-bounded covering walks are necessary and sufficient to achieve a constant factor approximate solution. Finally, Goemans et al. [14] have shown that in the weighted setting with linear delay functions, if at each step every player has the same probability to play a best response, after a polynomial number of steps the expected approximation ratio is constant.

The paper is organized as follows. In the next section we introduce some definition and notation. Section 3 is devoted to the analysis of the upper bound and lower bound to the approximation ratio achieved after a sequence of  $k$   $\beta$ -covering walks.

## 2 Definitions and Notation

A *weighted congestion game*  $\mathcal{G} = (N, E, (w_i)_{i \in N}, (\Sigma_i)_{i \in N}, (f_e)_{e \in E}, (c_i)_{i \in N})$  is a non-cooperative strategic game characterized by the existence of a set  $E$  of resources to be shared by the players in  $N = \{1, \dots, n\}$ .

Each player  $i$  has a weighted demand  $w_i \in \mathbb{R}^+$  and we denote by  $W$  the sum of the weights of all players, i.e.  $W = \sum_{i \in N} w_i$ . Moreover, let us denote by  $w_{max}$  the maximum weight.  $\Sigma_i$  is the strategy space of player  $i$ , and any strategy  $s_i \in \Sigma_i$  of player  $i$  is a subset of resources, i.e.  $\Sigma_i \subseteq 2^E$ . Given a strategy profile  $S = (s_1, \dots, s_n)$  and a resource  $e$ , we define the congestion  $\theta_e(S)$  on resource  $e$  by  $\theta_e(S) = \sum_{i \in N | e \in s_i} w_i$ . A delay function  $f_e : \mathbb{R}^+ \mapsto \mathbb{R}^+$  associates to resource  $e$  a delay depending on its congestion, so that the cost of player  $i$  for the pure strategy  $s_i$  is given by the weighted sum of the delays associated with resources in  $s_i$ , i.e.  $c_i(S) = \sum_{e \in s_i} w_i f_e(\theta_e(S))$ .

In this paper we will focus on congestion games with *polynomial* delay functions with maximum degree  $d$  and non-negative coefficients, that is for every resource  $e \in E$  the delay function is of the form  $f_e(x) = \sum_{j=1}^d a_{e,j} x^j$  with  $a_{e,j} \geq 0$  for all  $j = 0, \dots, d$ .

Given the strategy profile  $S = (s_1, \dots, s_n)$ , the social cost  $C(S)$  of a given state  $S$  is defined as the sum of all the players' costs, i.e.  $C(S) = \sum_{i \in N} c_i(S)$ . An optimal strategy profile  $S^* = (s_1^*, \dots, s_n^*)$  is one with minimum social cost, that we denote by OPT. We denote by  $E^* \subseteq E$  the set of resources used at a given optimal strategy profile  $S^*$ , i.e.  $E^* = \bigcup_{i \in N} s_i^*$ . The *approximation ratio* of state  $S$  is given by the ratio between the social cost of  $S$  and the social optimum, i.e.  $\frac{C(S)}{\text{OPT}}$ .

Each player acts selfishly and aims at choosing the strategy lowering her cost, given the strategy choices of other players. Given a strategy profile  $S$  and a strategy  $s'_i \in \Sigma_i$ , let  $(S \oplus s'_i) = (s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$ , i.e. the strategy profile

obtained from  $S$  if player  $i$  changes her strategy from  $s_i$  to  $s'_i$ . An *improvement move* of player  $i$  is a strategy  $s'_i$  such that  $c_i(S \oplus s'_i) < c_i(S)$ . Furthermore, a *best response* of player  $i$  in  $S$  is a strategy  $s_i^b \in \Sigma_i$  yielding the minimum possible cost, given the strategy choices of the other players, i.e.  $c_i(S \oplus s_i^b) \leq c_i(S \oplus s'_i)$  for any other strategy  $s'_i \in \Sigma_i$ . Notice that a best response corresponding to the strategy currently played in  $S$  by the involved player is not an improvement move.

The selfish behavior of players performing best responses can be modelled by the (*Best Response*) *Nash Dynamics Graph*. Formally the Nash Dynamics Graph associated to a congestion game  $\mathcal{G}$  is a directed graph  $\mathcal{B} = (V, A)$  where each vertex in  $V$  corresponds to a strategy profile and there is an edge  $(S, S') \in A$  with label  $i$ , where  $S' = S \oplus s'_i$  and  $s'_i \in \Sigma_i$ , if and only if both the following conditions are met: (i)  $s'_i$  is a best response of  $i$  in  $S$ ; (ii) if  $S \neq S'$ ,  $s'_i$  is also an improvement move of  $i$  in  $S$ . Observe that  $\mathcal{B}$  may contain loops, corresponding to best response in which a player maintains her current strategy. A *best response walk* is a directed walk in  $\mathcal{B}$ .

Given a best response walk in  $\mathcal{B}$  starting from an arbitrary state, we are interested in the social cost of its final state. We consider the following notions of best response walks, that are a refinement of the ones introduced in [8,16]:

**$\beta$ -bounded covering walk:** it's a best response walk  $R = ((S_R^0, S_R^1), (S_R^1, S_R^2), \dots, (S_R^i, S_R^{i+1}), \dots, (S_R^{\ell_R-1}, S_R^{\ell_R}))$  in  $\mathcal{B}$  of length  $\ell_R$ , where the edge  $(S_R^i, S_R^{i+1})$  has label  $\pi_R(i)$  for every  $0 \leq i \leq \ell_R - 1$ , i.e.  $\pi_R(i)$  is the player performing the  $i$ -th best response of the walk.  $\pi_R$  is such that every player performs at least a best response and at most  $\beta$  best responses in  $R$ .  $S_R^0$  is said the *initial* state of  $R$  and  $S_R^{\ell_R}$  its *final* state. For simplicity we denote  $R$  by a sequence of states, i.e.  $R = (S_R^0, \dots, S_R^{\ell_R})$ . When clear from the context, we will drop the index  $R$  from the notation, writing  $S^i$ ,  $\pi$  and  $\ell$  instead of  $S_R^i$ ,  $\pi_R$  and  $\ell_R$ , respectively.

**$\beta$ -bounded  $k$ -covering walk:** it's a best response walk  $P = (R_1, \dots, R_k)$  in  $\mathcal{B}$  corresponding to a sequence of  $k$   $\beta$ -bounded covering walks, i.e. such that each  $R_i$  is a  $\beta$ -bounded covering walk in  $\mathcal{B}$ .

Finally, we denote by  $\text{Apx}_k^\beta(\mathcal{G})$  the worst case approximation ratio of a state obtained after a  $\beta$ -bounded  $k$ -covering walk.

### 3 Upper and Lower Bounds

In this section we first provide an upper bound to the the social cost of the state achieved after a  $\beta$ -bounded  $k$ -covering walk starting from an arbitrary state, for any  $k \geq 1$  and  $\beta = O(1)$ . All the results hold for congestion games having polynomial delay functions with non-negative coefficients and maximum degree  $d$ , i.e. for every  $e \in E$ ,  $f_e(x) = \sum_{j=1}^d a_{e,j} x^j$  with  $a_{e,j} \geq 0$  for all  $j = 0, \dots, d$ . Without loss of generality, we can assume that for every  $e \in E$ ,  $f_e(x) = x^j$  with  $1 \leq j \leq d$ . In fact, given a congestion game  $\mathcal{G}$  having polynomial delays

with non-negative coefficients and maximum degree  $d$ , it is possible to obtain an equivalent congestion game  $\mathcal{G}'$ , having the same set of players and delay functions of the form  $f(x) = x^j$  with  $1 \leq j \leq d$  in the following way. For each resource  $e$  in  $\mathcal{G}$ , we include in  $\mathcal{G}'$  a set  $A_{e,j}$  of  $a_{e,j}$  resources for  $j = 1, \dots, d$  with delay function  $f_e(x) = x^j$  and  $n$  sets  $B_e^1, \dots, B_e^n$ , each containing  $a_{e,0}$  resources with delay  $f_e(x) = x$ ; moreover, given any strategy set  $s_i \in \Sigma_i$  in  $\mathcal{G}$ ,  $i = 1, \dots, n$ , we build a corresponding strategy set  $s'_i \in \Sigma'_i$  (in  $\mathcal{G}'$ ) by including in  $s'_i$ , for each  $e \in s_i$ , all the resources in the sets  $A_e = \cup_{j=1}^d A_{e,j}$  and  $B_e^i$ . If the coefficients  $a_{e,j}$  are not integers we can perform a similar reduction by exploiting a simple scaling argument.

For the sake of simplicity, we will only consider latencies of the form  $f_e(x) = x^d$ , but it can be verified that our proof works for the general case. All the details will be given in the full version of the paper.

The following technical lemma will be useful in the sequel.

**Lemma 1.** *For every pair of reals  $x, y \geq 0$  and every integer  $d \geq 1$ , it holds  $x^d \geq \frac{1}{2^{d-1}}(x+y)^d - y^d$ .*

*Proof.* Let us prove the statement by induction on  $d$ . When  $d = 1$  the inequality trivially holds. Now suppose that  $x^d \geq \frac{1}{2^{d-1}}(x+y)^d - y^d$  holds for every  $1 \leq d \leq t$ , it must be shown that the inequality holds even for  $d = t + 1$ . By the induction hypothesis we achieve that  $(x+y)^{t+1} \leq 2^{t-1}(x+y)(x^t + y^t) = 2^{t-1}(x^{t+1} + xy^t + y^{t+1} + yx^t)$ . The claim follows by observing that  $xy^t + yx^t \leq y^{t+1} + x^{t+1}$ .  $\square$

Let  $R = (S^0, \dots, S^\ell)$  be a  $\beta$ -bounded covering walk. Given the optimal strategy profile  $S^*$ , since the  $i$ -th moving player  $\pi(i)$  before moving can always select the strategy she would use in  $S^*$ ,  $c_{\pi(i)}(S^i)$  (that is the player's cost immediately after her best response) can be suitably upper bounded by  $\sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d$ . In order to state our results we define the following function

$$\Gamma(R) = \frac{1}{\text{OPT}} \sum_{i=1}^{\ell} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d,$$

which, by the same argument explained earlier, clearly represents an upper bound to  $\frac{1}{\text{OPT}} \sum_{i=1}^{\ell} c_{\pi(i)}(S^i)$ .

Lemmas 2 and 3 provide a lower and an upper bound to  $\Gamma(R)$ , respectively. From such Lemmas, we can easily derive the approximation achieved after a  $\beta$ -bounded covering walk.

**Lemma 2.** *For any  $\beta \geq 1$ , given a  $\beta$ -bounded covering walk  $R$  ending in  $S^\ell$ , it holds  $\Gamma(R) \geq \frac{1}{(d+1)} \frac{C(S^\ell)}{\text{OPT}}$ .*

*Proof.* Since the players perform best responses, inequality (I) below holds. In order to justify inequality 2, let us consider a resource  $e$ . Recall that the cost  $c_{\pi(i)}(S^i)$  incurred by a player  $\pi(i)$  on  $e$  is  $w_{\pi(i)} f_e(\theta_e(S^i))$ ; since  $f_e$  is a non

decreasing function, the sum of all the cost that players using  $e$  incur on  $e$  can be lower bounded by considering the resource used by many players each having infinitesimal weight; thus, the summation  $\sum_{e \in E} \sum_{i \in N | e \in s_{\pi(i)}} w_{\pi(i)} \theta_e^d(S^i)$  can be replaced by the integral  $\int_{x=0}^{\theta_e(S^\ell)} x^d dx$  because in  $R$  each player performs at least a best response. Therefore, by recalling the definition of  $\Gamma(R)$ ,

$$\begin{aligned}
 \Gamma(R) &= \frac{1}{\text{OPT}} \sum_{i=1}^{\ell} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d \\
 &\geq \frac{1}{\text{OPT}} \sum_{i=1}^{\ell} c_{\pi(i)}(S^i) \\
 &\geq \frac{1}{\text{OPT}} \sum_{e \in E} \sum_{\substack{i \in N \\ e \in s_{\pi(i)}}} w_{\pi(i)} \theta_e^d(S^i) \\
 &\geq \frac{1}{\text{OPT}} \sum_{e \in E} \int_{x=0}^{\theta_e(S^\ell)} x^d dx \\
 &\geq \frac{1}{(d+1)\text{OPT}} \sum_{e \in E} \theta_e^{d+1}(S^\ell) = \frac{C(S^\ell)}{(d+1)\text{OPT}}. \quad \square
 \end{aligned} \tag{1}$$

**Lemma 3.** *For any  $\beta \geq 1$ , given a  $\beta$ -bounded covering walk  $R$ , it holds  $\Gamma(R) \leq \beta(W + w_{max})^d$ .*

*Proof.* By the definition of  $\Gamma(R)$ , it holds

$$\begin{aligned}
 \Gamma(R) &= \frac{1}{\text{OPT}} \sum_{i=1}^{\ell} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d \\
 &\leq \frac{1}{\text{OPT}} \sum_{i=1}^{\ell} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (W + w_{max})^d \\
 &= \frac{(W + w_{max})^d}{\text{OPT}} \sum_{i=1}^{\ell} |s_{\pi(i)}^*| w_{\pi(i)} \\
 &\leq \beta(W + w_{max})^d, \tag{3}
 \end{aligned}$$

where [3](#) holds by observing that  $\text{OPT} \geq \frac{1}{\beta} \sum_{i=1}^{\ell} |s_{\pi(i)}^*| w_{\pi(i)}$ .  $\square$

As an immediate consequence of Lemmas [2](#) and [3](#),  $\text{Apx}_1^\beta(\mathcal{G}) \leq \beta(d+1)(W + w_{max})^d$ . Lemmas [4](#) and [5](#) will be useful to extend such result to a  $\beta$ -bounded  $k$ -covering walk  $P = \langle R_1, \dots, R_k \rangle$  by exploiting the relationship among two consecutive walks. To this aim we define the following function

$$H(S) = \sum_{e \in E^*} \theta_e^d(S) x_e,$$

where  $x_e = \sum_{j \in X_e} w_{\pi(j)}$  and  $X_e = \{i \in \{1, \dots, \ell\} | e \in s_{\pi(i)}^*\}$ .

The following lemma (resp. Lemma 5) will provide an upper bound (resp. a lower bound) to  $H(S^\ell)$  (resp.  $H(S^0)$ ), where  $S^\ell$  (resp.  $S^0$ ) is the last (resp. initial) state of a  $\beta$ -bounded covering walk. Since the final state of each walk  $R_t$  with  $t = 1, \dots, k - 1$  corresponds to the initial state of  $R_{t+1}$ , by combining  $k - 1$  times the results of lemmas 4 and 5, Theorem 1 finally derives an upper bound to  $\text{Apx}_k^\beta(\mathcal{G})$ .

**Lemma 4.** *For any  $\beta \geq 1$ , given a  $\beta$ -bounded covering walk  $R$  ending in  $S^\ell$ , it holds*

$$H(S^\ell) \leq \beta ((d + 1)\Gamma(R))^{\frac{d}{d+1}} \text{OPT}.$$

*Proof.*

$$\begin{aligned} H(S^\ell) &\leq \beta \sum_{e \in E^*} \theta_e^d(S^\ell) \theta_e(S^*) \\ &\leq \beta \left( \sum_{e \in E^*} (\theta_e^d(S^\ell))^{\frac{d+1}{d}} \right)^{\frac{d}{d+1}} \left( \sum_{e \in E^*} (\theta_e^{d+1}(S^*)) \right)^{\frac{1}{d+1}} \quad (4) \\ &= \beta \left( \sum_{e \in E^*} (\theta_e^{d+1}(S^\ell)) \right)^{\frac{d}{d+1}} \left( \sum_{e \in E^*} (\theta_e^{d+1}(S^*)) \right)^{\frac{1}{d+1}} \\ &= \beta (C(S^\ell))^{\frac{d}{d+1}} \text{OPT}^{\frac{1}{d+1}} \\ &\leq \beta ((d + 1)\Gamma(R)\text{OPT})^{\frac{d}{d+1}} \text{OPT}^{\frac{1}{d+1}} \quad (5) \\ &= \beta ((d + 1)\Gamma(R))^{\frac{d}{d+1}} \text{OPT}. \end{aligned}$$

where (4) follows from Hölder’s inequality

$$\sum_{j=1}^q a_j b_j \leq \left( \sum_{j=1}^q a_j^r \right)^{1/r} \left( \sum_{j=1}^q b_j^s \right)^{1/s},$$

by replacing  $r$  with  $(\frac{d+1}{d})$  and  $s$  with  $(d + 1)$ , and (5) follows from Lemma 2.  $\square$

**Lemma 5.** *For any  $\beta \geq 1$ , given a  $\beta$ -bounded covering walk  $R$  starting from  $S^0$ , it holds*

$$H(S^0) \geq \left( \frac{1}{2^{d-1}} - \frac{d}{\alpha} \right) \Gamma(R)\text{OPT} - \beta ((\alpha\beta)^d + 1) \text{OPT},$$

for any  $\alpha > d 2^{d-1}$ .

*Proof.* In order to lower bound  $H(S^0)$  with respect to  $\Gamma(R)$ , we define the following suitable potential function  $h_i(R) = \sum_{e \in E^*} g_e(S^i) x_e^{>i}$  for  $i \in \{0, \dots, \ell\}$ , where for a generic state  $S$ ,  $g_e(S) = \max\{0, f_e(\theta_e(S)) - f_e(\alpha\beta\theta_e(S^*))\}$  and  $x_e^{>k} = \sum_{j \in X_e^{>k}} w_{\pi(j)}$  where  $X_e^{>k} = \{i \in \{k + 1, \dots, \ell\} | e \in s_{\pi(i)}^*\}$ . Informally

speaking, such a potential function takes into account the delay due to the congestion of the not yet moving players during walk  $R$  above a “virtual” congestion frontier given by all the values  $\alpha\beta n_e(S^*)$ . Let  $\Delta_i(R) = h_{i-1}(R) - h_i(R)$  for  $i \in \{1, \dots, \ell\}$ . Notice that by the definition of the potential function  $h_i(R)$ , since  $h_\ell(R) = 0$ ,  $\sum_{i=1}^{\ell} \Delta_i(R) = h_0(R) \leq H(S^0)$ , that is a lower bound for  $\sum_{i=1}^{\ell} \Delta_i(R)$  is also a lower bound for  $H(S^0)$ ; therefore, in the following we focus on lower bounding  $\sum_{i=1}^{\ell} \Delta_i(R)$ .

Consider a generic step  $i$  in walk  $R$ , in which player  $\pi(i)$  performs a best response by selecting resources in  $s_{\pi(i)}^i$  and let us bound from below the value of  $\Delta_i(R)$  by evaluating how much player  $\pi(i)$  removes from  $h_{i-1}(R)$  and how much she adds to  $h_i(R)$ . Player  $\pi(i)$  in order to obtain  $h_i(R)$  removes at least  $\sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} g_e(S^{i-1})$  from  $h_{i-1}(R)$ , due to the decrease of the coefficients  $x_e^{>(i-1)}$  to  $x_e^{>i}$ . Let us evaluate how much player  $\pi(i)$  adds to  $h_i(R)$ . Player  $\pi(i)$  increases the value of  $h_i(R)$  only by resources whose congestion is above the virtual frontier after player  $\pi(i)$  plays her best response. Thus for each resource  $e \in s_{\pi(i)}^i$  such that  $\theta_e(S^i) > \alpha\beta\theta_e(S^*)$ , the increase of  $h_i(R)$  is equal to  $(g_e(S^i) - g_e(S^{i-1}))x_e^{>i}$  which, by the definition of  $g_e$ , is equal to  $(f_e(\theta_e(S^i)) - f_e(\theta_e(S^{i-1})))x_e^{>i}$ . Since  $f_e$  is convex, such quantity is at most  $(\theta_e(S^i) - \theta_e(S^{i-1}))f'_e(\theta_e(S^i))x_e^{>i} = w_{\pi(i)}f'_e(\theta_e(S^i))x_e^{>i}$ . Moreover since  $x_e^{>i} \leq x_e \leq \beta\theta_e(S^*) \leq \theta_e(S^i)/\alpha$ , we obtain that the increase for each resource  $e$  is at most  $w_{\pi(i)}f'_e(\theta_e(S^i))\theta_e(S^i)/\alpha$ . Thus, considering the previous quantity as an upper bound of the increase for all the resources in  $s_{\pi(i)}^i$ , player  $\pi(i)$  in order to obtain  $h_i(R)$  adds at most  $\frac{1}{\alpha} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)}f'_e(\theta_e(S^i))\theta_e(S^i)$  to  $h_{i-1}(R)$ . Therefore,

$$\Delta_i(R) \geq \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)}g_e(S^{i-1}) - \frac{1}{\alpha} \sum_{e \in s_{\pi(i)}^i \cap E^*} w_{\pi(i)}f'_e(\theta_e(S^i))\theta_e(S^i).$$

Finally, since  $g_e(S^{i-1}) \geq f_e(\theta_e(S^{i-1})) - f_e(\alpha\beta\theta_e(S^*))$  for every  $e \in E^*$ , it follows that

$$\begin{aligned} \Delta_i(R) &\geq \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (f_e(\theta_e(S^{i-1})) - f_e(\alpha\beta\theta_e(S^*))) \\ &\quad - \frac{1}{\alpha} \sum_{e \in s_{\pi(i)}^i \cap E^*} w_{\pi(i)}f'_e(\theta_e(S^i))\theta_e(S^i). \end{aligned} \quad (6)$$

Since  $f_e(x) = x^d$  for  $d \geq 1$ , we obtain

$$\begin{aligned} \Delta_i(R) &\geq \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e^d(S^{i-1}) - (\alpha\beta)^d \theta_e^d(S^*)) - \frac{d}{\alpha} \sum_{e \in s_{\pi(i)}^i \cap E^*} w_{\pi(i)} \theta_e^d(S^i) \\ &\geq \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e^d(S^{i-1}) - (\alpha\beta)^d \theta_e^d(S^*)) - \frac{d}{\alpha} c_{\pi(i)}(S^i). \end{aligned} \quad (7)$$

Since  $c_{\pi(i)}(S^i) \leq \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d$  and  $w_{\pi(i)} \leq \theta_e(S^*)$  for every  $e \in s_{\pi(i)}^*$ , by using Lemma [4](#), it follows that

$$\begin{aligned}
\Delta_i(R) &\geq \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e^d(S^{i-1}) - (\alpha\beta)^d \theta_e^d(S^*)) \\
&\quad - \frac{d}{\alpha} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d \\
&\geq \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} \left( \frac{1}{2^{d-1}} (\theta_e(S^{i-1}) + w_{\pi(i)})^d - w_{\pi(i)}^d - (\alpha\beta)^d \theta_e^d(S^*) \right) \\
&\quad - \frac{d}{\alpha} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d \\
&\geq \left( \frac{1}{2^{d-1}} - \frac{d}{\alpha} \right) \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d \\
&\quad - ((\alpha\beta)^d + 1) \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} \theta_e^d(S^*).
\end{aligned}$$

By summing up the values  $\Delta_i(R)$ , we obtain

$$\begin{aligned}
\sum_{i=1}^{\ell} \Delta_i(R) &\geq \left( \frac{1}{2^{d-1}} - \frac{d}{\alpha} \right) \sum_{i=1}^{\ell} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} (\theta_e(S^{i-1}) + w_{\pi(i)})^d \\
&\quad - ((\alpha\beta)^d + 1) \sum_{i=1}^{\ell} \sum_{e \in s_{\pi(i)}^*} w_{\pi(i)} \theta_e^d(S^*) \\
&\geq \left( \frac{1}{2^{d-1}} - \frac{d}{\alpha} \right) \Gamma(R) \text{OPT} - \beta ((\alpha\beta)^d + 1) \text{OPT},
\end{aligned}$$

and thus, since  $H(S^0) \geq \sum_{i=1}^{\ell} \Delta_i(R)$ , the claim follows.  $\square$

**Theorem 1.** *For any  $k \geq 1$  and any given  $d \geq 1$ , in every polynomial congestion game  $\mathcal{G}$  with delay functions having maximum degree  $d$ , it holds  $\text{ApX}_k^{O(1)}(\mathcal{G}) = O\left(W^{d\left(\frac{d}{d+1}\right)^{k-1}}\right)$ .*

*Proof.* Let  $P = \langle R_1, \dots, R_k \rangle$  be a  $\beta$ -bounded  $k$ -covering walk, where each  $R_j$  is a  $\beta$ -bounded covering walk. Let  $S_j^0$  the initial state of  $R_j$ .

From Lemma [5](#) we obtain that for each  $R_j$  with  $j = 2, \dots, k$  it holds that for any  $\alpha > d \cdot 2^{d-1}$

$$H(S_j^0) \geq \left( \frac{1}{2^{d-1}} - \frac{d}{\alpha} \right) \Gamma(R_j) \text{OPT} - \beta ((\alpha\beta)^d + 1) \text{OPT}. \quad (8)$$



Furthermore, since the final state of each walk  $R_t$  with  $t = 1, \dots, k - 1$  corresponds to the initial state of  $R_{t+1}$ , by applying Lemma 4 we obtain that for each  $R_j$  with  $j = 2, \dots, k$  it holds

$$H(S_j^0) \leq \beta ((d + 1)\Gamma(R_{j-1}))^{\frac{d}{d+1}} \text{OPT}. \tag{9}$$

By combining (8) and (9) we achieve a relation between  $\Gamma(R_j)$  and  $\Gamma(R_{j-1})$  for every  $j = 2, \dots, k$

$$\Gamma(R_j) \leq \beta \left( \frac{\alpha 2^{d-1}}{\alpha - d 2^{d-1}} \right) \left( ((d + 1)\Gamma(R_{j-1}))^{\frac{d}{d+1}} + ((\alpha\beta)^d + 1) \right) \tag{10}$$

for any  $\alpha > d 2^{d-1}$ .

From the previous inequalities (10), since  $\beta = O(1)$ , we obtain that for constant values of  $\alpha$  and  $d$  it holds

$$\Gamma(R_k) = O \left( (\Gamma(R_1))^{\left(\frac{d}{d+1}\right)^{k-1}} \right). \tag{11}$$

By applying Lemma 2 to  $\Gamma(R_k)$  and Lemma 3 to  $\Gamma(R_1)$  in (11), since  $\beta = O(1)$ , by constant value of  $d$  we obtain that the cost of the final state of walk  $P$  is

$$O \left( ((W + w_{max})^d)^{\left(\frac{d}{d+1}\right)^{k-1}} \text{OPT} \right),$$

and thus the approximation is

$$\text{Apx}_k^{O(1)}(\mathcal{G}) = O \left( W^{d \left(\frac{d}{d+1}\right)^{k-1}} \right). \quad \square$$

The following corollary is an immediate consequence of Theorem 1

**Corollary 1.** *For any polynomial congestion game  $\mathcal{G}$  with  $d = O(1)$ ,  $\text{Apx}_{\log \log W}^{O(1)}(\mathcal{G}) = O(1)$ .*

Finally, we are able to provide an almost matching lower bound to the approximation ratio achieved after a  $\beta$ -bounded  $k$ -covering walk. Due to space limitations, the proof is omitted and will appear in the full version of the paper.

**Theorem 2.** *For any  $d \geq 1$  there exists a congestion game  $\mathcal{G}$  with polynomial delay functions having maximum degree  $d$  such that  $\text{Apx}_k^1(\mathcal{G}) = \Omega \left( \frac{W^{d \left(\frac{d}{d+1}\right)^{k-1}}}{k} \right)$ .*

## References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. *J. ACM* 55(6) (2008)
2. Ackermann, H., Röglin, H., Vöcking, B.: Pure Nash equilibria in player-specific and weighted congestion games. *Theor. Comput. Sci.* 410(17), 1552–1563 (2009)

3. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact price of anarchy for polynomial congestion games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 218–229. Springer, Heidelberg (2006)
4. Awerbuch, B., Azar, Y., Epstein, A.: Large the price of routing unsplittable flow. In: STOC, pp. 57–66. ACM, New York (2005)
5. Awerbuch, B., Azar, Y., Epstein, A., Mirrokni, V.S., Skopalik, A.: Fast convergence to nearly optimal solutions in potential games. In: ACM Conference on Electronic Commerce, pp. 264–273. ACM, New York (2008)
6. Chien, S., Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. In: SODA, pp. 169–178. SIAM, Philadelphia (2007)
7. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: STOC, pp. 67–73. ACM, New York (2005)
8. Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and approximation in potential games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 349–360. Springer, Heidelberg (2006)
9. Yannakakis, M., Johnson, D.S., Papadimitriou, C.H.: How easy is local search? *Journal of Computer and System Sciences* 37, 79–100 (1988)
10. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The complexity of pure Nash equilibria. In: STOC, pp. 604–612. ACM, New York (2004)
11. Fanelli, A., Flammini, M., Moscardelli, L.: The speed of convergence in congestion games under best-response dynamics. In: Aceto, L., Damgaard, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part I. LNCS, vol. 5125, pp. 796–807. Springer, Heidelberg (2008)
12. Fotakis, D., 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)
13. Fotakis, D., Kontogiannis, S.C., Spirakis, P.G.: Selfish unsplittable flows. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 593–605. Springer, Heidelberg (2004)
14. Goemans, M.X., Mirrokni, V.S.: Adrian Vetta. Sink equilibria and convergence. In: FOCS, pp. 142–154 (2005)
15. Milchtaich, I.: Congestion games with player-specific payoff functions. *Games and Economic Behavior* 13(1), 111–124 (1996)
16. Mirrokni, V.S., Vetta, A.: Convergence issues in competitive games. In: Jansen, K., Khanna, S., Rolim, J.D.P., Ron, D. (eds.) RANDOM 2004 and APPROX 2004. LNCS, vol. 3122, pp. 183–194. Springer, Heidelberg (2004)
17. Nash, J.F.: Equilibrium points in  $n$ -person games. In: *Proceedings of the National Academy of Sciences*, vol. 36, pp. 48–49 (1950)
18. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
19. Skopalik, A., Vöcking, B.: Inapproximability of pure Nash equilibria. In: STOC, pp. 355–364. ACM, New York (2008)

# Parametric Packing of Selfish Items and the Subset Sum Algorithm

Leah Epstein<sup>1</sup>, Elena Kleiman<sup>1</sup>, and Julián Mestre<sup>2,\*</sup>

<sup>1</sup> Department of Mathematics, University of Haifa, 31905 Haifa, Israel  
lea@math.haifa.ac.il, elena.kleiman@gmail.com

<sup>2</sup> Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany  
jmestre@mpi-inf.mpg.de

**Abstract.** The subset sum algorithm is a natural heuristic for the classical Bin Packing problem: In each iteration, the algorithm finds among the unpacked items, a maximum size set of items that fits into a new bin. More than 35 years after its first mention in the literature, establishing the worst-case performance of this heuristic remains, surprisingly, an open problem.

Due to their simplicity and intuitive appeal, greedy algorithms are the heuristics of choice of many practitioners. Therefore, better understanding simple greedy heuristics is, in general, an interesting topic in its own right. Very recently, Epstein and Kleiman (*Proc. ESA 2008, pages 368-380*) provided another incentive to study the subset sum algorithm by showing that the Strong Price of Anarchy of the game theoretic version of the Bin Packing problem is *precisely* the approximation ratio of this heuristic.

In this paper we establish the exact approximation ratio of the subset sum algorithm, thus settling a long standing open problem. We generalize this result to the parametric variant of the Bin Packing problem where item sizes lie on the interval  $(0, \alpha]$  for some  $\alpha \leq 1$ , yielding tight bounds for the Strong Price of Anarchy for all  $\alpha \leq 1$ . Finally, we study the pure Price of Anarchy of the parametric Bin Packing game for which we show nearly tight upper and lower bounds for all  $\alpha \leq 1$ .

## 1 Introduction

**Motivation and framework.** The emergence of the Internet and its rapidly gained status as the predominant communication platform has brought up to the surface new algorithmic challenges that arise from the interaction of the multiple self-interested entities that manage and use the network. Due to the nature of the Internet, these interactions are characterized by the (sometimes complete) lack of coordination between those entities. Algorithm and network designers are interested in analyzing the outcomes of these interactions. An interesting and topical question is how much performance is lost due to the selfishness and unwillingness of network participants to cooperate. A formal framework for studying

---

\* Research supported by an Alexander von Humboldt Fellowship.

interactions between multiple rational participants is provided by the discipline of Game Theory. This is achieved by modeling the network problems as strategic games, and considering the quality of the Nash equilibria of these games. In this paper we consider *pure* Nash equilibria and strong equilibria. These equilibria are the result of the pure strategies of the participants of the game, where they choose to play an action in a deterministic, non-aleatory manner.

In this paper, we consider game theoretic variants of the well-known Bin Packing problem and its parametric version; see [5,6] for surveys on these problems. In the classic Bin Packing problem, we are given a set of items  $I = \{1, 2, \dots, n\}$ . The  $i$ th item in  $I$  has size  $s_i \in (0, 1]$ . The objective is to pack the items into unit capacity bins so as to minimize the number of bins used. In the parametric case, the sizes of items are bounded from above by a given value. More precisely, given a parameter  $\alpha \leq 1$  we consider inputs in which the item sizes are taken from the interval  $(0, \alpha]$ . Setting  $\alpha$  to 1 gives us the standard Bin Packing problem.

As discussed in [8], Bin Packing is met in a great variety of networking problems, such as the problem of packing a given set of packets into a minimum number of time slots for fairness provisioning and the problem of packing data for Internet phone calls into ATM packets, filling fixed-size frames to maximize the amount of data that they carry. This fact motivates the study of Bin Packing from a game theoretic perspective. The Parametric Bin Packing problem also models the problem of efficient routing in networks that consist of parallel links of same bounded bandwidth between two terminal nodes—similar to the ones considered in [2,8,14]. As Internet Service Providers often impose a policy which restricts the amount of data that can be downloaded/uploaded by each user, placing a restriction on the size of the items allowed to transfer makes the model more realistic.

**The model.** In this paper we study the Parametric Bin Packing problem both in cooperative and non-cooperative versions. In each case the problem is specified by a given parameter  $\alpha$ . The Parametric Bin Packing game is defined by a tuple  $BP(\alpha) = \langle N, (B_i)_{i \in N}, (c_i)_{i \in N} \rangle$ . Where  $N$  is the set of the items, whose size is at most  $\alpha$ . Each item is associated with a selfish player—we sometimes consider the items themselves to be the players. The set of strategies  $B_i$  for each player  $i \in N$  is the set of all bins. Each item can be assigned to one bin only. The outcome of the game is a particular assignment  $b = (b_j)_{j \in N} \in \times_{j \in N} B_j$  of items to bins. All the bins have unit cost. The cost function  $c_i$  of player  $i \in N$  is defined as follows. A player pays  $\infty$  if it requests to be packed in an invalid way, that is, a bin which is occupied by a total size of items which exceeds 1. Otherwise, the set of players whose items are packed into a common bin share its unit cost proportionally to their sizes. That is, if an item  $i$  of size  $s_i$  is packed into a bin which contains the set of items  $B$  then  $i$ 's payment is  $c_i = s_i / \sum_{k \in B} s_k$ . Notice that since  $\sum_{k \in B} s_k \leq 1$  the cost  $c_i$  is always greater or equal than  $s_i$ . The social cost function that we want to minimize is the number of used bins.

Clearly, a selfish item prefers to be packed into a bin which is as full as possible. In the non-cooperative version, an item will perform an improving step if there is a strictly more loaded bin in which it fits. At a Nash equilibrium, no item

can unilaterally reduce its cost by moving to a different bin. We call a packing that admits the Nash conditions *NE* packing. We denote the set of the Nash equilibria of an instance of the Parametric Bin Packing game  $G \in BP(\alpha)$  by  $NE(G)$ .

In the cooperative version of the Parametric Bin Packing game, we consider all (non-empty) subgroups of items from  $N$ . The cost functions of the players are defined the same as in the non-cooperative case. Each group of items is interested to be packed in a way so as to minimize the costs for all group members. Thus, given a particular assignment, all members of a group will perform a joint improving step (not necessarily into a same bin) if there is an assignment in which, for each member, the new bin will admit a strictly greater load than the bin of origin. The costs of the non-members may be enlarged as a result of this improving step. At a strong Nash equilibrium, no group of items can reduce the costs of all its members by moving to different bins. We denote the set of the strong Nash equilibria of an instance  $G$  of the Parametric Bin Packing game by  $SNE(G)$ . As a group can contain a single item,  $SNE(G) \subseteq NE(G)$  holds.

To measure the extent of deterioration in the quality of Nash packing due to the effect of selfish and uncoordinated behavior of the players (items) in the worst-case we use the Price of Anarchy (*PoA*) and the Price of Stability (*PoS*). These are the standard measures of the quality of the equilibria reached in uncoordinated selfish setting [14,17]. The *PoA/PoS* of an instance  $G$  of the Parametric Bin Packing game are defined to be the ratio between the social cost of the worst/best Nash equilibrium and the social optimum, respectively. As packing problems are usually studied via asymptotic measures, we consider asymptotic *PoA* and *PoS* of the Parametric Bin Packing game  $BP(\alpha)$ , that are defined by taking a supremum over the *PoA* and *PoS* of all instances of the Parametric Bin Packing game, for large sets  $N$ .

Recent research [11,9] initiated a study of measures that separate the effect of the lack of coordination between players from the effect of their selfishness. The measures considered are the Strong Price of Anarchy (*SPoA*) and the Strong Price of Stability (*SPoS*). These measures are defined similarly to the *PoA* and the *PoS*, but only strong equilibria are considered.

These measures are well defined only when the sets  $NE(G)$  and  $SNE(G)$  are not empty for any  $G \in BP(\alpha)$ . Even though pure Nash equilibria are not guaranteed to exist for general games, they always exist for the Bin Packing game: The existence of pure Nash equilibria was proved in [2] and the existence of strong Nash equilibria was proved in [8].

As we study the *SPoA/SPoS* measures in terms of the worst-case approximation ratio of a greedy algorithm for Bin Packing, we define here the *parametric worst-case ratio*  $R_A^\infty(\alpha)$  of algorithm  $A$  by

$$R_A^\infty(\alpha) = \lim_{k \rightarrow \infty} \sup_{I \in V_\alpha} \left\langle \frac{A(I)}{OPT(I)} \mid OPT(I) = k \right\rangle,$$

where  $A(I)$  denotes the number of bins used by algorithm  $A$  to pack the set  $I$ ,  $OPT(I)$  denotes the number of bins used in the optimal packing of  $I$  and  $V_\alpha$  is

the set of all list  $I$  for which the maximum size of the items is bounded from above by  $\alpha$ .

**Related work.** The first problems that were studied from a game theoretic point of view were job scheduling [7,14,16] and routing [17,18] problems. Since then, many other problems have been considered in this setting.

The classic Bin Packing problem was introduced in the early 70's [13,19]. This problem and its variants are often met in various real-life applications and it has a special place in theoretical computer science as one of the first problems to which approximation algorithms were suggested and analyzed with comparison to the optimal algorithm. Bilò [2] was the first to study the Bin Packing problem from a game theoretic perspective. He proved that the Bin Packing game admits a pure Nash equilibrium and provided non-tight bounds on the Price of Anarchy. He also proved that the Bin Packing game converges to a pure Nash equilibrium in a finite sequence of selfish improving steps, starting from any initial configuration of the items; however, the number of steps may be exponential. The quality of pure equilibria was further investigated by Epstein and Kleiman [8]. They proved that the Price of Stability of the Bin Packing game equals to 1, and showed almost tight bounds for the  $PoA$ ; namely, an upper bound of 1.6428 and a lower bound of 1.6416. Interestingly, this implies that the Price of Anarchy is not equal to the approximation ratio of any natural algorithm for Bin Packing. Yu and Zhang [20] later designed a polynomial time algorithm to compute a packing that is a pure Nash equilibrium. Finally, the  $SPoA$  was analyzed in [8].

A natural algorithm for the Bin Packing problem is the *Subset Sum* algorithm (or SS algorithm for short). In each iteration, the algorithm finds among the unpacked items, a maximum size set of items that fits into a new bin. The first mention of the *Subset Sum* algorithm in the literature is by Graham [10] who showed that its worst-case approximation ratio  $R_{SS}^\infty$  is at least  $\sum_{i=1}^{\infty} \frac{1}{2^i-1} \approx 1.6067$ . He also conjectured that this was indeed the true approximation ratio of this algorithm. The SS algorithm can be regarded as a refinement of the First-Fit algorithm [13], whose approximation ratio is known to be 1.7. Caprara and Pferschy [3] gave the first non-trivial bound on the worst-case performance of the SS algorithm, by showing that  $R_{SS}^\infty(1)$  is at most  $\frac{4}{3} + \ln \frac{4}{3} \approx 1.6210$ . They also generalized their results to the parametric case, giving lower and upper bounds on  $R_{SS}^\infty(\alpha)$  for  $\alpha < 1$ . In a follow-up paper, Caprara and Pferschy [4] considered two variants of the SS algorithm, called LSS and LRSS, that give preference to large items. They proved a nontrivial upper bound of  $\frac{13}{9} < 1.6067$  and gave lower bounds of 1.3643 and 1.30333 on the performances of LSS and LRSS, respectively. (Our results actually provide an improved lower bound of 1.3766 for LSS, which behaves like SS in the parametric case  $\alpha = 1/2$ .)

Surprisingly, the approximation ratio of the *Subset Sum* is deeply related to the Strong Price of Anarchy of the Bin Packing game. Indeed, the two concepts are equivalent [8]: Every output of the SS algorithm is a strong Nash equilibrium, and every strong Nash equilibrium is the output of some execution of the SS algorithm. Epstein and Kleiman [8] used this fact to show the existence of strong

equilibria for the Bin Packing game and to characterize the  $SPoA/SPoS$  in terms of this approximation ratio.

**Our results.** In this paper, we fully resolve the long standing open problem of finding the exact approximation ratio of the Subset Sum algorithm, proving Graham’s conjecture to be true. This in turn implies a tight bound on the Strong Price of Anarchy of the Bin Packing game. Then we extend this result to the parametric variant of Bin Packing where item sizes are all in an interval  $(0, \alpha]$  for some  $\alpha < 1$ . Interestingly, the ratio  $R_{SS}^\infty(\alpha)$  lies strictly between the upper and lower bounds of Caprara and Pferschy [3] for all  $\alpha \leq \frac{1}{2}$ . Finally, we study the pure Price of Anarchy for the parametric variant and show nearly tight upper bounds and lower bounds on it for any  $\alpha < 1$ . The tight bound of 1 on the Price of Stability proved in [8] for the general unrestricted Bin Packing game trivially carries over to the parametric case.

The main analytical tool we use to derive the claimed upper bounds is *weighting functions*—a technique widely used for the analysis of algorithms for various packing problems [13,15,19] and other greedy heuristics [11,12]. The idea of such weights is simple. Each item receives a weight according to its size and its assignment in some fixed  $NE$  packing. The weights are assigned in a way that the cost of the packing (the number of the bins used) is close to the total sum of weights. In order to complete the analysis, it is usually necessary to bound the total weight that can be packed into a single bin of an optimal solution.

Due to lack of space some proof are omitted [1].

## 2 Tight Worst-Case Analysis of the Subset Sum Algorithm

In this section we prove tight bounds for the worst-case performance ratio of the *Subset Sum* (SS) algorithm for any  $\alpha$ . It was proved in [8] that the strong equilibria coincide with the packings produced by the SS algorithm for Bin Packing. The equivalence for the  $SPoA$ ,  $SPoS$  and the worst-case performance ratio of the *Subset Sum* algorithm which was also proved in [8] still applies for the Parametric Bin Packing game; indeed, it holds for all possible lists of items (players), and in particular to lists where all items have size at most  $\alpha$ . This allows us to characterize the  $SPoA/SPoS$  in terms of  $R_{SS}^\infty(\alpha)$ .

First we focus on the unrestricted case, that is,  $\alpha = 1$ . Let  $\mathcal{B}_I$  be the set of bins used by our algorithm and  $\mathcal{O}_I$  be the optimal packing for some instance  $I$ . We are interested in the asymptotic worst-case performance of SS; namely, we want to identify constants  $\rho_{SS}$  and  $\delta_{SS}$  such that

$$|\mathcal{B}_I| \leq \rho_{SS} |\mathcal{O}_I| + \delta_{SS}. \quad (1)$$

Using the weighting functions technique, we charge the “cost” of the packing to individual items and then show for each bin in  $\mathcal{O}_I$  that the overall charge (weight) to items in the bin is not larger than  $\rho_{SS}$ .

<sup>1</sup> A full version of the paper is available at <http://arxiv.org/abs/0907.4311>

Let  $B \subseteq I$  be a bin in  $\mathcal{B}_I$ . We use the following short-hand notation  $s(B) = \sum_{j \in B} s_j$  and  $\min(B) = \min_{j \in B} s_j$ . Let  $s_{\min}$  be the size of the smallest yet-unpacked item just before opening  $B$ . For every  $i \in B$  we will charge item  $i$  a share  $w_i$  of the cost of opening the bin, where

$$w_i = \begin{cases} \frac{s_i}{s(B)} & \text{if } 1 - s_{\min} \leq s(B), \\ s_i & \text{otherwise.} \end{cases} \quad (2)$$

These weights are very much related to the payments of selfish players (items) in the Bin Packing game.

Let  $w(B)$  denote the total weight of items in a bin  $B$ . Note that if the size of items packed in  $B$  is large enough ( $s(B) \geq 1 - s_{\min}$ ) then  $w(B) = 1$  and thus the charged amount is enough to pay for  $B$ . Otherwise the charged amount only pays for a  $s(B)$  fraction of the cost. Let  $\hat{B}_1, \dots, \hat{B}_r$  be the bins that are underpaid listed in the order they are opened by the algorithm and let  $s_{\min}^i$  be the smallest item available when  $\hat{B}_i$  was opened. Notice that  $s_{\min}^i$  must belong to  $\hat{B}_i$  otherwise we could safely add the item to the bin. Also note that we cannot add  $s_{\min}^{i+1}$  to  $s(\hat{B}_i)$ , so we get

$$s(\hat{B}_i) + s_{\min}^{i+1} > 1 \implies s_{\min}^{i+1} > s_{\min}^i.$$

Therefore, because of the definition of the SS heuristic, for all  $i < r$ , it must be case that swapping  $s_{\min}^i$  with  $s_{\min}^{i+1}$  in  $\hat{B}_i$  must yield a set that cannot be packed into a single bin, so we get

$$s(\hat{B}_i) - s_{\min}^i + s_{\min}^{i+1} > 1 \implies 1 - s(\hat{B}_i) < s_{\min}^{i+1} - s_{\min}^i.$$

The total amount that is underpaid by all the  $\hat{B}_i$  bins can be bounded as follows

$$\sum_{i=1}^r (1 - s(\hat{B}_i)) \leq \sum_{i=1}^{r-1} (s_{\min}^{i+1} - s_{\min}^i) + (1 - s_{\min}^r) \leq 1.$$

This amount will be absorbed by the additive constant term  $\delta_{SS}$  in our asymptotic bound [\(II\)](#).

Let  $O$  be a set of items that can fit in a single bin, that is  $s(O) \leq 1$ , and denote with  $s_1, s_2, \dots, s_r$  the items contained in  $O$ , listed in *reverse order* of how our algorithm packs them. Our goal is to show that  $\sum_{i \in O} w_i$  is not too big. To that end, we first establish some properties that these values must have and then set up a mathematical program to find the sizes  $s_1, \dots, s_r$  obeying these properties and maximizing  $w(O)$ . Consider the point in time when our algorithm packs  $s_i$ . Let  $B$  be the bin the algorithm uses to pack  $s_i$  and let  $O_i = \{1, \dots, i\}$ .

Because  $O_i$  is a candidate bin for our algorithm we get  $s(B) \geq s(O_i)$ . Therefore, by [\(2\)](#), we have

$$w_i \leq \frac{s_i}{s(O_i)}. \quad (3)$$

Notice that if  $s(B) < 1 - \min(O_i)$  then  $i$ 's share is  $s_i$ . Therefore, we always have

$$w_i \leq \frac{s_i}{1 - \min(O_i)}. \quad (4)$$



Our job now is to find sizes  $s_1, \dots, s_r$  maximizing  $w(O)$  such that the weights obey (3) and (4). Equivalently, we are to determine the value of the following mathematical program

$$\text{maximize } \sum_{i=1}^r \frac{s_i}{\max \left\{ \sum_{j=1}^i s_j, 1 - \min_{1 \leq j \leq i} s_j \right\}} \quad (\text{MP}_r)$$

subject to

$$\begin{aligned} \sum_{i=1}^r s_i &\leq 1 \\ s_i &\geq 0 \quad \forall i \in [r] \end{aligned}$$

Let  $\lambda_r$  be the value of (MP<sub>r</sub>) and let  $\lambda = \sup_r \lambda_r$ . The following theorem shows that the worst-case approximation ratio of the SS algorithm is exactly  $\lambda$ .

**Theorem 1.** *For every instance  $I$ , we have  $|\mathcal{B}_I| \leq \lambda |\mathcal{O}_I| + 1$ . Furthermore, for every  $\delta > 0$ , there exists an instance  $I$  such that  $|\mathcal{B}_I| \geq (\lambda - \delta) |\mathcal{O}_I|$ .*

The necessary tools for proving the upper bound have been laid out above, we just need to put everything together:

$$|\mathcal{B}_I| \leq \sum_{B \in \mathcal{B}_I} \sum_{i \in B} w_i + 1 = \sum_{O \in \mathcal{O}_I} \sum_{i \in O} w_i + 1 \leq \sum_{O \in \mathcal{O}_I} \lambda_{|O|} + 1 \leq \lambda |\mathcal{O}_I| + 1.$$

To be able to prove the claimed lower bound, we first need to study some properties of (MP<sub>r</sub>). The following lemma fully characterizes the optimal solutions of (MP<sub>r</sub>). The proof of this lemma is omitted.

**Lemma 1.** *The optimal solution to (MP<sub>r</sub>) is*

$$s_i^* = \begin{cases} 2^{-i} & \text{if } i < r, \\ 2^{-r+1} & \text{if } i = r. \end{cases}$$

It follows that the optimal value of (MP<sub>r</sub>) is  $\lambda_r = \sum_{i=1}^{r-1} \frac{1}{2^i - 1} + \frac{1}{2^r - 1}$ . This expression increases as  $r$  grows. Therefore, the value is always at most

$$\lambda = \sum_{i=1}^{\infty} \frac{1}{2^i - 1}.$$

To lower bound the performance of the SS algorithm we use a construction based on Graham's original paper: The instance  $I$  has for each  $i \in [r - 1]$ ,  $N$  items of size  $2^{-i} + \varepsilon$ , and for  $i = r$ ,  $N$  items of size  $2^{-r+1} - r\varepsilon$ , where  $\varepsilon = 2^{-2r}$  and  $N$  is large enough so that  $N/s_i$  is integral for all  $i$ . The SS algorithm first packs the smallest items into  $N/2^{r-1}$  bins, then it packs the next smallest items into  $N/(2^{r-1} - 1)$  bins, the next items into  $N/(2^{r-2} - 1)$  bins, and so on. On the

other hand, the optimal solution uses only  $N$  bins. If we choose  $r$  to be such that  $2^r - 1 \geq \delta^{-1}$  then we get

$$|\mathcal{B}_I| = \lambda_r |\mathcal{O}_I| \geq \left( \lambda - \frac{1}{2^r - 1} \right) |\mathcal{O}_I| \geq (\lambda - \delta) |\mathcal{O}_I|.$$

Note that this lower bound example, for the case where there are  $r$  distinct item sizes, gives exactly the upper bound we found for  $\text{MP}_r$ .

**Corollary 1.** *For  $\alpha \in (\frac{1}{2}, 1]$ , the approximation ratio of the SS algorithm is  $R_{SS}^\infty(\alpha) = \sum_{i=1}^\infty \frac{1}{2^{i-1}} \approx 1.6067$ . Furthermore, the SPoA/SPoS of the BP( $\alpha$ ) game has the same value.*

**Parametric case.** To get a better picture of the performance of SS, we generalize Theorem 1 to instances where the size of the largest item is bounded by a parameter  $\alpha$ . Our goal is to establish the worst-case performance of the SS algorithm for instances in  $V_\alpha$  for all  $\alpha < 1$ .

Let  $t$  be the smallest integer such that  $\alpha \leq \frac{1}{t}$ . We proceed as we did before but with a slightly different weighting function:

$$w_i = \begin{cases} \frac{s_i}{s(\hat{B})} & \text{if } \max \{1 - s_{\min}, \frac{t}{t+1}\} \leq s(B), \\ s_i & \text{otherwise.} \end{cases} \tag{5}$$

As before there will be some bins that are underpaid. Let  $\hat{B}_1, \dots, \hat{B}_r$  be these bins and let  $s_{\min}^i$  be a smallest yet-unpacked item when the algorithm opened  $\hat{B}_i$ . These bins only pay for a  $s(\hat{B}_i)$  fraction of their cost. Even though we now have a more restrictive charging rule, the total amount underpaid is still at most 1. For all  $i < r$ , when  $s(\hat{B}_i) < 1 - s_{\min}^i$ , the same argument used above yields

$$1 - s(\hat{B}_i) < s_{\min}^{i+1} - s_{\min}^i.$$

Suppose that for some  $i$  we have  $s(\hat{B}_i) < \frac{t}{t+1}$  but  $s(\hat{B}_i) > 1 - s_{\min}^i$ . Note that this implies  $s_{\min}^i > 1/(t+1)$ . Since at this point every item has size in  $(\frac{1}{t+1}, \frac{1}{t}]$ , if there were at least  $t$  items left just before  $\hat{B}_i$  was opened, we could pack a bin with total size greater than  $\frac{t}{t+1}$ . Therefore,  $\hat{B}_i$  must be the last bin packed by the algorithm. Regardless whether such a bin exists or not, we always have  $1 - s(\hat{B}_r) \leq 1 - s_{\min}^r$ . Hence, the total amount underpaid is

$$\sum_{i=1}^r 1 - s(\hat{B}_i) \leq \sum_{i=1}^{r-1} (s_{\min}^{i+1} - s_{\min}^i) + (1 - s_{\min}^r) \leq 1.$$

The new weighting function (5) leads to the following mathematical program

$$\text{maximize } \sum_{i=1}^r \frac{s_i}{\max \left\{ \sum_{j=0}^i s_j, 1 - \min_{1 \leq j \leq i} s_j, t/(t+1) \right\}} \tag{MP_r^t}$$

subject to

$$\begin{aligned} \sum_{i=0}^r s_i &\leq 1 \\ s_i &\geq 0 && \forall i \in [r] \\ s_i &\leq 1/t && \forall i \in [r-1] \end{aligned}$$

Notice that  $s_r$  is allowed to be greater than  $1/t$ . This relaxation does not affect the value of the optimal solution, but it helps to simplify our analysis. From now on, we assume that  $r \geq t$ ; for otherwise the program becomes trivial. Define  $\lambda_r^t$  to be the value of  $(\text{MP}_r^t)$  and  $\lambda^t = \sup_r \lambda_r^t$ .

**Theorem 2.** *Let  $t \geq 2$  be an integer and  $\alpha \in (\frac{1}{t+1}, \frac{1}{t}]$ . For every instance  $I \in V_\alpha$ , we have  $|\mathcal{B}_I| \leq \lambda^t |\mathcal{O}_I| + 1$ . Furthermore, for every  $\delta > 0$ , there exist an instance  $I \in V_\alpha$  such that  $|\mathcal{B}_I| \geq (\lambda^t - \delta) |\mathcal{O}_I|$ .*

The proof of the upper bound is identical to that of Theorem 1. We only need to derive the counterpart of Lemma 1 for  $(\text{MP}_r^t)$ . Unlike its predecessor, Lemma 2 does not fully characterize the structure of the optimal solution of  $(\text{MP}_r^t)$ . Rather, we define an optimal solution  $s^*$  as a function of a parameter  $x$ . The proof this lemma is omitted.

**Lemma 2.** *An optimal solution to  $(\text{MP}_r^t)$  has the form*

$$s_i^* = \begin{cases} x & \text{if } i < t, \\ \frac{1-x(t-1)}{2^{i-t+1}} & \text{if } t \leq i < r, \\ \frac{1-x(t-1)}{2^{r-t}} & \text{if } i = r, \end{cases}$$

for some  $x \in [\frac{1}{t+1}, \frac{1}{t}]$ .

For any  $x \in [\frac{1}{t+1}, \frac{1}{t}]$ , we can construct a solution  $s^*$  for  $(\text{MP}_r^t)$  as described in Lemma 2. Let  $\lambda_r^t(x)$  be the value of this solution, that is,

$$\lambda_r^t(x) = x(t-1) \frac{t+1}{t} + \sum_{i=1}^{r-t} \frac{1}{\frac{2^i}{1-(t-1)x} - 1} + \frac{1}{\frac{2^{r-t}}{1-(t-1)x}}.$$

For any fixed  $x$ , the quantity  $\lambda_r^t(x)$  increases as  $r \rightarrow \infty$ . Therefore, it is enough to look at its limit value, which we denote by  $\lambda^t(x)$ :

$$\lambda^t(x) = \lim_{r \rightarrow \infty} \lambda_r^t(x) = x(t-1) \frac{t+1}{t} + \sum_{i=1}^{\infty} \frac{1}{\frac{2^i}{1-(t-1)x} - 1}.$$

It only remains to identify the value  $x \in [\frac{1}{t+1}, \frac{1}{t}]$  maximizing  $\lambda^t(x)$ . The answer is given by the following lemma, which we state without proof.

**Lemma 3.** *For every  $t \geq 2$ , the function  $\lambda^t(x)$  in the domain  $[\frac{1}{t+1}, \frac{1}{t}]$  attains its maximum at  $x = \frac{1}{t+1}$ .*

It follows that  $\lambda^t = \lambda^t(\frac{1}{t+1})$ , that is,

$$\lambda^t = 1 + \sum_{i=1}^{\infty} \frac{1}{(t+1)2^i - 1}.$$

Note that for a specific value of  $r$ ,

$$\lambda_r^t(\frac{1}{t+1}) = 1 + \sum_{i=1}^{r-t-1} \frac{1}{(t+1)2^i - 1} + \frac{1}{(t+1)2^{r-t-1}}.$$

For the lower bound on the performance of the SS algorithm, consider the instance  $I$  that for each  $i \in [t]$  has  $N$  items of size  $\frac{1}{t+1} + \varepsilon$ , for each  $i \in (t, r)$ , it has  $N$  items of size  $\frac{1}{(t+1)2^{i-t}} + \varepsilon$ , and for  $i = r$ , there are  $N$  items of size  $\frac{1}{(t+1)2^{r-1-t}} - r\varepsilon$ , where  $\varepsilon = \frac{1}{(t+1)2^{2-2r}}$  and  $N$  is large enough so that  $N/s_i$  is integral for all  $i$ . The SS algorithm first packs the smallest items into  $\frac{N}{(t+1)2^{r-t-1}}$  bins, then it packs the next smallest items into  $\frac{N}{(t+1)2^{r-1-t-1}}$  bins, and so on until reaching the items of size  $\frac{1}{t+1} + \varepsilon$  which are packed into  $N$  bins. The optimal solution uses  $N$  bins. If we choose  $r$  to be such that  $(t+1)2^{r-t} - 1 \geq \delta^{-1}$  then

$$|\mathcal{B}_I| = \lambda_r^t(\frac{1}{t+1}) |\mathcal{O}_I| \geq \left( \lambda^t - \frac{1}{(t+1)2^{r-t} - 1} \right) |\mathcal{O}_I| \geq (\lambda^t - \delta) |\mathcal{O}_I|.$$

**Corollary 2.** For each integer  $t \geq 1$  and  $\alpha \in (\frac{1}{t+1}, \frac{1}{t}]$ , the SS algorithm has an approximation ratio of  $R_{SS}^\infty(\alpha) = 1 + \sum_{i=1}^\infty \frac{1}{(t+1)2^i - 1}$ . Furthermore, the SPoA/SPoS of the BP( $\alpha$ ) game has the same value.

Table 1 compares our bound with the previously known upper bounds and lower bounds of Caprara and Pferschy [3]. Note that the true ratio lies strictly between previous bounds.

**Table 1.** Comparison of the worst-case ratio of FFD, SS, FF and PoA as a function of  $\alpha$  when  $\alpha \leq \frac{1}{t}$ , for  $t = 1, \dots, 10$

	$R_{FFD}(\alpha)$ [13]	CP lb [8]	$R_{SS}(\alpha)$	CP ub [3]	PoA( $\alpha$ )	$R_{FF}(\alpha)$ [13]
$t = 1$	1.222222	1.606695 [10]	1.606695	1.621015	[1.641632, 1.642857] [8]	1.700000
$t = 2$	1.183333	1.364307	1.376643	1.398793	[1.464571, 1.466667]	1.500000
$t = 3$	1.166667	1.263293	1.273361	1.287682	[1.326180, 1.326530]	1.333333
$t = 4$	1.150000	1.206935	1.214594	1.223143	[1.247771, 1.247863]	1.250000
$t = 5$	1.138095	1.170745	1.176643	1.182321	[1.199102, 1.199134]	1.200000
$t = 6$	1.119048	1.145460	1.150106	1.154150	[1.166239, 1.166253]	1.166667
$t = 7$	1.109127	1.126763	1.130504	1.133531	[1.142629, 1.142635]	1.142857
$t = 8$	1.097222	1.112360	1.115433	1.117783	[1.124867, 1.124871]	1.125000
$t = 9$	1.089899	1.100918	1.103483	1.105360	[1.111029, 1.111031]	1.111111
$t = 10$	1.081818	1.091603	1.093776	1.095310	[1.099946, 1.099947]	1.100000

### 3 Analysis of the Price of Anarchy

We now provide a lower bound for the Price of Anarchy of the parametric Bin Packing game with bounded size items and, in addition, prove a very close upper bound for each value of  $\frac{1}{t+1} < \alpha \leq \frac{1}{t}$  for a positive integer  $t \geq 2$ , that is, for all  $0 < \alpha \leq \frac{1}{2}$ . The case  $\frac{1}{2} < \alpha < 1$  ( $t = 1$ ) was extensively studied in [8]. Due to lack of space, all proofs in this section are omitted.

#### A construction of lower bound on the $PoA$ of parametric Bin Packing.

In this section we give the construction of a lower bound on  $PoA(\alpha)$ . For each value of  $t \geq 2$  we present a set of items which consists of multiple item lists. This construction is related to the one that we gave in [8] for  $\frac{1}{2} < \alpha \leq 1$ , though it is not a generalization of the former, which strongly relies on the fact that each item of size larger than  $\frac{1}{2}$  can be packed alone in a bin of the  $NE$  solution, whereas in the parametric case there are no such items. It is based upon techniques that are often used to design lower bounds on Bin Packing algorithms (see for example [15]), but it differs from these constructions in the notion of order in which packed bins are created (which does not exist here) and the demand that each bin satisfies the Nash stability property. Our lower bound is given by the following theorem.

**Theorem 3.** *For each integer  $t \geq 2$  and  $\alpha \in (\frac{1}{t+1}, \frac{1}{t}]$ , the  $PoA$  of the  $BP(\alpha)$*

*game is at least* 
$$\frac{t^2 + \sum_{j=1}^{\infty} (t+1)^{-j} \cdot 2^{-j(j-1)/2}}{t(t-1)+1}.$$

**An upper bound on the  $PoA$  of parametric Bin Packing.** We now provide a close upper bound on  $PoA(\alpha)$  for a positive integer  $t \geq 2$ . The technique used in [8] can be considered as a refinement of the one we use here. However, here we are required to use additional combinatorial properties of the  $NE$  packing. To bound the  $PoA$  from above, we prove the following theorem.

**Theorem 4.** *For each integer  $t \geq 2$  and  $\alpha \in (\frac{1}{t+1}, \frac{1}{t}]$ , the  $PoA$  of the parametric Bin Packing game  $BP(\alpha)$  is at most* 
$$\frac{2t^3 + t^2 + 2}{(2t+1)(t^2 - t + 1)}.$$

### 4 Concluding Remarks

In order to illustrate the results in the paper, we report in Table [ ] the values for the worst-case ratio of the  $SS$  algorithm for various values of  $\alpha$  along with previously known upper and lower bounds of Caprara and Pferschy [3], and the worst-case approximation ratios of  $FF$  and  $FFD$  algorithm Bin Packing. We also include the range of possible values for the  $PoA$  for different values of  $\alpha$ . We conjecture that the true value of the  $PoA$  equals our lower bound from Theorem [3].

## References

1. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: Proc. of SODA, pp. 189–198 (2007)
2. Bilò, V.: On the packing of selfish items. In: Proc. of IPDPS. IEEE, Los Alamitos (2006)
3. Caprara, A., Pferschy, U.: Worst-case analysis of the subset sum algorithm for bin packing. *Oper. Res. Lett.* 32(2), 159–166 (2004)
4. Caprara, A., Pferschy, U.: Modified subset sum heuristics for bin packing. *Inf. Process. Lett.* 96(1), 18–23 (2005)
5. Coffman Jr., E., Csirik, J.: Performance guarantees for one-dimensional bin packing. In: Gonzalez, T.F. (ed.) *Handbook of Approximation Algorithms and Metaheuristics*, ch. 32, pp. (32-1)–(32-18). Chapman & Hall/CRC, Boca Raton (2007)
6. Csirik, J., Leung, J.Y.-T.: Variants of classical one-dimensional bin packing. In: Gonzalez, T.F. (ed.) *Handbook of Approximation Algorithms and Metaheuristics*, ch. 33. Chapman & Hall/CRC, Boca Raton (2007)
7. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms* 3(1) (2007)
8. Epstein, L., Kleiman, E.: Selfish bin packing. In: Halperin, D., Mehlhorn, K. (eds.) *Esa 2008*. LNCS, vol. 5193, pp. 368–380. Springer, Heidelberg (2008)
9. 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)
10. Graham, R.L.: Bounds on multiprocessing anomalies and related packing algorithms. In: Proc. of the 1972 Spring Joint Computer Conference, pp. 205–217 (1972)
11. Immorlica, N., Mahdian, M., Mirrokni, V.S.: Cycle cover with short cycles. In: Proc. of STACS, pp. 641–653 (2005)
12. Jain, K., Mahdian, M., Markakis, E., Saberi, A., Vazirani, V.V.: Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *Journal of the ACM* 50(6), 795–824 (2003)
13. Johnson, D.S., Demers, A.J., Ullman, J.D., Garey, M.R., Graham, R.L.: Worst-case performance bounds for simple one-dimensional packing algorithms. *SIAM J. Comput.* 3(4), 299–325 (1974)
14. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: Proc. of STACS, pp. 404–413 (1999)
15. Lee, C.C., Lee, D.T.: A simple online bin packing algorithm. *J. ACM* 32, 562–572 (1985)
16. Mavronicolas, M., Spirakis, P.G.: The price of selfish routing. In: Proc. of STOC, pp. 510–519 (2001)
17. Roughgarden, T.: *Selfish routing and the price of anarchy*. MIT Press, Cambridge (2005)
18. Roughgarden, T., Tardos, É.: How bad is selfish routing? In: Proc. of FOCS, pp. 93–102 (2000)
19. Ullman, J.D.: The performance of a memory allocation algorithm. Technical Report 100. Princeton University, Princeton (1971)
20. Yu, G., Zhang, G.: Bin packing of selfish items. In: Papadimitriou, C., Zhang, S. (eds.) *WINE 2008*. LNCS, vol. 5385, pp. 446–453. Springer, Heidelberg (2008)

# An Online Multi-unit Auction with Improved Competitive Ratio<sup>\*</sup>

Sourav Chakraborty<sup>1</sup> and Nikhil Devanur<sup>2</sup>

<sup>1</sup> CWI, Amsterdam, Netherlands  
sourav.chakraborty@cwi.nl

<sup>2</sup> Microsoft Research, Redmond  
nikdev@microsoft.com

**Abstract.** We improve the best known competitive ratio (from  $1/4$  to  $1/2$ ), for the online multi-unit allocation problem, where the objective is to maximize the single-price revenue. Moreover, the competitive ratio of our algorithm tends to 1, as the bid-profile tends to “smoothen”. This algorithm is used as a subroutine in designing truthful auctions for the same setting: the allocation has to be done online, while the payments can be decided at the end of the day. Earlier, a reduction from the auction design problem to the allocation problem was known only for the unit-demand case. We give a reduction for the general case when the bidders have decreasing marginal utilities. The problem is inspired by sponsored search auctions.

## 1 Introduction

It is fairly common that a mechanism has to work in a dynamic environment, where there is an uncertainty in either the demand, or the supply, or both. This has led to the study of *online mechanism design* [7,4,13] and has presented significant new challenges compared to the traditional static setting. Most of the research has focused on dynamic demand case: the uncertainty is in the number and types of the bidders, their arrival and departure time, etc, such as airline tickets. On the other hand, very little is known for dynamic supply case: the uncertainty is in the number of items to be allocated, or more generally the set of feasible allocations, such as sponsored search. Mahdian and Saberi [14] initiated the study of the dynamic supply case by giving a constant competitive ratio algorithm for auctioning multiple copies of a single item with unit-demand bidders. We improve their competitive ratio by a factor of 2 by giving an alternate and simple algorithm, and also extend their results to handle bidders with multiple demand.

A bidder with *unit-demand* has a value  $u_i$  for one copy of the item, and his utility is  $u_i - p$  if he is allocated the item at price  $p$ , and 0 otherwise.

---

<sup>\*</sup> Part of the work was done when both the authors were visiting Microsoft Research, India in Summer 2007 and when the first author was a graduate student at the University of Chicago and the second author was in Toyota Technological Institute, Chicago.

**Definition 1 (Online Multi-unit Auction Problem, unit-demand ([14])).** *At the beginning of the auction, each bidder (with unit-demand) bids a value  $b_i$ . At each (discrete) time unit, a new copy arrives which must be allocated to a bidder immediately, or else it perishes. When there are no more copies left, the auction determines the prices charged to the winning bidders.*

Note that the auction has no prior knowledge of how many copies of the item will be produced.

**Definition 2.** *An auction is truthful if bidding  $b_i = u_i$  is a dominant strategy for each bidder.*

The goal of the auction is to maximize the revenue of the auctioneer, which is the sum of the prices charged to the winning bidders. The main motivation behind the work of [14] was sponsored search auctions, which are a major source of revenue for search engines like Google, Yahoo and MSN. The bidders are the advertisers and the items correspond to search queries. The queries arrive online and have to be allocated immediately, while the advertisers stay for the entire duration of the auction and present their bids ahead of time. The advertisers are only charged at the end of the day.

An alternate model is to ask that the prices are also determined online. For the sponsored search auction setting, charging at the end is closer to reality. Also, charging online seems to be considerably restrictive, as there are strong lower bounds for this model [1]. For the sponsored search auction setting, a more realistic model is when the bidders have multiple demand. We present an auction for this case as well, and our results for this case are of significant interest.

The auction problem considered here is also a natural extension of the line of work on digital goods auction: from unlimited supply ([9],[10],[11],[12]) to limited supply ([13],[5]), to unknown supply ([14] and this paper).

As is standard in the literature on digital goods auction, we give a competitive analysis of the auction, by comparing the revenue of the auction to a benchmark. The benchmark we use is once again a standard in digital goods auction, it is the *optimal single-price revenue* on hindsight:  $OPT := \max_p p \cdot |\{i : b_i \geq p\}|$ . The auction itself is allowed to charge different prices to different bidders, although our auction charges only two different prices.

**Definition 3 (Competitive Ratio).** *An auction is said to have a competitive ratio of  $\alpha$  if the expected revenue of the auction is at least  $\alpha OPT$ .*

[14] gave a reduction from the auction problem to the following algorithmic problem, with only a constant factor lost in the competitive ratio.

**Definition 4 (Online Multi-unit Allocation Problem, unit-demand).** *The algorithm is given the utility  $u_i$  of each bidder. At each (discrete) time*

---

<sup>1</sup> The lower bounds [2] are for a related problem, that of maximizing social welfare. It is an interesting open question if these lower bounds also hold for maximizing revenue.



unit, a new copy arrives which must be allocated to a bidder immediately, or else it perishes. When there are no more copies left, the algorithm charges all the winning bidders with a single price, that is smaller than their utilities.

There is no requirement of truthfulness in the allocation problem. Also, the algorithm itself has to charge the same price to all the bidders, unlike the auction which was allowed to charge different prices to different bidders. Also note that the revenue-maximizing single price is determined by the allocations made by the algorithm. It is simply the smallest winning utility. As with the auction problem, we compare the revenue of the algorithm with the optimal single price revenue on hindsight; the competitive ratio is defined analogously.

**Theorem 1.** ([14]) *There is a truthful mechanism for the online multi-unit auction problem with unit-demand bidders with a competitive ratio of  $O(\alpha)$  given an algorithm for the allocation problem with competitive ratio  $\alpha$ .*

## 1.1 Main Result

It can be easily seen that the competitive ratio of any deterministic algorithm for the allocation problem is arbitrarily small. So it is actually surprising that a randomized algorithm can even get a constant competitive ratio. The reason for this difficulty is that the revenue of the algorithm, as a function of the number of copies allocated can have many “peaks” and “valleys”. For any deterministic algorithm, an adversary can make sure that the algorithm either ends up in a valley, or is stuck on a small peak while the optimum is at a larger peak elsewhere. The key decision for an algorithm is when it is at a peak, it has to decide if it has to stay at the peak, or try to get to the next one. What our algorithm does is to simply wait at the current peak for a period of time chosen uniformly at random between 1 and the maximum distance between peaks seen so far. The simplicity of our algorithm is quite appealing. This improves the best known competitive ratio (from  $1/4$  to  $1/2$ ), for the online multi-unit allocation problem, which in turn gives a factor of 2 improvement for the online multi-unit auction problem.

**Theorem 2.** *There is an algorithm for the online multi-unit allocation problem for unit demand, that achieves a competitive ration of  $1/2$ .*

The proof of the competitive ratio relies on case analysis since the optimal revenue and the expected revenue of the algorithm vary depending on the total number of copies seen. A good idea of how the analysis goes can be had by considering the following instance: suppose there is one bid of 1 and many bids of  $\epsilon \ll 1$ . In this case the algorithm waits for a time chosen u.a.r between 1 and  $1/\epsilon$ . If the number of copies seen is  $m \leq 1/\epsilon$ , then the optimal revenue is 1, while the expected revenue is  $1 - x + \frac{x^2}{2}$  (where  $x = \epsilon m$ ), which is at least  $1/2$  when  $x \leq 1$ . If  $m \geq 1/\epsilon$  then the optimal revenue is  $\epsilon m$ , while the expected revenue is  $\epsilon m - 1/2 \geq \frac{\epsilon m}{2}$ .

Moreover, the competitive ratio of our algorithm tends to 1, as the bid-profile tends to “smoothen”. [14] also showed an upper bound of  $e/(e+1)$  for the allocation problem and closing the gap is an open problem. See Section 4 for a more detailed discussion on this.

Mahdian and Saberi [14] showed that using an algorithm for the online multi-unit allocation problem for unit demand with competitive ratio  $\rho$  one can construct a truthful auction for the online multi-unit auction problem with competitive ratio  $\rho/20$ . Thus,

**Corollary 1.** *There is a truthful auction for the online multi-unit auction problem, that achieves a constant competitive ratio.*

## 1.2 Extensions

A more realistic case in the context of sponsored search auction is when the bidders have multiple demand: bidders have decreasing marginal utilities for multiple copies of the item, and submit multiple bids. The optimum and the competitive ratio are defined analogous to the unit-demand case. The allocation problem remains the same even with multiple demands, since the problem does not really depend on the identity of the bidders. Hence, our algorithm for the allocation problem gives a competitive ratio of  $1/2$  even for this case.

However, the auction problem is harder with multiple demands, since it provides more ways for the bidders to lie and benefit. In particular, the auction obtained by using the reduction in [14] is not truthful for multiple demands. The reduction in [14] is based on random sampling with computing optimal “price offers”. But when run in an online setting, the prices offered decrease over time, due to which a bidder might regret not getting a copy earlier as the price decreased at a later time. The reduction in [14] takes care of this situation by a clever implementation that works only when all bidders want only one copy. It is not truthful when the bidders can submit multiple bids. We circumvent this difficulty by combining the random sampling technique with the VCG auction. However, we only get an asymptotic competitive ratio, that is the ratio tends to  $1/2$ , as a certain bidder dominance parameter tends to 0. The bidder dominance parameter is defined to be the maximum fraction of the optimum revenue that can be obtained from any single bidder. A small bidder dominance parameter indicates that the revenue from any one bidder is small compared to the optimal revenue.

**Definition 5.** *For any price  $p$  and any bidder  $i$  we denote by  $n(i, p)$  the number of bids of bidder  $i$  that are more than  $p$ . The bidder dominance parameter is*

$$\eta := \frac{\max_{i,p} n(i, p)p}{OPT}.$$

**Theorem 3.** *There is a truthful mechanism for the online multi-unit auction problem with multiple-demand bidders, that with probability more than  $(1 - \delta)$*

guarantees a revenue of at least  $\alpha OPT(1 - \epsilon)$  on expectation, where  $\alpha$  is the competitive ratio of the allocation algorithm that we use as the subroutine, if

$$\eta = O\left(\epsilon^2 / \log\left(\frac{n}{\delta}\right)\right),$$

where  $n$  is the number of distinct bid values.

The problem considered here is perhaps the simplest non-trivial case of the actual problem in sponsored search auctions. There are many extensions of which we have little understanding, for instance, one could consider multiple slots for every query. Another interesting extension is when the bidders have constant marginal utilities for the copies, but have daily budgets. [5][1] gave an auction for this case with known supply (the offline problem). Extending it to the online setting is an important open problem. The introduction of budgets also makes the multiple items case interesting. (Otherwise, assuming additive utilities, the auctions for different items are independent of each other.) Even the offline case of this problem is open.

**Subsequent Related Work:** Subsequent to our result, Devanur and Hartline [6] gave an alternate auction for the Online Multi-Unit Auction problem with a competitive ratio that is better than this paper. This auction does not use the reduction to the allocation problem. However, the auction in [6] is only for unit-demand bidders, so our results for the multiple-demand bidders are still the best. Also the online allocation problem, and the algorithm for it are interesting in their own right.

**Organization:** We present our algorithm for the Online Multi-unit Allocation problem in Section 2. Theorem 2. For lack of space we are unable to present the proof of the competitive ratio of the algorithm in this extended abstract. The auction for the multiple demands case and a sketch of the proof of Theorem 3 is given in Section 3. Section 4 contains a discussion on future work and open problems.

## 2 Algorithm for the Online Multi-unit Allocation Problem

Without loss of generality assuming that the utilities are  $u_1 \geq u_2 \geq \dots \geq u_n$ , the revenue obtained by allocating  $l$  units of the item is  $lu_l$ . Let  $1 = a_1 < b_1 < a_2 < b_2 < a_3 < b_3 < \dots$  be the critical points of the function  $lu_l$ , that is, the function  $lu_l$  is non-decreasing as  $l$  increases from  $a_i$  to  $b_i$  and for all  $b_i < l < a_{i+1}$  we have  $b_i u_{b_i} > lu_l$  and  $b_i u_{b_i} \leq a_{i+1} u_{a_{i+1}}$ .

The algorithm is in one of two states, ALLOCATE or WAIT. When it is in ALLOCATE, it allocates the next copy of the item. When it is in WAIT, it discards the next copy. The description of the algorithm is completed by specifying when it transits from one state to the other.

The algorithm is initially in ALLOCATE. It transits from ALLOCATE to WAIT when the number of copies allocated ( $X$ ) is equal to  $b_i$  for some  $i$ . It

transits from WAIT to ALLOCATE when the number of copies discarded till then ( $Y$ ) is equal to a random variable,  $T$ , for waiting time.  $T$  is reset every time the algorithm transits to WAIT.  $T$  is picked so that it is distributed uniformly between 0 and  $D_i$ , where  $D_0 = 0$  and for all  $i \geq 1$

$$D_i = \max_{j \leq i} (a_{j+1} - b_j)$$

(recall that  $X = b_i$ ). We further want to maintain the invariant that  $Y$  never exceeds  $T$ . Equivalently, the value of  $T$  can only increase during a run of the algorithm.

We still have to specify how  $T$  is picked. Because of the condition that  $T$  can only increase, we cannot pick  $T$  independently every time we transit to WAIT. If  $D_i \leq D_{i-1}$ , then we don't have to change  $T$  at all. If  $D_i > D_{i-1}$ , then

- w.p.  $\frac{D_{i-1}}{D_i}$  don't change  $T$ ,
- with the remaining probability pick  $T$  uniformly at random from the interval  $[D_{i-1}, D_i]$ .

It is easy to see that the resulting  $T$  is distributed uniformly in  $[0, D_i]$ . Note that in case  $T$  is not changed, then  $Y$  is already equal to  $T$ , and we transit back to ALLOCATE immediately. Equivalently, we don't transit to WAIT at all.

### Pseudocode for the Algorithm

1. initialize STATE = ALLOCATE,  $i=1$ ,  $X=Y=T=0$ ;
2. when a new copy is produced
3.     If (STATE = ALLOCATE)
4.         Allocate the copy to the next bidder;
5.          $X++$ ;
6.         If ( $X = b_i$ )
7.             If ( $D_i > D_{i-1}$ )
8.                 With prob  $1 - \frac{D_{i-1}}{D_i}$
9.                     set  $T$  to a random number from the interval  $[D_{i-1}, D_i]$ ;
10.                     STATE = WAIT;
11.              $i++$ ;
12.         If (STATE = WAIT)
13.             Discard the copy;
14.              $Y++$ ;
15.             If ( $Y = T$ )
16.                 STATE = ALLOCATE
17.         GO TO line 2.

Because of shortage of space we cannot present the analysis for competitive ratio in this extended abstract.

## 3 Bidders with Multiple Demand

Let  $\mathbf{B} = \{1, 2, \dots, n\}$  be the set of bidders. Each bidder can make multiple bids. We will design a truthful mechanism which has good competitive ratio. Our

mechanism will use an online multi-unit allocation algorithm as a sub-routine. Under a bidder-dominance assumption, the competitive ratio of our mechanism will be  $(1 - \epsilon)\alpha$  where  $\alpha$  is the competitive ratio of the allocation algorithm we use as our subroutine.

**The Mechanism:** We divide the set of bidders into two groups  $S$  and  $T$  by placing each bidder randomly into either of the groups. On each set of bidders  $S$  and  $T$  we will have fictitious runs of the allocation algorithm. Let the fictitious run of the allocation algorithm on the set  $S$  (respectively  $T$ ) allocates  $x(S, k)$  (respectively  $x(T, k)$ ) copies when  $k$  copies are produced.

Now when the  $j$ -th copy is produced, if  $j$  is even we compute  $x(S, j/2)$ . If at that time the number of copies allocated to bidders in  $T$  is less than  $x(S, j/2)(1 - 6\gamma)$  then we allocate the  $j$ -th copy to  $T$  otherwise discard the copy. Similarly, if  $j$  is odd we compute  $x(T, (j + 1)/2)$  and if the number of copies allocated to bidders in  $S$  is less than  $x(T, (j + 1)/2)(1 - 6\gamma)$  then we allocate the  $j$ -th copy to  $S$  otherwise discard the copy.

Finally let  $x_{final}(S)$  and  $x_{final}(T)$  copies are allocated to bidders in  $S$  and  $T$  respectively. The prices charged are the VCG payments, that is, as if we ran a VCG auction to sell  $x_{final}(S)$  copies to bidders in  $S$ .

Note that the even indexed copies will be allocated only to bidders in  $T$  and the odd-indexed copies will be allocated only to bidders in  $S$ . But the bids of bidders in  $S$  decides how many (odd-indexed) copies will be allocated to bidders in  $T$  and vice versa. This mechanism is similar to that in [11] on digital good auction with unlimited supplies except that in [11] the bids of bidders in  $S$  decides the cut off price for bidders in  $T$  and vice-versa.

If  $M$  is the number of copies of the item that are finally produced we denote by  $OPT = OPT(\mathbf{B}, M)$  the revenue obtained by the optimal single price allocation algorithm.

**Definition 6.** For any price  $p$  and any bidder  $i$  we denote by  $n(i, p)$  the number of bids of bidder  $i$  that are more than  $p$ .

We define the bidder dominance parameter  $\eta$  as

$$\eta = \frac{\max_{i,p} n(i, p)p}{OPT}.$$

**Theorem 4.** The above mechanism is a truthful mechanism. If all the bids are from a finite set of prices (say  $Q$ ) and if

$$\frac{1}{\eta} = \Omega \left( \log \left( \frac{|Q|}{\delta} \right) \left( \frac{1}{\epsilon^2} \right) \right)$$

and if we set  $\gamma = \epsilon/8$  then with probability more than  $(1 - \delta)$  our mechanism guarantees a revenue of at least  $\alpha OPT(1 - \epsilon)$  on expectation, where  $\alpha$  is the competitive ratio of the allocation algorithm that we use as the subroutine.

In the rest of this section we will give a sketch of the proof of the theorem. The detailed proof of the theorem is in the Appendix. The proof is similar to that in [11].

The proof that the mechanism is truthful follows from the facts that the number of copies allocated to each half is independent of the number of the bids of the bidders in that half and the fact that pricing is determined by the VCG auction.

The proof of the competitive ratio has two main parts: The first thing is that since the bidders are split randomly into two sets so with high probability the optimal revenue we can obtain from either of the sets is nearly half of what we can obtain from the whole set.

The second thing is that the discounting factor of  $(1 - 6\gamma)$  ensures that with high probability the eventual winners in  $S$  (respectively  $T$ ) are charged at least as much as our allocation algorithm charges during its fictitious run on the set  $T$  (respectively  $S$ ).

Note that the bound on the bidder dominance gives us an upper bound on  $n(i, p)$  that is the number of bids on any bidders that is more than  $p$ . This is essential for our analysis.

Let a fictitious run of the optimal single price allocation algorithm on  $S$  generates a revenue of  $OPT(S, j)$  after  $j$  copies are produced. By McDiarmid's Inequality and the bound on the bidder dominance parameter, with probability at least  $(1 - O(\delta))$  we have  $OPT(S, \lceil M/2 \rceil) > (1/2 - \gamma)OPT$ , where  $M$  is the final number of copies produced. Similarly we have  $OPT(T, \lfloor M/2 \rfloor) > (1/2 - \gamma)OPT$ .

For the second stage we again notice that since the set of bidders was partitioned randomly so with high probability the set of bids that are more than  $p$  is also evenly divided among the two sets  $S$  and  $T$ . From the McDiarmid's Inequality and from the bound on the bidder dominance parameter we see that with high probability the number of bids in  $S$  that are more than  $p$  is much more than  $(1 - 6\gamma)$  times the number of bids in  $T$  that are more than  $p$  (and vice versa).

Let  $ALG(S, j)$  and  $ALG(T, j)$  be the revenue is generated by the fictitious run of our allocation algorithm on  $S$  and  $T$  respectively after  $j$  items are produced. Now since the allocation algorithm is  $\alpha$  competitive we have that on expectation  $ALG(S, j) > \alpha OPT(S, j)$ . Thus with probability at least  $(1 - O(\delta))$  the revenue we earned on expectation is more than

$$ALG(S, \lceil M/2 \rceil)(1 - 6\gamma) + ALG(T, \lfloor M/2 \rfloor)(1 - 6\gamma) > \alpha(1 - 6\gamma)(1 - 2\gamma)OPT$$

which is greater than  $\alpha(1 - 8\gamma)OPT$ .

## 4 Conclusion and Open Problems

The optimal competitive ratio for the allocation problem is open. [14] showed an upper bound of  $e/(e + 1)$  for any randomized algorithm. The instance for which

they show this upper bound is when there is one bid of 1 and many bids of  $\epsilon$ . For this particular instance, the following algorithm gets a competitive ratio of  $2/3$ : with probability  $1/3$ , allocate just one copy and get a revenue of 1, and with probability  $2/3$ , run our algorithm. We conjecture that this algorithm can be generalized to get a  $2/3$  competitive ratio. Also, a better upper bound proof will probably have to consider instances with multiple peaks, where the ratio of the  $D_i$ 's to the  $a_{i+1}$ 's is large.

For the auction problem, the competitive ratio for the unit-demand case is quite small, and that for the multiple demand case holds only asymptotically. Getting it to a reasonably large constant (or proving that it is impossible) is an important open problem.

The most common scenario in sponsored search auctions is that the bidders have a constant utility for multiple copies of the item, but with a daily budget. Our allocation algorithm works for this case as well, but the reduction from the auction problem is not truthful. Borgs et al [5] give a truthful auction for the offline case with budgets, using the standard random sampling techniques with price offers. However, it is not clear how to extend their auction to the online case. The difficulty is the same as that for the multiple-demand case, that the price offers are decreasing over time. But unlike the multiple-demand case, there is no VCG auction for the budgets case, so our reduction does not work.

## Acknowledgements

The authors would like to thank Jason Hartline for useful discussions about the digital goods auction, and Deeparnab Chakrabarty for comments on an earlier draft of the paper.

## References

1. Abrams, Z.: Revenue maximization when bidders have budgets. In: SODA 2006, pp. 1074–1082 (2006)
2. Babaioff, M., Blumrosen, L., Roth, A.L.: Auctions with Online Supply (manuscript)
3. Balcan, M.F., Devanur, N.R., Hartline, J., Talwar, K.: Random Sampling Auctions for Limited Supply (manuscript) (submitted)
4. Blum, A., Hartline, J.D.: Near-optimal online auctions. In: SODA 2005, pp. 1156–1163 (2005)
5. Borgs, C., Chayes, J.T., Immorlica, N., Mahdian, M., Saberi, A.: Multi-unit auctions with budget-constrained bidders. In: ACM Conference on Electronic Commerce 2005, pp. 44–51 (2005)
6. Devanur, N.R., Hartline, J.: Limited and Online Supply and the Bayesian foundations of prior-free mechanism design. To appear in ACM Conference on Electronic Commerce 2009 (2009)
7. Parkes, D.: Online mechanisms. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*, ch. 16. CUP (2007)

8. Feige, U., Flaxman, A., Hartline, J.D., Kleinberg, R.D.: On the Competitive Ratio of the Random Sampling Auction. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 878–886. Springer, Heidelberg (2005)
9. Goldberg, A.V., Hartline, J.D.: Envy-free auctions for digital goods. In: ACM Conference on Electronic Commerce 2003, pp. 29–35 (2003)
10. Goldberg, A.V., Hartline, J.D.: Competitiveness via consensus. In: SODA 2003, pp. 215–222 (2003)
11. Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M.E., Wright, A.: Competitive auctions. *Games and Economic Behavior* (55), 242–269 (2006)
12. Goldberg, A.V., Hartline, J.D., Wright, A.: Competitive auctions and digital goods. In: SODA 2001, pp. 735–744 (2001)
13. Hajiaghayi, M.T., Kleinberg, R.D., Parkes, D.C.: Adaptive limited-supply online auctions. In: ACM Conference on Electronic Commerce 2004, pp. 71–80 (2004)
14. Mahdian, M., Saberi, A.: Multi-unit Auction with Unknown Supply. In: ACM Conference on Electronic Commerce 2006, pp. 243–249 (2006)



# Prediction Mechanisms That Do Not Incentivize Undesirable Actions<sup>\*</sup>

Peng Shi, Vincent Conitzer, and Mingyu Guo

Department of Computer Science, Duke University,  
Durham, NC 27708, USA  
ps46@duke.edu, {conitzer,mingyu}@cs.duke.edu

**Abstract.** A potential downside of prediction markets is that they may incentivize agents to take undesirable actions in the real world. For example, a prediction market for whether a terrorist attack will happen may incentivize terrorism, and an in-house prediction market for whether a product will be successfully released may incentivize sabotage. In this paper, we study *principal-aligned* prediction mechanisms—mechanisms that do not incentivize undesirable actions. We characterize all principal-aligned proper scoring rules, and we show an “overpayment” result, which roughly states that with  $n$  agents, any prediction mechanism that is principal-aligned will, in the worst case, require the principal to pay  $\Theta(n)$  times as much as a mechanism that is not. We extend our model to allow uncertainties about the principal’s utility and restrictions on agents’ actions, showing a richer characterization and a similar “overpayment” result.

**Keywords:** Prediction Markets, Proper Scoring Rules, Mechanism Design.

## 1 Introduction

Prediction markets reward agents for accurately assessing the probability of a future event.<sup>1</sup> Typically, agents buy or sell securities according to their beliefs, and they are rewarded based on the outcome that materializes. Empirical studies suggest that prediction markets make very accurate predictions, sometimes beating the best experts and polls [11, 5]. Currently, online markets such as NewsFutures and Intrade elicit public predictions about a wide variety of topics, and many technology companies, including HP [12], Google, Microsoft, and Yahoo!, use in-house prediction markets to elicit employees’ predictions on future products.<sup>2</sup>

A potential downside of prediction markets is that they may incentivize agents to take undesirable actions in the real world, if those actions affect the probability

---

<sup>\*</sup> This work is supported by NSF IIS-0812113, the Sloan Foundation, and a Yahoo! Faculty Research Grant. We thank the anonymous reviewers for helpful comments.

<sup>1</sup> For literature reviews on prediction markets, see [11, 4, 15].

<sup>2</sup> NewsFutures and InKlingMarkets both provide services to help companies run in-house prediction markets, thus making running these markets accessible to non-technology-based companies.

of the event. For example, the idea of organizing markets to predict terrorist activity, which was once seriously considered by the U.S. Department of Defense, has been dismissed in part based on the consideration that terrorists may stand to profit from such markets<sup>3</sup>. As another example, consider a software company that runs an in-house prediction market to assess whether a product will be released on time. The company may be concerned that the market provides an incentive to an employee to sabotage a timely release, if the employee predicts a late release.

On the other hand, not all real-world actions are necessarily undesirable. A terrorism prediction market may also incentivize an agent to *prevent* a terrorist attack, if she is predicting that no such attack will take place. Similarly, the software company’s in-house prediction market may incentivize an employee to work extra hard to finish her component of the product in time, if she is predicting that the product will be released on time.

The question that we study in this paper is the following: is it possible to design prediction mechanisms that do not incentivize undesirable actions? Here, an action is *undesirable* if it reduces the expected utility of the principal (center, organizer) of the prediction mechanism (*e.g.*, the Department of Defense or the software company in the above examples). We call such mechanisms *principal-aligned* since they, in some sense, align the agents’ incentives with those of the principal<sup>4</sup>.

The rest of this paper is organized as follows. In Section 2, we study proper scoring rules, which incentivize a single agent to truthfully report her subjective probabilities about an event. After reviewing proper scoring rules and a known characterization theorem, we give a complete characterization of principal-aligned proper scoring rules. In Section 3, we consider settings with  $n$  agents. We show a negative “overpayment” result that indicates (roughly stated) that a principal-aligned prediction mechanism will, in the worst case, require the principal to pay  $n$  times as much as a mechanism that is not principal-aligned.

---

<sup>3</sup> In July 2003, reports surfaced about a DARPA project to use prediction markets to guide policy decisions, and a possible topic was terrorist attacks. This ignited a political uproar and the proposal was quickly dropped. The arguments against it include creating incentives for a person to perform a violent act, as well as the distasteful thought of any person benefiting from such an attack. A survey of the proposal, the debate, and the aftermath can be found in [9].

<sup>4</sup> We note that it *is* possible for the prediction mechanisms in this paper to incentivize *desirable* actions. However, in this paper we will not have an explicit model for the agent’s costs for taking desirable actions (such as putting in extra effort so that a product is released on time), and as a result we will not be able to solve for the agent’s optimal action. All we can say is that the agent’s optimal action will be no less desirable than it would have been without the existence of the prediction mechanism. Moreover, whatever action the agent decides to take, she will be incentivized to report the true distribution conditional on that action.

We also assume that the prediction itself does not have secondary effects on the agent’s utility (an example of this would be the case where the agent’s manager observes the prediction of an early release date, and as a result will punish the agent if the product is not released early).

In Section 4, we extend our model to allow uncertainties about the principal’s utility, as well as restrictions on how agents’ actions can affect the underlying probabilities. We only want to disincentivize actions that are plausible under these restrictions and are definitely undesirable. We show how this provides richer structure to the class of principal-aligned proper scoring rules, and show that given sufficient uncertainty, any proper scoring rule can be transformed into a principal-aligned proper scoring rule by adding constant bonuses. We also show a similar overpayment result under the extended model. All omitted proofs can be found in the full version of this paper.

## 2 Principal-Aligned Proper Scoring Rules

### 2.1 Review of Proper Scoring Rules

Let  $\Omega = \{1, 2, \dots, m\}$  be the outcome space, with  $m$  possible outcomes. Let  $\mathbb{P} = \{\mathbf{p} \in \mathbb{R}^m : 0 < p_i < 1, \sum_{i=1}^m p_i = 1\}$  be the set of probability distributions over the outcomes. We define the standard basis  $\mathbf{O}_i$  ( $i = 1, 2, \dots, m$ ) as follows:  $\mathbf{O}_i$  is the vector in which the  $i$ -th element equals 1 and all other elements equal 0. Note that while  $\mathbf{O}_i$  is not in  $\mathbb{P}$ , all vectors in  $\mathbb{P}$  are in the span of the  $\mathbf{O}_i$ ’s.

**Definition 1.** A scoring rule is a function  $S : \mathbb{P} \times \Omega \rightarrow \mathbb{R}$ . For each report  $\mathbf{r} \in \mathbb{P}$  (on the underlying distribution) and each outcome  $i \in \Omega$ , it specifies a payment  $S(\mathbf{r}, i)$ . The expected payment  $\tilde{S}$  under the scoring rule  $S$  depends on both the report  $\mathbf{r}$  and the true probability distribution  $\mathbf{p}$  over the outcomes.  $\tilde{S}$  can be written as  $\tilde{S}(\mathbf{r}, \mathbf{p}) = \sum_{i=1}^m S(\mathbf{r}, i)p_i$ .

**Definition 2.** A scoring rule  $S : \mathbb{P} \times \Omega \rightarrow \mathbb{R}$  is (weakly) proper if  $\forall \mathbf{p}, \mathbf{r} \in \mathbb{P}$ ,  $\tilde{S}(\mathbf{p}, \mathbf{p}) \geq \tilde{S}(\mathbf{r}, \mathbf{p})$ . It is strictly proper if equality occurs if and only if  $\mathbf{r} = \mathbf{p}$ .

**Definition 3.** Given convex function  $G : \mathbb{P} \rightarrow \mathbb{R}$ , a subgradient is a vector function  $G^* : \mathbb{P} \rightarrow \mathbb{R}^m$  such that  $\forall \mathbf{x}, \mathbf{y} \in \mathbb{P}, G(\mathbf{y}) \geq G(\mathbf{x}) + G^*(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x})$ . Given convex function  $G : \mathbb{P} \rightarrow \mathbb{R}$ , a subtangent hyperplane at  $\mathbf{x} \in \mathbb{P}$  is a linear function  $H_{\mathbf{x}} : \mathbb{P} \rightarrow \mathbb{R}$  such that  $H_{\mathbf{x}}(\mathbf{x}) = G(\mathbf{x})$ , and  $\forall \mathbf{y} \in \mathbb{P}, H_{\mathbf{x}}(\mathbf{y}) \leq G(\mathbf{y})$ .

The following characterization of proper scoring rules was discovered first by Savage [13], but the version shown here is due to Gneiting and Raftery [5]. The intuition behind the characterization is illustrated in Figure 1.

**Theorem 1 (Gneiting & Raftery).** Given a convex function  $G : \mathbb{P} \rightarrow \mathbb{R}$  and a subgradient  $G^*$ , setting  $H_{\mathbf{r}}(\mathbf{p}) = G(\mathbf{r}) + G^*(\mathbf{r}) \cdot (\mathbf{p} - \mathbf{r})$  defines a family of

<sup>5</sup> This is available at

<http://www.cs.duke.edu/~pengshi/papers/2009-07-society-aligned.pdf>

<sup>6</sup> We assume that the probability of any outcome is positive. This assumption helps us handle peculiar cases with discontinuities at the boundary, and makes the ensuing math more elegant. We can handle the edge cases by taking limits.

<sup>7</sup>  $G^*$  always exists. If  $G$  is differentiable at  $\mathbf{x} \in \mathbb{P}$ , then the subgradient at  $\mathbf{x}$  is the gradient:  $G^*(\mathbf{x}) = (\nabla G)(\mathbf{x})$ . Otherwise, there may be many choices of  $G^*(\mathbf{x})$ .


subtangent hyperplanes such that  $H_r$  is subtangent at  $\mathbf{r}$ . Setting  $S(\mathbf{r}, i) = H_r(\mathbf{O}_i)$  defines a proper scoring rule. i.e.,

$$S(\mathbf{r}, i) = G(\mathbf{r}) - G^*(\mathbf{r}) \cdot \mathbf{r} + G^*(\mathbf{r}) \cdot \mathbf{O}_i = G(\mathbf{r}) - G^*(\mathbf{r}) \cdot \mathbf{r} + G_i^*(\mathbf{r})$$


Conversely, any proper scoring rule can be written in terms of a subgradient of some convex function  $G$  in the above fashion. We call  $G$  the cost function for the rule.

Any proper scoring rule  $S$  corresponds to a unique convex cost function  $G$  where  $G(\mathbf{p}) = \tilde{S}(\mathbf{p}, \mathbf{p})$  is the maximum expected payment the agent can obtain if the true probability is  $\mathbf{p}$ . Conversely, any convex function corresponds to some proper scoring rule.<sup>8</sup>

Currently, the cost function  $G$  is defined only in the open set  $\mathbb{P}$ ; we define  $G$  on the boundary of  $\mathbb{P}$  by taking limits. Since  $G$  is already continuous in  $\mathbb{P}$ ,<sup>9</sup> this makes  $G$  continuous everywhere.



**Fig. 1.** Geometric Intuition for Theorem 1. Every proper scoring rule corresponds to a convex cost function  $G$ . Suppose an agent reports  $\mathbf{r}$ ; then, the rule’s payment for each outcome is the corresponding intercept of the subtangent hyperplane at  $\mathbf{r}$ . Hence, for a given report, her possible expected payments are represented by the corresponding subtangent hyperplane, which is always below the upper envelope  $G$ . When the agent reports the true distribution  $\mathbf{p}$ , she attains the upper envelope  $G$ .



**Fig. 2.** Intuition behind Lemma 1. Consider the hyperplane normal to  $\mathbf{u}$  through  $\mathbf{x}$ . On one side of this hyperplane is  $\mathbb{P}_x$ , and for any  $\mathbf{y}$  in the region  $\mathbb{P}_x$ ,  $G(\mathbf{x}) \geq G(\mathbf{y})$ . By continuity of  $G$ ,  $G(\mathbf{x})$  must be constant along the hyperplane (the boundary of  $\mathbb{P}_x$ ).

## 2.2 Principal-Aligned Proper Scoring Rules

We now develop the concept of “principal-aligned” proper scoring rules. Suppose that agents can take actions in the real world to change the underlying distribution of the actual outcome. A principal-aligned rule disincentivizes any action that harms the principal in expectation.

<sup>8</sup> If  $G$  is not differentiable, then many families of subtangent hyperplanes can be specified, each of which corresponds to a proper scoring rule.

<sup>9</sup> A well-known fact about convex functions is that they are everywhere continuous in the interior of their domain.

Formally, let  $\mathbf{u} \in \mathbb{R}^m$  be a vector whose  $i$ th component is the principal's utility for outcome  $i$ . Note that given the true distribution  $\mathbf{p} \in \mathbb{P}$ , the principal's expected utility is  $\mathbf{p} \cdot \mathbf{u}$ .

**Definition 4.** A proper scoring rule  $S$  is aligned with vector  $\mathbf{u}$  if the cost function  $G$  satisfies:  $\forall \mathbf{p}_1, \mathbf{p}_2 \in \mathbb{P}$ , if  $(\mathbf{p}_2 - \mathbf{p}_1) \cdot \mathbf{u} > 0$ , then  $G(\mathbf{p}_2) \geq G(\mathbf{p}_1)$ . We call  $S$  strictly aligned if the inequality is always strict.

Since we have  $\tilde{S}(\mathbf{p}_2, \mathbf{p}_2) = G(\mathbf{p}_2)$  and  $\tilde{S}(\mathbf{r}, \mathbf{p}_1) \leq G(\mathbf{p}_1)$ , this definition says that if the true probability is  $\mathbf{p}_2$ , then the agent prefers reporting  $\mathbf{p}_2$  over changing the true probability to  $\mathbf{p}_1$  and reporting some  $\mathbf{r}$ .

We call  $\mathbf{u}$  uniform if  $\mathbf{u} = \alpha \mathbf{1}$  for some  $\alpha$ . Note that because  $\mathbf{1} \perp (\mathbf{p}_2 - \mathbf{p}_1) \forall \mathbf{p}_1, \mathbf{p}_2 \in \mathbb{P}$ , the above definition does not say anything when  $\mathbf{u}$  is uniform (that is, when the principal is indifferent among all outcomes).

**Definition 5.** A convex function  $G : \mathbb{P} \rightarrow \mathbb{R}$  is non-decreasing with respect to direction  $\mathbf{u} \in \mathbb{R}^m$  if  $\forall \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{P}$ ,  $(\mathbf{x}_2 - \mathbf{x}_1) \cdot \mathbf{u} > 0$  implies  $G(\mathbf{x}_2) \geq G(\mathbf{x}_1)$ . It is strictly increasing if the above inequality is strict.

**Lemma 1.** If convex function  $G : \mathbb{P} \rightarrow \mathbb{R}$  is non-decreasing with respect to non-uniform direction  $\mathbf{u}$ , then  $G(\mathbf{x}) = g(\mathbf{x} \cdot \mathbf{u})$  for some single-variable non-decreasing convex function  $g$ . The statement remains true when non-decreasing is replaced by strictly increasing.

*Proof.* For each  $\mathbf{x} \in \mathbb{P}$ , define the set  $\mathbb{P}_{\mathbf{x}} = \{\mathbf{y} | \mathbf{y} \in \mathbb{P}, (\mathbf{x} - \mathbf{y}) \cdot \mathbf{u} > 0\}$ .  $\mathbb{P}_{\mathbf{x}}$  is non-empty because  $\mathbf{u}$  is not uniform<sup>10</sup>. Moreover,  $\mathbb{P}_{\mathbf{x}}$  is open because it is the intersection of open sets  $\mathbb{P}$  and  $\{\mathbf{y} | (\mathbf{x} - \mathbf{y}) \cdot \mathbf{u} > 0\}$ .

Since  $G$  is non-decreasing with respect to  $\mathbf{u}$ ,  $G(\mathbf{x}) \geq \sup_{\mathbf{y} \in \mathbb{P}_{\mathbf{x}}} G(\mathbf{y})$ . But  $\mathbf{x}$  lies in the closure  $\bar{\mathbb{P}}_{\mathbf{x}}$ , so by continuity of  $G$ ,  $G(\mathbf{x}) \leq \sup_{\mathbf{y} \in \bar{\mathbb{P}}_{\mathbf{x}}} G(\mathbf{y})$ . This means that for all  $\mathbf{x} \in \mathbb{P}$ ,  $G(\mathbf{x}) = \sup_{\mathbf{y} \in \bar{\mathbb{P}}_{\mathbf{x}}} G(\mathbf{y})$ .

Note now that  $\forall \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{P}$ , whenever  $\mathbf{x}_1 \cdot \mathbf{u} = \mathbf{x}_2 \cdot \mathbf{u}$ , we have  $\mathbb{P}_{\mathbf{x}_1} \equiv \mathbb{P}_{\mathbf{x}_2}$ , which implies that  $G(\mathbf{x}_1) = G(\mathbf{x}_2)$ . Hence,  $G(\mathbf{x}) = g(\mathbf{x} \cdot \mathbf{u})$  for some single-variable function  $g$ . Moreover,  $g$  must be convex because  $G$  is convex, and  $g$  must be non-decreasing because  $G$  is non-decreasing w.r.t.  $\mathbf{u}$ . The above proof still holds if non-decreasing is replaced by strictly increasing. Figure 2 illustrates the intuition behind this lemma.

Combining Definition 4 and Lemma 1, we get the following characterization of principal-aligned proper scoring rules.

**Theorem 2.** Given any non-uniform principal utility vector  $\mathbf{u} \in \mathbb{R}^m$ , every proper scoring rule  $S$  aligned with  $\mathbf{u}$  corresponds to a cost function of the form  $G(\mathbf{p}) = g(\mathbf{p} \cdot \mathbf{u})$ , where  $g$  is a single-variable non-decreasing convex function. (By Theorem 1, this implies  $S(\mathbf{r}, i) = G(\mathbf{r}) - G^*(\mathbf{r}) \cdot \mathbf{r} + G_i^*(\mathbf{r})$ ,<sup>11</sup> where  $G^*$  is a subgradient of  $G$ .) Conversely, all such rules are aligned with  $\mathbf{u}$ .

<sup>10</sup> This is because  $\exists \mathbf{y} \in \mathbb{P}$  such that  $(\mathbf{x} - \mathbf{y}) \cdot \mathbf{u} \neq 0$ , in which case either  $\mathbf{y}$  or  $2\mathbf{x} - \mathbf{y}$  is in  $\mathbb{P}_{\mathbf{x}}$ .

<sup>11</sup> When  $g$  is differentiable,  $S(\mathbf{r}, i) = g(\mathbf{r} \cdot \mathbf{u}) + (u_i - \mathbf{r} \cdot \mathbf{u})g'(\mathbf{r} \cdot \mathbf{u})$ .

*The above statement remains true when aligned is replaced by strictly aligned, and non-decreasing is replaced by strictly increasing.*

Note that  $\mathbf{p} \cdot \mathbf{u}$  is the principal’s expected utility. The above theorem implies that for a fixed report, the agent’s expected reward is a non-decreasing function of the principal’s expected utility. Hence the interests of the agent and the principal are aligned.

### 3 Principal-Aligned Prediction Mechanisms

Now, suppose that we want to elicit predictions from  $n$  agents in a principal-aligned way. One possible method is to use a principal-aligned proper scoring rule for each agent, and allow agents to see previous agents’ reports. However, to incentivize agents to respond thoughtfully, each proper scoring rule requires some subsidy to implement, and hence this method requires  $\Theta(n)$  subsidy. On the other hand, if we do not require principal-alignment, we can use a *market scoring rule* and implement this with  $\Theta(1)$  subsidy.<sup>12</sup> In fact, we will show that this  $\Theta(n)$  gap always exists between the “cheapest” principal-aligned prediction mechanism and the “cheapest” non-principal-aligned mechanism. We formalize this in the next section.

#### 3.1 One-Round Prediction Mechanisms

We first make a technical note: As shown in [24], in all current implementations of prediction markets, agents may try to deceive others by giving false signals, to their own later profit. This makes analyzing incentives in a prediction market difficult without strong assumptions such as myopic agents.<sup>13</sup> Since our aim is to show a negative result about the minimum subsidy required in principal-aligned mechanisms, rather than to resolve such strategic issues, we focus on a *one-round* model in which agents can participate at most once, hence ruling out such strategic play. A negative result in this restricted model will carry over to general multi-agent prediction mechanisms.<sup>14</sup>

<sup>12</sup> Under a market scoring rule, each agent is paid according to a proper scoring rule, but must pay the proper-scoring rule payment of the previous agent (*i.e.*, the  $i$ th agent’s expected payment is  $\tilde{S}(\mathbf{r}_i, \mathbf{p}) - \tilde{S}(\mathbf{r}_{i-1}, \mathbf{p})$ ). As a result, from the perspective of the principal, almost all the payments cancel out, and the total amount that the principal must pay ( $\sum_i \tilde{S}(\mathbf{r}_i, \mathbf{p}) - \tilde{S}(\mathbf{r}_{i-1}, \mathbf{p}) = \tilde{S}(\mathbf{r}_n, \mathbf{p}) - \tilde{S}(\mathbf{r}_0, \mathbf{p})$ ) depends only on the last prediction and not on the number of agents. See [6] for further explanations of market scoring rules.

<sup>13</sup> [3] studies creating prediction markets that are fully incentive-compatible, without assuming myopic agents. However, the mechanisms proposed are quite different from typical prediction markets and it is not yet clear that they can be made to work in practice.

<sup>14</sup> This is because in any general prediction mechanism the agents may choose to act sequentially in a one-round fashion. Hence, relaxing the one-round requirement can only increase the worst-case subsidy required. The notions of “one-round” and “subsidy required” are made precise later.

**Definition 6.** We define a one-round prediction mechanism as follows. There is an event with a finite set of disjoint outcomes  $\Omega$ , and  $n$  agents are asked to give a probability for each outcome. There is some rule which decides the order in which agents report their predictions. Define  $E_j$  as the set of agents whose reports agent  $j$  cannot influence. When agent  $j$  reports prediction  $\mathbf{r}_j$  and event  $i \in \Omega$  occurs, she receives a payment of  $f_j(\{\mathbf{r}_k | k \in E_j\}, \mathbf{r}_j, i)$ .

**Definition 7.** A one-round prediction mechanism is truthful if regardless of others' reports, an agent's expected utility is maximized when she reports her true subjective probability vector  $\mathbf{p}$ . The mechanism satisfies voluntary participation if reporting the true  $\mathbf{p}$  always yields non-negative expected payment. The mechanism is feasible if it is truthful and satisfies voluntary participation.

Almost all current prediction mechanisms, when we add the restriction that each agent participates at most once, fit into this framework.<sup>15</sup>

**Lemma 2.** A one-round prediction mechanism is feasible if and only if for each agent  $j$ , the payment function  $f_j$ , holding fixed the reports  $\mathbf{r}_{-j} = \{\mathbf{r}_k | k \neq j\}$ , is a proper scoring rule  $S$  with a non-negative cost function. We call  $S$  the corresponding proper scoring rule in this situation.

**Definition 8.** A one-round prediction mechanism is aligned with principal utility vector  $\mathbf{u}$  if for every agent  $j$  and every combination of other agents' reports ( $\mathbf{r}_{-j} = \{\mathbf{r}_k | k \neq j\}$ ), the corresponding proper scoring rule is aligned with  $\mathbf{u}$ .

### 3.2 Minimum Subsidy for Principal-Aligned Prediction Mechanisms

We now formalize the notion that any principal-aligned prediction mechanism requires  $\Theta(n)$  subsidy to implement. Here, the subsidy is the minimum amount the principal must have to be solvent in expectation, no matter what the true probability and agents' reports are.

**Definition 9.** A one-round prediction mechanism requires subsidy  $M$  if  $\forall \epsilon > 0$ , for some true probability  $\mathbf{p}$  and some reports  $\mathbf{r} = \{\mathbf{r}_j\}$ , the total expected payment

$$\sum_{i,j} p_i f_j(\{\mathbf{r}_k | k \in E_j\}, \mathbf{r}_j, i) \geq M - \epsilon$$

Before deriving our minimum subsidy result, we introduce a notion of minimum incentives. In order to elicit useful predictions, we cannot simply offer the trivial scoring rule  $S \equiv 0$ . We assume that each agent will give a thoughtful report only if she can gain  $c > 0$  by reporting accurately. More precisely, to be meaningful, a prediction mechanism cannot allow agents to always obtain within  $c$  of the optimal expected payment by giving a constant report.

<sup>15</sup> Models that fit this framework include market scoring rules, Hanson's Market Maker [746], Pennock's Dynamic Parimutuel Market [1110], and the weighted-score mechanism in [8].



**Definition 10.** We say a proper scoring rule provides incentive  $c$  if the difference between the greatest attainable expected payment and the greatest expected payment the agent can guarantee with some constant report  $r$  is at least  $c$ .<sup>16</sup> A one-round prediction mechanism guarantees incentive  $c$  if for each agent  $j$  and each combination of others' reports  $\mathbf{r}_{-j}$ , the corresponding proper scoring rule provides incentive  $c$ .

**Lemma 3.** Given a non-uniform principal utility vector  $\mathbf{u}$ , let  $\mathbf{O}_{max}$  correspond to an optimal outcome for the principal (i.e., it maximizes  $O_i \cdot \mathbf{u}$  among all  $i$ ). If a proper scoring rule aligned with  $\mathbf{u}$  provides incentive  $c$ , and its cost function  $G(\mathbf{p})$  is non-negative, then  $G(\mathbf{O}_{max}) \geq c$ .

*Proof.* Suppose on the contrary that  $G(\mathbf{O}_{max}) < c$ . Let  $\mathbf{O}_{min}$  correspond to the principal's worst outcome (i.e., it minimizes  $O_i \cdot \mathbf{u}$  among all  $i$ ), and let  $d = \|\mathbf{O}_{max} - \mathbf{O}_{min}\|$ . Since  $\mathbf{u}$  is not uniform,  $d > 0$ . For  $0 < \epsilon < d$ , define  $r = \frac{\epsilon}{d}\mathbf{O}_{max} + \frac{d-\epsilon}{d}\mathbf{O}_{min}$ . Let  $G^*$  be the subgradient of  $G$  corresponding to  $S$  (in the sense of Theorem 1).

First note that  $\tilde{S}(r, \mathbf{O}_{max}) \geq 0$ . This is because by the definition of subgradients and by  $\mathbf{u} \cdot (\mathbf{O}_{min} - r) < 0$ , we have  $G^*(r) \cdot (\mathbf{O}_{min} - r) \leq G(\mathbf{O}_{min}) - G(r) \leq 0$ . Using Theorem 1 and the fact that  $(\mathbf{O}_{min} - r)$  is collinear with  $(r - \mathbf{O}_{max})$ , we have  $\tilde{S}(r, r) - \tilde{S}(r, \mathbf{O}_{max}) = G^*(r) \cdot (r - \mathbf{O}_{max}) \leq 0$ , which implies that  $\tilde{S}(r, \mathbf{O}_{max}) \geq G(r) \geq 0$ .

Because the scoring rule provides incentive  $c$  and because of convexity, either  $G(\mathbf{O}_{max}) - \tilde{S}(r, \mathbf{O}_{max}) \geq c$  or  $G(\mathbf{O}_{min}) - \tilde{S}(r, \mathbf{O}_{min}) \geq c$ .

However, the first inequality cannot hold, because  $G(\mathbf{O}_{max}) < c$  by assumption and  $\tilde{S}(r, \mathbf{O}_{max}) \geq 0$ .

Moreover, the second inequality cannot hold for sufficiently small  $\epsilon$ . This is because  $c - 0 > G(\mathbf{O}_{max}) - G(r) \geq G^*(r) \cdot (\mathbf{O}_{max} - r) = G^*(r) \cdot (r - \mathbf{O}_{min}) \left(\frac{d-\epsilon}{\epsilon}\right)$ . So as  $\epsilon \rightarrow 0$ , we need  $G^*(r) \cdot (r - \mathbf{O}_{min}) \rightarrow 0$ . This along with the identity  $\tilde{S}(r, \mathbf{O}_{min}) - \tilde{S}(r, r) = G^*(r) \cdot (\mathbf{O}_{min} - r)$  implies that as  $\epsilon \rightarrow 0$ ,  $\tilde{S}(r, \mathbf{O}_{min}) \rightarrow G(r)$ . Because  $G$  is continuous and bounded below by 0, we have that as  $\epsilon \rightarrow 0$ ,  $G(r) \rightarrow G(\mathbf{O}_{min}) \geq 0$ , so  $\tilde{S}(r, \mathbf{O}_{min}) \rightarrow G(\mathbf{O}_{min})$ . Hence, for sufficiently small  $\epsilon$ , the second inequality also fails. Contradiction.

Therefore,  $G(\mathbf{O}_{max}) \geq c$ .

**Theorem 3.** Let  $\mathbf{u}$  be a non-uniform principal utility vector and let  $n$  be the number of agents. Any feasible one-round prediction mechanism that guarantees incentive  $c$  and is aligned with  $\mathbf{u}$  requires subsidy  $cn$ .

*Proof.* For each agent  $j$  and each combination of others' reports, the corresponding proper scoring rule must have a non-negative cost function, must be aligned with  $\mathbf{u}$ , and must provide incentive  $c$ . By Lemma 3,  $G(\mathbf{O}_{max}) \geq c$ . Hence, if all agents report some  $\mathbf{r}$  that is arbitrarily close to  $\mathbf{O}_{max}$ , we get, by the continuity of convex function  $G$ , that the total expected payment can be arbitrarily close to  $nG(\mathbf{O}_{max}) \geq cn$ .

<sup>16</sup> In mathematical language, this states that  $\sup_{\mathbf{q}} \tilde{S}(\mathbf{q}, \mathbf{q}) - \sup_{\mathbf{r}} \left( \inf_{\mathbf{p}} \tilde{S}(\mathbf{r}, \mathbf{p}) \right) \geq c$ .



*Remark 1.* Suppose we sacrifice principal-alignment; then, we can implement the market scoring rule, based on, say, the quadratic scoring rule with cost function  $G(\mathbf{p}) = c \frac{m}{m-1} (\sum_i p_i^2 - \frac{1}{m})$ . This is a one-round prediction mechanism guaranteeing incentive  $c$  and requiring only subsidy  $c$ , which is the minimum possible.<sup>17</sup> This yields the following implication of Theorem 3: *Suppose we want a prediction mechanism that guarantees incentive  $c$ ; the cheapest principal-aligned mechanism requires  $n$  times as much subsidy as the cheapest non-principal-aligned mechanism in the worst case.*

## 4 Extensions: Uncertain Utilities and Restricted Actions

In practice, there may be uncertainty about the principal's utility vector  $\mathbf{u}$ , and there may be restrictions on the change in underlying probabilities that the agents' actions can bring about. In this section, we show that adding these features to our model provides a richer characterization. Moreover, if the agents can perform actions that are certainly undesirable, then principal-aligned prediction mechanisms still require  $\Theta(n)$  subsidy. We formalize these concepts via the following definitions.

**Definition 11.** *An action model is a function  $A : \mathbb{P} \rightarrow 2^{\mathbb{P}}$  such that  $\forall \mathbf{p} \in \mathbb{P}$ ,  $A(\mathbf{p})$  is a convex set satisfying  $\mathbf{p} \in A(\mathbf{p})$ .<sup>18</sup> Intuitively, if  $\mathbf{p}$  is the initial underlying probability vector, then for any  $\mathbf{p}' \in A(\mathbf{p})$ , there is some action that the agent can perform to change the probability vector to  $\mathbf{p}'$  (and the agent is not able to change it to any probability vector outside of  $A(\mathbf{p})$ ).*

**Definition 12.** *Let  $T$  be a set of possible utility vectors for the principal and let  $A$  be an action model. A proper scoring rule is aligned with  $T$  under  $A$  if  $\forall \mathbf{p} \in \mathbb{P}$  and  $\forall \mathbf{p}' \in A(\mathbf{p})$  such that  $\mathbf{u} \cdot \mathbf{p}' < \mathbf{u} \cdot \mathbf{p} \forall \mathbf{u} \in T$ ,<sup>19</sup> the cost function  $G(\mathbf{p}) \geq G(\mathbf{p}')$ . A one-round prediction mechanism is aligned with  $T$  under  $A$  if the corresponding proper scoring rule is always aligned with  $T$ . In all following references to  $T$ , we assume that  $T$  is not always uniform:  $T \setminus \{\alpha \mathbf{1}\} \neq \emptyset$ . (The principal is not definitely indifferent among all outcomes.)*

**Definition 13.** *Given a set  $T$  of possible utilities for the principal and an action model  $A$  for the agents, the corresponding bad direction function is  $B : \mathbb{P} \rightarrow 2^{\mathbb{R}^m}$ ,*

<sup>17</sup> By definition, guaranteeing incentive  $c$  requires subsidy  $c$ .

<sup>18</sup> We require  $\mathbf{p} \in A(\mathbf{p})$  because the agent can always choose to do nothing. Moreover,  $A(\mathbf{p})$  is convex because if by committing action  $a$ , the agent can change the underlying probability to  $\mathbf{q} \in A(\mathbf{p})$ , and by committing  $a'$ , the agent can change it to  $\mathbf{q}'$ , then by committing  $a$  with probability  $\lambda$  and  $a'$  with probability  $(1 - \lambda)$ , the agent can change the underlying probability to  $\lambda \mathbf{q} + (1 - \lambda) \mathbf{q}'$ .

<sup>19</sup> Here, we only guard against actions that harm the principal's expected utility for all possible vectors  $\mathbf{u} \in T$ . Alternatively we might have guarded against actions that are bad for any setting of  $\mathbf{u}$ , but this case is not interesting because the proof of Lemma 11 implies that the only proper scoring rule aligned with two linearly independent utility vectors is the trivial rule with  $G \equiv 1$ .

$$B(\mathbf{p}) = \{\mathbf{p}' - \mathbf{p} \mid \mathbf{p}' \in A(\mathbf{p})\} \cap \left( \bigcap_{\mathbf{u} \in T} \{\mathbf{x} \mid \mathbf{x} \cdot \mathbf{u} < 0\} \right)$$

We say that  $B$  is contained in a strict cone if  $\exists \epsilon > 0$  and  $\exists \mathbf{y}$  such that  $\forall \mathbf{p} \in \mathbb{P}$  and  $\forall \mathbf{x} \in B(\mathbf{p}), \mathbf{y} \cdot \mathbf{x} \geq \epsilon \|\mathbf{x}\| \|\mathbf{y}\|$ .

Define the magnitude of the worst action as  $d = \sup\{\|\mathbf{x}\| \mid \mathbf{x} \in B(\mathbf{p}), \mathbf{p} \in \mathbb{P}\}$ . Note that  $d = 0$  implies that agents cannot perform any action that is certainly bad.

### 4.1 Principal-Aligned Proper Scoring Rules in the Extended Model

The following characterization follows immediately from Definition 13.

**Theorem 4.** *Given a set  $T$  of possible utility vectors and an action model  $A$ , let the corresponding bad direction function be  $B$ . A proper scoring rule is aligned with  $T$  under  $A$  if and only if its cost function  $G$  satisfies  $\forall \mathbf{p} \in \mathbb{P}$ , and  $\forall \mathbf{p}'$  with  $\mathbf{p}' - \mathbf{p} \in B(\mathbf{p}), G(\mathbf{p}') \leq G(\mathbf{p})$ .*

**Definition 14.** *A proper scoring rule is bounded if  $S(\mathbf{r}, i)$  is bounded, that is,  $\exists C$  s.t.  $|S(\mathbf{r}, i)| \leq C \forall \mathbf{r} \in \mathbb{P}, i \in \Omega$ . Equivalently (by Theorem 7), a proper scoring rule is bounded if and only if both the cost function  $G(\mathbf{p})$  and the corresponding subgradient  $G^*(\mathbf{p})$  are bounded.*

The following theorem shows that under sufficient uncertainty, any proper scoring rule can be modified to be principal-aligned by adding constant bonuses.

**Theorem 5.** *Suppose that the bad direction function  $B$  is contained in a strict cone; then, any bounded<sup>20</sup> proper scoring rule  $S$  can be modified to be principal-aligned by adding constants  $\{k_i\}$  so that  $S'(\mathbf{r}, i) = S(\mathbf{r}, i) + k_i$ . Moreover,  $\|\mathbf{k}\| = \frac{M}{\epsilon}$ , where  $M$  is an upper bound<sup>21</sup> on the norm of the subgradient  $\|G^*(\mathbf{p})\|$ , and  $\epsilon$  is as in Definition 13.*

### 4.2 Principal-Aligned Prediction Mechanisms in the Extended Model


We now generalize the minimum subsidy result in Section 3.2 to the context of uncertain utilities and restricted actions. First, we note that the overpayment result does not hold if we allow the proper scoring rule to not be strict: as in the example shown in Figure 3, sometimes we can treat all reports for which we have relative preferences<sup>22</sup> to be the same, and implement a market scoring rule on the collapsed classes. In this case, the agents can obtain optimal payments while ignoring the directions for which the principal has preferences.

In practice it seems desirable to always strictly incentivize the agents to report as close to the true probability as possible. One way to formalize this is the notion of *uniform incentivization*.

<sup>20</sup> In practice, we can make any unbounded proper scoring rule bounded by restricting reports away from the boundary.

<sup>21</sup>  $M$  exists by the assumption that the scoring rule  $S$  is bounded. See Definition 14.

<sup>22</sup> By “relative preference” between reports  $\mathbf{r}'$  and  $\mathbf{r}$ , we mean that either  $\mathbf{r}' - \mathbf{r} \in B(\mathbf{r})$  or  $\mathbf{r} - \mathbf{r}' \in B(\mathbf{r}')$ .



**Fig. 3.** Suppose that there are 3 outcomes and no restrictions on actions.  $T$  is such that the only bad direction is  $b$ . Consider a market scoring rule on the relative probability of outcomes 1 and 2, guaranteeing incentive  $c$ . For every probability distribution  $r \in \mathbb{P}$  we follow the contour lines in the diagram and treat this as some  $r'$  on the segment connecting  $O_1$  and  $O_2$ . This makes the cost function  $G$  (for the corresponding proper scoring rule) in every situation constant along the contour lines. The resultant mechanism is principal-aligned, guarantees incentive  $c$ , but requires only  $O(c)$  subsidy.

**Definition 15.** Suppose that  $f : [0, \infty) \rightarrow \mathbb{R}$  is convex<sup>23</sup> with  $f(0) = 0$ ,  $f(x) > 0 \forall x > 0$ . A proper scoring rule provides uniform incentives according to  $f$  if  $\forall r, p \in \mathbb{P}$ ,  $\tilde{S}(p, p) - \tilde{S}(r, p) \geq f(\|p - r\|)$ . A one-round prediction mechanism guarantees uniform incentives according to  $f$  if for every agent  $j$  and every combination of others' reports, the corresponding proper scoring rule provides uniform incentives according to  $f$ .

Intuitively,  $f$  can be thought of as a measure of incentivization locally at  $p \in \mathbb{P}$ , and providing uniform incentives according to  $f$  guarantees a certain level of incentivization at all  $p \in \mathbb{P}$ . For a one-round prediction mechanism, guaranteeing uniform incentives according to some  $f$  corresponds to maintaining a minimal standard of incentivization for all agents in all situations. Conventional scoring rules such as the quadratic, the logarithmic, and the spherical scoring rules all guarantee uniform incentives according to some  $f$ .

**Theorem 6.** Suppose that a feasible one-round prediction mechanism with  $n$  agents is principal-aligned with set  $T$  of possible utility vectors, under action model  $A$  (assume that the magnitude of the worst action  $d > 0$ <sup>24</sup>); and guarantees uniform incentives according to some  $f$ . Then, the mechanism requires subsidy  $n \cdot f(\frac{d}{2})$  (which is  $\Theta(n)$ ).

In other words, even when the principal's utilities are uncertain and when agents' actions might be limited, suppose that some surely undesirable action exists. Then, under the requirements of feasibility and uniform incentivization, the cheapest principal-aligned mechanism requires  $\Theta(n)$  times as much subsidy as the cheapest non-principal-aligned mechanism.

One practical implication of our  $\Theta(n)$  subsidy results is that to run a useful prediction market without incentivizing undesirable actions, it may be impractical to let agents join freely. This is because agents may join just to get the

<sup>23</sup> The requirement that  $f$  is convex is natural since cost function  $G$  is convex.

<sup>24</sup> That is, it is possible for agents to perform a certainly undesirable action.

subsidy, without providing any useful additional information. However, it may be practical for an organization to run a principal-aligned in-house prediction market, because in this context the number of agents is naturally limited.

## References

1. Berg, J., Forsythe, R., Nelson, F., Rietz, T.: Results from a Dozen Years of Election Futures Markets Research. In: *Handbook of Experimental Economics Results*, ch. 80, vol. 1, pp. 742–751. Elsevier, Amsterdam (2008)
2. 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)
3. Conitzer, V.: Prediction markets, mechanism design, and cooperative game theory. In: *UAI 2009: Conference on Uncertainty in Artificial Intelligence* (2009)
4. Dimitrov, S., Sami, R.: Non-myopic strategies in prediction markets. In: *EC 2008: Proceedings of the 9th ACM conference on Electronic Commerce*, pp. 200–209. ACM, New York (2008)
5. Gneiting, T., Raftery, A.E.: Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477), 359–378 (2007)
6. Hanson, R.: Combinatorial information market design. *Information Systems Frontiers* 5(1), 107–119 (2003)
7. Hanson, R.: Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* 1(1), 1–15 (2007)
8. Lambert, N.S., Langford, J., Wortman, J., Chen, Y., Reeves, D., Shoham, Y., Pennock, D.M.: Self-financed wagering mechanisms for forecasting. In: *EC 2008: Proceedings of the 9th ACM conference on Electronic Commerce*, pp. 170–179. ACM, New York (2008)
9. Looney, R.: DARPA’s policy analysis market for intelligence: Outside the box or off the wall. *Strategic Insights* 2(9) (September 2003)
10. Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: *EC 2004: Proceedings of the 5th ACM conference on Electronic Commerce*, pp. 170–179. ACM, New York (2004)
11. Pennock, D.M., Sami, R.: Computational aspects of prediction markets. In: *Algorithmic Game Theory*, ch. 26, pp. 651–674. Cambridge University Press, Cambridge (2007)
12. Plott, C.R., Chen, K.-Y.: Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working Paper 1131, California Institute of Technology, Division of the Humanities and Social Sciences (March 2002)
13. Savage, L.J.: Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66(336), 783–801 (1971)
14. Wolfers, J., Zitzewitz, E.: Prediction markets. *The Journal of Economic Perspectives* 18(2), 107–126 (2004)
15. Wolfers, J., Zitzewitz, E.: Prediction markets in theory and practice. Research Papers 1927, Stanford University, Graduate School of Business (February 2006)

# Pricing Strategies for Viral Marketing on Social Networks

David Arthur<sup>1</sup>, Rajeev Motwani<sup>1</sup>, Aneesh Sharma<sup>2</sup>, and Ying Xu<sup>1</sup>

<sup>1</sup> Department of Computer Science, Stanford University

<sup>2</sup> Institute for Computational and Mathematical Engineering, Stanford University  
{darthur,rajeev,aneeshs,xuying}@cs.stanford.edu

**Abstract.** We study the use of viral marketing strategies on social networks that seek to maximize revenue from the sale of a single product. We propose a model in which the decision of a buyer to buy the product is influenced by friends that own the product and the price at which the product is offered. The influence model we analyze is quite general, naturally extending both the Linear Threshold model and the Independent Cascade model, while also incorporating price information. We consider sales proceeding in a cascading manner through the network, i.e. a buyer is offered the product via recommendations from its neighbors who own the product. In this setting, the seller influences events by offering a cash-back to recommenders and by setting prices (via coupons or discounts) for each buyer in the social network. This choice of prices for the buyers is termed as the seller's strategy.

Finding a seller strategy which maximizes the expected revenue in this setting turns out to be NP-hard. However, we propose a seller strategy that generates revenue guaranteed to be within a constant factor of the optimal strategy in a wide variety of models. The strategy is based on an *influence-and-exploit* idea, and it consists of finding the right trade-off at each time step between: generating revenue from the current user versus offering the product for free and using the influence generated from this sale later in the process.

## 1 Introduction

Social networks such as Facebook, Orkut and MySpace are free to join, and they attract tens of millions of users. Maintaining these websites for such a large group of users requires substantial investment from the host companies. To help recoup these investments, these companies often turn to monetizing the information that their users provide for free on these websites. This information includes both detailed profiles of users and also the network of social connections between the users. Not surprisingly, there is a widespread belief that this information could be a gold mine for targeted advertising and other online businesses. Nonetheless, much of this potential still remains untapped today. Facebook, for example, was valued at \$15 billion by Microsoft in 2007 [12], but its estimated revenue in 2008 was only \$300 million [15]. Thus, any monetization technology that can help bridge this gap is of paramount interest. Of particular interest are large-scale

monetization technologies that can effectively leverage the networked nature of the online social networks, and move away from the currently used paradigm of contextual advertising borrowed from sponsored search.

Recently, people have begun to consider such a monetization approach that is based on selling products through the spread of influence. Often, users can be convinced to purchase a product if many of their friends are already using it, even if these same users would be hard to convince through direct advertising. This is often a result of personal recommendations – a friend’s opinion can carry far more weight than an impersonal advertisement. In some cases, however, adoption among friends is important for even more practical reasons. For example, instant messenger users and cell phone users will want a product that allows them to talk easily and cheaply with their friends. Usually, this encourages them to adopt the same instant messenger program and the same cell phone carrier that their friends have. We refer the reader to previous work and the references therein for further explanations behind the motivation of the influence model [5,7].

In fact, many sellers already do try to utilize influence-and-exploit strategies that are based on these tendencies. In the advertising world, this has recently led to the adoption of *viral marketing*, where a seller attempts to artificially create word-of-mouth advertising among potential customers [9,10,13]. A more powerful but riskier technique has been in use much longer: the seller gives out free samples or coupons to a limited set of people, hoping to convince these people to try out the product and then recommend it to their friends. Without any extra data, however, this forces sellers to make some very difficult decisions. Who do they give the free samples to? How many free samples do they need to give out? What incentives can they afford to give to recommenders without jeopardizing the overall profit too much?

In this paper, we are interested in finding systematic answers to these questions. In general terms, we can model the spread of a product as a process on a social network. Each node represents a single person, and each edge represents a friendship. Initially, one or more nodes is “active”, meaning that person already has the product. This could either be a large set of nodes representing an established customer base, or it could be just one node – the seller – whose neighbors consist of people who independently trust the seller, or who are otherwise likely to be interested in early adoption.

At this point, the seller can encourage the spread of influences in two ways. First of all, it can offer cashback rewards to individuals who recommend the product to their friends. This is often seen in practice with “referral bonuses” – each buyer can optionally name the person who referred them, and this person then receives a cash reward. This gives existing buyers an incentive to recommend the product to their friends. Secondly, a seller can offer discounts to specific people in order to encourage them to buy the product, above and beyond any recommendations they receive. It is important to choose a good discount from the beginning here. If the price is not acceptable when a prospective buyer first receives recommendations, they might not bother to reconsider even if the price is lowered later.

After receiving discount offers and some set of recommendations, it is up to the prospective buyers to decide whether to actually go through with a purchase. In general, they will do so with some probability that is influenced by the discount and by the set of recommendations they have received. The form of this probability is a parameter of the model and it is determined by external factors, for instance, the quality of the product and various exogenous market conditions. While it is impossible for a seller to calculate the form of these probability exactly, they can estimate it from empirical observations, and use that estimate to inform their policies. One could interpret the probabilities according to a number of different models that have been proposed in the literature (for instance, the Independent Cascade and Linear Threshold models), and hence it is desirable for the seller to be able to come up with a strategy that is applicable to a wide variety of models.

Now let us suppose that a seller has access to data from a social network such as Facebook, Orkut, or MySpace (thus, the seller can estimate the underlying friendship structure). With this information in hand, a seller can model the spread of influence quite accurately, and the formerly inscrutable problems of who to offer discounts to, and at what price, become algorithmic questions that one can legitimately hope to solve. For example, if a seller knows the structure of the network, she can locate individuals that are particularly well connected and do everything possible to ensure they adopt the product and exert their considerable influence.

In this paper, we are interested in the algorithmic side of this question: Given the network structure and a model of the purchase probabilities, how should the seller decide to offer discounts and cashback rewards?

## 1.1 Our Contributions

We investigate seller strategies that address the above questions in the context of expected revenue maximization. We will focus much of our attention on *non-adaptive* strategies for the seller: the seller chooses and commits to a discount coupon and cashback offer for each potential buyer *before* the cascade starts. If a recommendation is given to this node at any time, the price offered will be the one that the seller committed to initially, irrespective of the current state of the cascade.

A wider class of strategies that one could consider are *adaptive strategies*, which do not have this restriction. For example, in an adaptive strategy, the seller could choose to observe the outcome of the (random) cascading process up until the last minute before making very well informed pricing decisions for each node. One might imagine that this additional flexibility could allow for potentially large improvements over non-adaptive strategies. Unfortunately, there is a price to be paid, in that good adaptive strategies are likely to be very complicated, and thus difficult and expensive to implement. The ratio of the revenue generated from the optimal adaptive strategy to the revenue generated from the optimal non-adaptive strategy is termed the “adaptivity gap”.



Our main theoretical contribution is a very efficient non-adaptive strategy whose expected revenue is within a constant factor of the optimal revenue from an *adaptive* strategy. This guarantee holds for a wide variety of probability functions, including natural extensions of both the Linear Threshold and Independent Cascade models<sup>1</sup>. Note that a surprising consequence of this result is that the adaptivity gap is constant, so one can make the case that not much is lost by restricting our attention to non-adaptive policies. We also show that the problem of finding an optimal non-adaptive strategy is NP-hard, which means an efficient approximation algorithm is the best theoretical result that one could hope for.

Intuitively, the seller strategy we propose is based on an *influence-and-exploit* idea, and it consists of categorizing each potential buyer as either an *influencer* or a *revenue source*. The influencers are offered the product for free and the revenue sources are offered the product at a pre-determined price, chosen based on the exact probability model. Briefly, the categorization is done by finding a spanning tree of the social network with as many leaves as possible, and then marking the leaves as revenue sources and the internal nodes as influencers. We can find such a tree in near-linear time [8, 11]. Cashback amounts are chosen to be a fixed fraction of the total revenue expected from this process. The full details are presented in section 3.

In practice, we propose using this approach to find a strategy that has good global properties, and then using local search to improve it further. This kind of combination has been effective in the past, for example on the k-means problem [2]. Indeed, experiments show that combining local search with the above influence-and-exploit strategy is more effective than using either approach on its own. See the full version of the paper for details [1].

## 1.2 Related Work

The problem of *social contagion* or spread of influence was first formulated by the sociological community, and introduced to the computer science community by Domingos and Richardson [3]. An influential paper by Kempe, Kleinberg and Tardos [6] solved the *target set selection* problem posed by [3] and sparked interest in this area from a theoretical perspective (see [7]). This work has mostly been limited to the *influence maximization* paradigm, where influence has been taken to be a proxy for the revenue generated through a sale. Although similar to our work in spirit, there is no notion of price in this model, and therefore, our central problem of setting prices to encourage influence spread requires a more complicated model.

A recent work by Hartline, Mirrokni and Sundararajan [5] is similar in flavor to our work, and also considers extending social contagion ideas with pricing information, but the model they examine differs from our model in a several aspects. The main difference is that they assume that the seller is allowed to

---

<sup>1</sup> More precisely, the strategy achieves a constant-factor approximation for any *fixed* model, independent of the social network. If one changes the model, the approximation factor does vary, as made precise in Section 3.



approach arbitrary nodes in the network at any time and offer their product at a price chosen by the seller, while in our model the cascade of recommendations determines the timing of an offer and this cannot be directly manipulated. In essence, the model proposed in [5] is akin to advertising the product to arbitrary nodes, bypassing the network structure to encourage a desired set of early adopters. Our model restricts such direct advertising as it is likely to be much less effective than a direct recommendation from a friend, especially when the recommender has an incentive to convince the potential buyer to purchase the product (for instance, the recommender might personalize the recommendation, increasing its effectiveness). Despite the different models, the algorithms proposed by us and [5] are similar in spirit and are based on an *influence-and-exploit* strategy.

This work has been inspired by a direction mentioned by Kleinberg [7], and is our interpretation of the informal problem posed there (of moving influence to edges). Finally, we point out that the idea of cashbacks has been implemented in practice, and new retailers are also embracing the idea [9,10,13]. We note that some of the systems being implemented by retailers are quite close to the model that we propose, and hence this problem is relevant in practice.

## 2 The Formal Model

Let us start by formalizing the setting stated above. We represent the social network as an undirected graph  $G(V, E)$ , and denote the initial set of adopters by  $S^0 \subseteq V$ . We also denote the *active* set at time  $t$  by  $S^{t-1}$  (we call a node *active* if it has purchased the product and *inactive* otherwise). Given this setting, the recommendations cascade through the network as follows: at each time step  $t \geq 1$ , the nodes that became active at time  $t - 1$  (i.e.  $S^0$  for  $t = 1$ , and  $u \in S^{t-1} \setminus S^{t-2}$  for  $t \geq 2$ ) send recommendations to their currently inactive friends in the network:  $N^{t-1} = \{v \in V \setminus S^{t-1} \mid (u, v) \in E, u \in S^{t-1} \setminus S^{t-2}\}$ . Each such node  $v \in N^{t-1}$  is also given a price  $c_{v,t} \in \mathbb{R}$  at which it can purchase the product. This price is chosen by the seller to either be full price or some discounted fraction thereof.

The node  $v$  must then decide whether to purchase the product or not (we discuss this aspect in the next section). If  $v$  does accept the offer, a fixed cashback  $r > 0$  is given to a recommender  $u \in S^{t-1}$  (note that we are fixing the cashback to be a positive constant for all the nodes as the nodes are assumed to be non-strategic and any positive cashback provides incentive for them to provide recommendations). If there are multiple recommenders, the buyer must choose only one of them to receive the cashback; this is a system that is quite standard in practice. In this way, offers are made to all nodes  $v \in N^{t-1}$  through the recommendations at time  $t$  and these nodes make a decision at the end of this time period. The set of active nodes is then updated and the same process is repeated until the process quiesces, which it must do in finite time since any step with no purchases ends the process.

In the model described above, the only degree of freedom that the seller has is in choosing the prices and the cashback amounts. It wants to do this in a way

that maximizes its own expected revenue (the expectation is over randomness in the buyer strategies). Since the seller may not have any control over the seed set, we are looking for a strategy that can maximize the expected revenue starting from any seed set on any graph. In most online scenarios, producing extra copies of the product has negligible cost, so maximizing expected revenue will also maximize expected profit.

Now we can formally state the problem of finding a revenue maximizing strategy as follows:

*Problem 1.* Given a connected undirected graph  $G(V, E)$ , a seed set  $S^0$ , a fixed cashback amount  $r$ , and a model  $M$  for determining when nodes will purchase a product, find a strategy that maximizes the expected revenue from the cascading process described above.

In this work, we restrict attention to policies that choose a single price for a node (as against a price per edge) as otherwise this would lead to strategizing by the buyers. We are particularly interested in non-adaptive policies, which correspond to choosing a price for each node in advance, making the price independent of both the time of the recommendation and also the state of the cascade at the time of the offer. Our goal will be threefold: (1) to show that this problem is NP-hard even for simple models  $M$ , (2) to construct a constant-factor approximation algorithm for a wide variety of models, and (3) to show that restricting to non-adaptive policies results in at most a constant factor loss of profit.

To simplify the exposition, we will assume the cashback  $r = 0$  for now. At the end of Section 4, we will show how the results can be generalized to work for positive  $r$ , which should be sufficient incentive for buyers to pass on recommendations.

## 2.1 Buyer Decisions

In this section, we discuss how to model the probability that a node will actually buy the product given a set of recommendations and a price. We use a very general model in this work that naturally extends the most popular traditional models proposed in the influence maximization literature, including both Independent Cascade and Linear Threshold.

Consider an abstract model  $M$  for determining the probability that a node will buy a product given a price and what recommendations it has received. To model the lack of information about individual node preferences, we will assume that each node's value for the product is drawn from a known distribution. Further, the distribution is dependent on the neighbors of the node that already own the product. Thus, we will assume that the value of the product for a node with degree  $d$ ,  $k \geq 1$  of whose neighbors already have the product, is drawn at random from a distribution specified by the cumulative distribution function  $F_{k,d}(\cdot)$ , i.e.  $F_{k,d}(t)$  is the probability that the node's value is less than  $t$ . Note that such a node will buy the product at an offer price  $x$  with probability  $1 - F_{k,d}(x)$ . For technical reasons, it is convenient to work with the inverse of  $1 - F_{k,d}(x)$ , which

we call  $C_{k,d}(x)$ , i.e. if a node with degree  $d$  and  $k$  active neighbors is offered the product at price  $C_{k,d}(x)$ , it buys with probability  $x$ <sup>2</sup>. We allow the model  $M$  to be fairly general, imposing only the following conditions:

1. We assume that the seller has full information about the different functions  $C_{k,d}(\cdot)$  for all  $k \in \{1, 2, \dots, d\}$  and  $d \in \{1, 2, \dots, n\}$ <sup>3</sup>.
2. The functions  $C_{k,d}(\cdot)$  have bounded domain, which we assume to be  $[0, 1]$  for all  $k, d$  w.l.o.g. This only means that a node will never pay more than full price for the product, where the full price is assumed to be 1. Without an assumption like this, the seller could potentially achieve unbounded revenue on a network, which makes the problem degenerate.
3. We also assume that  $C_{k,d}(1) = 0$  and  $C_{k,d}(0) = 1$  for any  $k, d$ . In other words, a node will always accept the product and recommend it to friends if it receives a recommendation with price 0. Since nodes are given positive cash rewards for making recommendations, this condition is true for any rational buyer. Also, we assume that each node has non-trivial interest in the item and hence a node's value will never exceed the full price of the product.
4. All the functions  $C_{k,d}(\cdot)$  are assumed to be differentiable at 0 and 1.

Now, given a model  $M$  that satisfies the stated assumptions, we make the following observations (all proofs have been moved to the full version of the paper [1] due to lack of space):

**Lemma 1.** *For a cascading process proceeding according to a model  $M$  that satisfies the above assumptions, the following properties hold:*

1. *If the social network is a single line graph with  $S^0$  being the two endpoints, the maximum expected revenue for the optimal choice of prices is at most a constant  $L$ .*
2. *There exist constants  $f, c, q$  so that if more than fraction  $f$  of a given node's neighbors recommend the product to the node at cost  $c$ , the node will purchase the product with probability  $q$ .*

Intuitively, the first property states that each prospective buyer on a social network should have some chance of rejecting the product (unless it's given to them for free), and therefore the maximum revenue on a line is bounded by a geometric series, and is therefore constant. The second property rules out extreme inertia, for example the case where no buyer will consider purchasing a product unless almost all of its neighbors have already done so.

Next, we define a parameter that captures how complicated the model  $M$  is, and our final approximation bound will be in terms of this parameter. Since

<sup>2</sup> It is sometimes useful to consider functions  $p(\cdot)$  that are not one-to-one. These functions have no formal inverse, but in this case,  $c$  can still be formally defined as  $C(x) = \max_y |1 - F_{k,d}(y)| \geq x$ .

<sup>3</sup> The different functions can be approximated in practice by running experiments and observing people's behavior. Also, this general model encompasses a wide variety of models, many of which require estimating a much smaller number of parameters.

the assumptions stated above are quite general, model complexity enables us to parametrize the model in terms of the profit/long-term investment trade-off (property 1) and “nice” myopic prices (property 2). Specific values of model complexity for instances of well-known cascade models are also provided below.

**Definition 1.** *The “model complexity” of a model  $M$  is defined to be  $\frac{L}{(1-f)_{cq}}$ .*

The following corollary is immediately obtained from the proof of lemma [□](#)

**Corollary 1.** *A model  $M$  that satisfies all the above stated assumptions has constant model complexity.*

While it may not be obvious that all the above assumptions are met in general, we will show that they are for both the Independent Cascade and Linear Threshold models, and indeed, the arguments there extend naturally to many other cases as well.

In the traditional Independent Cascade model, there is a fixed probability  $p$  that a node will purchase a product each time it is recommended to them. These decisions are made independently for each recommendation, but each node will buy the product at most once. To generalize this to multiple prices, it is natural to make  $p$  a function  $[0, 1] \rightarrow [0, 1]$  where  $p(x)$  represents the probability that a node will buy the product at price  $x$ . We can express this in terms of our cumulative distribution function as:  $F_{k,d}(x) = 1 - (1 - p(x))^k$ . Again, we can work with the inverse of  $F_{k,d}(x)$ . In fact, since  $F_{k,d}(\cdot)$  has a *decoupled* form we find it easier to work with the inverse of  $p$  which we call  $C'$ . Our general conditions on the model reduce to setting  $C'(0) = 1$  and  $C'(1) = 0$  in this case. The last two assumptions can now be stated along with the model as follows:

**Definition 2.** *Fix a cost function  $C' : [0, 1] \rightarrow [0, 1]$  with  $C'(0) = 1, C'(1) = 0$  and with  $C'$  differentiable at 0 and 1. We define the Independent Cascade Model  $ICM_C$  as follows:*

*Every time a node receives a recommendation at price  $C'(x)$ , it buys the product with probability  $x$  and does nothing otherwise. If a node receives multiple recommendations, it performs this check independently for each recommendation but it never purchases the product more than once.*

The following lemma follows immediately from the proof of lemma [□](#)

**Lemma 2.** *Fix a cost function  $C'$ . Then:*

1.  $ICM_C$  has bounded (model) complexity.
2. If  $C'$  has maximum slope  $m$  (i.e.  $|C'(x) - C'(y)| \leq m|x - y|$  for all  $x, y$ ), then  $ICM_C$  has  $O(m^2)$  complexity.
3. If  $C'$  is a step function with  $n$  regularly spaced steps (i.e.  $C'(x) = C'(y)$  if  $\lfloor \frac{x}{n} \rfloor = \lfloor \frac{y}{n} \rfloor$ ), then  $ICM_C$  has  $O(n^2)$  complexity.

In the traditional Linear Threshold model, there are fixed influences  $b_{v,w}$  on each directed edge  $(v, w)$  in the network. Each node independently chooses a threshold  $\theta$  uniformly at random from  $[0, 1]$ , and then purchases the product

if and when the total influence on it from nodes that have recommended the product exceeds  $\theta$ .

To generalize this to multiple prices, it is natural to make  $b_{v,w}$  a function  $[0, 1] \rightarrow [0, 1]$  where  $b_{v,w}(x)$  indicates the influence  $v$  exerts on  $w$  as a result of recommending the product at price  $x$ . To simplify the exposition, we will focus on the case where a node is equally influenced by all its neighbors. (This is not strictly necessary but removing this assumptions requires rephrasing the definition of  $f$  to be a *weighted* fraction of a node's neighbors.) Finally, we assume for all  $v, w$  that  $b_{v,w}(0) = 1$  to satisfy the second general condition for models.

**Definition 3.** Fix a max influence function  $B : (0, 1] \rightarrow [0, 1]$ , not uniformly 0. We define the Linear Threshold Model  $LTM_B$  as follows: Every node independently chose a threshold  $\theta$  uniformly at random from  $[0, 1]$ . A node will buy the product at price  $x > 0$  only if  $B(x) \geq \frac{\theta}{\alpha}$  where  $\alpha$  denotes the fraction of the node's neighbors that have recommended the product. A node will always accept a recommendation if the product is offered for free.

**Lemma 3.** Fix a max influence function  $B$  and let  $K = \max_x x \cdot B(x)$ . Then  $LTM_B$  has complexity  $O(\frac{1}{K})$ .

We omit the proof since it is similar to that of Lemma 1. In fact, it is simpler since, on a line graph, a node either gets the product for free or it has probability at most  $\frac{1}{2}$  of buying the product and passing on a recommendation.

### 3 Approximating the Optimal Revenue

In this section, we present our main theoretical contribution: a non-adaptive seller strategy that achieves expected revenue within a constant factor of the revenue from the optimal *adaptive* strategy. We show the problem of finding the exact optimal strategy is NP-hard (see the full version of the paper [1] for a proof), so this kind of result is the best we can hope for. Note that our approximation guarantee is against the strongest possible optimum, which is perhaps surprising: it is unclear a priori whether such a strategy should even exist.

The strategy we propose is based on computing a *maximum-leaf spanning tree* (MAXLEAF) of the underlying social network graph, i.e., computing a spanning tree of the graph with the maximum number of leaf nodes. The MAXLEAF problem is known to be NP-Hard, and it is in fact also MAX SNP-Complete, but there are several constant-factor approximation algorithms known for the problem [4, 8, 11, 14]. In particular, one of these is nearly linear-time [11], making it practical to apply on large online social network graphs. The seller strategy we attain through this is an *influence-and-exploit* strategy that offers the product to all of the interior nodes of the spanning tree for free, and charges a fixed price from the leaves. Note that this strategy works for all the buyer decision models discussed above, including multi-price generalizations of both Independent Cascade and Linear Threshold.

We consider the setting of Problem [1](#), where we are given an undirected social network graph  $G(V, E)$ , a seed set  $S^0 \subseteq V$  and a buyer decision model  $M$ . Throughout this section, we will let  $L$ ,  $f$ ,  $c$  and  $q$  denote the quantities that parametrize the model complexity, as described in Section [2.1](#). To simplify the exposition, we will assume that the seed set is a singleton node (i.e.,  $|S^0| = 1$ ). If this is not the case, the seed nodes can be merged into a single node, and we can make much the same argument in that case. We will also ignore cashbacks for now, and return to address them at the end of the section.

The exact algorithm we will use is stated below:

- Use the MAXLEAF algorithm [11](#) to compute an approximate max-leaf spanning tree  $T$  for  $G$  that is rooted at  $S_0$ .
- Offer the product to each internal node of  $T$  for free.
- For each leaf of  $T$  (excluding  $S_0$ ), independently flip a biased coin. With probability  $\frac{1+f}{2}$ , offer the product to the node for free. With probability  $\frac{1-f}{2}$ , offer the product to the node at cost  $c$ .

We henceforth refer to this strategy as STRATEGYMAXLEAF. Our analysis will revolve around what we term as “good” vertices, defined formally as follows:

**Definition 4.** *Given a graph  $G(V, E)$ , we define the good vertices to be the vertices with degree at least 3 and their neighbors.*

On the one hand, we show that if  $G$  has  $g$  good vertices, then the MAXLEAF algorithm will find a spanning tree with  $\Omega(g)$  leaves. We then show that each leaf of this tree leads to  $\Omega(1)$  revenue, implying STRATEGYMAXLEAF gives  $\Omega(g)$  revenue overall. Conversely, we can decompose  $G$  into at most  $g$  line-graphs joining high-degree vertices, and the total revenue from these is bounded by  $gL = O(g)$  for all policies, which gives the constant-factor approximation we need.

We begin by bounding the number of leaves in a max-leaf spanning tree. For dense graphs, we can rely on the following fact [8,11](#):

**Fact 1.** *The max-leaf spanning tree of a graph with minimum degree at least 3 has at least  $n/4 + 2$  leaves [8,11](#).*

In general graphs, we cannot apply this result directly. However, we can make any graph have minimum degree 3 by replacing degree-1 vertices with small, complete graphs and by contracting along edges to remove degree-2 vertices. We can then apply Fact [1](#) to analyze this auxiliary graph, which leads to the following result (all the proofs for this section have been moved to the full version of the paper [11](#) due to lack of space):

**Lemma 4.** *Suppose a connected graph  $G$  has  $n_3$  vertices with degree at least 3. Then  $G$  has a spanning tree with at least  $\frac{n_3}{8} + 1$  leaves.*

We must further extend this to be in terms of the number of good vertices  $g$ , rather than being in terms of  $n_3$ :

**Lemma 5.** *Given an undirected graph  $G$  with  $g$  good vertices, the  $\text{MAXLEAF}$  algorithm [11] will construct a spanning tree with  $\max(\frac{g}{50} + 0.5, 2)$  leaves.*

We can now use this to prove a guarantee on the performance of  $\text{STRATEGYMAXLEAF}$  in terms of the number of good vertices on an arbitrary graph:

**Lemma 6.** *Given a social network  $G$  with  $g$  good vertices,  $\text{STRATEGYMAXLEAF}$  guarantees an expected revenue of  $\Omega((1 - f)cq \cdot g)$ .*

Now that we have computed the expected revenue from  $\text{STRATEGYMAXLEAF}$ , we need to characterize the optimal revenue to bound the approximation ratio. This bound is given by the following lemma.

**Lemma 7.** *The maximum expected revenue achievable by any strategy (adaptive or not) on a social network  $G$  with  $g$  good vertices is  $O(L \cdot g)$ .*

Now, we can combine the above lemmas to state the main theorem of the paper, which states that  $\text{STRATEGYMAXLEAF}$  provides a constant factor approximation guarantee for the revenue.

**Theorem 1.** *Let  $K$  denote the complexity of our buyer decision model  $M$ . Then, the expected revenue generated by  $\text{STRATEGYMAXLEAF}$  on an arbitrary social network is  $O(K)$ -competitive with the expected revenue generated by the optimal (adaptive or not) strategy.*

*Proof.* This follows immediately from Lemmas 6 and 7, as well as the fact that  $K = \frac{L}{(1-f)cq}$ .

As a corollary, we get the fact that the adaptivity gap is also constant:

**Corollary 2.** *Let  $K$  denote the complexity of our buyer decision model  $M$ . Then the adaptivity gap is  $O(K)$ .*

Now we briefly address the issue of cashbacks that were ignored in this exposition. We set the cashback  $r$  to be a small fraction of our expected revenue from each individual  $r_0$ , i.e.  $r = z \cdot r_0$ , where  $z < 1$ . Then, our total profit will be  $n \cdot r_0 \cdot (1 - z)$ . Adding this cashback decreases our total profit by a constant factor that depends on  $z$ , but otherwise the argument now carries through as before, and nodes now have a positive incentive to pass on recommendations.

In light of Corollary 2, one might ask whether the adaptivity gap is not just 1. In other words, is there any benefit at all to be gained from using non-adaptive strategies? In fact, there is. For example, consider a social network consisting of 4 nodes  $\{v_1, v_2, v_3, v_4\}$  in a cycle, with  $v_3$  connected to two other isolated vertices. Suppose furthermore that a node will accept a recommendation with probability 0.5 unless the price is 0, in which case the node will accept it with probability 1. On this network, with seed set  $S^0 = \{v_1\}$ , the optimal adaptive strategy is to always demand full price unless exactly one of  $v_2$  and  $v_4$  purchases the product initially, in which case  $v_3$  should be offered the product for free. This beats the optimal non-adaptive strategy by a factor of 1.0625.



## Acknowledgments

This research has been supported in part by NSF Grant ITR-0331640, TRUST (NSF award number CCF-0424422), and grants from Cisco, Google, KAUST, Lightspeed, and Microsoft. The third author is grateful to Jason Hartline and Mukund Sundararajan for useful discussions.

## References

1. Arthur, D., Motwani, R., Sharma, A., Xu, Y.: Pricing strategies for viral marketing on social networks (2009), <http://arxiv.org/abs/0902.3485>
2. Arthur, D., Vassilvitskii, S.: k-means++: the advantages of careful seeding. In: SODA 2007: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms, Philadelphia, PA, USA, pp. 1027–1035. Society for Industrial and Applied Mathematics (2007)
3. Domingos, P., Richardson, M.: Mining the network value of customers. In: Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 57–66 (2001)
4. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. WH Freeman & Co., New York (1979)
5. Hartline, J., Mirrokni, V., Sundararajan, M.: Optimal Marketing Strategies over Social Networks. In: Proceedings of the 17th international conference on World Wide Web (2008)
6. Kempe, D., Kleinberg, J., Tardos, É.: Maximizing the spread of influence through a social network. In: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 137–146 (2003)
7. Kleinberg, J.: Cascading Behavior in Networks: Algorithmic and Economic Issues. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.) Algorithmic Game Theory. Cambridge University Press, New York (2007)
8. Kleitman, D.J., West, D.B.: Spanning Trees with Many Leaves. SIAM Journal on Discrete Mathematics 4, 99 (1991)
9. Leskovec, J., Singh, A., Kleinberg, J.: Patterns of influence in a recommendation network. In: Ng, W.-K., Kitsuregawa, M., Li, J., Chang, K. (eds.) PAKDD 2006. LNCS (LNAI), vol. 3918, pp. 380–389. Springer, Heidelberg (2006)
10. Leskovec, J., Adamic, L.A., Huberman, B.A.: The dynamics of viral marketing. ACM Trans. Web 1(1), 5 (2007)
11. Lu, H.I., Ravi, R.: Approximating Maximum Leaf Spanning Trees in Almost Linear Time. Journal of Algorithms 29(1), 132–141 (1998)
12. BBC News. Facebook valued at \$15 billion (2007), <http://news.bbc.co.uk/2/hi/business/7061042.stm>
13. Schonfeld, E.: Amiando makes tickets go viral and widgetizes event management (2008), <http://www.techcrunch.com/2008/07/17/amiando-makes-tickets-go-viral-and-widgetizes-event-event-management-200-discount-for-techcrunch-readers/>
14. Solis-Oba, R.: 2-Approximation Algorithm for finding a Spanning Tree with Maximum Number of leaves. In: Proceedings of the Sixth European Symposium on Algorithms, pp. 441–452 (1998)
15. Wikipedia. Facebook revenue in 2008 (2008), <http://en.wikipedia.org/wiki/Facebook>



# Consistent Continuous Trust-Based Recommendation Systems

Ola Rozenfeld and Moshe Tennenholtz

Technion – Israel Institute of Technology, Haifa 32000, Israel,  
Microsoft R&D Labs, Israel  
ola.rozenfeld@gmail.com,  
moshet@ie.technion.ac.il

**Abstract.** The goal of a trust-based recommendation system is to generate personalized recommendations from known opinions and trust relationships. Prior work introduced the axiomatic approach to trust-based recommendation systems, but has been extremely limited by considering binary systems, while allowing these systems to be inconsistent. In this work we introduce an axiomatic approach to deal with consistent continuous trust-based recommendation systems. We introduce the model, discuss some basic axioms, and provide a characterization of a class of systems satisfying a set of basic axioms. In addition, as it turns out, relaxing some of the axioms leads to additional interesting systems, which we examine.

## 1 Introduction

Many online systems offer their users access to a vast variety of products or services, which drives the need for high quality personalized recommendations. Often, these systems utilize social network structures of their users, as well as the users' opinions of one another and the products, in order to improve the quality of recommendations. Google's page ranking system [17] uses links to represent "voting" for web pages, Amazon and eBay's reputation systems (e.g. [18]) aggregate feedbacks that users leave for transactions, and the Epinions trust/reputation system (e.g. [16]) aggregates explicit trust/distrust links between its users. In recent years, recommendation and reputation systems became the focus of intensive research (e.g. [4,19,8,22,11]).

In this work we focus on the setting where there is one item of interest, and various users have rated this item. A user wishes to "predict" her own rating of the item by consulting her friends, who, in turn, might consult their friends and so on. There exist many automated recommendation systems that fit this general framework; however, this raises the question of comparing the relative merits of these systems. There exist two main approaches to studying and comparing recommendation/reputation systems: the experimental approach and the axiomatic approach. The experimental approach evaluates the performance of a given system on a particular set of users and ratings (along various performance metrics); the obvious advantage of this approach is that it provides a practical estimate of the real accuracy of the system. However, the results of a particular system on different data might be different; we need a general understanding of the properties of different systems *before* we decide which one we want to implement for

a particular setting. The axiomatic approach, which has a long history in social choice theory, aims to achieve these goals.

In [6] the authors use the axiomatic approach in the restricted case of recommendation systems in which the setting is an annotated directed graph, where some of the nodes are labeled by votes of  $+$  and  $-$ . In that model a node represents an agent, an edge directed from  $a$  to  $b$  represents the fact that agent  $a$  trusts agent  $b$ , and a subset of the nodes are labeled by  $+$  or  $-$ , indicating that these nodes have already formed positive or negative opinions about the item under question. Based on this input, a recommendation system must output a recommendation for each unlabeled node. In this paper we extend this study to the general case, where trust values and votes are continuous, allowing users to express a range of recommendations, instead of just a binary “like/dislike”. This extension is not only a technical one, since it allows considering *consistent* trust-based recommendation systems, as explained below.

One of the most desired properties of a recommendation system is that it should be consistent with its own recommendations. Namely, if a system comes with a particular recommendation to agent  $p$ , based on other agents’ observations and the trust network, then if this recommendation turns out to be correct, this should not change any recommendations for the other agents. It is easy to see that the simplified binary setting does not allow for such consistency. Hence, we view the axiomatic study of consistent continuous trust-based recommendation systems as essential for the understanding/analysis of desirable realistic systems.

## 1.1 Overview of Results

In section 2 we introduce our model of continuous trust-based recommendation systems, where weights are non-negative real numbers, and votes are normalized to the interval  $[-1, 1]$ . In this framework we define our model axioms, which are minimal requirements we may wish to have, as well as more elaborated axioms dealing with the notion of consistency; these are further extended to several monotonicity axioms. In section 3 we introduce an axiom termed Independence of Irrelevant Stuff [IIS], and show that this axiom simplifies matters by allowing to consider only the way votes are locally aggregated. We show a natural recommendation system that satisfies our axioms, the Random Walk system, which is an extension of the system studied in previous work. Much emphasis in this paper is given to transformations of recommendation systems which preserve their desired properties; in particular, in section 4 we introduce two transformations that allow to tweak an existing recommendation system in order to give more/less weight to radical votes, and to define the meaning of different trust values. This brings us in section 5 to a theorem which fully characterizes the set of recommendation systems which satisfy a set of five natural axioms; namely, this characterization shows that all systems satisfying these axioms must be a modification of the Random Walk system under the two transformations above, and that each such modified Random Walk system satisfies the axioms. We also show that all five axioms are essential. By relaxing the axioms, we get some interesting new recommendation systems: systems which incorporate discount factors and systems which incorporate the aggregated network vote. Due to the tight space constraints of the proceedings, these do not appear here, but they can be seen online in the full version

of the paper at the authors' web sites: <http://technion.ac.il/~olga/> or <http://ie.technion.ac.il/Home/Users/Moshet0.html>

Similarly, most proofs are omitted here, but they all appear online.

### 1.2 Additional Related Work


There are several ways to study recommendation systems. Standard evaluation tools include simulations and field experiments (e.g. [7,18,13]). In addition, one may also consider computational properties of suggested systems. As far as axiomatic studies are concerned, our work builds on previous work on axiomatizations of ranking systems. The literature on the axiomatic approach to ranking systems deals with both global ranking systems [1,2,21,9,22,4,5] and personalized ranking systems [7,10,3,15]. Personalized ranking systems are very close to trust-based recommendation systems. In such systems, agents rank some of the other agents. Then an aggregated ranking of agents, personalized to the perspective of a particular agent, is generated based on that information. However, previous studies on the axiomatic approach have not been concerned with situations where the participants share reviews or opinions on items of interest which are external to the system. Many existing recommendation system are based on collaborative filtering (CF), which is a completely different approach than the trust-based systems considered in this paper. Combining trust-based and CF approaches is a direction of recent research [20].

## 2 Model and Axioms

The trust-based recommendation setting can be modeled formally as follows:

A *voting network* is a directed annotated graph  $G = (N, V, E, v, w)$  where  $N$  is a set of agents,  $V \subseteq N$  is a set of voters (the agents who have an informed opinion on the item of interest),  $v : V \rightarrow [-1, 1]$  gives the vote of each voter, which represents his opinion on the item (normalized to  $[-1, 1]$ ),  $w : E \rightarrow \mathbb{R}^{\geq 0}$ , where  $w(x, y)$  represents the trust of agent  $x$  in agent  $y$ . We will denote by  $\bar{V}$  the set of non-voters  $N \setminus V$ .

Example of a voting network appears in Fig.1



**Fig. 1.** Here, voters are designated by a double circle ( $V = \{C, D\}$ ). The votes  $v$  appear inside the circle; the trust weights appear near the appropriate edge.

A *recommendation system*  $R$  takes as input a voting network  $G$  and has as output the recommendations  $r(u) \in [-1, 1]$  for every  $u \in \bar{V}$  ( $R(G) : \bar{V} \rightarrow [-1, 1]$ ). We denote  $R(G)$  by  $R_G$ ; for notation simplicity, we extend the definition of  $R_G$  to  $N$  by setting  $R_G(u) = v(u)$  for all  $v \in V$  (so that  $R_G(u)$  now means recommendation or vote of  $u$ ).

This model is the natural extension of the model in [6] for continuous trust and vote values. First, we restrict ourselves to recommendation systems that satisfy the following basic conditions:

1. **Anonymity:** Isomorphic graphs correspond to isomorphic recommendations. Formally, for a permutation  $\pi$  of  $N$  such that  $\pi(G) = G'$  (where  $\pi(G)$  stands for applying  $\pi$  on  $v$  and  $w$  appropriately), it holds that  $R(G') = \pi(R_G)$
2. **Neutrality:** The system is a priori indifferent towards positive or negative opinions – switching the signs on votes will cause switched signs of recommendations. Formally,  $R(-G) = -R_G$ , where  $-G = (N, V, E, -v, w)$ .
3. **No-edge equals zero trust:** Adding edges with weight 0 does not affect recommendations.
4. **No-node equals orphan non-voter:** Adding non-voters with no incoming or outgoing edges does not affect any existing recommendations.
5. **Continuity:** The recommendation of a node  $v$  is a *continuous* function of the votes and the trust in every point, except possibly for points where all  $v$ 's outgoing trust equals 0.

We call this set of conditions **Model Axioms (MA)**. In essence, we want these requirements as part of our model. Requirements 1 and 2 are natural, and appeared in the same form in the binary model of [6] as well (here, we renamed Symmetry to Anonymity in order to be consistent with social choice literature). In 3 we wanted to express that trust value of 0 is equivalent to no trust at all, and in 4 we wanted to express our implicit assumption that  $N$  represents only those agents who are in some sense informed (either know other agents or tried the product); implicitly, there is an infinite amount of agents in the system, but the recommendation system should only take the informed ones into account. (An alternative way to model the requirements 3 and 4 was to fix the set of agents to the set of natural numbers,  $\mathbb{N}$ , and require the input trust weights to induce a finite graph). Finally, in 5, we want the system to be “stable” in a sense: we don’t want small changes in the input to cause big changes in the output. However, the case where a node  $v$  is a sink is a reasonable exception to this rule, since trusting no one at all and having trust in someone (no matter how small) are fundamentally different situations. For example, in many natural systems it holds that, if a node trusts a single voter with a vote of  $v$ , then its recommendation will also be  $v$ , no matter what is the weight of the trust link. However, if a node does not trust anyone at all, its recommendation is 0. We don’t want these cases to violate our continuity requirement.

In the following we implicitly assume MA; in particular, when we speak of paths in a voting network  $G$  we will refer only to edges of positive weight; parents (or predecessors) of a node  $u$  ( $Pred(u)$ ) are nodes  $z$  for which  $w(z, u) > 0$ ; similarly, children (or successors) of a node  $u$  ( $Succ(u)$ ) are nodes  $z$  for which  $w(u, z) > 0$ . However, we admit that the requirement MA[4] is debatable; we will show some natural recommendation systems that violate it.


In this work, we want to concentrate on the consistency requirement:

**Node Consistency (NC):** Let  $u$  be a non-voter. If we turn  $u$  into a voter with the vote  $R_G(u)$ , no recommendations will change. Formally, for any voting network  $G$  and agent

$u \in \bar{V}$ , let  $G' = (N, V \cup \{u\}, E, v \cup \{(u, R_G(u))\}, w)$ . Then, for all  $z \in \bar{V}$ ,  $R_{G'}(z) = R_G(z)$ .

Intuitively, consistent recommendation systems rely on their own predictions – they don’t change their prediction for an agent if their predictions for other agents turned out to be correct. Their output is stable, in a sense: if the output is added to the input, nothing changes.

Consistency seems a very natural requirement; however, many systems used in real life are not consistent. Consider, for example, the Majority system: the recommendation of a node is 1 if the majority of his neighbors cast a positive vote, -1 if the majority of his neighbors cast a negative vote, and 0 otherwise. This system does not rely on its own predictions, since only voter neighbors are counted (it violates node consistency). A similar problem occurs in all binary systems – the Random Walk and Personalized Page Rank systems, as defined in [6], are not consistent because of similar considerations: intuitively, a binary system which has to commit to recommendations of 0, 1 or -1, cannot be very sensitive to nuances in the input (in the sense that many different inputs map to the same output), and therefore has little chances of being consistent. [1] However, this particular problem is somewhat artificial: we can redefine the systems to fit appropriately into the continuous setting. For example, the Personalized Page Rank system (PPR) [12] can be adapted as follows: the recommendation of a node  $v$  is the expected vote value of a voter that can be reached from a random walk starting from  $v$ , with a restarting probability  $\alpha$  (in the binary setting, it was defined as the sign of the above value). It is easy to see that even this continuous system is not consistent. Consider the graph in Fig 2: the PPR recommendation of node B is 1; but if B becomes a voter with a vote of 1, the PPR recommendation of A will become 0.5 instead of 0.2.



**Fig. 2.** Here, the votes and recommendations of PPR appear inside the circles;  $\alpha = 0.5$ ; the trust weight is 1 for all edges

Thus far, we spoke of consistency in terms of node values (votes). However, our system has two kinds of input: nodes and edges (votes and trust). Below is one intuitive requirement concerning trust values:

**Edge Consistency (EC):** If we increase trust in an agent with a value (recommendation or vote) equal to ours, our value will not change. Formally, for any voting network  $G$ , agent  $u \in \bar{V}$ , edge  $(u, z) \in E$ , and  $w' > w(u, z)$  let  $G' = (N, V, E, v, w \setminus \{(u, z), w(u, z)\} \cup \{(u, z), w'\})$ . Then,  $R_G(u) = R_G(z) \Rightarrow R_{G'}(u) = R_G(u)$ .

Consistency in itself is too weak a requirement: it only tells us how the system should behave if its output is fed back into the input, without any changes. For example, a recommendation system that always gives a recommendation of 0 to all agents is consistent, as well as the AVE system that gives the recommendation of the average vote to

<sup>1</sup> It would be interesting to translate this intuition into a formal impossibility proof.

all agents. Now we would like to express the intuition about how the output of a recommendation system should change if the input changes. The input of the system comes in two parts: votes and trust. Therefore, as before, we formulate two requirements: what happens when a vote changes, and what happens when trust changes.

**Node Monotonicity (NM):** For any voting network  $G$  and agent  $u \in N$ , let  $G' = (N, V \cup \{u\}, E, v \setminus \{(u, R_G(u))\} \cup \{(u, r')\}, w)$ . Then:

1. If  $r' = R_G(u)$ , then for every  $z \in N$ ,  $R_{G'}(z) = R_G(z)$
2. For every parent  $z$  of  $u$ ,  
 $sgn(R_{G'}(z) - R_G(z)) = sgn(r' - R_G(u))$

**Edge Monotonicity (EM):** For any voting network  $G$ , agent  $u \in \bar{V}$ , edge  $(u, z) \in E$ , and  $w' > w(u, z)$  let  $G' = (N, V, E, v, w \setminus \{(u, z), w(u, z)\} \cup \{(u, z), w'\})$ . Then,  $sgn(R_{G'}(u) - R_G(u)) = sgn(R_G(z) - R_G(u))$ .

Here,  $sgn(x)$  is defined for  $x \in \mathbb{R}$  as:

$$sgn(x) = \begin{cases} \frac{x}{|x|} & x \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

Note that NM implies NC and EM implies EC (follows from the definitions). Intuitively, by node monotonicity we demand the following: if we change a parent's vote, the child's vote must change appropriately (this is what makes the monotonicity requirement strict). The trivial system of recommending 0 to everyone satisfies NC, but not NM. Edge monotonicity states that increasing trust in a given agent, all other things being equal, should bring our recommendation strictly closer to the vote of that agent<sup>2</sup> (note that due to NC, it doesn't matter if we restrict  $z$  to belong to  $V$  in the definition of EM). Note also that AVE system satisfies NM and EC, but not EM (since it ignores the trust network).


The first question of interest for us is: which recommendation systems satisfy NM and EM? The Random Walk system (RW) is defined as follows: intuitively, to get a recommendation for a non-voter  $u$ , we perform a random walk in  $G$  starting from  $u$ , where in each step we move from node  $z$  to a random neighbor  $z'$ , choosing  $z'$  with a probability proportionate to the trust of  $z$  in  $z'$ , relatively to the overall trust of  $z$ . If the walk reaches a voter, the value of the walk is the vote value. Otherwise, the value of the walk is 0. The recommendation is the expected value of a walk over all random walks. Formally, for a voting network  $G$ , let  $S \subseteq \bar{V}$  be the set of non-voters that cannot reach any voter. For each  $u \in N$ , create a variable  $r_u \in [-1, 1]$ . Solve the following from  $r_u$ :

$$r_u = \begin{cases} 0 & u \in S \\ v(u) & u \in V \\ \frac{\sum_{z \in Succ(u)} w(u, z) r_z}{\sum_{z \in Succ(u)} w(u, z)} & \text{otherwise} \end{cases}$$

The recommendations  $RW(G)$  are defined as  $RW(G)(u) = r_u$ .

Fig 3 shows the result of RW on the example network from Fig 1.

<sup>2</sup> The Positive Response axiom in the binary setting of [6] can be derived from EM in our setting.



**Fig. 3.** Here, the votes and the recommendations of RW appear inside the circles

This system is a direct adaptation of the Random Walk system from [6] to the continuous setting (in the binary setting, the recommendation of  $u$  was -1, 0 or 1 according to the sign of  $r_u$ , which is neither continuous nor consistent). It is easy to see that the RW recommendation system satisfies MA, NM and EM.

### 3 IIS Axiom

The following axiom makes things easier, because it allows us to restrict attention to “local” functions:

**Independence of Irrelevant Stuff (IIS)** encompasses the following two requirements:

1. Let  $z \in N$  be a node not reachable from node  $u$ . Then for the subgraph  $G'$  in which node  $z$  and all its associated edges were removed,  $R_{G'}(u) = R_G(u)$ .
2. Let  $G = (N, V, E, v, w)$  and  $e \in V \times N$  an edge leaving a voter. Then for the subgraph  $G' = (N, V, E \setminus \{e\}, v, w \setminus \{(e, w(e))\})$  in which  $e$  has been removed,  $R_{G'} = R_G$ .

The first requirement demands that a path of trust exist between a node  $u$  and a given voter in order for  $u$  to be influenced in any way by that voter – a condition consistent with what we expect from trust-based recommendation systems: we want to ignore non-trusted opinions. The second requirement captures the intuition that an agent trusts his own opinion infinitely more than opinions of others. It is easy to see that RW satisfies IIS; in the full version of the paper (see <http://technion.ac.il/~olga>) we show some systems that satisfy MA+NM+EM, but not IIS[1]<sup>3</sup>

**Proposition 1.** *A recommendation system  $R$  that satisfies MA, IIS and NC can be written as:  $R_G(u) = F \begin{pmatrix} v_1, \dots, v_n \\ w_1, \dots, w_n \end{pmatrix}$  for each non-voter  $u$ , where the  $v_i$ 's are the recommendations/votes of children of  $u$ ,  $w_i$ 's are the appropriate weights (formally: for each  $z_i$  s.t.  $(u, z_i) \in E$ ,  $v_i = R_G(z_i)$ ,  $w_i = w(u, z_i)$ ), and  $F$  is an aggregate function of  $v_1, \dots, v_n, w_1, \dots, w_n$  (also denoted  $F(\vec{v}, \vec{w})$ ) that satisfies the following:*

1.  $F(\vec{v}, \vec{w}) \in [-1, 1]$  for all  $v_1, \dots, v_n \in [-1, 1]$  and  $w_1, \dots, w_n \in \mathbb{R}^{\geq 0}$ .
2.  $F \begin{pmatrix} v_1, \dots, v_n \\ w_1, \dots, w_n \end{pmatrix} = F \begin{pmatrix} v_{\pi(1)}, \dots, v_{\pi(n)} \\ w_{\pi(1)}, \dots, w_{\pi(n)} \end{pmatrix}$ , where  $\pi$  is any permutation of  $\{1, \dots, n\}$ .

<sup>3</sup> We conjecture that IIS[2] can be derived from MA+NM+EM, but we have yet to find a proof.



3.  $F(\vec{v}, \vec{w}) = -F(-\vec{v}, \vec{w})$
4.  $F(\vec{v}, \vec{0}) = 0$
5.  $F\left(\begin{matrix} v_1, v_2, \dots, v_n \\ 0, w_2, \dots, w_n \end{matrix}\right) = F\left(\begin{matrix} v_2, \dots, v_n \\ w_2, \dots, w_n \end{matrix}\right)$
6.  $F$  is continuous everywhere except possibly at points where  $\vec{w} = \vec{0}$

## 4 Scaling Votes and Trust

We have already seen one recommendation system that satisfies MA, NM, EM and IIS; the following propositions will enable us to create many more such systems. All of them introduce transformations on recommendation systems that preserve (some of) their desired properties.

More formally, let  $T$  be a transformation on the space of recommendation systems. We say that  $T$  *preserves a set of properties*  $X$  if, for every recommendation system  $R$  satisfying *all* the properties in  $X$ , the recommendation system  $T(R)$  satisfies all these properties as well.

The following transformation can be used to give more (or less) weight to “radical” opinions (those closer to the ends of the  $[-1, 1]$  scale):

**Proposition 2.** *Let  $f : [-1, 1] \rightarrow [-1, 1]$  be a continuous monotone anti-symmetric onto function. Then, the transformation  $T_f(R) = f^{-1} \circ R \circ f$  (meaning: apply  $f$  to all the votes in  $G$ , then apply  $R$ , then apply  $f^{-1}$  on the resulting recommendations) preserves  $\{MA, NM, EM\}$  and  $\{MA, NM, EM, IIS\}$*

Note that a continuous monotone function is also one-to-one. If we want to use a function  $f$  which is not onto, we need an additional assumption in order for the transformation to preserve the desired properties (without an additional assumption,  $f^{-1}$  might not be well defined). We call this assumption Conservativeness:

**Conservativeness:** A recommendation system  $R$  is said to satisfy Conservativeness if for every voting network  $G$  and non-voter  $u$  with a single positive weighted outgoing edge to a voter  $z$  (w.l.o.g.  $v(z) \geq 0$ ),  $0 \leq R_G(u) \leq v(z)$ .

Conservativeness states that a recommendation for an agent cannot be more radical than the vote of someone he trusts – the value can only be altered in the direction of uncertainty. We can show that conservative systems satisfy the natural betweenness property – the recommendation of a node falls between the votes of his neighbors, possibly skewed towards uncertainty:

**Proposition 3.** *Let  $R$  be a recommendation system that satisfies MA, NC, EM, IIS and Conservativeness. For a voting network  $G$  and  $u \in N$ , we define  $R_{Succ(u)} = \{R_G(z) | z \in Succ(u)\}$ . Then, for every voting network  $G$  and every  $u \in \overline{V}$ ,*

$$\min\{R_{Succ(u)} \cup \{0\}\} \leq R_G(u) \leq \max\{R_{Succ(u)} \cup \{0\}\}$$

*Moreover, when  $|R_{Succ(u)}| > 1$ , the inequalities are strict:*

$$\min\{R_{Succ(u)} \cup \{0\}\} < R_G(u) < \max\{R_{Succ(u)} \cup \{0\}\}$$

The RW system obviously satisfies Conservativeness. If we begin with a conservative system, we can show a wider family of transformations that preserve our desired properties:



**Proposition 4.** *Let  $f : [-1, 1] \rightarrow [-1, 1]$  be a continuous monotone anti-symmetric function. Then, the transformation  $T_f(R) = f^{-1} \circ R \circ f$  preserves  $\{MA, NM, EM, IIS, Conservativeness\}$ .*

*Proof (Sketch).* From IIS and prop. 3 above,  $T_f(R)$  is well defined, even if  $f$  is not onto; therefore, MA, NM, EM and IIS are preserved by prop. 2. Preservation of Conservativeness follows directly from monotonicity of  $f$ .  $\square$

It is an interesting open question whether the Conservativeness requirement is indeed necessary for the proof – meaning, whether there exist recommendation systems satisfying MA, NM, EM and IIS, but not Conservativeness.

The following transformation can be used to specify what contributes more to the decision – many neighbors with small trust values or few neighbors with larger trust values (given that the total trust is the same in both cases):

**Proposition 5.** *Let  $g : \mathbb{R}^+ \rightarrow \mathbb{R}^+$  be a continuous monotone increasing function, with  $g(0) = 0$ . Then, the transformation  $T_g(R) = R \circ g$  (meaning: apply  $g$  to all the weights in  $G$ , then apply  $R$ ) preserves  $\{MA, NM, EM\}$  and  $\{MA, NM, EM, IIS\}$ .*

*Proof (Sketch).* The proof follows directly from properties of  $R$  and  $g$ ; the assumption  $g(0) = 0$  is required in order for MA[3] to hold.  $\square$

## 5 Axiomatic Characterization of RW Systems

In this section we wish to characterize the family of recommendation systems that results from the closure of RW under the transformations  $T_f$  and  $T_g$ , for any  $f$  and  $g$ . In order to do that, we need to define some additional properties of recommendation systems. We start with some intuition:

One desirable requirement is the transitivity of trust relations (appeared originally in [6]). Let  $G$  be a graph,  $v$  a node, and  $A, B$  two disjoint sets of nodes. We say that  $v$  trusts  $A$  more than  $B$  if the recommendation of  $v$  in a voting network on  $G$  where all nodes in  $A$  vote 1, all nodes in  $B$  vote -1, and no other nodes vote, is positive.

**Transitivity (TR):** We say that a recommendation system  $R$  is transitive if, for every graph  $G$ , node  $v$  and disjoint sets  $A, B, C$ , if  $v$  trusts  $A$  more than  $B$  and  $v$  trusts  $B$  more than  $C$ , then  $v$  trusts  $A$  more than  $C$ .

[6] showed, for the binary setting, that Transitivity is impossible to reconcile with IIS and their versions of monotonicity. This still holds in the continuous case:

**Proposition 6.** *The requirements MA[1,2,3], IIS, either one of NM or EM, and Transitivity are inconsistent.*

So the Transitivity requirement needs to be weakened, if we want any hope of accommodating it. In the binary setting, [6] had the axiom of Trust Propagation: if  $u$  trusts  $v$  with a weight  $k$ , and  $v$  trusts (only) nodes  $v_1, \dots, v_k$  with a weight of 1, then the edge  $(u, v)$  can be replaced with  $k$  edges  $(u, v_i)$ , with no recommendations being affected.


The direct translation of this axiom into the continuous setting would have the weight on the edge  $(u, v)$  equal to the sum of the weights on  $(v, v_i)$ , but we found it to be too

restrictive. We did not want to limit the way the trust weights are interpreted (does a trust link of weight 2 “matter” exactly as much as two trust links of weight 1?). Therefore, the axiom we chose retains the spirit of Trust Propagation, while avoiding to specify the exact way of translating the trust weights:

**Separability:** Let  $\oplus$  be an operator  $\oplus : \mathbb{R}^{\geq 0} \times \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}^{\geq 0}$  satisfying:

1.  $w \oplus 0 = w$
2. **Associativity:**  $(w_1 \oplus w_2) \oplus w_3 = w_1 \oplus (w_2 \oplus w_3)$
3.  $\oplus$  is a continuous, strictly monotone increasing function in both variables

We say that a recommendation system  $R$  is *separable* if, for any voting network  $G$ , if a node  $v$  trusts only nodes  $v_1, v_2$  with respective trust weights of  $w_1, w_2$ , and a node  $u$  trusts  $v$  with a weight of  $(w_1 \oplus w_2)$ , then the edge  $(u, v)$  can be replaced with edges  $(u, v_1), (u, v_2)$  with weights  $w_1, w_2$ , and no recommendations will change (see Fig. 4).



**Fig. 4.** Separability axiom

The RW system obviously satisfies Separability, with  $\oplus = +$ .

Next, we have a more strict variant of Conservativeness:

**Neighborhood Consensus (NCS):** A recommendation system  $R$  is said to satisfy Neighborhood Consensus if for every voting network  $G$  and non-voter  $u$  with a single positive weighted outgoing edge to a voter  $z$ ,  $R_G(u) = v(z)$ .

Note that NCS+EM imply that if all neighbors of a node agree on their votes, the recommendation for that node equals these votes (hence the name). Despite being more strict than Conservativeness, this property is very natural and appears in many social choice systems. The RW system also satisfies this requirement.

Now we can formulate the characterization result:

**Theorem 1.** *A recommendation system  $R$  can be defined as the unique solution of the system of equations:*

$$r_u = \begin{cases} 0 & \text{no path exists from } u \text{ to a voter} \\ v(u) & u \in V \\ f^{-1} \left[ \frac{\sum_{z \in Succ(u)} g(w(u,z)) f(r_z)}{\sum_{z \in Succ(u)} g(w(u,z))} \right] & \text{otherwise} \end{cases}$$

where  $f$  is a continuous monotone anti-symmetric function, and  $g$  a continuous monotone increasing function, with  $g(0) = 0$ , if and only if  $R$  satisfies the following conditions: MA, IIS, NM, NCS and Separability.

We note that each axiom in the above theorem is necessary for the result to hold; namely, there exist examples of systems satisfying all but one of these axioms which could not be written in the above form. Below we shortly discuss each example.

Relaxing MA is not interesting – in particular, it makes no sense to consider systems which are not anonymous. Relaxing NCS is the focus of the section on discount factors (see full paper online) – in particular, applying prop. 7 there on RW can be used to create examples of systems satisfying MA, NM, IIS and Separability, but not NCS. It is also possible to relax IIS (see section 7 in the full paper), resulting in systems satisfying MA, NM, NCS and Separability, but not IIS. Relaxing the NM requirement is trivial: a system that always gives recommendations of 0 will do. Relaxing Separability – we apply the following transformation on the graph  $G$ : for each node  $v$  with outgoing edges  $e_1, \dots, e_k$  and appropriate weights  $w_1, \dots, w_k$ , let  $t_v = w_1 + \dots + w_k$ . We change the weights of  $e_1, \dots, e_k$  to  $w_1^{t_v}, \dots, w_k^{t_v}$ , respectively (we raise all weights to the power of  $t_v$ ). Let  $G'$  be the graph resulting from the transformation. Our system returns the result of RW on  $G'$ . It is easy to see that the system satisfies MA, NM, NCS and IIS, but not Separability.

## 6 Conclusions and Further Work

This paper explored continuous trust based recommendation systems. Our main requirement was consistency - we focused on recommendation systems which rely on their own predictions. First, we showed the “basic” such system – Random Walk – and then we introduced a few transformations that customize recommendation systems according to specific needs: give more or less weight to radical opinions, tweak the importance of having multiple trust links of small weights vs. few trust links of large weights, tweak the “damping factor” (discount) of trusted opinions, and consider the aggregated opinion of the network. All along we used the axiomatic approach – we started by formalizing our requirements from recommendation systems, and then identified families of recommendation systems that conform to these requirements.

All the recommendation systems in this paper were some version of Random Walk applied to a transformed voting network. A central question is whether this focus is justified. In section 5 we showed that if we want our system to conform to some additional requirements, then, indeed, we have no choice but to use a generalized version of Random Walk. But what if we consider the most general possible recommendation system? Are there any continuous consistent recommendation systems *at all* that are not fundamentally based on Random Walk? We strongly suspect that the answer is negative; however, the concept of being “fundamentally based on Random Walk” is not easy to formalize. Attempts of full characterization of consistent continuous recommendation systems are in the focus of our current research.

## References

1. Altman, A., Tennenholtz, M.: On the axiomatic foundations of ranking systems. In: Proc. 19th International Joint Conference on Artificial Intelligence, pp. 917–922 (2005)
2. Altman, A., Tennenholtz, M.: The pagerank axioms. In: Proceedings of the 6th ACM conference on Electronic Commerce (2005)

3. Altman, A., Tennenholtz, M.: An axiomatic approach to personalized ranking systems. In: Proc. 20th International Joint Conference on Artificial Intelligence (2007)
4. Altman, A., Tennenholtz, M.: Quantifying incentive compatibility of ranking systems. In: Proc. of AAAI 2006 (2006)
5. Altman, A., Tennenholtz, M.: Incentive compatible ranking systems. In: Proceedings of the 2007 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2007 (2007)
6. Andersen, R., Borgs, C., Chayes, J., Feige, U., Flaxman, A., Kalai, A., Mirrokni, V., Tennenholtz, M.: Trust-based recommendation systems: an axiomatic approach. In: Proceedings of WWW 2008 (2008)
7. Avesani, P., Massa, P., Tiella, R.: A trust-enhanced recommender system application: Moleskiing. In: SAC 2005: Proceedings of the, ACM symposium on Applied computing, pp. 1589–1593. ACM Press, New York (2005)
8. Bakos, Y., Dellarocas, C.N.: Cooperation without enforcement? a comparative analysis of litigation and online reputation as quality assurance mechanisms. MIT Sloan School of Management Working Paper No. 4295-03 (2003)
9. Borodin, A., Roberts, G.O., Rosenthal, J.S., Tsaparas, P.: Link analysis ranking: algorithms, theory, and experiments. *ACM Trans. Inter. Tech.* 5(1), 231–297 (2005)
10. Cheng, A., Friedman, E.: Sybilproof reputation mechanisms. In: P2PECON 2005: Proceeding of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems, pp. 128–132. ACM Press, New York (2005)
11. Dash, R., Ramchurn, S., Jennings, N.: Trust-based mechanism design. In: Proceedings of the Third International Joint Conference on Autonomous Agents and MultiAgent Systems, pp. 748–755 (2004)
12. Haveliwala, T.: Topic-sensitive pagerank. In: WWW 2002: Proceedings of the 11th international conference on World Wide Web, pp. 517–526 (2002)
13. Houser, D., Wooders, J.: Reputation in auctions: Theory and evidence from ebay. Working Paper, University of Arizona (2000)
14. Kleinberg, J.M.: Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)* 46(5), 604–632 (1999)
15. Levien, R.: Attack Resistant Trust Metrics. PhD thesis, University of California, Berkeley (2002)
16. Massa, P., Avesani, P.: Controversial users demand local trust metrics: An experimental study on epinions.com community. In: Proc. of AAAI 2005, pp. 121–126 (2005)
17. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: Bringing order to the web. Technical Report, Stanford University (1998)
18. Resnick, P., Zeckhauser, R.: Trust among strangers in internet transactions: Empirical analysis of ebay’s reputation system. *The Economics of the Internet and E-Commerce*, M.R. Baye (2002)
19. Resnick, P., Zeckhauser, R., Friedman, R., Kuwabara, E.: Reputation systems. *Communications of the ACM* 43(12), 45–48 (2000)
20. Singh, A.P., Gunawardana, A., Meek, C., Surendran, A.C.: Recommendations using absorbing random walks. In: North East Student Colloquium on Artificial Intelligence, NESCAI (2007)
21. Slutzki, G., Volij, O.: Ranking participants in generalized tournaments. *International Journal of Game Theory* 33(2), 255–270 (2005)
22. Tennenholtz, M.: Reputation systems: An axiomatic approach. In: Proceedings of UAI 2004 (2004)

# The Stackelberg Minimum Spanning Tree Game on Planar and Bounded-Treewidth Graphs

Jean Cardinal<sup>1</sup>, Erik D. Demaine<sup>2</sup>, Samuel Fiorini<sup>1</sup>,  
Gwenaël Joret<sup>1,\*</sup>, Ilan Newman<sup>3</sup>, and Oren Weimann<sup>4</sup>

<sup>1</sup> Université Libre de Bruxelles (ULB), B-1050 Brussels, Belgium  
{jcardin,sfiorini,gjoret}@ulb.ac.be

<sup>2</sup> MIT, Cambridge, MA 02139, USA  
edemaine@mit.edu

<sup>3</sup> University of Haifa, Haifa 31905, Israel  
ilan@cs.haifa.ac.il

<sup>4</sup> Weizmann Institute of Science, Rehovot 76100, Israel  
oren.weimann@weizmann.ac.il

**Abstract.** The Stackelberg Minimum Spanning Tree Game is a two-level combinatorial pricing problem introduced at WADS'07. The game is played on a graph, whose edges are colored either red or blue, and where the red edges have a given fixed cost. The first player chooses an assignment of prices to the blue edges, and the second player then buys the cheapest spanning tree, using any combination of red and blue edges. The goal of the first player is to maximize the total price of purchased blue edges. We study this problem in the cases of planar and bounded-treewidth graphs. We show that the problem is NP-hard on planar graphs but can be solved in polynomial time on graphs of bounded treewidth.


## 1 Introduction

A young startup company has just acquired a collection of point-to-point tubes between various sites on the Interweb. The company's goal is to sell the use of these tubes to a particularly stingy client, who will buy a minimum-cost spanning tree of the network. Unfortunately, the company has a direct competitor: the government sells the use of a different collection of point-to-point tubes at publicly known prices. Our goal is to set the company's tube prices to maximize the company's income, given the government's prices and the knowledge that the client will buy a minimum spanning tree made from any combination of company and government tubes. Naturally, if we set the prices too high, the client will rather buy the government's tubes, while if we set the prices too low, we unnecessarily reduce the company's income.

This problem is called the *Stackelberg Minimum Spanning Tree Game* [CDF<sup>+</sup>07], and is an example in the growing family of algorithmic game-theoretic problems about combinatorial optimization in graphs [BHK08, GvLSU09, LMS98, RSM05, BGPW08]. More formally, we are given an undirected graph  $G$  (possibly with parallel edges, but no loops), whose edge set  $E(G)$  is partitioned into

---

\* Aspirant F.R.S. – FNRS.



**Fig. 1.** A sample instance of the STACKMST problem. The goal is to assign prices to the blue edges to maximize the total price of the blue edges purchased in a minimum spanning tree.

a red edge set  $R(G)$  and a blue edge set  $B(G)$ . We are also given a cost function  $c : R(G) \rightarrow \mathbb{R}^+$  assigning a positive cost to each red edge. The STACKMST problem is to assign a price  $p(e)$  to each blue edge  $e$ , resulting in a weighted graph  $(G, c \cup p)$ , to maximize the total price of blue edges in a minimum spanning tree. We assume that, if there is more than one minimum spanning tree, we obtain the maximum possible income. (Otherwise, we could decrease the prices slightly and get arbitrarily close to the same income.) Figure 1 shows an example.

This problem is thus a two-player two-level optimization problem, in which the leader (the company) chooses a strategy (a price assignment), taking into account the strategy of the follower (the client), which is determined by a second-level optimization problem (the minimum spanning tree problem). Such a game is known as a *Stackelberg game* in economics.

The complexity and approximability of the STACKMST problem has been studied in a previous paper [CDF<sup>+</sup>07], which shows the following results. The problem is APX-hard, but can be approximated within a logarithmic factor. Constant-factor approximation exist for the special cases in which the given costs are bounded or take a bounded number of distinct values. Finally, an integer programming formulation has an integrality gap corresponding to the best known approximation factors.

Instead of restricting the edge weights, we can restrict the class of allowed graphs, with the hope of obtaining better approximation algorithms. One natural class of graphs is *planar graphs*, on which many important problems admit polynomial-time approximation schemes. Many of these results use Baker’s technique [Bak94] or more modern variations [Kle06, DHM07, DHK09], which ultimately rely on the ability to efficiently solve the problem in graphs of bounded treewidth in polynomial time. Such algorithms generally use dynamic programming, using a textbook technique for well-behaved problems. In particular, the problem of checking a graph-theoretic property expressible in monadic second-order logic is fixed-parameter tractable with respect to the treewidth of the graph; see [Com08] for a survey. However, few if any such dynamic programs have been developed for a two-level optimization problem such as STACKMST, and standard techniques do not seem to apply.

In this paper, we consider the STACKMST problem in these two graph classes: planar graphs and bounded-treewidth graphs. We prove in Section 2 that STACKMST remains NP-hard when restricted to planar graphs. We develop in Section 4 a polynomial-time dynamic programming algorithm for STACKMST in graphs of bounded treewidth. Along the way, we develop in Section 3 a dynamic programming algorithm for series-parallel graphs, or equivalently, biconnected graphs of treewidth at most 2, which are also planar.

To our knowledge, our algorithms are the first examples of a two-level pricing problem solved by dynamic programming on a graph decomposition tree. We believe that this result provides insight into the structure of the problem, and could be a stepping stone toward a polynomial-time approximation scheme for planar graphs. More generally, we believe that our techniques may be useful in the design of dynamic programming algorithms for other pricing problems in graphs, including pricing problems with many followers [BHK08, GvLSU09], and Stackelberg problems involving shortest paths [RSM05] or shortest path trees [BGPW08].


## 2 Planar Graphs

We consider the STACKMST problem on planar graphs. We strengthen the hardness result given in [CDF<sup>+</sup>07] by showing that the problem remains NP-hard in this special case. The reduction is from the *minimum connected vertex cover problem*, which is known to be NP-hard, even in planar graphs [GJ79]. The minimum connected vertex cover problem involves finding a minimum-size subset  $C$  of the vertices of a graph, such that every edge has at least one endpoint in  $C$ , and  $C$  induces a connected graph.

**Theorem 1.** *The STACKMST problem is NP-hard, even when restricted to planar graphs.*

The reduction is the following. Given a planar graph  $G = (V, E)$ , with  $|V| = n$  and  $|E| = m$ , we construct an instance of STACKMST with red costs in  $\{1, 2\}$ . Let  $G' = (V', R \cup B)$  be the graph for this instance, with  $(R, B)$  a bipartition of the edge set. We first let  $V' = V \cup E$ . The set of blue edges  $B$  is the set  $\{ve : e \in E, v \in e\}$ . Thus the blue subgraph is the vertex-edge incidence graph of  $G$ , which is clearly planar. Given a planar embedding of the blue subgraph, we connect all vertices  $e \in E$  of  $G'$  by a tree, all edges of which are red and have cost 1. The graph can be kept planar by letting those red edges be nonintersecting chords of the faces of the embedding. Finally, we double all blue edges by red edges of cost 2. The whole construction is illustrated in figure 2(a). Let  $t$  be a positive integer. It can be shown that the revenue for an optimal price function for  $G'$  is at least  $m + 2n - t - 1$  if and only if there exists a connected vertex cover of  $G$  of size at most  $t$ . The proof can be found in the long version of the paper [1].

<sup>1</sup> <http://arxiv.org/abs/0909.3221>



**Fig. 2.** Illustration of the reduction in Theorem 1

### 3 Series-Parallel Graphs

We now describe a polynomial-time dynamic programming algorithm for solving the STACKMST problem on series-parallel graphs.

We use the following inductive definition of (connected) series-parallel graphs. Consider a connected graph  $G$  with two distinguished vertices  $s$  and  $t$ . The graph  $(G, s, t)$  is a *series-parallel* graph if either  $G$  is a single edge  $(s, t)$ , or  $G$  is a *series* or *parallel* composition of two series-parallel graphs  $(G_1, s_1, t_1)$  and  $(G_2, s_2, t_2)$ . The series composition of  $G_1$  and  $G_2$  is formed by setting  $s = s_1, t = t_2$  and identifying  $t_1 = s_2$ ; the parallel composition is formed by identifying  $s = s_1 = s_2$  and  $t = t_1 = t_2$ .

**Theorem 2.** *The STACKMST problem can be solved in  $O(m^4)$  time on series-parallel graphs.*

#### 3.1 Definitions

Let us fix an instance of STACKMST, that is, a graph  $G$  with  $E(G) = R(G) \cup B(G)$  endowed with a cost function  $c : R(G) \rightarrow \mathbb{R}_+$ . Denote by  $c_1, c_2, \dots, c_k$  the different values taken by  $c$ , in increasing order. Let also  $c_0 := 0$ .

For two distinct vertices  $s, t \in V(G)$  of  $G$  and a subset  $F \subseteq B(G)$  of blue edges, define  $\mathcal{P}(G, F, s, t)$  as the set of  $st$ -paths in the graph  $(V(G), R(G) \cup F)$ . Let also  $\tilde{\mathcal{P}}(G, F, s, t)$  denote the subset of paths in  $\mathcal{P}(G, F, s, t)$  that contain at least one red edge. Cardinal *et al.* [CDF<sup>+</sup>07] proved the following.

**Lemma 1** ([CDF<sup>+</sup>07]). *Suppose that  $G$  contains a red spanning tree, and let  $F \subseteq B(G)$  be an acyclic subset of blue edges. Then, the maximum revenue achievable by the leader, over solutions where the set of blue edges bought by the follower is exactly  $F$ , is obtained by setting the price of each edge  $st \notin F$  to  $+\infty$ , and the price of each edge  $st \in F$  to  $\min \left\{ \max_{e \in P \cap R(G)} c(e) \mid P \in \tilde{\mathcal{P}}(G, F, s, t) \right\}$ .*

This lemma states that if we know the set of blue edges that will eventually be bought, the price of a selected blue edge  $st$  is given by the minimum, over the paths from  $s$  to  $t$ , of the largest red cost on this path.



Motivated by this result, we introduce some more notations. For a subset  $Z \subseteq E(G)$  of edges, we define  $\text{mc}(Z)$  as the maximum cost of a red edge in  $Z$  if  $Z \cap R(G) \neq \emptyset$ , as  $c_0 = 0$  otherwise. (The two letters  $\text{mc}$  stand for “max cost”.) We define  $w(G, F, s, t)$  as

$$w(G, F, s, t) := \begin{cases} \min \{ \text{mc}(P) \mid P \in \mathcal{P}(G, F, s, t) \} & \text{if } \mathcal{P}(G, F, s, t) \neq \emptyset; \\ c_k & \text{otherwise.} \end{cases}$$

Similarly,

$$\tilde{w}(G, F, s, t) := \begin{cases} \min \{ \text{mc}(P) \mid P \in \tilde{\mathcal{P}}(G, F, s, t) \} & \text{if } \tilde{\mathcal{P}}(G, F, s, t) \neq \emptyset; \\ c_k & \text{otherwise.} \end{cases}$$

Thus, the price assigned to the edge  $st \in F$  in Lemma 11 is  $\tilde{w}(G, F, s, t)$ . Also, we will consider graphs that do not necessarily contain a red spanning tree; this is why we need to treat the case where  $\mathcal{P}(G, F, s, t)$  or  $\tilde{\mathcal{P}}(G, F, s, t)$  is empty in the above definitions.

In what follows, we let  $[k] := \{0, 1, \dots, k\}$ . Our algorithm for series-parallel graphs associates a value to each pair  $(H, q)$ , where  $q \in [k]^2$ , and  $H$  is a graph appearing in the series-parallel decomposition of  $G$ .

A subset  $F \subseteq B(G)$  of blue edges realizes  $q = (i, j) \in [k]^2$  in  $(G, s, t)$  if  $F$  is acyclic and  $w(G, F, s, t) = c_i$ . Although this property does not depend on  $j$ , the formulation will appear to be convenient. Similarly, we say that  $q$  is *realizable* in  $(G, s, t)$  if there exists such a subset  $F$ .

For  $j \in [k]$  and distinct vertices  $s, t \in V(G)$ , let  $G^+$  denote the graph  $G$  with an additional red edge between  $s$  and  $t$  of cost  $c_j$ . We define

$$\text{OPT}_{(i,j)}(G, s, t) := \max \left\{ \sum_{uv \in F} \tilde{w}(G^+, F, u, v) \mid F \text{ realizes } (i, j) \text{ in } (G, s, t) \right\},$$

if such a subset  $F \subseteq B(G)$  exists, and set  $\text{OPT}_{(i,j)}(G, s, t) := -\infty$  otherwise.

Intuitively, we want to keep track of optimal acyclic subsets of blue edges for every graph  $G$  obtained during the construction of a series-parallel graph. The problem is, that the weights of the blue edges in the optimal solution might change as we compose graphs in the series-parallel decomposition. However, the weights of edges depend only on the maximum red costs, or *bottlenecks*, of the new  $st$ -paths that will be added to  $G$ . We can thus prepare  $\text{OPT}(G, s, t)$  for every possible set of bottlenecks. These bottlenecks are the values  $j$  in what precedes.

Note that by Lemma 11, if  $G$  has a red spanning tree, then the maximum revenue achievable by the leader on instance  $G$  equals  $\max_{i \in [k]} \text{OPT}_{(i,k)}(G, s, t)$ . This will be the result returned by the algorithm.

### 3.2 Series Compositions

Let  $q = (i, j)$ ,  $q_1 = (i_1, j_1)$ , and  $q_2 = (i_2, j_2)$ , with  $q, q_1, q_2 \in [k]^2$ . We say that the pair  $(q_1, q_2)$  is *series-compatible* with  $q$  if

- (S1)  $\max\{i_1, i_2\} = i$ ;
- (S2)  $\max\{j, i_2\} = j_1$ , and
- (S3)  $\max\{j, i_1\} = j_2$ ,

Our dynamic program uses the following recursion.

**Lemma 2.** *Suppose that  $(G, s, t)$  is a series composition of  $(G_1, s_1, t_1)$  and  $(G_2, s_2, t_2)$ , and that  $q \in [k]^2$  is realizable in  $(G, s, t)$ . Then*

$$\text{OPT}_q(G, s, t) = \max\{\text{OPT}_{q_1}(G_1, s_1, t_1) + \text{OPT}_{q_2}(G_2, s_2, t_2)\}$$

where the maximum is taken over all pairs  $(q_1, q_2)$  that are series-compatible with  $q$ .

We now prove that the recursion is valid. We need the following lemmas. In what follows,  $(G, s, t)$  is a series composition of  $(G_1, s_1, t_1)$  and  $(G_2, s_2, t_2)$ ;  $q, q_1, q_2 \in [k]^2$  with  $q = (i, j)$ ,  $q_1 = (i_1, j_1)$ , and  $q_2 = (i_2, j_2)$  are such that  $(q_1, q_2)$  is series-compatible with  $q$ ; and  $F_\ell \subseteq B(G_\ell)$  realizes  $q_\ell$  in  $(G_\ell, s, t)$ , for  $\ell = 1, 2$ .

We first observe that  $F := F_1 \cup F_2$  realizes  $q$ .

**Lemma 3.**  *$F$  realizes  $q$  in  $(G, s, t)$ .*

*Proof.* Since  $V(G_1) \cap V(G_2) = \{t_1\} (= \{s_2\})$ , the set  $F$  is clearly acyclic. It remains to show  $w(G, F, s, t) = c_i$ . Every  $st$ -path in  $\mathcal{P}(G, F, s, t)$  is the combination of an  $s_1t_1$ -path of  $\mathcal{P}(G_1, F_1, s_1, t_1)$  with an  $s_2t_2$ -path of  $\mathcal{P}(G_2, F_2, s_2, t_2)$ . It follows

$$w(G, F, s, t) = \max\{w(G_1, F_1, s_1, t_1), w(G_2, F_2, s_2, t_2)\} = \max\{c_{i_1}, c_{i_2}\} = c_i,$$

where the last equality is from (S1). □

The next lemma motivates the definition of series-compatibility.

**Lemma 4.** *Let  $G^+$  be the graph  $G$  augmented with a red edge  $st$  of cost  $c_j$ , and  $G_\ell^+$  (for  $\ell = 1, 2$ ) the graph  $G_\ell$  augmented with a red edge  $s_\ell t_\ell$  of cost  $c_{j_\ell}$ . Then for  $\ell = 1, 2$  and every edge  $uv \in F_\ell$ ,*

$$\tilde{w}(G^+, F, u, v) = \tilde{w}(G_\ell^+, F_\ell, u, v).$$

*Proof.* We prove the statement for  $\ell = 1$ , the case  $\ell = 2$  follows by symmetry. Let  $uv \in F_1$ , and let  $e = st$  and  $e_1 = s_1t_1$  be the additional red edges in  $G^+$  and  $G_1^+$ , respectively.

*Claim.*  $\tilde{w}(G^+, F, u, v) \geq \tilde{w}(G_1^+, F_1, u, v)$ .

*Proof.* The claim is true if  $\tilde{\mathcal{P}}(G^+, F, u, v) = \emptyset$ , since then  $\tilde{w}(G^+, F, u, v) = c_k \geq \tilde{w}(G_1^+, F_1, u, v)$ . Suppose thus  $\tilde{\mathcal{P}}(G^+, F, u, v) \neq \emptyset$ , and let  $P \in \tilde{\mathcal{P}}(G^+, F, u, v)$ . It is enough to show that  $\text{mc}(P) \geq \tilde{w}(G_1^+, F_1, u, v)$ . This clearly holds if  $e \notin E(P)$ , as  $P$  belongs then also to  $\tilde{\mathcal{P}}(G_1^+, F_1, u, v)$  (recall that  $|V(G_1) \cap V(G_2)| = 1$ ). Hence, we may assume  $e \in E(P)$ . It follows  $s_1, t_1 \in V(P)$ .

Let  $P_1$  denote the path of  $\tilde{\mathcal{P}}(G_1^+, F_1, u, v)$  obtained by replacing the subpath  $s_1Pt_1$  of  $P$  with the edge  $e_1$ . Using (S2), we obtain

$$\text{mc}(s_1Pt_1) = \max\{c_j, \text{mc}(t_2Pt_1)\} \geq \max\{c_j, c_{i_2}\} = c_{j_1},$$

implying  $\text{mc}(P) \geq \text{mc}(P_1) \geq \tilde{w}(G_1^+, F_1, u, v)$ . □

*Claim.* Conversely,  $\tilde{w}(G^+, F, u, v) \leq \tilde{w}(G_1^+, F_1, u, v)$ .

*Proof.* Again, this trivially holds if  $\tilde{\mathcal{P}}(G_1^+, F_1, u, v)$  is empty. Suppose thus  $\tilde{\mathcal{P}}(G_1^+, F_1, u, v) \neq \emptyset$ , and let  $P_1 \in \tilde{\mathcal{P}}(G_1^+, F_1, u, v)$ . Similarly as before, it is enough to show that  $\tilde{w}(G^+, F, u, v) \leq \text{mc}(P_1)$ . This is true if  $e_1 \notin E(P_1)$ , since then  $P_1 \in \tilde{\mathcal{P}}(G^+, F, u, v)$ . Assume thus  $e_1 \in E(P_1)$ .

If  $\mathcal{P}(G_2, F_2, s_2, t_2) = \emptyset$ , then  $i_2 = k$  and  $\text{mc}(P_1) \geq c_{j_1} = \max\{c_j, c_{i_2}\} = c_k \geq \tilde{w}(G^+, F, u, v)$  by (S2). We may thus assume that  $\mathcal{P}(G_2, F_2, s_2, t_2)$  contains a path  $P_2$ ; we choose  $P_2$  such that  $\text{mc}(P_2) = c_{i_2}$ .

Denote by  $P$  the path obtained from  $P_1$  by replacing the edge  $e_1$  with the combination of edge  $e$  and path  $P_2$ . Since  $P \in \tilde{\mathcal{P}}(G^+, F, u, v)$ , (S2) yields

$$\begin{aligned} \text{mc}(P_1) &= \max\{c_{j_1}, \text{mc}(P_1 - e_1)\} = \max\{c_j, c_{i_2}, \text{mc}(P_1 - e_1)\} \\ &= \max\{c_j, \text{mc}(P_2), \text{mc}(P_1 - e_1)\} = \text{mc}(P) \geq \tilde{w}(G^+, F, u, v). \quad \square \end{aligned}$$

The lemma follows from Claims 3.2 and 3.2 □

We are now ready to prove the correctness of the recursion step in Lemma 2.

*Proof (Lemma 2).* Let  $q$  and  $G^+$  be defined as before. We first show:

*Claim.* There exist  $q_1, q_2 \in [k]^2$  such that  $(q_1, q_2)$  is series-compatible with  $q$  and  $\text{OPT}_q(G, s, t) \leq \text{OPT}_{q_1}(G_1, s, t) + \text{OPT}_{q_2}(G_2, s, t)$ .

*Proof.* Let  $F \subseteq B(G)$  be a subset of blue edges realizing  $q$  in  $(G, s, t)$  such that

$$\text{OPT}_q(G, s, t) = \sum_{uv \in F} \tilde{w}(G^+, F, u, v).$$

For  $\ell = 1, 2$ , let also  $F_\ell := F \cap E(G_\ell)$  and  $q_\ell := (i_\ell, j_\ell)$ , with  $i_\ell$  the index such that  $c_{i_\ell} = w(G_\ell, F_\ell, s_\ell, t_\ell)$ , and  $j_\ell := \max\{j, i_{\ell+1}\}$  (indices are taken modulo 2).  $F_\ell$  ( $\ell = 1, 2$ ) clearly realizes  $q_\ell$  in  $(G_\ell, s_\ell, t_\ell)$ . It is also easily verified that  $(q_1, q_2)$  is series-compatible with  $q$ . Hence we can apply Lemma 4

$$\begin{aligned} \text{OPT}_q(G, s, t) &= \sum_{uv \in F_1} \tilde{w}(G_1^+, F_1, u, v) + \sum_{uv \in F_2} \tilde{w}(G_2^+, F_2, u, v) \\ &\leq \text{OPT}_{q_1}(G_1, s_1, t_1) + \text{OPT}_{q_2}(G_2, s_2, t_2), \end{aligned}$$

as claimed. □

*Claim.*  $\text{OPT}_q(G, s, t) \geq \text{OPT}_{q_1}(G_1, s_1, t_1) + \text{OPT}_{q_2}(G_2, s_2, t_2)$  holds for every  $q_1, q_2 \in [k]^2$  such that  $(q_1, q_2)$  is series-compatible with  $q$ .

*Proof.* Suppose that  $(q_1, q_2)$  is series-compatible with  $q$ . Let  $F_\ell \subseteq B(G_\ell)$  ( $\ell = 1, 2$ ) be a subset of blue edges of  $G_\ell$  such that

$$\text{OPT}_{q_\ell}(G_\ell, s_\ell, t_\ell) = \sum_{uv \in F_\ell} \tilde{w}(G_\ell^+, F_\ell, u, v).$$

By Lemma 3,  $F := F_1 \cup F_2$  realizes  $q$  in  $(G, s, t)$ . Using again Lemma 4, we have:

$$\begin{aligned} \text{OPT}_q(G, s, t) &\geq \sum_{uv \in F} \tilde{w}(G^+, F, u, v) \\ &= \sum_{uv \in F_1} \tilde{w}(G_1^+, F_1, u, v) + \sum_{uv \in F_2} \tilde{w}(G_2^+, F_2, u, v) \\ &= \text{OPT}_{q_1}(G_1, s_1, t_1) + \text{OPT}_{q_2}(G_2, s_2, t_2), \end{aligned}$$

and the claim follows. □

The lemma follows from Claims 3.2 and 3.2. □

### 3.3 Parallel Compositions

The recursion step for parallel compositions follows a similar scheme. Let  $q, q_1, q_2 \in [k]^2$  with  $q = (i, j)$ ,  $q_1 = (i_1, j_1)$ , and  $q_2 = (i_2, j_2)$ . We say that the pair  $(q_1, q_2)$  is *parallel-compatible* with  $q$  if

- (P1) at least one of  $i_1, i_2$  is non-zero;
- (P2)  $\min\{i_1, i_2\} = i$ ;
- (P3)  $\min\{j, i_2\} = j_1$ , and
- (P4)  $\min\{j, i_1\} = j_2$ ,

The recursion step for parallel composition is as follows. The proof is omitted, due to the space limitation. We refer the reader to the long version of the paper.

**Lemma 5.** *Suppose that  $(G, s, t)$  is a parallel composition of  $(G_1, s, t)$  and  $(G_2, s, t)$ , and that  $q \in [k]^2$  is realizable in  $(G, s, t)$ . Then*

$$\text{OPT}_q(G, s, t) = \max\{\text{OPT}_{q_1}(G_1, s, t) + \text{OPT}_{q_2}(G_2, s, t)\}$$

where the maximum is taken over all pairs  $(q_1, q_2)$  that are parallel-compatible with  $q$ .

### 3.4 The Algorithm

A series-parallel decomposition of a connected series-parallel graph can be computed in linear time [VTL82]. Given such a decomposition, Lemmas 2 and 5 yield the following algorithm: consider each graph  $(H, s, t)$  in the decomposition tree in a bottom-up fashion. If  $H$  is a single edge, compute  $\text{OPT}_q(H, s, t)$  for every  $q \in [k]^2$ . If  $(H, s, t)$  is a series or parallel composition of  $(H_1, s_1, t_1)$  and  $(H_2, s_2, t_2)$ , compute  $\text{OPT}_q(H, s, t)$  for every  $q \in [k]^2$  based on the previously computed values for  $(H_1, s_1, t_1)$  and  $(H_2, s_2, t_2)$ , relying on Lemmas 2 and 5.

For every  $q = (i, j) \in [k]^2$ , there are  $O(k)$  possible values for either series-compatible or parallel-compatible pairs  $(q_1, q_2)$ . Hence every step costs  $O(k)$  times. Since there are  $O(k^2)$  possible values for  $q$ , and  $O(m)$  graphs in the decomposition of  $G$ , the overall complexity is  $O(k^3m) = O(m^4)$ . Furthermore, it is not difficult to keep track of a witness at each step.

## 4 Bounded-Treewidth Graphs

In the previous section, we gave a polynomial-time algorithm for solving the STACKMST problem on series-parallel graphs, which are biconnected graphs of treewidth 2. In this section, we extend the algorithm to graphs of bounded treewidth.

**Theorem 3.** *The STACKMST problem can be solved in  $m^{O(\omega^2)}$  time on graphs of treewidth  $\omega$ .*

We follow Abrahamson and Fellows [AF93] and characterize a graph of treewidth  $\omega$  as an  $\omega$ -*boundaried graph*. An  $\omega$ -boundaried graph is a graph with  $\omega$  distinguished vertices (called *boundary vertices*), each uniquely labeled by a label in  $\{1, \dots, \omega\}$ .  $\omega$ -boundaried graphs are formed recursively by the following composition operators:

1. The null operator  $\emptyset$  creates a boundaried graph which has only boundary vertices, and they are all isolated.
2. The binary operator  $\oplus$  takes the disjoint union of two  $\omega$ -boundaried graphs by identifying the  $i$ th boundary vertex of the first graph with the  $i$ th boundary vertex of the second graph. If there are only two boundary vertices  $s$  and  $t$  then this is exactly a parallel-composition.
3. The unary operator  $\eta$  introduces a new isolated vertex and makes this the new vertex with label 1 in the boundary. The previous vertex that was labeled 1 is removed from the boundary but not from the graph.
4. The unary operator  $\epsilon$  adds an edge between the vertices labeled 1 and 2 in the boundary.
5. Unary operators that permute the labels of the boundary vertices.

Any  $\omega$ -boundaried graph (and hence any graph of treewidth  $\omega$ ) can be constructed by applying  $O(\omega n)$  compositions according to the above five operators. This construction as well as the boundary vertices can be found in linear time [Bod96].

### 4.1 Definitions

Given an  $\omega$ -boundaried graph  $G = (V, E)$  and two distinct boundary vertices  $a, b \in \{1, 2, \dots, \omega\}$ , we call an  $ab$ -path *internal* if the only boundary vertices it passes through are  $a$  and  $b$ . We slightly modify the definition of  $\mathcal{P}(G, F, a, b)$ ,  $\widehat{\mathcal{P}}(G, F, a, b)$ ,  $w(G, F, a, b)$ , and  $\widetilde{w}(G, F, a, b)$  from the previous section to only include internal  $ab$ -paths.

As in the series-parallel case, we want to keep track of optimal acyclic subsets of blue edges for every graph  $G$  obtained during the construction of a bounded-treewidth graph. Notice that the weights of edges in an optimal solution depend only on the bottlenecks of the new internal paths that will be added to  $G$ . We thus prepare  $OPT(G)$ s for every possible set of bottlenecks (the  $j_{abs}$  in what follows) between any two boundary vertices  $a$  and  $b$ .

Let  $I_{\omega \times \omega}$  be a matrix of pairs where  $I[a, b] = I[b, a] = (i_{ab}, j_{ab})$  for some  $i_{ab}, j_{ab} \in \{0, 1, \dots, k\}$ . Let  $OPT_I(G)$  be the optimal solution to STACKMST (that is, an acyclic subset of blue edges  $F$ ) on the graph  $G^+$  obtained from  $G$  by adding a red edge connecting  $a$  and  $b$  of cost  $c_{j_{ab}}$  for every pair of distinct boundary vertices  $a, b \in \{1, 2, \dots, \omega\}$ , subject to the conditions that for every distinct  $a, b \in \{1, 2, \dots, \omega\}$  we have  $w(G, F, a, b) = c_{i_{ab}}$ .

During the construction, we store for every graph  $G$  the partial solutions  $OPT_I(G)$  for every possible  $I$ . In cases where  $OPT_I(G)$  is undefined (no proper  $F$  exists), we set  $OPT_I(G) = -\infty$ . Also, we abuse the notation  $OPT_I(G)$  for denoting both the acyclic subset  $F$  and its revenue.

## 4.2 The Algorithm

We now describe our bounded-treewidth algorithm by showing how to maintain the  $OPT_I$  information as  $G$  is constructed by the five operators. We present the algorithm along with a proof sketch of its correction. The formal proof of correctness follows similar lines to the series-parallel case. We use  $(i_{ab}, j_{ab})$  to denote  $I$ 's pairs,  $(i_{ab}^1, j_{ab}^1)$  for  $I_1$ 's pairs,  $(i_{ab}^2, j_{ab}^2)$  for  $I_2$ 's pairs, and  $(i'_{ab}, j'_{ab})$  for  $I'$ 's pairs.

We begin with the null operator  $\emptyset$  that creates a new graph  $G$  with isolated vertices labeled  $1, \dots, \omega$ . Therefore, we set  $OPT_I(G) = 0$  (associated with  $F = \emptyset$ ) for every  $I$  whose entries are all of the form  $(k, j_{ab})$ . The value  $i_{ab}$  is required to be  $k$  as there are no internal paths at all; the  $j_{ab}$  can be arbitrary values in  $\{0, 1, \dots, k\}$ . For all other  $I$  we set  $OPT_I(G) = -\infty$ .

If  $G = G_1 \oplus G_2$ , then  $OPT_I(G) = \max\{OPT_{I_1}(G) \cup OPT_{I_2}(G)\}$ . This operator is a lot like a parallel-composition of series-parallel graphs. Indeed, in the following conditions on the compatibility of  $I, I_1, I_2$ , the first four conditions are exactly the same as in a parallel-composition, only they must hold for every pair of boundary vertices (whereas in the series-parallel case there was only one pair). The fifth condition makes sure that the blue edges that will be purchased do not form a cycle. Therefore, if  $G = G_1 \oplus G_2$ , we require that

1. at least one of  $i_{ab}^1, i_{ab}^2$  is non-zero,
2.  $i_{ab} = \min\{i_{ab}^1, i_{ab}^2\}$ ,
3.  $j_{ab}^1 = \min\{j_{ab}, i_{ab}^2\}$ ,
4.  $j_{ab}^2 = \min\{j_{ab}, i_{ab}^1\}$ ,

for every distinct  $a, b \in \{1, \dots, \omega\}$ , and moreover that

5. the graph  $H$  is acyclic, where  $V(H) = \{1, \dots, \omega\}$  and distinct  $a, b \in V(H)$  are adjacent in  $H$  if  $i_{ab} = 0$ .

If  $G = \eta(G')$ , then a new isolated boundary vertex  $v$  with label 1 is created, and the old 1-labeled vertex  $u$  is now no longer a boundary vertex. Since no edges are modified, the optimal solutions for  $G'$  and  $G$  are the same and we set  $OPT_I(G) = OPT_{I'}(G')$ . Notice that an  $ab$ -path between two distinct boundary vertices  $a, b \neq u$  that go through  $u$  is not an *internal* path in  $G'$ , but could be in

$G$  (if the path does not contain any other boundary vertex). This fact is captured by the first two of the following four conditions. If  $G = \eta(G')$  we require that, for every distinct  $a, b \in \{2, \dots, \omega\}$ ,

1.  $i_{ab} = \min \{i'_{ab}, \max\{i'_{a1}, i'_{1b}\}\}$ ,
2.  $j'_{ab} = \min \{j_{ab}, \max\{j_{a1}, j_{1b}\}\}$ ,
3.  $i_{a1} = k$ , (there is no internal path incident to  $v$  as it is isolated)
4.  $j'_{a1} = k$ . (there will be no new internal paths originating from  $u$ )

If  $G = \epsilon(G')$ , then  $G$  is obtained from  $G'$  by adding an edge  $e$  between vertices labeled 1 and 2. Notice that  $\epsilon(G') = G' \oplus G''$  where  $G''$  is the boundary graph which has only boundary vertices, and its only edge is  $e$ . Thus, instead of dealing with the  $\epsilon$  operator we can introduce two new null-like operators that create a graph  $G$  isomorphic to  $G''$  with the edge  $e$  being red and blue, respectively.

- If  $e$  is red with cost  $c(e)$  then we set  $OPT_I(G) = 0$  (associated with  $F = \emptyset$ ) for every  $I$  whose entries are all of the form  $(k, j_{ab})$  (the  $j_{ab}$ s can be anything in  $\{0, 1, \dots, k\}$ ) except that  $i_{12} = c(e)$ . For all other  $I$  we set  $OPT_I(G) = -\infty$ .
- If  $e$  is blue then we set  $OPT_I(G) = 0$  (associated with  $F = \emptyset$ ) for every  $I$  whose entries are all of the form  $(k, j_{ab})$  (the  $j_{ab}$ s can be anything in  $\{0, 1, \dots, k\}$ ).

In addition, we set  $OPT_I(G) = p(e)$  for every  $I$  whose entries are all of the form  $(k, j_{ab})$  except that  $i_{12} = 0$ . This corresponds to setting  $F = \{e\}$ . For such a  $I$ , the price  $p(e)$  assigned to  $e$  is determined as follows. Let  $P_{12}$  be the set of all ordered sequences of the form  $a_1-a_2-\dots-a_t$  where every  $a_i \in \{1, 2, \dots, \omega\}$  is a unique boundary vertex,  $a_1 = 1$ ,  $a_t = 2$ , and  $2 \leq t \leq \omega$  (if  $t = 2$  then the path is simply 1-2). Every such sequence, together with  $e$  could close a cycle when new internal paths will be added. We therefore set

$$p(e) = \min_{a_1-\dots-a_t \in P_{12}} \max_{2 \leq i \leq t} j_{a_{i-1}a_i}$$

For all other  $I$  we set  $OPT_I(G) = -\infty$ .

Unary operators that permute the labels of the boundary vertices are trivial to handle. They merely represent a permutation of  $I$ .

*Time Complexity.* The composition operators require us to check every combination of at most three different  $I$  for compatibility. There are  $k^{\omega^2}$  possible  $I$ , so we need to check  $O(k^{3\omega^2})$  combinations. Each check requires at least  $O(\omega^2)$  time to read the  $I$ . The most time-consuming check is the one of the  $\epsilon$  operator when it adds a blue edge  $e$ . This might require figuring out  $p(e)$ . Notice that  $|P_{12}| < \omega!$ , so we can perform the check in  $O(\omega!)$  time. The total complexity of the above algorithm is therefore bounded by  $O(k^{3\omega^2} \cdot \omega!) = m^{O(\omega^2)}$ .

Although the problem is polynomial for every constant value of  $\omega$ , it is unclear whether there exists a fixed-parameter algorithm of complexity  $O(f(\omega)n^c)$  for some function  $f$  of  $\omega$  and a constant  $c$ . We conjecture that under reasonable complexity-theoretic assumptions, such an algorithm does not exist.

## References

- [AF93] Abrahamson, K.R., Fellows, M.R.: Finite automata, bounded treewidth, and well-quasiordering. In: Robertson, N., Seymour, P. (eds.) *Graph Structure Theory*, pp. 539–564 (1993)
- [Bak94] Baker, B.S.: Approximation algorithms for NP-complete problems on planar graphs. *J. ACM* 41(1), 153–180 (1994)
- [BGPW08] Bilò, D., Gualà, L., Proietti, G., Widmayer, P.: Computational aspects of a 2-player Stackelberg shortest paths tree game. In: Papadimitriou, C., Zhang, S. (eds.) *WINE 2008*. LNCS, vol. 5385, pp. 251–262. Springer, Heidelberg (2008)
- [BHK08] Briest, P., Hoefer, M., Krysta, P.: Stackelberg network pricing games. In: *Proc. 25th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 133–142 (2008)
- [Bod96] Bodlaender, H.L.: A linear time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.* 25, 1305–1317 (1996)
- [CDF<sup>+</sup>07] 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) (to appear in *Algorithmica*)
- [Cou08] Courcelle, B.: Graph structure and monadic second-order logic: Language theoretical aspects. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) *ICALP 2008, Part I*. LNCS, vol. 5125, pp. 1–13. Springer, Heidelberg (2008)
- [DHK09] Demaine, E.D., Hajiaghayi, M., Kawarabayashi, K.: Approximation algorithms via structural results for apex-minor-free graphs. In: *Proc. 36th International Colloquium on Automata, Languages and Programming, ICALP (2009)*
- [DHM07] Demaine, E.D., Hajiaghayi, M., Mohar, B.: Approximation algorithms via contraction decomposition. In: *Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 278–287 (2007)
- [GJ79] Garey, M.R., Johnson, D.S.: *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, New York (1979)
- [GvLSU09] Grigoriev, A., van Loon, J., Sitters, R., Uetz, M.: Optimal pricing of capacitated networks. *Networks* 53(1), 79–87 (2009)
- [Kle06] Klein, P.N.: A subset spanner for planar graphs, with application to subset TSP. In: *Proc. 38th ACM Symposium on Theory of Computing (STOC)*, pp. 749–756 (2006)
- [LMS98] Labbé, M., Marcotte, P., Savard, G.: A bilevel model of taxation and its application to optimal highway pricing. *Management Science* 44(12), 1608–1622 (1998)
- [RSM05] Roch, S., Savard, G., Marcotte, P.: An approximation algorithm for Stackelberg network pricing. *Networks* 46(1), 57–67 (2005)
- [VTL82] Valdes, J., Tarjan, R.E., Lawler, E.L.: The recognition of series parallel digraphs. *SIAM J. Comput.* 11(2), 298–313 (1982)



# Tighter Bounds for Facility Games

Pinyan Lu<sup>1</sup>, Yajun Wang<sup>1</sup>, and Yuan Zhou<sup>2</sup>

<sup>1</sup> Microsoft Research Asia

{pinyanl,yajunw}@microsoft.com

<sup>2</sup> Carnegie Mellon University

yuanzhou@cs.cmu.edu

**Abstract.** In one dimensional facility games, public facilities are placed based on the reported locations of the agents, where all the locations of agents and facilities are on a real line. The cost of an agent is measured by the distance from its location to the nearest facility.

We study the approximation ratio of social welfare for *strategy-proof* mechanisms, where no agent can benefit by misreporting its location. In this paper, we use the total cost of agents as social welfare function. We study two extensions of the simplest version as in [9]: two facilities and multiple locations per agent. In both cases, we analyze randomized *strategy-proof* mechanisms, and give the first lower bound of 1.045 and 1.33, respectively. The latter lower bound is obtained by solving a related linear programming problem, and we believe that this new technique of proving lower bounds for randomized mechanisms may find applications in other problems and is of independent interest.

We also improve several approximation bounds in [9], and confirm a conjecture in [9].

## 1 Introduction

In a facility game, a planner is building public facilities while agents (players) are submitting their locations. In this paper, we study the facility game in one dimension, i.e., the locations of the agents and the facilities are in the real line. Let the position reported by agent  $i$  be  $x_i \in \mathcal{R}_i \subseteq \mathcal{R}$ . Assume the number of agents is  $n$  and the number of public facilities available is  $k$ . A (deterministic) mechanism for the  $k$ -facility game is simply a function

$$f : \mathcal{R}_1 \times \mathcal{R}_2 \times \cdots \times \mathcal{R}_n \rightarrow \mathcal{R}^k.$$

In this paper, we assume  $\mathcal{R}_i = \mathcal{R}$  for all agents. The *cost* of an agent is the distance from its *true* location to the nearest facility. Let  $\{l_1, l_2, \dots, l_k\}$  be the set of locations of the facilities. The cost of agent  $i$  is  $\text{cost}(\{l_1, \dots, l_k\}, x_i) = \min_{1 \leq j \leq k} |x_i - l_j|$ . A randomized mechanism returns a distribution over  $\mathcal{R}^k$ . Then the cost of agent  $i$  is the *expected* cost over the distribution returned by the randomized mechanism.

An agent may misreport its location if it can reduce its own cost. A usual solution concept is *strategy-proofness*, which is also the focus of this paper. In

a *strategy-proof* mechanism, no agent can unilaterally misreport its location to reduce its own cost. For  $\mathbf{x} = \{x_1, x_2, \dots, x_i, \dots, x_n\} \in \mathcal{R}^n$ , we define  $\mathbf{x}_{-i} = \{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n\}$ . A mechanism is *strategy-proof* if for any  $x_i$  and  $x'_i \neq x_i$ ,  $\text{cost}(f(\mathbf{x}_{-i}, x_i), x_i) \leq \text{cost}(f(\mathbf{x}_{-i}, x'_i), x_i)$ . In other words, no matter what other agents' strategies are, one of the best strategies for agent  $i$  is reporting its true location. Our strategy-proof randomized mechanisms are defined by the *expected costs* of the agents.

The facility game problem has a rich history in social science literature. Consider the case that we are building one facility in a discrete set of locations (alternatives). Agents are reporting its preference for the alternatives. The renowned Gibbard-Satterthwaite theorem [6,10] showed that if the preference on the alternatives for each agent can be arbitrary, the only strategy-proof mechanisms are the dictatorships when the number of alternatives are greater than two.

In the facility game, however, the preferences on the facility locations are not arbitrary. In particular, agent  $i$  has a single preferred location  $x_i$ . When two locations are on the same side of  $x_i$ , agent  $i$  will always prefer the one closer to  $x_i$ . This kind of admissible individual preferences are defined as *single-peaked* preferences, which was first discussed by Black [3]. Since the Gibbard-Satterthwaite theorem does not hold with single-peaked preferences, the facility game admits a much richer set of strategy-proof mechanisms. Moulin [8] characterized the class of all strategy-proof mechanisms for *one-facility* game in the real line. (One unnecessary assumption in the proof is dropped by Barberà and Jakson [2], and Sprumont [12].) In particular, a *generalized median voter scheme* is sufficient to characterize all strategy-proof mechanisms. Interested readers may refer to the detailed survey by Barberà [1].

More recently, Procaccia and Tennenholtz [9] studied the facility game in a different perspective. They consider the facility game as a special case of the game theoretic optimization problems where the optimal social welfare solution is not strategy-proof. They treat the facility game in a broader concept of the games that payments are not allowed or infeasible. Such mechanism design problems without payments are rarely studied by computer scientists, except some special problems [11].

Procaccia and Tennenholtz studied strategy-proof mechanisms with provable approximation ratios on social welfare, when the optimal solution is not strategy-proof. For the simplest case of one facility, the median mechanism is both strategy-proof and optimal for social welfare. Then Procaccia and Tennenholtz studied two extensions: (1) there are two facilities; (2) each agent controls multiple locations (with one facility). In both cases, the optimal solutions are no longer strategy-proof in general. Therefore, it is interesting to study strategy-proof mechanisms with good approximation ratios for these extensions. This is also the focus of this paper. A strategy-proof mechanism has an approximation ratio of  $\alpha$  if for every input instance, the social cost for the output of the mechanism is always at most  $\alpha$  times the social cost for any solution.

We remark that, if payment is allowed, then the well-know Vickrey-Clarke-Groves (VCG) mechanism [13,47] will give both optimal and strategy-proof

solutions for both extensions. However, in many real world scenarios, payment is not available as noted by Schummer and Vohra [11]. We focus on the strategy-proof mechanisms without money in this paper.

## 1.1 Our Result

We study the approximation ratios of social welfare for the strategy-proof mechanisms in the facility game with one or more facilities. The social welfare function we use is the *social cost*, i.e., the total cost of all agents. We focus on the approximation ratios for *social cost* of the strategy-proof mechanisms, where we improve most results in [9]. Furthermore, we also provide several novel approximation bounds which are not previously available. Table 1 summarizes our contribution.

**Table 1.** Our results are in bold. The numbers in brackets are previous results in [9] unless stated otherwise. (N/A means no previous known bound.)

	Two Facilities	Multi-Location Per Agent (One Facility)
Deterministic	UB: $(n - 2)$ LB: <b>2</b> (1.5)	UB: (3 [5]) LB: (3 [5])
Randomized	UB: <b><math>n/2</math></b> $(n - 2)$ LB: <b>1.045</b> (N/A)	UB: <b>3</b> - $\frac{2 \min_{j \in N} w_j}{\sum_{j \in N} w_j} (2 + \frac{ w_1 - w_2 }{w_1 + w_2})$ for $n = 2$ only) LB: <b>1.33</b> (N/A)

The organization of the paper is as follows. In Section 2, we provide improved upper and lower bounds of both deterministic and randomized strategy-proof mechanisms for the two-facility game. In Section 3, we study the cases when each agent controls more than one location. We conclude our paper in Section 4 with several open problems.


## 2 The Two-Facility Game

In this section, we study strategy-proof mechanisms for the two-facility game. We first provide a better randomized mechanism achieving approximation ratio  $n/2$  for social cost. The only previously known upper bound is  $n - 2$ , which is from a deterministic mechanism. Then we study the lower bounds both for the deterministic and randomized cases. For deterministic mechanisms, the lower bound is improved to 2 from 1.5 in [9]. For randomized mechanisms, we provide the first non-trivial approximation ratio lower bound of 1.045.

### 2.1 A Better Randomized Mechanism

The following mechanism is inspired by Mechanism 2 from [9]. However, our proof is different and much simpler.

*Mechanism 1.* See Figure [1](#) for reference. Let  $\mathbf{x} = \{x_1, x_2, \dots, x_n\}$  be the reported locations of the agents. Define  $lt(\mathbf{x}) = \min\{x_i\}$ ,  $rt(\mathbf{x}) = \max\{x_i\}$  and  $mt(\mathbf{x}) = (lt(\mathbf{x}) + rt(\mathbf{x}))/2$ . We further define the *left boundary*  $lb(\mathbf{x}) = \max\{x_i : i \in N, x_i \leq mt(\mathbf{x})\}$  and the *right boundary*  $rb(\mathbf{x}) = \min\{x_i : i \in N, x_i \geq mt(\mathbf{x})\}$ . Let  $dist(\mathbf{x}) = \max\{rt(\mathbf{x}) - rb(\mathbf{x}), lb(\mathbf{x}) - lt(\mathbf{x})\}$ . We set  $\overline{lb}(\mathbf{x}) = lt(\mathbf{x}) + dist(\mathbf{x})$  and  $\overline{rb}(\mathbf{x}) = rt(\mathbf{x}) - dist(\mathbf{x})$ . The mechanism returns  $(lt(\mathbf{x}), rt(\mathbf{x}))$  or  $(\overline{lb}(\mathbf{x}), \overline{rb}(\mathbf{x}))$ , each with probability  $1/2$ .



**Fig. 1.** Mechanism 1 picks  $(lt(\mathbf{x}), rt(\mathbf{x}))$  or  $(\overline{lb}(\mathbf{x}), \overline{rb}(\mathbf{x}))$ , each with probability  $1/2$

**Theorem 1.** *Mechanism 1 is strategy-proof. The approximation ratio of Mechanism 1 is  $n/2$  for social cost.*

*Proof.* We first prove the approximation ratio assuming that all agents report their true locations. By symmetry, we assume  $rt(\mathbf{x}) - rb(\mathbf{x}) \geq lb(\mathbf{x}) - lt(\mathbf{x})$  as in Figure [1](#). Since we only have two facilities, either  $lt(\mathbf{x})$  and  $rb(\mathbf{x})$  or  $rb(\mathbf{x})$  and  $rt(\mathbf{x})$  are served by a same facility. Therefore the optimal solution is least  $\min\{|lt(\mathbf{x}) - rb(\mathbf{x})|, |rb(\mathbf{x}) - rt(\mathbf{x})|\} = dist(\mathbf{x})$ . On the other hand, for each agent, its expected cost is exactly  $dist(\mathbf{x})/2$  in this mechanism. So Mechanism 1 has an approximation ratio of  $\frac{n}{2}$ .

We then prove that Mechanism 1 is strategy-proof. We first show that any point other than the 3 points defining  $lt(\mathbf{x}), rt(\mathbf{x})$  and  $rb(\mathbf{x})$  cannot benefit by misreporting its location. Let the new configuration be  $\mathbf{x}'$ . Consider the 3 points defining the previous  $lt(\mathbf{x}), rt(\mathbf{x})$  and  $rb(\mathbf{x})$ . No matter how the 3 points are partitioned by the new  $mt(\mathbf{x}')$ ,  $dist(\mathbf{x}') \geq rt(\mathbf{x}) - rb(\mathbf{x})$ , where  $\mathbf{x}'$  is the new configuration. We know that the expected cost for any location in this configuration is at least  $dist(\mathbf{x}')/2$ , which is at least as large as the honest cost  $dist(\mathbf{x}) = rt(\mathbf{x}) - rb(\mathbf{x})$ . The same argument also shows  $lt(\mathbf{x})$  (resp.  $rt(\mathbf{x})$ ) does not have incentive of reporting positions on the left (resp. right).

Consider the point  $rb(\mathbf{x})$ . Its expected cost is  $\frac{rt(\mathbf{x}) - rb(\mathbf{x})}{2}$  if it reports its true location. By lying, it cannot move the left or right boundary towards itself, and as a result, its expected cost in any new configuration is at least  $\min\{|lt(\mathbf{x}) - rb(\mathbf{x})|, |rb(\mathbf{x}) - rt(\mathbf{x})|\}/2 = (rt(\mathbf{x}) - rb(\mathbf{x}))/2$ . Therefore, the point at  $rb(\mathbf{x})$  has no incentive to lie.


The only possible case left to analyze is that the agent at  $lt(\mathbf{x})$  (resp.  $rt(\mathbf{x})$ ) is reporting a location to the right (resp. left). Its expected cost is  $(rt(\mathbf{x}) - \overline{lb}(\mathbf{x}))/2$  if it reports its true location. Reporting a location on its right can only move  $\overline{lb}(\mathbf{x}')$  toward right, which will hurt itself. Therefore the agent at  $lt(\mathbf{x})$  has no

incentive to lie. Similar argument also holds if the agent at  $\text{rt}(\mathbf{x})$  reports its location on the left of  $\text{rt}(\mathbf{x})$ .

To sum up, no agent has incentive to lie. Therefore, **Mechanism 1** is strategy-proof.

### 2.2 Lower Bounds

In this section, we show the approximation ratio lower bounds both for deterministic and randomized strategy-proof mechanisms. Both bounds are proved by the following construction, which is similar to the 1.5 lower bound example in [9].



**Fig. 2.** Lower bound example for the two-facility game

**Theorem 2 (Lower bound for deterministic mechanisms).** *In the two-facility game, any deterministic strategy-proof mechanism  $f : \mathcal{R}^n \rightarrow \mathcal{R}^2$  has an approximation ratio of at least  $2 - \frac{4}{n-2}$  for social cost.*

*Proof.* See Figure 2 for the configuration. We have  $n - 2$  nodes at the origin and the left node at  $-1$  and the right node at  $1$ .

Assume to the contrary, there exists a strategy-proof mechanism with approximation ratio less than 2. Then this mechanism has to place one facility in the range  $(-\frac{2}{n-2}, \frac{2}{n-2})$ . Now consider the left node and the right node at  $-1$  and  $1$ . At least one of them is  $1 - 2/(n - 2)$  away from its closest facility. Without loss of generality, assume the right node at  $1$  is at least  $1 - \frac{2}{n-2}$  away from the facilities.

If there is one facility on the right of  $1$ , it must be placed at a position right to  $2 - 2/(n - 2)$  by our assumption. In this case, since the optimal cost is 1, the approximation ratio is at least  $2 - \frac{4}{n-2}$  as one facility is always close to the origin.

Now consider the case that the closest facility to the right node at  $1$  is on the left. Let  $I$  be the image set of the closest facility to the right node when the right node moves and all other nodes remain fixed. Clearly, by strategy-proofness,  $I \cap (\frac{2}{n-2}, 2 - \frac{2}{n-2}) = \emptyset$ . On the other hand,  $I \cap [2 - \frac{2}{n-2}, +\infty) \neq \emptyset$ , otherwise the approximation ratio is unbounded when the right node moves to the infinity.

Take  $p$  as the left most point of  $I \cap [2 - \frac{2}{n-2}, +\infty)$ . ( $p$  always exists, as  $I$  is a closed set.) If we place the right node at  $p - 1 + \frac{2}{n-2}$ , the closest facility to  $x$  is at  $p$ . Therefore, the cost of the mechanism for such a configuration is at least  $2 - \frac{4}{n-2}$ , as the other facility has to be close to the origin. Because the optimal cost is still 1, the approximation ratio is at least  $2 - \frac{4}{n-2}$ .

If the mechanism is randomized, the output is a distribution over  $\mathcal{R}^2$ . Notice that in a randomized mechanism, the cost of an agent is measured by the *expected distance* from its true location to the closest facility. We give the first non-trivial (greater than 1) approximation ratio lower bound of strategy-proof mechanisms for social cost in Theorem 3.

**Theorem 3 (Lower bound for randomized mechanisms).** *In a two-facility game, any randomized strategy-proof mechanism has an approximation ratio of at least  $1 + \frac{\sqrt{2}-1}{12-2\sqrt{2}} - \frac{1}{n-2} \geq 1.045 - \frac{1}{n-2}$  for the social cost for any  $n \geq 5$ .*

*Proof.* Again, we consider the point set as in Figure 2. Let the expected distance from  $-1, 0$  and  $1$  to the closest facility be  $e_1, e_2$  and  $e_3$  respectively. Clearly, we have  $e_1 + e_2 + e_3 \geq 1$ . For any randomized strategy-proof mechanism with approximation ratio at most 2,  $e_2 \leq \frac{2}{n-2}$ . Without loss of generality, we assume  $e_3 \geq \frac{1}{2} - \frac{1}{n-2}$ .

Now we place the right node at  $1$  to a new position at  $1 + \alpha$  for some  $\alpha \in (0, 1/2)$ . Let  $e'_3$  be the expected distance from  $1 + \alpha$  to the nearest facility at the new configuration by the same strategy-proof mechanism. Because of strategy-proofness,  $e'_3 \geq \frac{1}{2} - \alpha - \frac{1}{n-2}$ . (The condition  $n \geq 5$  guarantees  $e'_3 \geq 0$  for the optimal  $\alpha$  chosen later.)

Let  $p(x)$  be the probability density function of the probability that the closest facility to the right node at  $1 + \alpha$  is at  $x$  in the new configuration. When  $x \leq -\frac{1}{n-2}$ , the closest facility is at *weighted distance* at least 1 to nodes at 0. When  $x \geq \frac{1}{n-2}$ , for any placement of the other facility, the sum of the *weighted distances* to the closest facility for the nodes at  $-1$  and  $0$  is at least 1. In these two cases, the weighted distance to nodes at  $-1$  and  $0$  is at least 1. Denote  $P = \int_{-\frac{1}{n-2}}^{\frac{1}{n-2}} p(x) dx$ . Therefore, the total cost of the mechanism in the new configuration is at least:

$$\text{cost} \geq (1 - P) \cdot 1 + e'_3 \geq 1 + \frac{1}{2} - \alpha - \frac{1}{n-2} - P.$$

On the other hand, consider the distance to the node at  $1 + \alpha$ . When the closest facility to  $1 + \alpha$  is  $x \in (-\frac{1}{n-2}, \frac{1}{n-2})$ , the total weighted distance from the nodes to the closest facilities is at least  $1 + \alpha$ . Therefore, we have

$$\text{cost} \geq (1 - P) \cdot 1 + P \cdot (1 + \alpha) = 1 + \alpha \cdot P.$$

The optimal ratio is achieved when  $P = \frac{1/2 - \alpha - 1/(n-2)}{1 + \alpha}$  and the approximation ratio is at least

$$1 + \frac{1}{2} - \alpha - \frac{1}{n-2} - \frac{1/2 - \alpha - 1/(n-2)}{1 + \alpha} \geq 1 + \frac{1}{2} - \frac{1}{n-2} - \frac{\alpha^2 + 1/2}{1 + \alpha}.$$

Define  $g(\alpha) = \frac{\alpha^2 + 1/2}{1 + \alpha}$ . The maximum ratio is achieved when  $g'(\alpha) = 0$  with  $\alpha = \frac{2-\sqrt{2}}{4}$ , and the approximation ratio is at least  $1 + \frac{\sqrt{2}-1}{12-2\sqrt{2}} - \frac{1}{n-2}$ .

Both lower bounds for deterministic and randomized strategy-proof mechanisms can be generalized to  $k$  facilities for  $k \geq 3$ . (Consider the configuration that two nodes on the two sides, and  $k - 1$  group of nodes in between. Each group of nodes (including the two singletons) are at unit distance away.) We have a direct corollary.

**Corollary 1 (Lower bound for the  $k$ -facility game).** *In the  $k$ -facility game for  $k \geq 2$ , any deterministic strategy-proof mechanism has an approximation ratio of at least  $2 - \frac{4}{m}$  for the social cost, where  $m = \lfloor \frac{n-2}{k-1} \rfloor$ . Any randomized strategy-proof mechanism for the  $k$ -facility game has an approximation ratio of at least  $1 + \frac{\sqrt{2}-1}{12-2\sqrt{2}} - \frac{1}{m} \geq 1.045 - \frac{1}{m}$ .*

### 3 Multiple Locations Per Agent

In this section, we study the case that each agent controls multiple locations. Assume agent  $i$  controls  $w_i$  locations, i.e.,  $\mathbf{x}_i = \{x_{i1}, x_{i2}, \dots, x_{iw_i}\}$ . A (deterministic) mechanism with one facility in the multiple locations setting is a function  $f : \mathcal{R}^{w_1} \times \dots \times \mathcal{R}^{w_n} \rightarrow \mathcal{R}$  for  $n$  agents. Then, for agent  $i$ , its cost is defined as  $\text{cost}(l, \mathbf{x}_i) = \sum_{j=1}^{w_i} |l - x_{ij}|$ , where  $l$  is the location of the facility. As before, we are interested in minimizing the social cost of the agents, i.e.,  $\sum_{i \in N} \sum_{j=1}^{w_i} |l - x_{ij}|$ , where  $N = \{1, 2, \dots, n\}$ .

We first give a tight analysis of a randomized strategy-proof mechanism proposed in [9]. This in particular confirms a conjecture of [9]. Then we prove the first approximation ratio lower bound of 1.33 for any randomized truthful mechanism. This lower bound even holds for the simplest case that there are only two player and each controls the same number of locations. As pointed out by [9], our result here can be directly applied in the incentive compatible regression learning setting of Dekel et al. [5].

#### 3.1 A Tight Analysis of a Randomized Mechanism

In [9], Procaccia and Tennenholtz proposed the following randomized mechanism in the setting of multiple locations:

*Randomized Median Mechanism:* Given  $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ , return  $\text{med}(\mathbf{x}_i)$  with probability  $w_i / (\sum_{j \in N} w_j)$ .

If  $w_i$  is even,  $\text{med}(\mathbf{x}_i)$  can either report the  $\frac{w_i}{2}$ -th location or  $\frac{w_i}{2} + 1$ -th location of  $\mathbf{x}_i$ . In [9], Procaccia and Tennenholtz gave a tight analysis for the case of two players ( $n = 2$ ), which has an approximation ratio of  $2 + \frac{|w_1 - w_2|}{w_1 + w_2}$ . They proposed as an open question for the bound in the general setting. In this section, we give a tight analysis of this randomized mechanism in the general setting, which in particular confirms the conjecture. Notice that  $2 + \frac{|w_1 - w_2|}{w_1 + w_2} = 3 - \frac{2 \min_{j \in N} w_j}{\sum_{j \in N} w_j}$ , when  $n = 2$ .

**Theorem 4.** *The Randomized Median Mechanism has an approximate ratio of  $3 - \frac{2 \min_{j \in N} w_j}{\sum_{j \in N} w_j}$  for social cost.*

*Proof.* If  $n = 1$ ,  $\text{med}(x_1)$  is the optimal solution. So the mechanism has an approximate ratio of  $3 - 2w_1/w_1 = 1$ . Now we consider the case for  $n \geq 2$ .

Without loss of generality, we can reorder the players so that  $\text{med}(\mathbf{x}_1) \leq \text{med}(\mathbf{x}_2) \leq \dots \leq \text{med}(\mathbf{x}_n)$ . Then it must be the case that  $\text{med}(\mathbf{x}_1) \leq \text{med}(\mathbf{x}) \leq \text{med}(\mathbf{x}_n)$ . The idea here is to construct a worst case instance for this mechanism and then analyze the approximate ratio for the worst case. Let  $i'$  be the largest  $i$  such that  $\text{med}(\mathbf{x}_i) \leq \text{med}(\mathbf{x})$ .

*Claim.* We can assume that the worst case satisfies the following properties: (1)  $w_i$  is even for all  $i \in N$ ; (2) for all  $i \leq i'$ ,  $\text{med}(\mathbf{x}_i)$  returns the  $\frac{w_i}{2}$ -th point of  $\mathbf{x}_i$ ; (3) and for all  $i > i'$ ,  $\text{med}(\mathbf{x}_i)$  returns the  $(\frac{w_i}{2} + 1)$ -th point of  $\mathbf{x}_i$ .

We justify the claim as follows: if some  $w_i$  is odd, we can add one more point for agent  $i$  at the global median  $\text{med}(\mathbf{x})$ , then the original  $\text{med}(\mathbf{x}_i)$  is still one of  $i$ -th two medians after adding the new point. We still return that value when we need to return  $\text{med}(\mathbf{x}_i)$ . After the modification, the expected cost can only increase while the optimal cost remain the same. So we can assume all  $w_i$  are even in a worst case. The properties (2) and (3) are obvious because returning the other point only improves the performance of the mechanism.

Now we assume that our instance satisfies all properties in Claim 1. By symmetry, we can further assume  $\sum_{i=1}^{i'} w_i \geq \sum_{i=i'+1}^n w_i$ . Let  $W = \sum_{j \in N} w_j$  and  $R(\text{med}(\mathbf{x}_i))$  be the rank of  $\text{med}(\mathbf{x}_i)$  in the whole set  $\mathbf{x}$ . Let  $X$  be the ordered global set of  $\mathbf{x}$  and  $X_i$  be the  $i$ th location in  $X$ . We perturb the points so that  $X_i$  and  $R(\text{med}(\mathbf{x}_i))$  are well defined. Then for all  $i \leq i'$ ,  $R(\text{med}(\mathbf{x}_i)) \geq \sum_{j=1}^i \frac{w_j}{2}$ ; for all  $i > i'$ ,  $R(\text{med}(\mathbf{x}_i)) \leq W - \sum_{j=i}^n \frac{w_j}{2}$ . The worst case happens when the above two sets of inequalities all reach equalities.

We further make the two sides more symmetric as follows. If  $w_1 > w_n$ , previously, the mechanism returns  $X_{\frac{w_1}{2}}$  with probability  $\frac{w_1}{W}$  and returns  $X_{W+1-\frac{w_n}{2}}$  with probability  $\frac{w_n}{W}$ . We modify the mechanism by returning  $X_{\frac{w_n}{2}}$  and  $X_{W+1-\frac{w_n}{2}}$  both with probability  $\frac{w_n}{W}$  and returning  $X_{\frac{w_1}{2}}$  with probability  $\frac{w_1-w_n}{W}$ . We continue this process and finally we can get the following mechanism. There are  $0 = k_0 < k_1 < k_2 < \dots < k_m$  and  $l \leq m$ . The mechanism returns  $X_{k_i}$  and  $X_{W+1-k_i}$  both with probability  $\frac{k_i-k_{i-1}}{k_m+k_l}$  if  $1 \leq i \leq l$ ; returns  $X_{k_i}$  with probability  $\frac{k_i-k_{i-1}}{k_m+k_l}$  if  $l < i \leq m$ . (The meaning of  $k_j$ s are roughlyly  $k_j = \sum_{i=1}^j \frac{w_i}{2}$ . However due to the symmetrization process described above, we also have  $k_j = \sum_{i=n}^{n-j} \frac{w_i}{2}$  for  $j \leq l$ .) We have  $k_1 = \frac{\min\{w_1, w_n\}}{2} \geq \frac{\min_{i \in N} w_i}{2}$  and  $k_m + k_l = W/2$ .

To simply the notation, we define  $\bar{i} = W + 1 - i$  and  $K = k_m + k_l$ . The optimal solution is  $\text{OPT} = \sum_{i=1}^{W/2} (X_{\bar{i}} - X_i) \geq s = \sum_{j=1}^m a_j$ , where  $a_j = \sum_{i=k_{j-1}+1}^{k_j} (X_{\bar{i}} - X_i)$ . Now we can compute the expected cost for this mechanism. For  $1 \leq i \leq l$ , we calculate the cost for  $X_{k_i}$  and  $X_{\bar{k}_i}$  together. They both have probability  $\frac{k_i-k_{i-1}}{k_m+k_l}$ .

The cost for  $X_{k_i}$  is  $\sum_{j=1}^i a_j + \sum_{j=k_i+1}^{W/2} (|X_j - X_{k_i}| + |X_{\bar{j}} - X_{k_i}|)$ . And we write that the cost for  $X_{\bar{k}_i}$  as  $\sum_{j=1}^i a_j + \sum_{j=k_i+1}^{W/2} (|X_j - X_{\bar{k}_i}| + |X_{\bar{j}} - X_{\bar{k}_i}|)$ . We combine the cost of  $X_{k_i}$  and  $X_{\bar{k}_i}$  together.



$$\begin{aligned}
 & 2 \sum_{j=1}^i a_j + 2 \sum_{j=k_i+1}^{W/2} |X_{\bar{k}_i} - X_{k_i}| \\
 &= 2 \sum_{j=1}^i a_j + 2(K - k_i)|X_{\bar{k}_i} - X_{k_i}| \leq 2 \sum_{j=1}^i a_j + 2 \frac{a_i(K - k_i)}{k_j - k_{j-1}}
 \end{aligned}$$

Now consider the case for  $l + 1 \leq i \leq m$ . Similarly, the cost of  $X_{k_i}$  is

$$\sum_{j=1}^i a_j + \sum_{j=k_i+1}^{W/2} (|X_j - X_{k_i}| + |X_j - X_{k_i}|) \leq \sum_{j=1}^i a_j + 2 \frac{a_i(K - k_i)}{k_j - k_{j-1}}$$

Therefore the expected cost of the mechanism is no more than

$$\begin{aligned}
 & \sum_{j=1}^l \frac{k_j - k_{j-1}}{K} (2 \sum_{i=1}^j a_i + 2 \frac{a_j(K - k_j)}{k_j - k_{j-1}}) + \sum_{j=l+1}^m \frac{k_j - k_{j-1}}{K} (\sum_{i=1}^j a_i + 2 \frac{a_j(K - k_j)}{k_j - k_{j-1}}) \\
 & \leq \frac{1}{K} (2k_l \sum_{i=1}^l a_i + (k_m - k_l) \sum_{i=1}^l a_i + k_m \sum_{i=l+1}^m a_i - 2k_1 \sum_{j=1}^m a_j) + 2s \\
 & = \frac{1}{K} (K \sum_{i=1}^l a_i + k_m \sum_{i=l+1}^m a_i - 2k_1 \sum_{j=1}^m a_j) + 2s \\
 & \leq (3 - \frac{2k_1}{K})s \leq (3 - \frac{2 \min_{j \in N} w_j}{\sum_{j \in N} w_j}) \text{OPT}
 \end{aligned}$$

The following corollary confirms a conjecture of [9] regarding the case where each agent controls the same number of locations.

**Corollary 2.** *If all the players control the same number of locations, the approximate ratio of Randomized Median Mechanism is  $3 - \frac{2}{n}$  for social cost.*

### 3.2 Lower Bounds for Randomized Strategy-Proof Mechanisms

In this section, we consider the lower bound of the approximation ratios for randomized strategy-proof mechanisms in the multiple locations setting. We first give a 1.2 lower bound of the approximation ratio, based on a very simple instance. Then we extend to a more complicated instance, which we derive a lower bound of 1.33 by solving a linear programming instance.

**Theorem 5.** *Any randomized strategy-proof mechanism of the one-facility game has an approximation ratio at least 1.2 for social in the setting that each agent controls multiple locations.*

*Proof.* We assume to the contrary that there exists one strategy-proof mechanism  $M$  which has an approximate ratio  $c < 1.2$ . Consider the following three instances:

**Instance 1:** First player has 2 points at 0 and 1 point at 1; second player has 3 points at 1.

**Instance 2:** First player has 3 points at 0; second player has 3 points at 1.

**Instance 3:** First player has 3 points at 0; second player has 1 point at 0 and 2 points at 1.

Let  $P_1, P_2$  and  $P_3$  be the distribution of the facility the mechanism  $M$  gives for these three instances respectively. For all  $x \in R$  and a distribution  $P$  on  $R$ , we use  $\text{cost}(P, x)$  to denote  $E_{y \sim P} |y - x|$ . Then we have (for all  $i = 1, 2, 3$ )

$$\text{cost}(P_i, 0) + \text{cost}(P_i, 1) \geq 1.$$

We use  $p_1(x), p_2(x)$  and  $p_3(x)$  to denote the probability density function of  $P_1, P_2$  and  $P_3$  respectively. Let

$$\forall i \in \{1, 2, 3\}, L_i = \int_{-\infty}^0 -xp_i(x)dx \quad \text{and} \quad R_i = \int_1^{+\infty} (x - 1)p_i(x)dx.$$

Now, we compute the cost of the players in each distribution. For the first player in Instance 1, its cost in distribution  $P_i$  is

$$\begin{aligned} 2\text{cost}(P_i, 0) + \text{cost}(P_i, 1) &= \text{cost}(P_i, 0) + (\text{cost}(P_i, 0) + \text{cost}(P_i, 1)) \\ &= \text{cost}(P_i, 0) + \int_{-\infty}^{+\infty} (|x| + |x - 1|)p_i(x)dx = \text{cost}(P_i, 0) + 2L_i + 2R_i + 1 \end{aligned}$$

Since  $L_1, R_1 \geq 0$ , It's easy to see

$$\text{cost}(P_1, 0) \leq \text{cost}(P_1, 0) + 2(L_1 + R_1) \leq \text{cost}(P_2, 0) + 2(L_2 + R_2), \quad (1)$$

where the second inequality is because of the strategy-proofness (of the first player in Instance 1). By symmetry, we also have

$$\text{cost}(P_3, 1) \leq \text{cost}(P_2, 1) + 2(L_2 + R_2). \quad (2)$$

Using similar calculation as above, we can get the expected cost of Instance 1 as follows.

$$2\text{cost}(P_1, 0) + 4\text{cost}(P_1, 1) = 2\text{cost}(P_1, 1) + 2(2L_1 + 1 + 2R_1) \geq 2\text{cost}(P_1, 1) + 2.$$

Since the optimal cost is 2 and the approximate ratio is less than 1.2, we know that  $\text{cost}(P_1, 1) + 2 < 2 \times 1.2 = 2.4$ . Therefore, we have  $\text{cost}(P_1, 1) < 0.2$  and hence  $\text{cost}(P_1, 0) > 0.8$ . Substituting the above inequality into (1), we get  $\text{cost}(P_2, 0) + 2(L_2 + R_2) > 0.8$ . Again by symmetry, we also have  $\text{cost}(P_2, 1) + 2(L_2 + R_2) > 0.8$ . Adding these two inequalities together, we have  $\text{cost}(P_2, 0) + \text{cost}(P_2, 1) + 4(L_2 + R_2) > 1.6$ . We also have  $\text{cost}(P_2, 0) + \text{cost}(P_2, 1) = 1 + 2(L_2 + R_2)$ . Substituting this, we get  $L_2 + R_2 > 0.1$ . On the other hand, note the approximate ratio condition of Instance 2 requires that  $1 + 2(L_2 + R_2) < 1.2$ . Thus we reach a contradiction.

To prove the lower bound of 1.33, we extend the above instances as follows. We employ  $2K + 1$  ( $K \geq 1$  is an integer) instances (for  $K = 1$ , this is exactly the same set of instances as above):

**Instance  $i$  ( $1 \leq i \leq K$ ):** First player has  $K + i$  points at 0 and  $K + 1 - i$  points at 1; second player has all  $2K + 1$  points at 1.

**Instance  $K + 1$ :** First player has all  $2K + 1$  points at 0; second player has all  $2K + 1$  points at 1.

**Instance  $i$  ( $K + 2 \leq i \leq 2K + 1$ ):** First player has all  $2K + 1$  points at 0; second player has  $i - K - 1$  points at 0 and  $3K + 2 - i$  points at 1.

Again, let  $P_i$  be the distribution of output of the mechanism on Instance  $i$ . Define the variables as  $X_i = \text{cost}(P_i, 0)$  and  $Y_i = \text{cost}(P_i, 1)$ . Then, the strategy-proofness among the instances can be listed as linear constrains. Assume the approximation ratio is  $\alpha$ . We want to compute the minimal ratio  $\alpha$  so that all constrains are satisfied. It is then straightforward to formulate the following linear programming problem.

$$\begin{aligned}
 & \text{Minimize: } \alpha \\
 & \text{Subject to:} \\
 & (K + i)X_i + (3K + 2 - i)Y_i \leq (K + i)\alpha, & 1 \leq i \leq K + 1 \\
 & (K + i)X_i + (3K + 2 - i)Y_i \leq (3K + 2 - i)\alpha, & K + 2 \leq i \leq 2K + 1 \\
 & (K + i)X_i + (K + 1 - i)Y_i \\
 & \quad \leq (K + i)X_{i+1} + (K + 1 - i)Y_{i+1}, & 1 \leq i \leq K \\
 & (i - K - 1)X_i + (3K + 2 - i)Y_i \\
 & \quad \leq (i - K - 1)X_{i-1} + (3K + 2 - i)Y_{i-1}, & K + 2 \leq i \leq 2K + 1 \\
 & X_i \geq 0, Y_i \geq 0, X_i + Y_i \geq 1, & 1 \leq i \leq 2K + 1
 \end{aligned}$$

First two sets of constrains come from the approximate ratio constrain. The next two sets of constrains are enforced by strategy-proofness. And the last two sets of constrains are boundary conditions.

Choosing  $K = 500$ , we solve this LP problem by computer and the optimal value is greater than 1.33. Therefore, if we set the approximation ratio to 1.33, there is no feasible solution for the linear programming which implies no feasible strategy-proof mechanism for the instances. So we have an approximation lower bound of 1.33.

**Theorem 6.** *Any randomized strategy-proof mechanism of the one-facility game has an approximation ratio at least 1.33 in the setting that each agent controls multiple locations.*

The numerical computation suggests that the optimal value for this LP problem is close to  $\frac{4}{3}$  when  $K$  is large. It would be interesting to give an analytical proof for a lower bound of  $\frac{4}{3}$ . We leave it as an open question.

## 4 Conclusion

In this paper, we study the strategy-proof mechanisms in facility games. We derive approximation bounds for such mechanisms for social cost both in the two-facility game and the multiple location setting. Our results improves several bounds previously studied [9]. We also obtain some new approximation ratio lower bounds.

There are still a lot of interesting open questions. For example, in the two-facility game, the deterministic mechanism has an approximation ratio of  $n - 2$  for social cost, while the lower bound is only 2. In randomized case, there is also a huge gap between  $n/2$  and 1.045.

## References

1. Barberà, S.: An introduction to strategy-proof social choice functions. *Social Choice and Welfare* 18(4), 619–653 (2001)
2. Barberà, S., Jackson, M.: A characterization of strategy-proof social choice functions for economies with pure public goods. *Social Choice and Welfare* 11(3), 241–252 (1994)
3. Black, D.: On the rationale of group decision-making. *The Journal of Political Economy*, 23–34 (1948)
4. Clarke, E.H.: Multipart Pricing of Public Goods. *Public Choice* 11, 17–33 (1971)
5. Dekel, O., Fischer, F., Procaccia, A.D.: Incentive compatible regression learning. In: *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 884–893. Society for Industrial and Applied Mathematics, Philadelphia (2008)
6. Gibbard, A.: Manipulation of voting schemes: A general result. *Econometrica: Journal of the Econometric Society*, 587–601 (1973)
7. Groves, T.: Incentives in Teams. *Econometrica* 41, 617–631 (1973)
8. Moulin, H.: On strategy-proofness and single peakedness. *Public Choice* 35(4), 437–455 (1980)
9. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. In: *Proceedings of the 10th ACM Conference on Electronic Commerce* (2009)
10. Satterthwaite, M.A.: Strategy-proofness and Arrows conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10(2), 187–217 (1975)
11. Schummer, J., Vohra, R.V.: *Mechanism design without money*, ch. 10. Cambridge University Press, Cambridge (2007)
12. Sprumont, Y.: The division problem with single-peaked preferences: a characterization of the uniform allocation rule. *Econometrica: Journal of the Econometric Society*, 509–519 (1991)
13. Vickrey, W.: Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 8–37 (1961)

# Degrees of Guaranteed Envy-Freeness in Finite Bounded Cake-Cutting Protocols

Claudia Lindner and Jörg Rothe

Institut für Informatik, Heinrich-Heine-Universität Düsseldorf, 40225 Düsseldorf, Germany

**Abstract.** Fair allocation of goods or resources among various agents is a central task in multiagent systems and other fields. The specific setting where just one divisible resource is to be divided fairly is commonly referred to as cake-cutting, and agents are called players in this setting. Cake-cutting protocols aim at dividing a cake and assigning the resulting portions to several players in a way that each of the players, according to his or her valuation of these portions, feels to have received a “fair” amount of the cake. An important notion of fairness is envy-freeness: No player wishes to switch the portion of the cake received with another player’s portion. Despite intense efforts in the past, it is still an open question whether there is a *finite bounded* envy-free cake-cutting protocol for an arbitrary number of players, and even for four players. In this paper, we introduce the notion of degree of guaranteed envy-freeness (DGEF, for short) as a measure of how good a cake-cutting protocol can approximate the ideal of envy-freeness while keeping the protocol finite bounded. We propose a new finite bounded proportional protocol for any number  $n \geq 3$  of players, and show that this protocol has a DGEF of  $1 + \lceil n^2/2 \rceil$ . This is the currently best DGEF among known finite bounded cake-cutting protocols for an arbitrary number of players. We will make the case that improving the DGEF even further is a tough challenge, and determine, for comparison, the DGEF of selected known finite bounded cake-cutting protocols, among which the Last Diminisher protocol turned out to have the best DGEF, namely,  $2 + n(n-1)/2$ . Thus, the Last Diminisher protocol has  $\lceil n/2 \rceil - 1$  fewer guaranteed envy-free-relations than our protocol.

**Keywords:** Cake-cutting protocol, fair division, multiagent resource allocation.

## 1 Introduction

Research in the area of cake-cutting started off in the 1940s with the pioneering work of Steinhaus [18] who, to the best of our knowledge, was the first to introduce the problem of fair division. Dividing a good (or a resource) fairly among several players such that each of them is satisfied with the portion received is of central importance in many fields. In the last 60 years this research area has developed vividly, spreading out into various directions and with applications in areas as diverse as economics, mathematics, computer science, and psychology. While some lines of this research seek to find reasonable interpretations of what “fairness” really stands for and how to measure it [9,7], others study proofs of existence or impossibility theorems regarding fair division (see,

e.g., [1]), or design new cake-cutting procedures [3,21,17] and, relatedly, analyze their complexity with respect to both upper and lower bounds [14,23,15]. As cake-cutting procedures involve several players, they are also referred to as “protocols.”

Cake-cutting protocols aim at achieving a *fair* division of an infinitely divisible resource among  $n$  players, who each may have different valuations of different parts of the resource. We focus on the notion of envy-freeness in finite bounded cake-cutting protocols. Cake-cutting protocols are either finite or continuous. While a *finite* protocol always provides a solution after only a finite number of decisions, a *continuous* protocol could potentially run forever. Among finite protocols, one can further distinguish between bounded and unbounded ones. A finite *bounded* cake-cutting protocol is present if we know in advance that a certain number of steps (that may depend on the number of players) will suffice to divide the resource fairly—independently of how the players may value distinct parts of the resource in a particular case and independently of the strategies chosen by the players. In contrast, in finite *unbounded* cake-cutting protocols, we cannot predict an upper bound on how many steps will be required to achieve the same goal. Aiming to apply cake-cutting procedures to real-world scenarios, it is important to develop fair *finite bounded* cake-cutting protocols. In this context, “fairness” is often interpreted as meaning “envy-freeness.” A division is *envy-free* if no player has an incentive to switch his or her portion with the portion any other player received.

For the division of a divisible good among  $n$  players, Steinhaus [19] proved that an envy-free division always exists. However, the current state of the art—after six decades of intense research—is that for arbitrary  $n$ , and even for  $n = 4$ , the development of *finite bounded* envy-free cake-cutting protocols still appears to be out of reach, and a big challenge for future research. For  $n > 3$  players, hardly any envy-free cake-cutting protocol is known, and the ones that are known are either finite unbounded or continuous (see, e.g., [3,16,5]).

Our goal in this paper is to look for compromises that can be made with respect to envy-freeness while keeping the protocol finite bounded: We propose an approach to evaluate finite bounded (yet possibly non-envy-free) cake-cutting protocols with respect to their *degree of guaranteed envy-freeness* (DGEF), a notion to be formally introduced in Section 3. Informally put, this notion provides a measure of how good such a protocol can approximate the (possibly for this particular protocol unreachable) ideal of envy-freeness in terms of the number of envy-free-relations that exist even in the worst case. To put the DGEF approach into practice, we present a new finite bounded proportional cake-cutting protocol with a significantly enhanced degree of guaranteed envy-freeness in Section 4, and discuss its significance in Section 5. A comparison to related work is drawn in the full version of this paper [13].

## 2 Preliminaries and Notation

Cake-cutting is about dividing a cake into portions that are assigned to the players such that each of them feels, according to his or her valuation of the portions, to have received a fair amount of the cake (where “cake” is a metaphor for the resource or the good to be divided). The cake is assumed to be infinitely divisible and can be divided

into arbitrary pieces without losing any of its value. Given  $n$  players, cake  $C$  is to be divided into  $n$  portions that are to be distributed among the players so as to satisfy each of them. A portion is not necessarily a single piece of cake; it can be a collection of disjoint, possibly noncontiguous pieces of  $C$ . The players may have different individual valuations of the single pieces of the cake: One player may prefer the pieces with the chocolate on top, whereas another player may prefer the pieces with the cherry topping.

More formally, cake  $C$  is represented by the unit interval  $[0, 1]$  of real numbers. By performing cuts,  $C$  is divided into  $m$  pieces  $c_k$ ,  $1 \leq k \leq m$ , which are represented by subintervals of  $[0, 1]$ . Each player  $p_i$ ,  $1 \leq i \leq n$ , assigns value  $v_i(c_k) = v_i(x_k, y_k)$  to piece  $c_k \subseteq C$ , where  $c_k$  is represented by the subinterval  $[x_k, y_k] \subseteq [0, 1]$  and  $p_i$ 's valuation function  $v_i$  maps subintervals of  $[0, 1]$  to real numbers in  $[0, 1]$ . We require each function  $v_i$  to satisfy the following properties:

1. *Normalization:*  $v_i(0, 1) = 1$ .
2. *Positivity:*<sup>1</sup> For all  $c_k \subseteq C$  with  $c_k \neq \emptyset$  we have  $v_i(c_k) > 0$ .
3. *Additivity:* For all  $c_k, c_\ell \subseteq C$  with  $c_k \cap c_\ell = \emptyset$  we have  $v_i(c_k) + v_i(c_\ell) = v_i(c_k \cup c_\ell)$ .
4. *Divisibility:*<sup>2</sup> For all  $c_k \subseteq C$  and for each  $\alpha$ ,  $0 \leq \alpha \leq 1$ , there exists some  $c_\ell \subseteq c_k$  such that  $v_i(c_\ell) = \alpha \cdot v_i(c_k)$ .

For simplicity, we write  $v_i(x_k, y_k)$  instead of  $v_i([x_k, y_k])$  for intervals  $[x_k, y_k] \subseteq [0, 1]$ . Due to Footnote<sup>2</sup> no ambiguity can arise. For each  $[x, y] \subseteq [0, 1]$ , define  $\|[x, y]\| = y - x$ .

We assume  $C$  to be heterogeneous (i.e., subintervals of  $[0, 1]$  having equal size can be valued differently by the same player). Moreover, distinct players may value one and the same piece of  $C$  differently, i.e., their individual valuation functions will in general be distinct. Every player knows only the value of (arbitrary) pieces of  $C$  corresponding to his or her own valuation function. Players do not have any knowledge about the valuation functions of other players.

A *division of  $C$*  is an assignment of disjoint and nonempty portions  $C_i \subseteq C$ , where  $C = \bigcup_{i=1}^n C_i = \bigcup_{k=1}^m c_k$ , to the players such that each player  $p_i$  receives a portion  $C_i \subseteq C$  consisting of at least one nonempty piece  $c_k \subseteq C$ . The goal is to assign all portions in as fair a way as possible. There are different interpretations, though, of what ‘‘fair’’ might mean. To distinguish between different degrees of fairness, among others, the following two notions have been introduced in the literature (see, e.g., [17]):

**Definition 1.** Let  $v_1, v_2, \dots, v_n$  be the valuation functions of the  $n$  players. A division of cake  $C = \bigcup_{i=1}^n C_i$ , where  $C_i$  is the  $i$ th player's portion, is said to be: (i) simple fair (a.k.a. proportional) if for each  $i$ ,  $1 \leq i \leq n$ ,  $v_i(C_i) \geq 1/n$ ; (ii) envy-free if for each  $i$  and  $j$ ,  $1 \leq i, j \leq n$ ,  $v_i(C_i) \geq v_i(C_j)$ .

A cake-cutting protocol describes an interactive procedure for obtaining a division of a given cake. A protocol is characterized by a set of rules and a set of strategies (see, e.g., [4]), which have to be followed by all players for them to be guaranteed a fair

<sup>1</sup> The literature is a bit ambiguous regarding this assumption. Some papers require the players' values for nonempty pieces of cake to be *nonnegative* (i.e.,  $v_i(c_k) \geq 0$ ) instead of positive.

<sup>2</sup> Divisibility implies that for each  $x \in [0, 1]$ ,  $v_i(x, x) = 0$ . That is, isolated points are valued 0, and open intervals have the same value as the corresponding closed intervals.

portion of the cake. The rules determine the course of action, such as a request to cut the cake, whereas the strategies define how to achieve a certain degree of fairness, e.g., by advising the players where to cut the cake.

### 3 Degrees of Guaranteed Envy-Freeness

The design of envy-free cake-cutting protocols for any number  $n$  of players seems to be quite a challenge. For  $n \leq 3$  players, several protocols that always provide envy-free divisions have been published, both finite (bounded and unbounded) and continuous ones [21,4,17]. However, to the best of our knowledge, up to date no finite bounded cake-cutting protocol for  $n > 3$  players is known to always provide an envy-free division. For practical purposes, it would be most desirable to have *finite bounded* cake-cutting protocols that always provide divisions as fair as possible. We propose an approach that weakens the concept of envy-freeness for the purpose of keeping the protocols finite bounded.

On the one hand, in this section we study known simple fair (i.e., proportional) cake-cutting protocols that are finite bounded, and determine their “degree of guaranteed envy-freeness” (see Definition 3). On the other hand, in Section 4 we propose a new finite bounded proportional cake-cutting protocol that—compared with the known protocols—has an enhanced DGEF.

When investigating the degree of envy-freeness of a cake-cutting protocol for  $n$  players, for each player  $p_i$ ,  $1 \leq i \leq n$ , the value of his or her portion needs to be compared to the values of the  $n - 1$  other portions (according to the measure of player  $p_i$ )<sup>3</sup>. Thus,  $n(n - 1)$  pairwise relations need to be investigated in order to determine the degree of envy-freeness of a cake-cutting protocol for  $n$  players. A player  $p_i$  envies another player  $p_j$ ,  $1 \leq i, j \leq n$ ,  $i \neq j$ , when  $p_i$  prefers player  $p_j$ 's portion to his or her own. If  $p_i$  envies  $p_j$ , we call the relation between these two players an *envy-relation*; otherwise, we call it an *envy-free-relation*.

**Definition 2.** Let cake  $C = \bigcup_{i=1}^n C_i$  be divided among all players in  $P = \{p_1, p_2, \dots, p_n\}$ , where  $v_i$  is  $p_i$ 's valuation function and  $C_i$  is  $p_i$ 's portion. Let  $p_i, p_j \in P$  be any two distinct players. An *envy-relation* occurs in this division if  $p_i$  envies  $p_j$  (denoted by  $p_i \Vdash p_j$ ), i.e., if  $v_i(C_i) < v_i(C_j)$ ; an *envy-free-relation* occurs if  $p_i$  does not envy  $p_j$  (denoted by  $p_i \not\vdash p_j$ ), i.e., if  $v_i(C_i) \geq v_i(C_j)$ .

We mention the following properties of envy-relations and envy-free-relations<sup>4</sup>. No player can envy him- or herself, i.e., envy-relations are irreflexive: The inequality

<sup>3</sup> We will use “valuation” and “measure” interchangeably.

<sup>4</sup> Various analogs of envy-relations and envy-free-relations have also been studied, from an economic perspective, in the different context of multiagent allocation of indivisible resources. Feldman and Weiman [10] consider “non-envy relations,” which are similar to our notion of envy-free-relations, and Chauduri [6] introduces “envy-relations.” Despite some similarities, their notions differ from ours, both in their properties and in the way properties holding for their and our notions are proven. For example, Chauduri [6] notes that mutual envy cannot occur in a market equilibrium, i.e., in this case his “envy-relations” are asymmetric, which is in sharp contrast to two-way envy being allowed for our notion.



$v_i(C_i) < v_i(C_i)$  never holds. Thus,  $v_i(C_i) \geq v_i(C_i)$  always holds. However, when counting envy-free-relations for a given division, we disregard the trivial envy-free-relations  $p_i \not\ll p_i$ ,  $1 \leq i \leq n$ . Neither envy-relations nor envy-free-relations need to be transitive. These observations imply that envy-relations and envy-free-relations are either one-way or two-way, i.e., it is possible that: (a) two players envy each other ( $p_i \ll p_j$  and  $p_j \ll p_i$ ), (b) neither of two players envies the other ( $p_i \not\ll p_j$  and  $p_j \not\ll p_i$ ), (c) one player envies another player but is not envied by this other player ( $p_i \ll p_j$  and  $p_j \not\ll p_i$ ).

Assuming all players to follow the rules and strategies, some cake-cutting protocols always guarantee an envy-free division (i.e., they always find an envy-free division of the cake), whereas others do not. Only protocols that *guarantee* an envy-free division in *every* case, even in the worst case (in terms of the players' valuation functions), are considered to be envy-free. An envy-free division may be obtained by coincidence, just because the players have matching valuation functions that avoid envy, and not because envy-freeness is enforced by the rules and strategies of the cake-cutting protocol used. In the worst case, however, when the players have totally nonconforming valuation functions, an envy-free division would not just happen by coincidence, but needs to be enforced by the rules and strategies of the protocol. An envy-free-relation is said to be *guaranteed* if it exists even in the worst case.

**Definition 3.** For  $n \geq 1$  players, the *degree of guaranteed envy-freeness* (DGEF, for short) of a given proportional [\[1\]](#) cake-cutting protocol is defined to be the maximum number of envy-free-relations that exist in every division obtained by this protocol (provided that all players follow the rules and strategies of the protocol), i.e., the DGEF (which is expressed as a function of  $n$ ) is the number of envy-free-relations that can be guaranteed even in the worst case.

This definition is based on the idea of weakening the notion of fairness in terms of envy-freeness in order to obtain cake-cutting protocols that are fair (though perhaps not envy-free) *and* finite bounded, where the fairness of a protocol is given by its degree of guaranteed envy-freeness. The higher the degree of guaranteed envy-freeness the fairer the protocol.

Proposition [\[2\]](#) gives an upper and a lower bound on the degree of guaranteed envy-freeness for proportional cake-cutting protocols. Its proof can be found in the full version of this paper [\[13\]](#).

**Proposition 1.** *Let  $d(n)$  be the degree of guaranteed envy-freeness of a proportional cake-cutting protocol for  $n \geq 2$  players. It holds that  $n \leq d(n) \leq n(n - 1)$ .*

An envy-free cake-cutting protocol for  $n$  players guarantees that no player  $p_i$  envies any other player  $p_j$ , i.e., the DGEF of an envy-free protocol equals  $n(n - 1)$ , the upper bound in Proposition [\[2\]](#).

The degree of fairness of a division obtained by applying a proportional cake-cutting protocol highly depends on the rules of this protocol. Specifying and committing to appropriate rules often increases the degree of guaranteed envy-freeness, whereas the

<sup>5</sup> We restrict the notion of DGEF to proportional protocols only, since otherwise the DGEF may overstate the actual level of fairness, e.g., if all the cake is given to a single player.

lack of such rules jeopardizes it in the sense that the number of guaranteed envy-free-relations may be limited to the worst-case minimum of  $n$  as stated in Proposition 1. In this context, “appropriate rules” are those that involve the players’ evaluations of other players’ pieces and portions that still are to be assigned. Concerning a particular piece of cake, involving the evaluation of as many players as possible in the allocation process helps to keep the number of envy-relations to be created low, since this allows to determine early on whether a planned allocation later may turn out to be disadvantageous—and thus allows to take adequate countermeasures. In contrast, omitting mutual evaluations means to forego additional knowledge that could turn out to be most valuable later on. For example, say player  $p_i$  is going to get assigned piece  $c_j$ . If the protocol asks all other players to evaluate piece  $c_j$  according to their measures, all envy-relations to be created by the assignment of piece  $c_j$  to player  $p_i$  can be identified before the actual assignment and thus countermeasures (such as trimming piece  $c_j$ ) can be undertaken. However, if the protocol requires no evaluations on behalf of the other players, such potential envy-relations cannot be identified early enough to prevent them from happening.

**Lemma 1.** *A proportional cake-cutting protocol with  $n \geq 2$  players has a DGEF of  $n$  (i.e., each player is guaranteed only one envy-free-relation) if the rules of the protocol require none of the players to value any of the other players’ portions.*

Our next result shows the DGEF for a number of well-known finite bounded *proportional* cake-cutting protocols. Note that these protocols have been developed with a focus on achieving proportionality, and not on maximizing the DGEF. The proofs of Lemma 1 and Theorem 1 can be found in the full version of this paper [13].

**Theorem 1.** *For  $n \geq 3$  players, the proportional protocols in Table 1 have a DGEF as shown in the same table.*

**Table 1.** DGEF of selected finite bounded cake-cutting protocols

Protocol	DGEF
Last Diminisher [18]	$2 + n(n-1)/2$
Lone Chooser [11]	$n$
Lone Divider [12]	$2n - 2$
Cut Your Own Piece (no strategy) [20]	$n$
Cut Your Own Piece (left-right strategy)	$2n - 2$
Divide and Conquer [8]	$n \cdot \lfloor \log n \rfloor + 2n - 2^{\lfloor \log n \rfloor + 1}$
Minimal-Envy Divide and Conquer [2]	$n \cdot \lfloor \log n \rfloor + 2n - 2^{\lfloor \log n \rfloor + 1}$
Recursive Divide and Choose [22]	$n$

## 4 A Protocol with an Enhanced DGEF

Figure 1 shows a finite bounded proportional cake-cutting protocol with an enhanced DGEF for  $n$  players, where  $n \geq 3$  is arbitrary. Unless specified otherwise, ties in this

protocol can be broken arbitrarily. Regarding the DGEF results in Table 1 the Last Diminisher protocol [8] shows the best results for  $n \geq 6$ , whereas the best results for  $n < 6$  are achieved by the Last Diminisher protocol as well as both the Divide and Conquer protocols [8, 2]. The protocol in Figure 1 improves upon these degrees of guaranteed envy-freeness for all  $n \geq 3$  and improves upon the DGEF of the Last Diminisher protocol by  $\lceil n/2 \rceil - 1$  additional guaranteed envy-free-relations.

Before presenting our protocol in detail, let us give an intuitive, high-level explanation. Both the protocol in Figure 1 and the Last Diminisher protocol are, more or less, based on the same idea of determining a piece of minimal size that is valued exactly  $1/n$  by one of the players (who is still in the game), which guarantees that all other players (who are still in the game) will not envy this player for receiving this particular piece. However, the protocol in Figure 1 works in a more parallel way, which makes its enhanced DGEF of  $\lceil n^2/2 \rceil + 1$  possible (see Theorem 3). To ensure that the parallelization indeed pays off in terms of increasing the degree of guaranteed envy-freeness, the “inner loop” (Steps 4.1 through 4.3) of the protocol is decisive. In addition, the protocol in Figure 1 provides a proportional division in a finite bounded number of steps (see Theorem 2), just as the Last Diminisher protocol.

*Remark 1.* Some remarks on the protocol in Figure 1 are in order:

1. From a very high-level perspective the procedure is as follows: The protocol runs over several rounds in each of which it is to find a player  $p_j$  who takes a portion from the left side of the cake, and to find a player  $p_k$  who takes a disjoint portion from the right side of the cake, such that none of the players still in the game envy  $p_j$  or  $p_k$  (at this, appropriate “inner-loop handling” might be necessary, see Figure 1 for details). Thereafter,  $p_j$  and  $p_k$  are to drop out with their portions, and a new round is started with the remaining cake (which is being renormalized, see remarks 3 and 4 below) and the remaining players. Finally, the Selfridge–Conway protocol is applied to the last three players in the game.<sup>6</sup>
2. The trivial cases  $n = 1$  (where one player receives all the cake) and  $n = 2$  (where each proportional division is always envy-free) are ignored.
3. Regarding  $n \geq 5$  players, if at any stage of our protocol the same player marks both the leftmost smallest piece and the rightmost smallest piece, the cake may be split up into two pieces and later on merged again. To simplify matters, in such a case the interval boundaries are adapted as well, which is expressed in Step 8 of Figure 1. Simply put, the two parts of the cake are set next to each other again to ensure a seamless transition. This can be done without any loss in value due to additivity of the players’ valuation functions.
4. In Steps 1 and 9.1, the value of subcake  $C' \subseteq C$  is normalized such that  $v_i(C') = 1$  for each player  $p_i$ ,  $1 \leq i \leq s$ , for the sake of convenience. In more detail, each player  $p_i$  values  $C'$  at least  $s/n$  of  $C$ , i.e.,  $v_i(C') \geq (s/n) \cdot v_i(C)$ . Thus, by receiving a proportional share (valued  $1/s$ ) of  $C'$  each player  $p_i$  is guaranteed at least a proportional share (valued  $1/n$ ) of  $C$ .

<sup>6</sup> This protocol has been developed by Banach and Knaster and was first presented in Steinhaus [18].

<sup>7</sup> This protocol is known to be a finite bounded envy-free cake-cutting protocol for  $n = 3$  players (see Stromquist [21]).

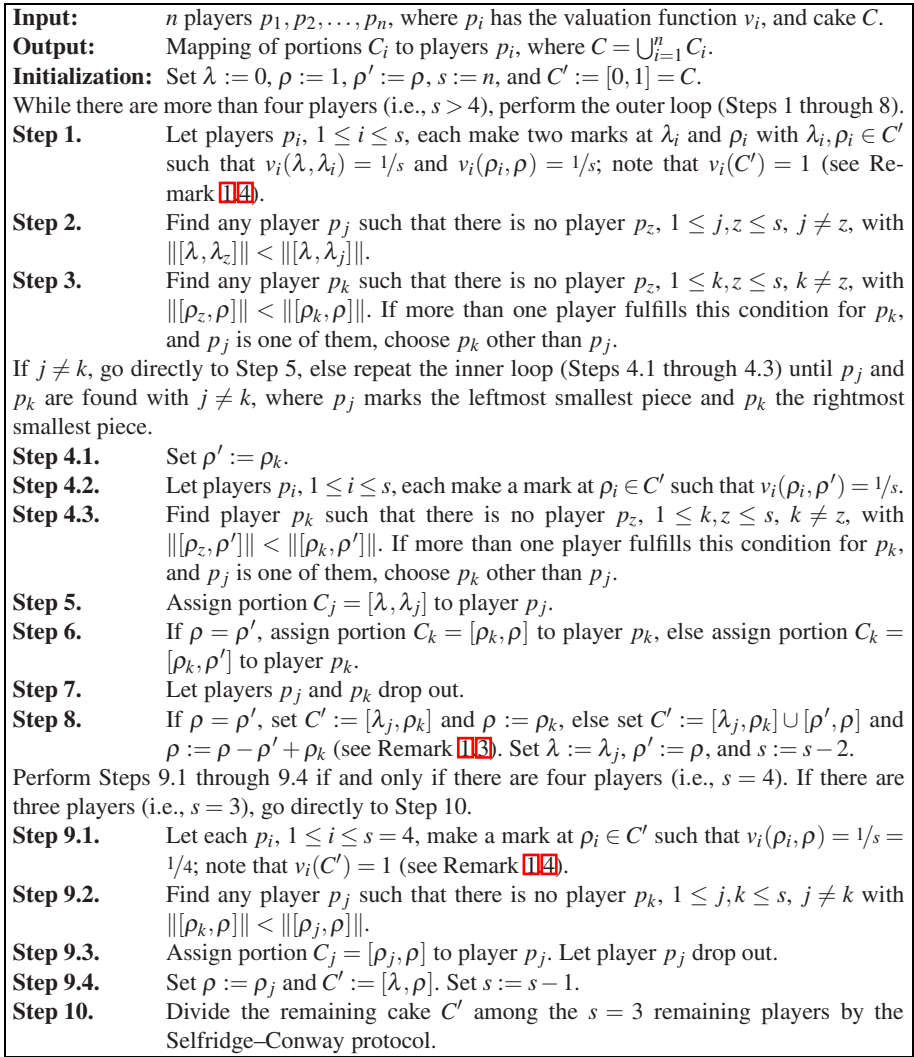


Fig. 1. A proportional protocol with an enhanced DGEF of  $\lceil n^2/2 \rceil + 1$  for  $n \geq 3$  players

**Theorem 2.** *The protocol in Figure 1 is finite bounded and proportional.*

The proofs of Theorems 2 and 3 are given in the full version of this paper [13], but we mention that the protocol is bounded by  $(7 \cdot \lceil (n-4)/2 \rceil) + (3 \cdot \sum_{i=1}^{\lceil (n-4)/2 \rceil} (n-2i)) + 4 + 9$  steps, which shows that it indeed is finite bounded. The proof of proportionality follows along the lines of the proof of Theorem 3 and in particular uses that each player is assigned a portion valued exactly  $1/s$  of a subcake that he or she values to be worth at least  $s/n$  of the given cake, according to his or her measure.

**Theorem 3.** For  $n \geq 5$  players, the cake-cutting protocol in Figure 1 has a DGEF of  $\lceil n^2/2 \rceil + 1$ .<sup>8</sup>

## 5 Discussion

It may be tempting to seek to decrease envy (and thus to increase the DGEF) via trading, aiming to get rid of potential circular envy-relations. Indeed, if the DGEF is *lower than*  $n(n-1)/2$ , the number of guaranteed envy-free-relations can be improved to this lower bound by resolving circular envy-relations (of which two-way envy-relations are a special case) by means of circular trades after the execution of the protocol.<sup>9</sup> Thus, in this case, involving subsequent trading actions adds on the number of guaranteed envy-free-relations. Furthermore, having  $n(n-1)/2$  guaranteed envy-free-relations after all circular envy-relations have been resolved, three more guaranteed envy-free-relations can be gained by applying an envy-free protocol (e.g., the Selfridge–Conway protocol) to the three most envied players, which yields to an overall lower bound of  $3 + n(n-1)/2$  guaranteed envy-free-relations. Note, though, that the DGEF is defined to make a statement on the performance of a particular protocol and not about all sorts of actions to be undertaken afterwards.

However, if the DGEF of a proportional cake-cutting protocol is  $n(n-1)/2$  or higher (such as the DGEF of the protocol presented in Figure 1) then circular envy-relations are not *guaranteed* to exist, and hence, in this case, trading has no impact on the number of guaranteed envy-free-relations.

Although the well-known protocols listed in Table 1 have not been developed with a focus on maximizing the DGEF,<sup>10</sup> linking their degrees of guaranteed envy-freeness to the lower bound provided by involving, e.g., the Selfridge–Conway protocol and guaranteed trading opportunities indicates that the development of cake-cutting protocols with a considerably higher DGEF or even with a DGEF close to the maximum of  $n(n-1)$  poses a true challenge. That is why we feel that the enhanced DGEF of the protocol presented in Figure 1 constitutes a significant improvement.

<sup>8</sup> Note that the same formula holds if  $n = 3$ , but for the special case of  $n = 4$  (see [13] for details) even one more envy-free-relation can be guaranteed (i.e., for  $n = 4$  players, the DGEF of the protocol in Figure 1 is  $(n^2/2) + 2$ ).

<sup>9</sup> To be specific here, all occurrences of “guaranteed envy-free-relations” in this and the next paragraph refer to those envy-free-relations that are guaranteed to exist after executing some cake-cutting protocol *and in addition, subsequently, performing trades that are guaranteed to be feasible*. This is in contrast with what we mean by this term anywhere else in the paper; “guaranteed envy-free-relations” usually refers to those envy-free-relations that are guaranteed to exist after executing the protocol only. As is common, we consider trading not to be part of a cake-cutting protocol, though it might be useful in certain cases (for example, Brams and Taylor mention that trading might be used “to obtain better allocations; however, this is not a procedure but an informal adjustment mechanism” [4] page 44]). In particular, the notion of DGEF refers to (proportional) cake-cutting protocols without additional trading.

<sup>10</sup> Quite remarkably, without any trading actions and without involving, e.g., the Selfridge–Conway protocol the Last Diminisher protocol achieves with its DGEF almost (being off only by one) the trading- and Selfridge–Conway-related bound of  $3 + n(n-1)/2$  mentioned above.

## 6 Conclusions

Finite bounded protocols that guarantee an envy-free division for  $n > 3$  players are still a mystery. However, finite bounded protocols are the ones we are looking for in terms of practical implementations. We propose to weaken the requirement of envy-freeness, while insisting on finite boundedness. To this end, we introduced the notion of degree of guaranteed envy-freeness for proportional cake-cutting protocols and determined the DGEF in existing finite bounded proportional cake-cutting protocols. We expect that the concept of DGEF is suitable to extend the scope for the development of new finite bounded cake-cutting protocols by allowing to approximate envy-freeness step by step. In this context, we proposed a new finite bounded proportional cake-cutting protocol, which provides a significantly enhanced DGEF compared with those in Table 11. In particular, our protocol has  $\lceil n/2 \rceil - 1$  more guaranteed envy-free-relations than the Last Diminisher protocol, which previously was the best finite bounded proportional protocol with respect to the DGEF. To achieve this significantly enhanced DGEF, our protocol makes use of parallelization with respect to the leftmost and the rightmost pieces. In this regard, adjusting the values of the pieces to be marked from  $1/n$  to  $1/s$  (with  $s$  players still in the game) and applying an appropriate inner-loop procedure is crucial to make the parallelization work.

**Acknowledgments.** We thank Ariel Procaccia for valuable comments and Jeff Rosenschein for hosting Jörg Rothe's visit to Jerusalem. This work was supported in part by DFG grants RO 1202/{11-1, 12-1}, the ESF's EUROCORES program LogICCC, and the Humboldt Foundation's TransCoop program.

## References

1. Akin, E.: Vilfredo Pareto cuts the cake. *Journal of Mathematical Economics* 24(1), 23–44 (1995)
2. Brams, S., Jones, M., Klamler, C.: Divide-and-Conquer: A proportional, minimal-envy cake-cutting procedure. In: Dagstuhl Seminar 07261. *Dagstuhl Seminar Proceedings* (2007)
3. Brams, S., Taylor, A.: An envy-free cake division protocol. *The American Mathematical Monthly* 102(1), 9–18 (1995)
4. Brams, S., Taylor, A.: *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, Cambridge (1996)
5. Brams, S., Taylor, A., Zwicker, W.: A moving-knife solution to the four-person envy-free cake-division problem. *Proceedings of the American Mathematical Society* 125(2), 547–554 (1997)
6. Chauduri, A.: Formal properties of interpersonal envy. *Theory and Decision* 18, 301–312 (1985)
7. Chauduri, A.: Some implications of an intensity measure of envy. *Social Choice and Welfare* 3(4), 255–270 (1986)
8. Even, S., Paz, A.: A note on cake cutting. *Discrete Applied Mathematics* 7, 285–296 (1984)
9. Feldman, A., Kirman, A.: Fairness and Envy. *The American Economic Review* 64(6), 995–1005 (1974)
10. Feldman, A., Weiman, D.: Envy, wealth, and class hierarchies. *Journal of Public Economics* 11(1), 81–91 (1979)

11. Fink, A.: A note on the fair division problem. *Mathematics Magazine* 37(5), 341–342 (1964)
12. Kuhn, H.: On games of fair division. In: Shubik, M. (ed.) *Essays in Mathematical Economics in Honor of Oskar Morgenstern*, Princeton University Press, Princeton (1967)
13. Lindner, C., Rothe, J.: Degrees of guaranteed envy-freeness in finite bounded cake-cutting protocols. Technical Report cs.GT/0902.0620, ACM Computing Research Repository (CoRR) (September 2009), <http://arxiv.org/abs/0902.0620>
14. Magdon-Ismail, M., Busch, C., Krishnamoorthy, M.: Cake-cutting is not a piece of cake. In: Alt, H., Habib, M. (eds.) *STACS 2003*. LNCS, vol. 2607, pp. 596–607. Springer, Heidelberg (2003)
15. Procaccia, A.: Thou shalt covet thy neighbor’s cake. In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI, July 2009*, pp. 239–244 (2009)
16. Robertson, J., Webb, W.: Near exact and envy free cake division. *Ars Combinatoria* 45, 97–108 (1997)
17. Robertson, J., Webb, W.: *Cake-Cutting Algorithms: Be Fair If You Can*. A K Peters, Stanford (1998)
18. Steinhaus, H.: The problem of fair division. *Econometrica* 16, 101–104 (1948)
19. Steinhaus, H.: Sur la division pragmatique. *Econometrica* 17, 315–319 (1949)
20. Steinhaus, H.: *Mathematical Snapshots*, 3rd edn. Oxford University Press, New York (1969)
21. Stromquist, W.: How to cut a cake fairly. *The American Mathematical Monthly* 87(8), 640–644 (1980)
22. Tasnádi, A.: A new proportional procedure for the  $n$ -person cake-cutting problem. *Economics Bulletin* 4(33), 1–3 (2003)
23. Woeginger, G., Sgall, J.: On the complexity of cake cutting. *Discrete Optimization* 4(2), 213–220 (2007)

# Approximate Pure Nash Equilibria via Lovász Local Lemma

Thành Nguyen\* and Éva Tardos\*\*

Cornell University

**Abstract.** In many types of games, mixed Nash equilibria is not a satisfying solution concept, as mixed actions are hard to interpret. However, pure Nash equilibria, which are more natural, may not exist in many games. In this paper we explore a class of graphical games, where each player has a set of possible decisions to make, and the decisions have bounded interaction with one another. In our class of games, we show that while pure Nash equilibria may not exist, there is always a pure approximate Nash equilibrium. We also show that such an approximate Nash equilibrium can be found in polynomial time. Our proof is based on the Lovász local lemma and Talagrand inequality, a proof technique that can be useful in showing similar existence results for pure (approximate) Nash equilibria also in other classes of games.

## 1 Introduction

In his Nobel prize winning work, John Nash proved the existence of mixed equilibria, which are now called Nash equilibria. A nice property of Nash equilibria is that they exist in (almost) any game. On the other hand, in many applications, for example in network design problem [3] and facility locations [20], agents need to make a decision that then will be visible to all other players, such as locating a facility (network router, a store, a server, etc.). In such games, there is no natural interpretation of a randomized action of the players. Before the decision is made, it makes sense to think of two locations as equally likely for being the selected choice. But once a decision is made, and assuming this decision is observable for all other players, these other players will react to the actual decision made, and the fact that *a priori* another decision was equally likely becomes irrelevant. This is especially true if, once the decision is made, it becomes hard to undo, such as locating a facility that requires significant investment cost. In such contexts pure Nash equilibria are much more natural, but unfortunately, they may not exist.

Our interest in this paper is to explore the existence of pure Nash equilibria (or approximate equilibria). Recent results on the hardness of computing Nash equilibria [8,5] inspired work on finding approximate equilibria, see for example [7]. In this paper, we attempt to start similar investigations for approximate pure Nash

---

\* Supported in part by ONR grant N00014-08-1-0031.

\*\* Supported in part by ITR grant CCR-0325453, CCR-0729006, and ONR grant N00014-08-1-0031.



equilibria. In a general model, we proved that there exists pure constant approximate Nash, and we can find it in polynomial time. In this paper, we also introduce a new technique for proving the existence of approximate Nash equilibria.

*Our model.* The class of games we consider is a variant of graphical games, where different players interact in limited ways. Graphical game is a general class of games that successfully captures and exploits the locality and sparsity of direct influences in games with large number of players. In a graphical game we are given an undirected graph, in which players are identified with vertices, and a player's payoff function is entirely determined by the action of the player and his/her neighbors. In this paper we consider a slightly different class with limited direct interactions. We will assume that each player makes many decisions, and assume that each decision is directly influenced by only a limited number of other decisions (while a player can influence a large set of other players through his many decisions).

More precisely, in the game we consider, each player  $i$  will be represented as a set of vertices  $S_i$ , one vertex for each of the decisions the player has to make. In each vertex, the player has two options, to play 0 or 1. We can think of the decision of playing 1 as representing the strategy to locate a service at this location, or start a business, develop a new product, etc. In this game, each player has a complex strategy set: the number of strategies is  $2^{|S_i|}$  corresponding to all possible subsets of  $S_i$ . We will describe the direct interactions of these decisions by an undirected graph  $G$  on all the vertices  $\cup_i S_i$ , with the set of edges connecting a vertex of a player to another player's vertex. We call this graph *interaction graph*.

We model the utility function of a player in two steps. First, on each possible location  $j$  (vertex in  $j \in S_i$ ) there is an *outcome function*, the outcome of the decision player  $i$  made on this vertex, which is a function mapping decisions made at all vertices connected to the vertex  $j$  to a multiple dimensional vector in  $\{0, 1\}^h$ . We'll think of the outcome as the level of success of the decision made in several criteria. For example, a product can be a success in one part of the market but not in the other, or building a factory at a location can create profit for the company, offer jobs for the locals but can cause environmental problems as well. A player in this model sometimes needs to make a decision that balances the trade-off between many factors. We model the payoff function of each player  $i$  as a MAXSAT formula of the whole outcome vector on  $S_i$ .

The proposed model of outcomes, and utility functions via a MAXSAT expression is very general, allowing us to model interests by players in many aspects. A term in the MAXSAT formula can model the success of each of the player's ventures, but other terms can express combinatorial goals. For example, in the case of companies developing products, one of the goals can be: at least one of the given products needs to be successful in a given market. The outcome function is an arbitrary function and therefore can capture some complex situations. For example, when developing a strategy whether to invest and build facility in various locations or to develop some new products, a company needs to take into account many factors: whether there will be too many competitors in the

location nearby, or too many similar products in the market. In some cases a product can only be successful if there are some other products available on the market. Our model captures some of these types of problems. A more precise description of the model is presented in Section 2.

*Our result and technique.* In this paper we will prove that if we assume the utility functions satisfy a Lipschitz condition, that is  $U_i(X_i) - U_i(X'_i) \leq \Delta$  whenever  $X_i, X'_i$  differ on one coordinate, then the game has a pure approximate Nash equilibrium, at which each player  $i$  has an payoff at least

$$OPT_i/2 - O\left(\Delta(\sqrt{OPT_i} + \Delta)\sqrt{\log n + \log d}\right),$$

where  $OPT_i$  is the optimal payoff he can get by a deviation,  $d$  is the maximum degree in the interaction graph  $G$  and  $n$  is the maximum size of a player’s action set  $S_i$ , that is  $n = \max_i |S_i|$ . Furthermore, such a configuration can be found in polynomial time.

Note that when  $OPT_i \gg \Delta, \log n, \log d$  then this is a constant approximate Nash. Furthermore, in the case of using MAXSAT to describe user utility functions, finding a better than constant approximate Nash is unlikely, as finding a better than constant approximation for even a single player’s optimization problem is NP hard [12].

In many cases, the assumption that  $OPT_i$  is relatively larger than  $\log n$  and  $\log d$  is reasonable, as the optimal solution of a MAXSAT formula is at least a constant times the number of its clauses. In our model if the number of clauses is much larger than  $\log n$  and the outcome functions are not all “constant”, that is, for a player by changing his decision on a vertex from 0 to 1 or from 1 to 0, the outcome of this decision also changes, then  $OPT_i$  satisfies the condition above.

At the heart of our proof is the use of Lovász local lemma and Talagrand’s inequality. The Lipschitz condition is used extensively for a concentration bound using Talagrand’s inequality. Lovász local lemma is used to prove the guarantee that every player has a payoff near the optimal. These techniques are nontrivial and general, and we hope that they will be useful in proving similar results about approximate Nash equilibria in many other settings.

*A bad example.* To illustrate the difficulties and some of the ideas in our techniques, let us start with the following example:




Fig. 1. An example

There are two players in the game. The first player plays on  $n$  vertices  $u_1, u_2, \dots, u_n$  and the second one plays on  $v_1, v_2, \dots, v_n$ . Let  $o(u_i)$  and  $o(v_i) \in \{0, 1\}$  be the outcome on vertex  $u_i$  and  $v_i$  respectively. The utility function of player 1 is a SAT formula consisting of the clauses:  $(o(u_1) \vee o(u_2)), (o(u_1) \vee o(u_3)), \dots, (o(u_1) \vee o(u_n))$ . Similarly, the utility function of player 2 is a SAT formula consisting of the clauses:  $(o(v_1) \vee o(v_2)), (o(v_1) \vee o(v_3)), \dots, (o(v_1) \vee o(v_n))$ . The underlying graph is a matching  $\{u_i v_i\}$ . Now the outcome functions are defined as follows: On all the vertices  $u_i, v_i$ , where  $i \geq 2$ , the outcome function is 0 no matter what the players' strategies are. On  $v_1$  and  $u_1$ , the outcome function is the same as the payoff function of the “matching penny” game described by the following table:

	0	1
0	(0, 1)	(1, 0)
1	(1, 0)	(0, 1)

It is not hard to see that it is a mixed strategy Nash if player 1 plays 0 or 1 with  $1/2$  probability on  $u_1$  and player 2 plays 0 or 1 with  $1/2$  probability on  $v_1$ . Given any deterministic configuration of the game, there is always a player that can improve his payoff from 0 to  $n - 1$  by changing his strategy on his first vertex. Thus, the game not only has no pure Nash equilibria but also has no “reasonable” approximate pure Nash equilibria. It turns out that the main obstacle in this example is the property that in the utility functions there is a variable such that by changing its value the utility function changes rapidly.

*Related works.* The complexity of mixed Nash equilibria is studied in [8,5]. In many applications, such as most of the network formation games, pure Nash Equilibria are usually considered a more realistic model of rationality. See the [2,3,20] for detail.

Potential games [17] is essentially the only class of games that is known to have pure Nash equilibria. The complexity of finding a pure Nash in potential games is proved to be PLS-complete [11]. Inspired by these hardness results, many researchers have been investigating pure approximate Nash equilibria in various games [2,18].

Graphical games were introduced in [14], and have been extensively studied. (See the survey [15].) The complexity of finding pure Nash equilibria in graphical games were studied in [6,9]. Most of the results concerning pure NE in graphical games are either negative or for a small class of graphs such as trees or graphs with bounded tree-width. Our model is a variant of graphical games. Here, instead of representing each player as a vertex in the graph, we consider each player as a set of vertices and thus, it models more complex strategy sets for players. In this paper, we provide a positive result for the existence of a pure approximate Nash in general graphs that can be found in polynomial time.

Lovász local lemma is proved in [10]. Lovász local lemma is a powerful tool in probability and has a wide range of applications, see the book of [1] for details and references. To the best of our knowledge, our paper is the first application of the local lemma in the area of algorithmic game theory.

*Structure of the paper.* In the next section, we will define our model more precisely. In Section 3 we will prove the existence of an approximate pure NE. Using an algorithmic version of the local lemma, we also obtain an algorithm to find such a solution.

## 2 The Model and Notations


Consider a game consisting  $N$  players, each player  $i$  has a set of possible *vertices*  $S_i$ , one vertex for each of the decisions that the player has to make. On each of his/her vertex, player  $i$  can choose one of two strategies 0 or 1. We call it the *decision* of player  $i$  on the vertex. In this game, each player  $i$  can have up to  $2^{|S_i|}$  strategies corresponding to all possible subsets of  $S_i$ .

We assume that the size of  $S_i$  is at most  $n$  for every  $i$ , and all the set  $S_i$  are disjoint. The latter assumption does not affect the generality of our model, as each player can have his/her own copy of a vertex.

Similar to graphical games [14], we model the direct interactions of the players' decisions by an undirected graph  $G$  on the set of all the vertices  $\cup_i S_i$ . The edges of  $G$  connects a vertex of a player to another player's vertex. We call  $G$  the *interaction graph*. See Figure 2 for an example. In the rest of the paper, we use  $d$  as the maximum degree of the interaction graph.

The utility function of a player is modeled in two steps. First, at each vertex  $j \in S_i$ , there is an *outcome function*. The outcome function on a vertex  $j$  is a function mapping the decisions of the players on  $j$  and neighbors of  $j$  to a finite set of outcomes. In this paper, we consider the case where each outcome is a multiple dimensional vector in  $\mathcal{O} = \{0, 1\}^h$  for a constant  $h$ . For a vertex  $j$ , we denote  $\Gamma(j)$  as the set of  $j$ ' neighbors, and  $o_j$  the outcome function on  $j$ , we have:

$$o_j : \{0, 1\}^{|\Gamma(j) \cup j|} \rightarrow \mathcal{O} = \{0, 1\}^h$$



**Fig. 2.** An example of our model

Now, given a configuration of the game, each player has an outcome on each of his vertices. The utility of player  $i$  denoted by  $U_i$  is a function mapping the outcome vector on  $S_i$  to a non negative number.

$$U_i : \{\mathcal{O}\}^{|S_i|} \rightarrow \mathbb{R}^+$$

One way to think about  $U_i$  is that it is function on  $m_i = h \cdot |S_i|$  boolean variables measuring the success of the decisions made in many aspects. In many cases, there are trade-offs between these decisions and some of the goals of a player is a combination of several outcomes. We model  $U_i$  as a *MAXSAT function*, that is for each  $U_i$  there is a Normal Conjunctive Formula (NCF)  $\Phi$  on boolean variables  $x_1, \dots, x_{m_i}$  such that the value of  $U_i(x_1, \dots, x_{m_i})$  is the number of satisfying clauses in  $\Phi$ . Note that our results can be naturally extended to a weighted version of MAXSAT.

As we have seen in the counter example in the introduction, we will need a *Lipschitz condition* for the utility functions. This condition assumes that by changing the decision at one vertex, the player’s utility does not change very much. More precisely, the utility function  $U_i$ , described above as a composition of a MAXSAT and several outcome functions, can be considered as a function of the player  $i$ ’s strategy  $X_i \in \{0, 1\}^{|S_i|}$  and all other players’  $X_{-i} \in \{0, 1\}^{|\cup_{j \neq i} S_j|}$ . We assume that  $\Delta$  is the Lipschitz constant for the class of utilities we consider. That is, for every  $i$ ,  $|U_i(X_i, X_{-i}) - U_i(X'_i, X_{-i})| \leq \Delta$  whenever  $X_i, X'_i$  differ in only one coordinate.

In the rest of the paper, we call the class of games defined above *MAXSAT games* with the *Lipschitz constant*  $\Delta$ .

### 3 Existence of Approximate Pure Nash

In this section we will prove the following result about the existence of a pure approximate Nash equilibrium for MAXSAT games.

**Theorem 1.** *In the MAX SAT game, there exists an approximate Nash, where each player  $i$  obtains a payoff at least*

$$OPT_i/2 - O\left(\Delta(\sqrt{OPT_i} + \Delta)\sqrt{\log n + \log d}\right),$$

where  $OPT_i$  is the optimum that player  $i$  can achieve assuming that other players do not change their strategies, and  $\Delta$  is the Lipschitz constant of the players’ utilities.


We now give some intuition before proving the theorem. Consider an arbitrary configuration of the game, and a player  $i$ . If player  $i$  assumes that other players do not change their strategies, then on a vertex  $j \in S_i$ , the outcome function will be a function mapping his decision on  $j$  to a vector in  $\{0, 1\}^h$ . Let  $x_j$  be the player’s decision on vertex  $j$ . The outcome  $o_j$  will be a vector whose coordinates can be a constant (independent of  $x_j$ ) or  $x_j$  or  $\neg x_j$ . Because the player  $i$ ’s utility can be expressed as a MAXSAT function on the outcome vector, it is also a MAXSAT function on  $\{x_j, j \in S_i\}$ . Therefore, if player  $i$  tries to find a strategy to maximize his own utility, he needs to solve a MAXSAT problem. It is, however, known that by assigning each variable to 0 and 1 with probability 1/2, we get at least a 2-approximate solution *in expectation*. We note that in the worst case, this simple algorithm gives the best possible approximation on some instances of MAXSAT, as shown by [12].

With this intuition, let us consider a simple strategy where each player’s decision on a location is to pick either 0 or 1 randomly with probability  $1/2$ . This strategy is an 2-approximate *mixed strategy Nash*. However, by selecting the strategies randomly, players might end up in a situation when they can deviate and significantly increase their payoff. Because the number of players  $N$  can be arbitrarily large, independent of  $n$  and  $d$ , the number of such players can also be arbitrarily large. Our idea is to bound the probability that such event happens and then use Lovász local lemma to prove the existence of approximate Nash equilibria. Lovász local lemma, proved in [10], can be stated as follows:

**Lemma 1 (Lovász).** *Let  $A_1, A_2, \dots, A_N$  be a series of events such that each event occurs with probability at most  $p$  and such that each event is independent of all the other events except for at most  $D$  of them. If  $e \cdot p \cdot (D + 1) < 1$  (where  $e = 2.718\dots$ ), then there is a nonzero probability that none of the events occur.  $\square$*

In order to use Lovász local lemma, we will define  $A_i$  as the event that player  $i$  can “significantly” improve his payoff by deviating. Thus, to prove that the probability that none of  $A_i$  occur is positive, we will need to bound the probability of  $A_i$  and to show that each  $A_i$  is independent of all but few other  $A_j$ s. We call the number of such events the *dependence number* of  $A_i$ . We first give a bound on the dependence numbers.

**Lemma 2.** *For every  $i$ , the dependence number of  $A_i$  is at most  $2nd^2$ .*



**Fig. 3.**  $A_1$  and  $A_3$  are independent,  $A_1$  and  $A_2$  are not independent

*Proof.*  $A_i$  is the event that a player  $i$  can find a “significantly” better move. This event depends on the decision of player  $i$  on his vertices and the decision of all other players on the neighboring vertices. Consider two events  $A_i$  and  $A_j$ . If each path connecting two vertices of player  $i$  and player  $j$  has a distance at least 3, then the set of neighboring vertices of player  $i$  and player  $j$  are disjoint. Hence  $A_i$  and  $A_j$  are independent of each other. See Figure 3 for an example. The dependence number of  $A_i$ , therefore is bounded by the number of  $A_j$  that has a path of length at most 2 connecting from a vertex of  $j$  to a vertex of  $i$ . Since each vertex can belong to at most one player, the number of such  $A_j$ ’s that might not be independent of  $A_i$  is at most the number of vertices having

a path of length 1 or 2 to any vertex of player  $i$ . Now the maximum degree of the interaction graph is bounded by  $d$ . Thus the number of vertices within a distance 2 from player  $i$ 's vertices is at most  $n \cdot d + n \cdot d \cdot d \leq 2nd^2$ .  $\square$

We now give a simple way to define and prove a bound on  $Pr(A_i)$ . This approach only gives a weaker result than Theorem [1](#). We then use a more powerful technique to prove our main theorem. The correct definition of  $A_i$  and the formal proof of Theorem [1](#) will be given at the end of this section.

For a player  $i$ , assume that other players do not change their strategies, the utility  $U_i$  is a function on the strategy set  $\{0, 1\}^{|S_i|}$ . As discussed at the beginning of this section, a random strategy gives at least a 2- approximate solution in expectation. Because  $U_i$  satisfies a Lipschitz condition, one can use a concentration inequality to bound the probability that the value of  $U_i$  is much smaller than the expected value. In particular we can use the following Hoeffding-Azuma inequality [\[3,4\]](#):

**Lemma 3 (Hoeffding-Azuma).** *Let  $U : \{0, 1\}^n \rightarrow R$  be a function such that  $|U(x) - U(x')| \leq \Delta$  whenever  $x, x'$  differ in only one coordinate, then if  $x_1, \dots, x_n$  are independent boolean variables, we have :  $Pr(|U(x_1, \dots, x_n) - E(U(x_1, \dots, x_n))| \geq \lambda \Delta \sqrt{n}) \leq 2e^{-\frac{\lambda^2}{2}}$ .  $\square$*

If we define  $A_i$  to be the event that:

$$U_i(x_1, \dots, x_{|S_i|}) - E(U_i(x_1, \dots, x_{|S_i|})) \geq \lambda \Delta \sqrt{|S_i|}$$

with a  $\lambda$  chosen such that

$$(2e^{-\frac{\lambda^2}{2}})e(2d^2n + 1) \leq 1$$

then we can apply the local lemma to prove that there the game has a configuration where each player  $i$  get at least  $\frac{1}{2}OPT_i - \Delta \lambda \sqrt{|S_i|}$ . Now:

$$2e^{-\frac{\lambda^2}{2}}e(2d^2n + 1) \leq 1 \Leftrightarrow 2e^{\frac{\lambda^2}{2}} \geq e(2d^2n + 1)$$

One can choose

$$\lambda = \sqrt{2 \log(2e(2d^2n + 1))} = O(\sqrt{\log n + \log d}),$$

and obtain the following result:

*Claim.* There exists a configuration such that each player obtains a payoff at least  $\frac{1}{2}OPT_i - O(\Delta \sqrt{|S_i|}(\log n + \log d))$ , where  $OPT_i$  is the optimum that that player can obtain assuming other players do not change their strategies.  $\square$

The additive error in the result above is of order  $O(\sqrt{|S_i|}(\log n + \log d))$ , thus when  $OPT_i$  is relatively small compared with  $\sqrt{|S_i|}$  the result is rather weak. To prove the stronger result in Theorem [1](#), we will apply a stronger concentration

technique developed by Talagrand [19]. We start by describing a class of functions that can be used in this approach.

**Definition 1.** We call a non negative function  $f$  defined on a set  $\Omega$  of  $n$  dimensional vectors a  $c$ -configuration function if it has the following property: for each  $x \in \Omega$  there is a non-negative unit  $n$ -dimensional vector  $\alpha$  such that for each  $y \in \Omega$  we have:

$$f(y) > f(x) - \sqrt{cf(x)}d_\alpha(x, y),$$

where

$$d_\alpha(x, y) = \sum_{i \in \{1, \dots, n\}; x_i \neq y_i} \alpha_i.$$

**Lemma 4 (Talagrand).** Let  $f$  be a  $c$ -configuration function,  $X$  be a random variable taking each coordinate independently from any distribution, and let  $m$  be a median for  $f(X)$ , that is,  $m : Pr(f(X) \leq m) = \frac{1}{2}$ . Then for any  $t > 0$

$$Pr(f(X) \leq m - t) \leq 2e^{-t^2/4cm}.$$

$$Pr(f(X) \geq m + t) \leq 2e^{-t^2/4c(m+t)}.$$

$$Pr(|f(X) - m| \geq t) \leq 2e^{-t^2/4cm} + 2e^{-t^2/4c(m+t)}.$$

□

The Talagrand’s inequality exploits the structure of the function  $f$ . The concentration error can be given as a function of the value of the median instead of the number of the variables. Therefore, Talagrand’s inequality gives a stronger result when the median is relatively small.

We will show later that the class of utility functions considered in this paper are  $\Delta^2$ -configuration functions, where  $\Delta$  is the Lipschitz constant of the utilities. We will then apply the Talagrand’s inequality. A technical problem here is that, the concentration is however bounded around the *median* of the variable. Some technical work will be needed to express the concentration inequality around the expected value. We will use the following result instead of Lemma 4:

**Lemma 5.** Let  $f$  be a  $c$ -configuration function, and let  $\mu$  be the expected value for  $f(X)$ , where  $X$  is a random variable taking each coordinate independently from any distribution, then for all  $\lambda > 10$

$$Pr(f(X) < \mu - 60\lambda(\sqrt{c\mu} + c)) \leq 2e^{-\lambda^2}$$

*Proof.* Before proving the lemma, let us remark that all the constants are chosen in the calculation for convenience. We did not attempt to optimize them. Now, recall that  $\mu$  and  $m$  is the expected value and median of  $f(X)$ . We have:

$$|\mu - m| = |E(f(X) - m)| \leq E(|f(X) - m|) = \int_0^\infty Pr(|f(X) - m| > t) dt$$



Because of the Lemma [4](#)

$$\begin{aligned} \int_0^\infty Pr(|f(X) - m| > t) dt &\leq 2 \int_0^\infty e^{-t^2/4cm} dt + 2 \int_0^\infty e^{-t^2/4c(m+t)} dt \\ &= 2\sqrt{\pi cm} + 2 \int_0^\infty e^{-t^2/4c(m+t)} dt \\ &\leq 2\sqrt{\pi cm} + 2 \int_0^m e^{-t^2/4c(m+m)} dt + 2 \int_m^\infty e^{-t^2/4c(t+t)} dt \\ &= 2\sqrt{\pi cm} + 2 \int_0^m e^{-t^2/8cm} dt + 2 \int_m^\infty e^{-t^2/8ct} dt \\ &< 2\sqrt{\pi cm} + 2\sqrt{2\pi cm} + 16c \leq 10\sqrt{cm} + 16c \end{aligned}$$

Thus we have

$$|\mu - m| < 10\sqrt{cm} + 16c \tag{1}$$

From here it is not hard to see that

$$15(\sqrt{\mu} + \sqrt{c}) \geq \sqrt{m} \tag{2}$$

Now, applying the first inequality in Lemma [4](#) for  $t = 2\lambda\sqrt{cm}$ , we have:

$$2e^{-\lambda^2} \geq Pr(f(X) < m - 2\lambda\sqrt{cm}) \tag{3}$$

Using the fact that  $m \geq \mu - |\mu - m|$ , we have:

$$Pr(f(X) < m - 2\lambda\sqrt{cm}) \geq Pr(f(X) < \mu - |\mu - m| - 2\lambda\sqrt{cm}) \tag{4}$$

Now, due to [\(1\)](#)  $|\mu - m| \leq 10\sqrt{c\pi m} + 16c$ , we have:

$$Pr(f(X) < \mu - |\mu - m| - 2\lambda\sqrt{cm}) \geq Pr(f(X) < \mu - 10\sqrt{cm} - 16c - 2\lambda\sqrt{cm}) \tag{5}$$

Combining [\(3\)](#), [\(4\)](#) and [\(5\)](#) we obtain

$$\begin{aligned} 2e^{-\lambda^2} &\geq Pr(f(X) < \mu - 10\sqrt{cm} - 16c - 2\lambda\sqrt{cm}) \\ &\Leftrightarrow 2e^{-\lambda^2} \geq Pr(f(X) < \mu - \sqrt{cm}(10 + 2\lambda) - 16c) \end{aligned}$$

Using the fact from [\(2\)](#) that  $\sqrt{m} < 15(\sqrt{\mu} + \sqrt{c})$  we obtain:

$$2e^{-\lambda^2} \geq Pr(f(X) < \mu - 15\sqrt{c}(\sqrt{\mu} + \sqrt{c})(10 + 2\lambda) - 16c)$$

Therefore if  $\lambda$  big enough ( $> 10$ ), we can simplify the last expression to get:

$$2e^{-\lambda^2} \geq Pr(f(X) < \mu - 60\lambda(\sqrt{c\mu} + c))$$

This is what we need to prove. □

We are now ready to prove the main theorem:

*Proof (Proof of Theorem 7)*

Given a MAXSAT utility function  $U$  with the Lipschitz constant  $\Delta$ , we first show that  $U$  is a  $\Delta^2$ -configuration function. And hence we can apply Lemma 5. Given a configuration of the game, let  $X$  be the strategy of a player, and  $k$  be the value of his payoff. In other words, by playing  $X$ , the player can satisfy  $k$  clauses in the corresponding Normal conjunctive formula. From each satisfying clause, pick a variable. It is possible that we pick the same variable for different clauses, thus the number of variables picked is  $k' \leq k$ . Each of these variables corresponds to a decision made on a vertex. We can consider these decisions as a set of “witnesses”. The reason is that if the player makes the same decisions as these witnesses, he is guaranteed to obtain a payoff of at least  $k$ . Furthermore, because of the Lipschitz condition, if at most  $l$  witness-decisions are made differently, then the player obtains a payoff of at least  $k - \Delta l$ .

Now let  $\alpha \in \mathbb{R}^n$  be a non negative vector that takes  $\frac{1}{\sqrt{k'}}$  on the coordinates of the witnesses and 0 elsewhere. Let  $Y$  be an alternative strategy such that  $d_\alpha(X, Y) = \frac{l}{\sqrt{k'}}$ . This means that  $X$  and  $Y$  differs on  $l$  witness. Thus the payoff of the player playing  $Y$  is at least  $k - \Delta l$ . This shows that:

$$U(Y) \geq U(X) - \Delta\sqrt{k'}d_\alpha(X, Y) \geq U(X) - \sqrt{\Delta^2 U(X)}d_\alpha(X, Y).$$

The last inequality is from the fact that  $U(X) = k \geq k'$ . This proves that the utility function is a  $\Delta^2$ -configuration function.

Let  $\mu = E(U(X))$  and  $\lambda > 10$  be a parameter, we now apply Lemma 5.

$$Pr(U(X) < \mu - 60\lambda(\Delta\sqrt{\mu} + \Delta^2)) \leq 2e^{-\lambda^2}$$

We need to choose  $\lambda$  such that  $e \cdot 2e^{-\lambda^2} \cdot (2nd^2 + 1) < 1$ . Again without optimizing the constant, let us choose  $\lambda = 2\sqrt{\log n + \log d} + 10$ . And we have:

$$Pr\left(U(X) \leq \mu - 60(2\sqrt{\log n + \log d} + 10)(\Delta\sqrt{\mu} + \Delta^2)\right) \leq \frac{1}{e(2nd^2 + 1)} \quad (6)$$

Note that  $\mu$  is the expected value of a MAXSAT function, therefore,  $\mu$  is at least  $1/2$  the optimal value. Therefore, we define  $A_i$  to be the event that the player  $i$  gets a payoff less than

$$\begin{aligned} & \frac{OPT_i}{2} - 60(2\sqrt{\log n + \log d} + 10)(\Delta\sqrt{\frac{OPT_i}{2}} + \Delta^2) \\ &= \frac{OPT_i}{2} - O\left(\Delta(\sqrt{OPT_i} + \Delta)\sqrt{\log n + \log d}\right) \end{aligned}$$

where  $OPT_i$  is the maximum payoff that player  $i$  can get assuming other players do not change their strategies. Because of (6), we have

$$Pr(A_i) \leq \frac{1}{e(2nd^2 + 1)}.$$

Thus, according to Lovász local lemma, there exists a configuration where none of  $A_i$  occur. This is what we need to prove.  $\square$

**Note:** Using algorithmic versions of Lovász local lemma, for example, [16], we can give a polynomial time algorithm for finding such an approximate equilibrium.

## References

1. Alon, N., Spencer, J.H.: The probabilistic method. In: Discrete Math and Optimization. Wiley-Interscience Ser. Wiley, New York (1992)
2. Anshelevich, E., Dasgupta, A., Tardos, E., Wexler, T.: Near-Optimal Network Design with Selfish Agents. *Theory of Computing* 4, 77–109 (2008)
3. 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 (2004)
4. Azuma, K.: Weighted sums of certain dependent random variables. *Tokoku Math. J.* 19, 357–367 (1967)
5. Chen, X., Deng, X.: Settling the Complexity of 2-Player Nash-Equilibrium. In: FOCS 2006 (2006)
6. Daskalakis, C., Papadimitriou, C.: Computing Nash in graphical games via Markov random field. In: EC 2006 (2006)
7. Daskalakis, C., Mehta, A., Papadimitriou, C.: Progress in Approximate Nash Equilibria. In: EC 2007 (2007)
8. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.: The Complexity of Computing a Nash Equilibrium. In: STOC 2006 (2006)
9. Elkind, E., Goldenberg, L.A., Goldenberg, P.W.: Nash Equilibria in Graphical Games on Tree Revisited. In: EC 2006 (2006)
10. Erdős, P., Lovász, L.: Problems and results on 3-chromatic hypergraphs and some related questions. In: Hajnal, A., Rado, R., Sós, V.T. (eds.) *Infinite and finite sets (to Paul Erdős on his 60th birthday)*, vol. II, pp. 609–627. North-Holland, Amsterdam (1975)
11. Fabrikant, A., Papadimitriou, C., Talwar, K.: The Complexity of Pure Nash Equilibria. In: STOC 2004 (2004)
12. Hastad, J.: Some optimal inapproximability results. *Journal of ACM* 48, 798–859 (2001)
13. Hoeffding, W.: Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* 58, 713–721 (1963)
14. Kearns, M., Littman, M., Singh, S.: Graphical models for game theory. In: UAI (2001)
15. Kearns, M.: Graphical Games. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
16. Moser, R.: A Constructive Proof of the Lovász Local Lemma. In: STOC 2009 (2009)
17. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International journal of Game Theory* 2, 65–67 (1973)
18. Skopalik, A., Vöcking, B.: Inapproximability of Pure Nash Equilibria. In: STOC 2008 (2008)
19. Talagrand, M.: Concentration of measure and isoperimetric inequalities in product spaces. *Publications Mathématiques de L’IHÉS* 81, 73–205 (1995)
20. Vetta, A.: Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. In: FOCS 2002 (2002)

# Externalities in Keyword Auctions: An Empirical and Theoretical Assessment\*

Renato Gomes<sup>1</sup>, Nicole Immorlica<sup>2</sup>, and Evangelos Markakis<sup>3</sup>

<sup>1</sup> Northwestern University, Department of Economics  
rdbgomes@gmail.com

<sup>2</sup> Northwestern University, Department of EECS  
nicimm@gmail.com

<sup>3</sup> Athens University of Economics and Business  
Department of Informatics, Athens, Greece  
markakis@gmail.com

**Abstract.** The value of acquiring a slot in a sponsored search list (that comes along with the organic links in a search engine’s result page) might depend on who else is shown in the other sponsored positions. To empirically evaluate this claim, we develop a model of ordered search applied to keyword advertising, in which users browse slots from the top to the bottom of the sponsored list and make their clicking decisions slot by slot. Our contribution is twofold: first, we use impression and click data from Microsoft Live to estimate the ordered search model. With these estimates in hand, we are able to assess how the click-through rate of an ad is affected by the user’s click history and by the other competing links. Our dataset suggests that externality effects are indeed economically and statistically significant. Second, we study Nash equilibria of the Generalized Second Price Auction (GSP) and characterize the scoring rule that produces greatest profits in a complete information setting.

## 1 Introduction

Sponsored search advertising is a booming industry that accounts for a significant part of the revenue made by search engines. For queries with most commercial interest, Google, Yahoo! and MSN Live make available to advertisers up to three links above the organic results (these are the *mainline slots*), up to eight links besides the organic results (*sidebar slots*) and, more recently, MSN Live even sells links below the organic results (*bottom slots*).

As such, an advertiser that bids for a sponsored position is seldom alone; and is usually joined by his fiercest competitors. Indeed, it is widely believed that the value of acquiring a sponsored slot highly depends on the identity and position of the other

---

\* Most of this work was done while all authors were at the Center for Math and Computer Science (CWI), Amsterdam. The project was partially funded by the Microsoft external research program “Beyond Search: Semantic Computing and Internet Economics”.

<sup>1</sup> A full version of our work along with missing proofs of the theoretical results is available at <http://pages.cs.aueb.gr/~markakis/research/pubs.html>

advertisers. Putting it differently, advertisers impose *externalities* on each other, which affect their click-through rates and might have consequences on their bidding behavior.

The literature on sponsored search auctions mostly assumes click-through rates are separable, i.e., the click through rate of a bidder is a product of two quantities, the first expressing the quality of the bidder and the second the quality of the slot she occupies. Such models cannot capture the externalities that one advertiser imposes on the others. To capture these externality effects, we depart from the separable model and study a model in which users perform *ordered search*, that is, (i) they browse the sponsored links from top to bottom and (ii) they take clicking decisions slot by slot. After reading each ad, users decide whether to click on it or not and, subsequently, decide whether to continue browsing the sponsored list or to simply skip it altogether (for a formal definition and motivation for this model, see Section 2). With this formulation, we are able to estimate *continuation probabilities* for each ad (which are simply the probabilities of continuing searching the sponsored list after clicking or not on some ad) and *conditional click-through rates* for each ad (which tell the probability of a click conditional on the user's previous clicking history). Continuation probabilities capture *position externalities*, that is, they capture the negative impact that top links impose on the click-through rates of bottom links (as users stop browsing either because their search needs were already fulfilled or because they got tired of previous bad matches). In turn, conditional click-through rates capture *information externalities*, as we can assess how the information collected by the user by clicking on one given link impacts the click-through rates of the other links he eventually reads.

We used three months of impression and clicking data from Microsoft Live to estimate the ordered search model. We report our findings from three selected search terms: *ipod*, *diet pill* and *avg antivirus*. For each of the selected keywords, we selected the logs in which the most clicked advertisers occupied the mainline slots. Our main empirical findings can be summarized as follows: first, our dataset suggests that both position and information externalities are economically and statistically significant - and the returns to keyword advertising (in terms of clicks) strongly depend on the identity of the other advertisers. Secondly, our estimates suggest that users roughly divide in two groups: the first group has a low clicking probability and usually drops the sponsored list without going through all the mainline slots. In contrast, the second group of users clicks more often and tends to read most of the sponsored links (price research behavior).

Inspired by our empirical findings, we set up an auction model in which advertisers submit their bids taking click-through rates as implied by ordered search. As prescribed by the rules of the Generalized Second Price auction (GSP), search engines then multiply each bid by a weight defined by a *scoring rule* (which solely depends on each advertiser's characteristics), producing a score for each advertiser. Advertisers are then ranked by their score; slots are assigned in decreasing order of scores and each advertiser pays per click the minimum bid necessary to keep his position.

In this framework, we characterize the revenue-maximizing complete information Nash equilibrium (under any scoring rule). We then show that this equilibrium can be implemented under any valuation profiles and advertiser's search parameters if and only if the search engine ranks bids using a particular weighting rule that combines click-through rates and continuation probabilities. Interestingly, this is the same ranking

rule derived in [10] for solving the efficient allocation problem (in a non-strategic environment). Finally we provide an impossibility theorem: there is no scoring rule that implements an efficient equilibrium where advertisers pay their VCG payments for all valuations and search parameters (this is by far the most analyzed equilibrium of the separable click-through rate model). This extends an observation first made by [7], who argue that such an equilibrium does not exist in the rank by revenue GSP.

**Comparisons with Related Work.** The issue of externalities in ad auctions has recently attracted quite a bit of attention from the research community [3,11,10,7,6]. Initial studies were largely theoretical, and involved proposing models for user-search behavior that would explain externalities. Athey and Ellison [3] proposed one of the first such models. In their work, they assume that users search in a top-down manner and that clicking is costly. They then derive the resulting equilibria. Closely related are the cascade models of Aggarwal et al [1] and Kempe and Mahdian [10]. These models associate with each ad a click-through-rate as well as a continuation probability representing the probability that a user continues the search after viewing the given ad. They then proceed to solve the winner determination problem in their models. Recently Giotis and Karlin [7] studied the equilibria of the cascade model in GSP auctions.

Our model of ordered search generalizes these previous models slightly by allowing click-through-rates and continuation probabilities to depend on the clicking history of the user. This enables us to model both position externalities as well as information externalities. Our empirical work shows that both effects are significant.

This is the first paper to empirically document externalities in sponsored search. In a subsequent work, [8] estimated a model of unordered search in which users read all advertisements before choosing a subset of them to click on. We believe this is a valid and worth-exploring model of users' behavior. We nevertheless think that ordered search is a more natural starting point. Indeed, it is hard to reconcile the assumption that users perform unordered search with the advertisers' competition to obtain the top positions (why pay more to get a top slot if users read the whole list anyway?). Moreover, unlike [8], we allow click-through rates to depend on the click history of users (this captures users' learning by browsing).

On the theoretical side, this work advances the equilibrium analysis of the GSP in the presence of externalities. The previous theoretical literature to study GSP equilibria [2,5,12,13,11] mostly focused on the separable click-through-rate model. The only exceptions are Athey and Ellison [3] and Giotis and Karlin [7], mentioned above.

## 2 The Ordered Search Model

In order to study externalities in sponsored advertising, we develop a model of users' behavior that assumes ordered search. The main elements of this model are, first, that users make their choices about clicking on sponsored links by analyzing one link at a time and, secondly, that they browse sponsored results from top to bottom. Our focus on such an ordered search model is motivated by various reasons. First, as the work of [4] demonstrates, position bias is present in organic search. In particular [4] compares a sequential search model with four other models (including the separable model) and concludes that sequential search provides the best fit to the click logs they have considered. Secondly, sequential search is further substantiated as a natural way to browse

through a list of ads by the eye-tracking experiments of Joachims *et al.* [9], where it is observed that users search and click in a top down manner. Moreover, as the value per click of each advertiser tends to be correlated with its relevance, ordered search is a good heuristic for users (see [3]).

Users typically do not click on all the ads of the list, as it is costly both in terms of time and cognitive effort to go through a website and assimilate its content. For this reason, users only click on a link if it looks good enough to compensate for its browsing cost. Moreover, users typically change their willingness to incur this browsing cost as they collect new information through their search, and hence the decision about whether to continue reading ads naturally depends on the click history of the user. To formalize these ideas, we denote the click history of users as they browse through the sponsored links by  $H = \{j : \text{link } j \text{ received a click}\}$ .<sup>2</sup> Here we will focus on two types of externalities:

**Information Externalities.** An ad imposes information externalities on others by providing a user who has clicked on his link with information regarding the search – e.g., prices or product reviews. This, in turn, affects the user’s willingness to click on all links displayed below in the sponsored search list. To make these points formally, let’s denote the expected quality of link  $j$  by  $u_j$ . In order to save on browsing costs, a searcher with click history  $H$  clicks on link  $j$  only if its perceived quality exceeds some optimal threshold, which we denote by  $T_H$ . We set  $H = \{\emptyset\}$  if no links were previously clicked (no extra information gathered through search),  $H = \{j\}$  if only link  $j$  was clicked and  $H = \{j, k\}$  if links  $j$  and  $k$  were clicked in this order. We let the clicking threshold  $T_H$  on the ad’s perceived quality depend on the information gathered by the searcher in his previous clicks, but assume that  $T_H$  is not affected by the precise order of clicks. That is, we impose  $T_{\{j,k\}} = T_{\{k,j\}}$ . In addition, we summarize any user specific bias towards a link by the random term  $\varepsilon^{ij}$ . Hence, a user with click history  $H$  that reaches the slot occupied by advertiser  $j$  clicks on it if and only if

$$u_j - \varepsilon^{ij} \geq T_H.$$

We assume that the idiosyncratic preference parameters  $\varepsilon^{ij}$  are independently and identically distributed across bidders and advertisers, with a cumulative distribution function  $F$ . Thus, the probability that a searcher  $i$  with click history  $H$  clicks on link  $j$  is:

$$F_j(H) \equiv \text{Prob}\{\varepsilon^i \leq u_j - T_H\} = F(u_j - T_H).$$

We call  $F_j(H)$  the *conditional click-through rate* of  $j$  given the click history  $H$ . By virtue of browsing from the top, users have no previous clicks when they analyze the first slot. Hence, if advertiser  $j$  occupies the first position, his chance of getting a click, which we call *click-through rate*, is  $F_j \equiv F_j(\{\emptyset\})$ .

The difference between advertiser  $j$ ’s click-through rate,  $F_j$ , and his conditional click-through rate,  $F_j(H)$ ,  $H \neq \emptyset$ , indicates the impact of information externalities.

**Position Externalities.** An ad additionally imposes externalities on other ads by virtue of its position in the ordered search list. This can happen in one of two manners: first,

<sup>2</sup> Note we abstract away order information; i.e., we assume a user’s behavior depends on past clicks, but not on the order in which the clicks were made.

the user may tire of the search if the ads he has read appear to be poorly related to the search term; second, the user may leave the search if an ad he has read and clicked on has satisfied his search need. We capture the first effect with a parameter,  $\lambda_j$ , that indicates the probability a user keeps browsing the sponsored links after reading ad  $j$  and *choosing not to click on it*. We capture the second effect with a parameter,  $\gamma_j$ , that indicates the probability a user keeps browsing the sponsored links after *clicking link  $j$* . The parameters  $\lambda_j$  and  $\gamma_j$  are referred to as the *continuation probabilities* of ad  $j$  and jointly capture its position externalities imposed on the ads that follow.

Note that, unlike many models in the literature, in our model the position externalities may depend on both the advertiser and clicking behavior of the user.

We model the user behavior for a given sponsored list using the above parameters as follows. She reads the first ad  $A_1$  in the list and clicks on it with probability  $F_{A_1}$ . Conditional on clicking on  $A_1$ , she reads the second ad  $A_2$  with probability  $\gamma_{A_1}$  and clicks on it with probability  $F_{A_2}(\{A_1\})$ . Conditional on not clicking on  $A_1$ , she reads the second ad with probability  $\lambda_{A_1}$  and clicks on it with probability  $F_{A_2}$ . Thus, the probability she clicks on ad  $A_1$  is simply  $F_{A_1}$  while the probability she clicks on  $A_2$  is

$$(1 - F_{A_1})\lambda_{A_1}F_{A_2} + F_{A_1}\gamma_{A_1}F_{A_2}(\{A_1\}).$$

This behavior extends to multiple advertisers in the natural way.

## 2.1 Data Description

Our data consists of impression and clicking records associated to queries that contained the keywords *ipods*, *diet pills* and *avg antivirus* in Microsoft’s Live Search. We chose these keywords because, first, a user that searches for any of them has a well defined objective and, second, because they are highly advertised. Within each of these keywords, we selected the three most popular advertisers (in number of clicks) and considered all impressions in which at least two of these advertisers are displayed.<sup>3</sup>

**Table 1.** Keywords and Advertisers

keyword	advertisers	# of obs.
ipod	(A): store.apple.com	8,398
	(B): cellphonestop.net	
	(C): nextag.com	
diet pill	(A): pricesexposed.net	4,652
	(B): dietpillvalueguide.com	
	(C): certiphene.com	
avg antivirus	(A): Avg-Hq.com	1,336
	(B): avg-for-free.com	
	(C): free-avg-download.com	

<sup>3</sup> Regarding the impressions that contain only two of the three selected advertisers, we only kept those logs which display our selected advertisers in the first two positions. By doing this, we can disregard the advertisers on slot 3 and below without biasing our estimates.



For the keyword *ipod*, the Apple Store ([www.store.apple.com](http://www.store.apple.com)) is the most important advertiser, followed by the online retailer Cell Phone Shop ([www.cellphoneshop.net](http://www.cellphoneshop.net)) and by the price research website Nextag ([www.nextag.com](http://www.nextag.com)). All the 8398 *ipod* observations in our sample refer to impressions that happened between August 1st and November 1st of 2007.

The most popular advertisers for *diet pills* are, first, the meta-search website Price Exposed ([pricesexposed.net](http://pricesexposed.net)), followed by the diet pills retailer [dietpillvalueguide.com](http://dietpillvalueguide.com) and then by [certiphene.com](http://certiphene.com) (which only sells the diet pill certiphene). All 4,652 impressions considered happened between August 1st and October 1st of 2007.

For *avg antivirus*, the most popular advertiser is the official AVG website, followed by the unofficial distributors of the AVG antivirus [avg-for-free.com](http://avg-for-free.com) and [free-avg-download.com](http://free-avg-download.com). The 1,336 observations range from September 1st to November 1st of 2007. The sample provided by Microsoft AdWords displays impressions associated to different keywords with varying intensities through time. This is why ranges differ across the selected keywords; and we have no reason to expect such differences might affect the estimates of our model.

**Table 2.** Distribution of Advertisers per Slot

slot	ipod	diet pill	antivirus
first	(A): 6,460 (76.92%)	(A): 1,912 (41.10%)	(A): 1,233 (92.29%)
	(B): 1,864 (22.20%)	(B): 908 (19.52%)	(B): 71 (5.31%)
	(C): 74 (0.88%)	(C): 1,832 (39.38%)	(C): 32 (2.40%)
second	(A): 1,438 (17.12%)	(A): 1,848 (39.72%)	(A): 88 (6.59%)
	(B): 5,826 (69.37%)	(B): 1,988 (42.73%)	(B): 674 (50.45%)
	(C): 1,134 (13.50%)	(C): 816 (17.54%)	(C): 574 (42.96%)
third	(A): 26 (0.31%)	(A): 472 (10.15%)	(A): 9 (0.67%)
	(B): 22 (0.26%)	(B): 692 (14.88%)	(B): 21 (1.57%)
	(C): 950 (11.31%)	(C): 668 (14.36%)	(C): 355 (26.57%)
	(other): 7,400 (88.12%)	(other): 2,820 (60.62%)	(other): 951 (71.18%)

All keywords possess a leading advertiser that occupies the first position in most of the observations. For *ipod*, the Apple Store occupies the first slot in roughly 77% of the cases, while the Cell Phone Shop appears in 22% of the observations. The situation is reversed when we look at the second slot: the Cell Phone Shop is there in almost 70% of the observations, while the Apple Store and Nextag appear respectively in 17% and 13% of the cases. As table 2 below makes clear, advertising for *diet pills* or *avg antivirus* display a similar pattern.

For all the keywords considered, approximately one out of four impressions got at least one click (25.26% for *ipods*, 24.24% for *diet pills* and 35.55% for *avg antivirus*). As one should expect, click-through rates are decreasing for most of the queries: among the clicks associated to *diet pill*, 56.73% occurred in the first slot, 34.04% in the second and 9.21% in the third. For *ipod*, the concentration of clicks in the first slot is even higher, as one can see from table 3. The keyword *avg antivirus* is an interesting exception, as most of the clicks happened in the second slot (54.5%).

**Table 3.** Distribution of Clicks per Slot

slot	ipod	diet pill	antivirus
first	1,572 (74.08%)	640 (56.73%)	205 (43.15%)
second	524 (24.69%)	384 (34.04%)	259 (54.52%)
third	30 (1.41%)	104 (9.21%)	11 (2.31%)
total	2,122 (100%)	1,128 (100%)	475 (100%)

## 2.2 Estimation Results

At this stage, it is not possible to tell whether a high click-through rate in the first slot is simply due to users' behavior or is the effect of very high quality advertisers. In the same vein, what explains the very low click-through rate in the third slot for *ipod*? Is it because advertisers are bad matches for the users' search or is it the result of search externalities imposed by the links in the first two slots?

In order to evaluate externalities, we must estimate the parameters of our model. We do this with the well-established *maximum likelihood method*, which selects values for the parameters that maximize the probability of the sample. First we must derive an expression, called the *log-likelihood*, for the (log of) probability of the sample given the parameters of the model.<sup>4</sup> Our log-likelihood function is:

$$\log L = \sum_n \log [\text{Prob} (\{j_n, k_n, l_n; c_n^1, c_n^2, c_n^3\})],$$

where the probability of observations  $\{j_n, k_n, l_n; c_n^1, c_n^2, c_n^3\}$  is derived from our empirical model.

Next we estimate the parameters to be those that maximize the log-likelihood. Before discussing our estimation results, we need to make one important observation. The conditional click-through rate of some advertiser  $j$ ,  $F_j(\{k\})$ , is the probability that a random user clicks on ad  $j$  given that this user clicked on advertiser  $k$ 's link and kept searching until he read  $j$ 's link. Note that  $F_j(\{k\})$  abstracts from position externalities, as this is the probability that a user that *read* the ad gives a click on it. We have three reasons to think that conditional click-through rates should differ from baseline click-through rates. First, link  $k$  may offer low prices for ipods, hence even if the user keeps browsing the sponsored list after clicking on  $k$  (an event of probability  $\gamma_k$ ), he will be less likely to click on  $j$ . This is the *negative externality* effect, which pushes, let's say  $F_j(\{k\})$ , to be less than  $F_j$ . Second, link  $k$  may increase the users' willingness to click on  $j$ , which may happen if, for example, link  $k$  is a meta-search website. In this case,  $F_j(\{k\})$  is greater than  $F_j$ , which corresponds to a *positive externality* effect.

These first two reasons for  $F_j$  to depart from  $F_j(\{k\})$  relate to information externalities. There is a third reason, though, not related to externalities but to the structure of our data, that may explain why  $F_j \neq F_j(\{k\})$ : the group of users that make at least one click may be fundamentally different from the total pool of users that perform

<sup>4</sup> It is common to use the log of the probability as opposed to the probability itself to simplify the algebra. As log is a monotone function, maximizing the log-likelihood corresponds to maximizing the likelihood.

searches on Microsoft Live. As such, the conditional click-through rate  $F_j(\{k\})$  reflects the probability of  $j$  getting a click among a quite selected group of users. It is natural to think that these users click more often on sponsored links than a common user; and this should push  $F_j(\{k\})$  to be higher than  $F_j$ . We call this the *selection* effect.

As a consequence, we can safely interpret estimates such that  $F_j > F_j(\{k\})$  as evidence that advertiser  $k$  imposes a negative externality on advertiser  $j$ . Nevertheless, if  $F_j < F_j(\{k\})$ , as we do not observe any users' characteristics, we cannot tell apart positive externalities from purely selection effects. We need to keep this in mind in order to interpret the estimation results.

One can directly test whether the selection effect is driving our estimates by looking at the continuation probabilities  $\lambda_j$  and  $\gamma_j$ . Clearly, absent any selection effect and granted  $j$  is not a meta-search website,  $\lambda_j$ , the probability that a user keeps browsing after *not* clicking on  $j$ , is expected to be higher than  $\gamma_j$ , the probability that a user keeps browsing after clicking on  $j$ . The reason for this is that users may only fulfil their search needs if they do click on  $j$ , in which case they are not expected to return to the results page. As a consequence, having  $\lambda_j$  significantly lower than  $\gamma_j$  is strong evidence in favor of the selection effect.

We are now able to discuss our estimation results, which are displayed at Table 4. We find that for the search terms we investigated, selection effects were ubiquitous. Nonetheless, we observed significant negative externalities in two of them (*ipod* and *avg antivirus*). For the third keyword (*diet pills*), we observed that conditional click-through-rates were higher than the base-line click-through-rates, although it is not possible to determine whether to attribute this to the selection effect or to positive externalities. In the following subsections, we discuss the results for each keyword in detail.

***ipod* Results:** For this keyword, the lead advertiser (the Apple Store) has a very high click-through rate: 21%. Its competitors, the Cell Phone Shop and Nextag, have 8.7% and 10.4%, respectively. These estimates can be interpreted as the probability that the first slot gets a click when it is occupied by one of these three advertisers. The difference between the Apple Store click-through rate and that of its competitors is significant at the 1% level. As such, the lead advertiser (who occupies the top position in 76% of the observations – see Table 2) is also the most effective in attracting clicks.

Our estimates detect that Apple Store imposes a negative externality on the Cell Phone Shop (as  $F_B = 0.08 > 0.04 = F_B(\{A\})$ , and the difference is significant at 5%) and on Nextag (as  $F_C = 0.10 > 0.04 = F_C(\{A\})$ , and the difference is significant at 5%). This means that the information provided by the Apple Store website reduced by half the appeal to a random user of the links to the Cell Phone Shop or the Nextag. The lack of observations in which users click on Nextag and then click on Apple Store or the Cell Phone Shop prevents us from being able to estimate  $\gamma_C, F_A(\{C\})$ ,  $F_A(\{B, C\})$ ,  $F_B(\{C\})$  and  $F_B(\{A, C\})$ .

The selection effect indeed seems to play a role in our estimates. Looking at the results, one can see that  $\gamma$ 's are higher than  $\lambda$ 's for at least two advertisers: for the Apple Store,  $\gamma_A = 0.94 > 0.76 = \lambda_A$  (although the difference is not significant) and for the Cell Phone Shop,  $\gamma_B = 1 > 0.62 = \lambda_B$  (significant at 15%). This suggests that users that click on a link are more likely to keep browsing the sponsored list. As the

results presented above point out, though, the selection effect wasn't strong enough to shadow the negative externalities that the Apple Store imposes on its competitors.

**diet pill Results:** Alike the *ipod* case, the leading advertiser for *diet pill* is also the most effective in terms of attracting users: the click-through rate of *pricesexposed.net*, roughly 21%, is significantly (at 1% level) higher than that of its competitors (15% for *dietpillvalueguide.com* and 5% for *certiphene.com*).

We didn't find evidence of negative information externalities among *diet pill* advertisers. For *pricesexposed.net*, the click-through rate jumps from roughly 21% to 31% if *certiphene.com* was previously clicked; and the difference is significant at 10%. The same happens with *dietpillvalueguide.com*: its click-through rate goes from 15% to either 66% (in case *certiphene.com* got a click) or to 33% (in case *certiphene.com* and *pricesexposed.net* had clicks); and both differences are significant at 5%.

Interestingly, the click-through rate of *certiphene.com* jumps from 5% to 8% (difference significant at 5%) if *dietpillvalueguide.com* was previously clicked by the user. Since *dietpillvalueguide.com* is a website specialized in comparing diet products, one can think that positive reviews of the Certiphene pills might explain this difference.

As discussed above, we cannot rule out that the selection effect explains this difference, though. Indeed, our estimates imply that users are more likely to keep browsing the sponsored links if they clicked on *certiphene.com*:  $\gamma_C = 1 > 0.57 = \lambda_C$  (significant at 5%).

**Table 4.** Estimates of the Ordered Search Model

keyword	ipod	diet pill avg			ipod	diet pill avg			ipod	diet pill avg		
$F_A$	0.210 (0.005)	0.210 (0.008)	0.151 (0.010)	$F_B$	0.087 (0.006)	0.146 (0.034)	0.364 (0.050)	$F_C$	0.104 (0.012)	0.051 (0.004)	0.215 (0.042)	
$F_{A(B)}$	0.250 (0.038)	0.232 (0.032)	0.00 (0.074)	$F_{B(A)}$	0.030 (0.022)	0.146 (0.034)	0.364 (0.050)	$F_{C(A)}$	0.040 (0.032)	0.052 (0.017)	0.242 (0.042)	
$F_{A(C)}$	—	0.317 (0.065)	—	$F_{B(C)}$	—	0.663 (0.080)	—	$F_{C(B)}$	0.095 (0.032)	0.088 (0.029)	0.121 (0.889)	
$F_{A(B,C)}$	—	0.664 (0.075)	—	$F_{B(A,C)}$	—	0.334 (0.083)	—	$F_{C(A,B)}$	0.327 (0.190)	0.664 (0.089)	0.125 (0.699)	
$\lambda_A$	0.676 (0.056)	0.760 (0.064)	1.0 (0.217)	$\lambda_B$	0.627 (0.042)	0.673 (0.057)	0.183 (0.049)	$\lambda_C$	1.00 (0.057)	0.579 (0.037)	0.424 (0.201)	
$\gamma_A$	1.00 (0.777)	0.940 (0.195)	1.00 (0.231)	$\gamma_B$	1.00 (0.820)	1.00 (0.743)	0.686 (0.902)	$\gamma_C$	—	1.00 (0.892)	—	

**avg antivirus Results:** Unlike the previous keywords, the leading advertiser for *avg antivirus* is not the one with highest CTR. In fact, *Avg-Hq.com* has the lowest CTR (15%), while *avg-for-free.com* and *free-avg-download.com* have a 20% and 21% CTRs, respectively (higher than *Avg-Hq.com*'s CTR at a 15% confidence level).

Our estimates detect that *avg-for-free.com* imposes a negative externality on *Avg-Hq.com*, as  $F_A = 0.15 > 0 = F_A(\{B\})$  (significant at 5%). As in the *ipod* case, the lack of observations in which users click on *free-avg-download.com* and then click on *Avg-Hq.com* or *avg-for-free.com* makes it impossible to estimate  $\gamma_C$ ,  $F_A(\{C\})$ ,  $F_A(\{B, C\})$ ,  $F_B(\{C\})$  and  $F_B(\{A, C\})$ .

### 3 Equilibrium Analysis

We'll now analyze how advertisers bid given that users do ordered search. We return to a model with a set of  $N$  advertisers denoted by  $A_j, j \in \{1, \dots, N\}$  and  $K$  slots. Each advertiser  $A_j$  has a value of  $v_{A_j}$  per click. Search engines use the following generalization of the second-price auction to sell sponsored links: first, each advertiser  $A_j$  submits a bid  $b_{A_j}$  representing his willingness to pay per click. Then each advertiser's bid is multiplied by a weight  $w_{A_j}$  that solely depends on his characteristics, producing a score  $s_{A_j} = w_{A_j} \cdot b_{A_j}$ . Next, advertisers are ranked in decreasing order of their scores and the  $j^{\text{th}}$  highest ranked advertiser gets the  $j^{\text{th}}$  slot. When an advertiser receives a click, he is charged a price equal to the smallest bid he could have submitted that would have allowed him to maintain his position in the list. Labeling advertisers such that  $A_i$  denotes the advertiser ranked in the  $i$ 'th slot, we see that advertiser  $A_j$  pays  $p_{A_j}$  where:

$$p_{A_j} \cdot w_{A_j} = b_{A_{j+1}} \cdot w_{A_{j+1}} \quad \text{which gives} \quad p_{A_j} = \frac{b_{A_{j+1}} \cdot w_{A_{j+1}}}{w_{A_j}}.$$

The total payment of advertiser  $A_j$  is then  $p_{A_j} \cdot q^j$ , where  $q^j$  is the total number of clicks of slot  $j$ . To simplify the analysis, we'll take the ordered search model of the previous section and assume that baseline and conditional click-through rates are the same for each advertiser, that is,  $F_{A_j} = F_{A_j}(H)$  for any history  $H$ . Although our empirical exercise suggests that baseline and conditional click-through rates indeed differ, this assumption is necessary to bring tractability to our theoretical model of bidding. Further, our main theoretical conclusions remain valid under the more general ordered search model of the previous section.

With this assumption in hand, the total number of clicks of the  $j^{\text{th}}$  slot is given by:

$$q^j = F_{A_j} \cdot \prod_{k=1}^{j-1} c_{A_k}, \quad \text{where} \quad c_{A_k} = F_{A_k} \gamma_{A_k} + (1 - F_{A_k}) \lambda_{A_k}.$$

Each term  $c_{A_k}$  accounts for the fraction of users that continue browsing the sponsored list after coming across advertiser  $A_k$ . As such, the total number of clicks of slot  $j$  is the product of advertiser  $A_j$ 's click-through rate ( $F_{A_j}$ ) and the total number of users that reach that position ( $\prod_{k=1}^{j-1} c_{A_k}$ ). Advertiser  $A_j$ 's payoff is then  $(v_{A_j} - p_{A_j})q^j$ .

We are interested in analyzing the Nash equilibria and the resulting efficiency of various scoring rules. A *complete information Nash equilibrium* is a vector of bids such that no advertiser can unilaterally change his bid and improve his payoff.

#### 3.1 Can Scoring Rules Help?

Search engines have often changed their auction rules for keyword advertising in order to increase revenue. Yahoo! first dropped a generalized first-price auction and adopted the rank-by-bid GSP in early 1997. Ten years later, Yahoo! opted for a less drastic change and simply altered its scoring rule from rank-by-bid to rank-by revenue (in which case  $w_{A_j} = F_{A_j}$ ). Microsoft's Live Search followed the same path and also in

2007 moved from the rank-by-bid to the rank-by-revenue GSP. Recently, Google also changed its scoring rule, although its precise functional form was not made public.

We will focus on a very interesting, but so far neglected, equilibrium of the GSP: the one that maximizes the search engine's revenue among all pure strategy Nash equilibria. The next lemma derives the bid profile that maximizes revenue for the search engine:

**Lemma 1.** *Consider the GSP with scoring rule  $w_{A_j}$ , selling  $K$  slots to  $N > K$  advertisers. Let advertisers  $A_1, \dots, A_K$  be the efficient assignees of slots 1 to  $K$  and assume advertisers submit bids according to:*

$$b_{A_j} = (1 - c_{A_j}) \frac{w_{A_{j-1}}}{w_{A_j}} v_{A_{j-1}} + c_{A_j} \frac{w_{A_{j+1}}}{w_{A_j}} b_{A_{j+1}}$$

$$\text{for } j \in \{2, \dots, K\}, b_{A_{K+1}} = \frac{w_{A_K}}{w_{A_{K+1}}} v_{A_K}, \quad b_{A_1} > b_{A_2} \quad (1)$$

$$\text{and } b_{A_j} < b_{A_{K+1}} \text{ for } j > K + 1. \quad (2)$$

If this bid profile constitutes a Nash equilibrium, than it maximizes the search engine's revenue among all pure strategy Nash equilibria. We call it the greedy bid profile.

As the next proposition shows, such a bid profile is an equilibrium for all  $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$  if and only if weights are given by:

$$w_{A_j} = \frac{F_{A_j}}{1 - c_{A_j}} = \frac{F_{A_j}}{1 - (F_{A_j} \gamma_{A_j} + (1 - F_{A_j}) \lambda_{A_j})}.$$

Although at first awkward, the scoring rule above is a quite natural one. Indeed, as first proved by [10], advertiser  $j$  comes on top of advertiser  $k$  in the efficient allocation if and only if  $v_{A_j} \cdot w_{A_j} \geq v_k \cdot w_{A_k}$ .

**Proposition 1.** *Consider the GSP with scoring rule  $w_{A_j}$ , selling  $K$  slots to  $N > K$  advertisers. The greedy bid profile constitutes a complete information Nash equilibrium for all valuations and search parameters  $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$  if and only if  $w_{A_j} = \frac{F_{A_j}}{1 - c_{A_j}}$  (up to a multiplicative constant). In this case, the equilibrium allocation is efficient and the search engines's revenue is maximal.*

Our next proposition brings a pessimistic message about what scoring can achieve in the GSP. It shows that there is no scoring rule for which an efficient equilibrium where each advertiser pays his Vickrey-Clark-Groves payments exists for all profiles of valuations and search parameters. This extends a result by [7], who shows that the GSP equipped with the "rank-by-revenue" scoring function ( $w_{A_K} = F_{A_K}$ ) does not possess an efficient equilibrium that implements VCG payments. Recall the VCG payments charge each advertiser the welfare difference imposed on the others:

$$p_{A_j}^V = W(N - \{A_j\}) - (W(N) - q^j v_{A_j})$$

where for a set  $S$ ,  $W(S)$  is the optimal social welfare of the agents in  $S$ .

**Proposition 2.** *Consider the GSP selling  $K$  slots to  $N > K$  advertisers. There is no scoring rule  $w_{A_j}$  which depends solely on advertiser  $A_j$ 's search parameters  $(F_{A_j}, \gamma_{A_j}, \lambda_{A_j})$  that implements an efficient equilibrium with VCG payments for all valuations and search parameters  $\{(v_{A_j}, F_{A_j}, \gamma_{A_j}, \lambda_{A_j})\}_{j=1}^N$ .*

## 4 Conclusion

This work documents information and position externalities among sponsored search advertisers. Our results bring suggestive evidence that part of the population of users perform price research through the sponsored list (as a user that clicks on a link is more likely to keep browsing the sponsored list than users that don't make clicks at all).

Our approach relies on the assumption that users browse from the top to the bottom of the sponsored list and take clicking decisions link by link. It would be interesting to extend the analysis (both empirical and theoretical) to allow users to perform other search procedures.

## References

1. Aggarwal, G., Feldman, J., Muthukrishnan, S., Pal, M.: Sponsored search auctions with markovian users. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 621–628. Springer, Heidelberg (2008)
2. Aggarwal, G., Goel, A., Motwani, R.: Truthful auctions for pricing search keywords. In: Proc. the Seventh ACM Conference on Electronic Commerce, Ann Arbor, Michigan, USA (2006)
3. Athey, S., Ellison, G.: Position auctions with consumer search. working paper (2008)
4. Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: WSDM 2008, pp. 87–94 (2008)
5. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth in keywords. *American Economic Review* 97(1), 242–259 (2007)
6. Ghosh, A., Mahdian, M.: Externalities in online advertising. In: International World Wide Web Conference, WWW (2008)
7. Giotis, I., Karlin, A.R.: On the equilibria and efficiency of the gsp mechanism in keyword auctions with externalities. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 629–638. Springer, Heidelberg (2008)
8. Jeziorskiy, P., Segal, I.: What makes them click: Empirical analysis of consumer demand for search advertising (working paper) (2009)
9. Joachims, T., Granka, L., Pan, B., Hembrooke, H., Gay, G.: Accurately predicting click-through data as implicit feedback. In: SIGIR, pp. 154–161 (2005)
10. Kempe, D., Mahdian, M.: A cascade model for externalities in sponsored search. In: Workshop on Internet and Network Economics (WINE), Shanghai, China (2008)
11. Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: Proc. the Seventh ACM Conference on Electronic Commerce, Ann Arbor, Michigan, USA (2006)
12. Lahaie, S., Pennock, D., Saberi, A., Vohra, R.: 28: Sponsored search auctions. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
13. Varian, H.: Position auctions. *International Journal of Industrial Organization* 25, 1163–1178 (2007)



# Covering Games: Approximation through Non-cooperation<sup>\*</sup>

Martin Gairing

Department of Computer Science, University of Liverpool, U.K.  
m.gairing@liverpool.ac.uk

**Abstract.** We propose approximation algorithms under game-theoretic considerations. We introduce and study the *general covering problem* which is a natural generalization of the well-studied *max-n-cover problem*. In the general covering problem, we are given a universal set of weighted elements  $E$  and  $n$  collections of subsets of the elements. The task is to choose one subset from each collection such that the total weight of their union is as large as possible. In our game-theoretic setting, the choice in each collection is made by an independent player. For covering an element, the players receive a payoff defined by a non-increasing *utility sharing function*. This function defines the fraction that each covering player receives from the weight of the elements.

We show how to construct a utility sharing function such that every Nash Equilibrium approximates the optimal solution by a factor of  $1 - \frac{1}{e}$ . We also prove that any sequence of unilateral improving steps is polynomially bounded. This gives rise to a polynomial-time local search approximation algorithm whose approximation ratio is best possible.

## 1 Introduction

**Motivation and Framework.** Large scale distributed systems, like the Internet, usually lack a centralized control authority. Instead, they are operated and controlled in a distributed fashion by competing entities – modeled as *players* – which make their decisions in order to optimize their own private *utility*. Such systems are assumed to end up in a *Nash equilibrium* [20] – a state in which no player wishes to unilaterally leave her own strategy in order to improve the value of her private *utility*. However, Nash equilibria are often suboptimal solutions with respect to the *social objective function*. The *price of anarchy* [18] is a measure for the performance degradation. It is defined as the worst-case ratio between the values of a social objective function in a Nash equilibrium and in an optimum solution.

As the designer of a distributed system we are faced with the main challenge of how to design the distributed system in order to optimize this social objective function even in the presence of myopic players. However, even if all players

---

<sup>\*</sup> This work was supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD).



adhere to some centralized authority, finding the optimum solution with respect to the social objective might be an  $\mathcal{NP}$ -hard optimization problem. Approximation algorithms [23] are a powerful tool for coping with intractable optimization problems. In general, they compute (in polynomial time) suboptimal solutions but with a provable performance guarantee. However, approximation algorithms usually presume a centralized authority. In this paper we propose to consider approximation algorithms that take the selfish user behavior into account. More precisely, we propose to design the distributed system in a way that the price of anarchy is optimized and the system is guaranteed to converge to a Nash equilibrium in polynomial time.

For our study of such approximation algorithms we consider a very general covering problem. In covering problems a finite set of elements has to be covered with subsets of the elements. Such problems arise in many contexts: Covering problems can be used to model service installation problems in distributed systems. Moreover, many packing problems and fixed parameter optimization problems can be modeled as a covering problem (see [13]). A well-studied representative is the *max- $n$ -cover* problem (see e.g. [9,14]): Given a finite set of weighted elements  $E$ , choose  $n$  subsets of the elements (from a given collection of subsets) such that the total weight of their union is as large as possible. We consider a generalization of the *max- $n$ -cover* problem that we call *general covering problem*. Here, we are given not 1 but  $n$  collections of subsets and we have to choose one subset from each collection. Although this generalization seems natural, we are not aware of any previous work on it.

We study the general covering problem as a *covering game*, where the choice for the subset in each of the  $n$  collections is made by an independent player. Covering games are a subclass of the *congestion games* introduced by Rosenthal [21]. For covering an element, the players receive a payoff defined by a *utility sharing function*. This function defines the fraction that each covering player receives from the weight of the element and depends only on local parameters. Those parameters are the *number of players* covering an element and the *cardinality* of an element (i.e. the number of players that have this element in at least one of their strategies). We only make two natural assumptions on the utility sharing function: First, we assume that it is non-increasing in the number of players covering the element. And second, we want that the payoff to the players for covering an element does not exceed the weight of the element.

The focus of this paper is to design utility sharing functions, such that:

1. In any Nash equilibrium the total weight of the covered elements is as large as possible. Or more precisely, the price of anarchy is maximized.
2. A Nash equilibrium is reached in polynomial time.

Obviously, each utility sharing function that fulfills both of these properties gives rise to a local search approximation algorithm. In fact, we will show that this approach yields essentially the best possible approximation ratio.

**Contribution.** In this paper, we introduce and study *covering games*. Such games are congestion games that have the *general covering problem* as underlying structure.

In the first part of the paper, we focus on the design of *utility sharing functions* that maximize the price of anarchy. In particular, we construct a utility sharing function which achieves a price of anarchy of  $1 - \frac{1}{e}$  (Theorem 4) and show that no utility sharing function performs better (Theorem 1). To show this, we first prove a corresponding result on the price of anarchy, that depends on the maximum cardinality  $k$  of the elements (Theorem 3). Surprisingly, we get matching bounds for each fixed  $k$ . All our results on the price of anarchy hold for pure and mixed Nash equilibria.

In the second part of the paper, we show how to use our results on the price of anarchy to construct a *local search approximation algorithm* for the general covering problem (Theorem 6), which runs in polynomial time, if the weights of the elements are polynomially bounded. Our hardness result in Theorem 7 shows that this restriction on the weights is necessary. For the general case, we also present a (centralized) approximation algorithm, which is based on LP-rounding and generalizes an algorithm for MAXSAT [12].

**Related Work.** Congestion games and variants thereof have long been used to model non-cooperative resource sharing among selfish players. Rosenthal [21] showed that congestion games always possess pure Nash equilibria. However, computing such a pure Nash equilibrium is  $\mathcal{PLS}$ -complete [8]. The price of anarchy in congestion games has been studied extensively (see e.g. [1,6,10,22]).

The *general covering problem* is a natural generalization of the well-studied *max- $n$ -cover* problem. For the *max- $n$ -cover* problem, the greedy approach yields a  $(1 - \frac{1}{e})$ -approximation [14] and no polynomial time algorithm can do better, unless  $NP \subseteq TIME(n^{O(\log \log n)})$  [9]. Applying the greedy approach to our more general problem guarantees only a  $\frac{1}{2}$ -approximation. For an overview on approximation algorithms for covering problems, we refer to [13, Chapter 3].

The MAXSAT problem is a special case of the *generalized covering* problem, where each of the  $n$  collections consists of at most 2 subsets (corresponding to true/false). The power of local search for approximating MAXSAT has been studied in [2,17]. MAXSAT has also been considered in a game-theoretic setting as a SAT-game [4,11], which is itself a special case of our covering games. Bilò [4] mainly focuses on the expressiveness of SAT-games. Mavronicolas et al. [19] concentrate on structural properties and complexity questions for a generalization of SAT-games, called *weighted boolean formula games*. Here, each player controls a set of variables and aims to maximize the total weight of his satisfied formulas. Giannakos et al. [11] study the price of anarchy (for pure Nash equilibria) of SAT-games under different utility sharing functions and point out the relation to approximation algorithms. Our work generalizes their results in two perspectives. First, we consider a far more general class of games, and second, we allow for mixed Nash equilibria. Moreover, in Example 1 we show that their main result [11, Thm. 5] is incorrect.

Under certain conditions on the utility sharing functions, our games fall in the class of *valid utility games* [24]. For such games, Vetta [24] shows that each Nash equilibrium is a  $\frac{1}{2}$ -approximation. Our result improves this ratio to  $1 - \frac{1}{e}$ .

Coordination mechanisms have been introduced in [7] as a notion to improve the price of anarchy. The idea is to define local policies such that the corresponding price of anarchy is as small as possible. A few other papers follow this approach, e.g. [3,5,15]. Our task of designing utility sharing functions can be seen as such a coordination mechanism. However, we take the idea of coordination mechanisms one step further. In the design of our utility sharing function, the goal is not only to optimize the price of anarchy, but also to ensure that the system converges to a Nash equilibrium in polynomial time. Azar et al. [3] pursue a similar approach for the unrelated scheduling problem. However, the price of anarchy of their best coordination mechanism increases significantly by requiring a polynomial convergence.

**Roadmap.** The rest of the paper is organized as follows. In Section 2, we introduce *covering games*. Section 3 comprises our results on the price of anarchy, while Section 4 presents our approximation algorithms. We conclude in Section 5. Due to lack of space, some proofs are omitted.

## 2 Model

For any two integers  $l \leq m$ , denote  $[m] = \{1, \dots, m\}$  and  $[l, m] = \{l, \dots, m\}$ . For a vector  $\mathbf{v} = (v_1, \dots, v_n)$ , let  $\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$  and  $(\mathbf{v}_{-i}, v'_i) = (v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_n)$ .

**The general covering problem.** In the *general covering problem* we are given a finite set of elements  $E$  and a weight function  $w : E \mapsto \mathbb{N}$  that assigns a positive integer weight  $w_e$  to each element  $e \in E$ . Moreover, we are given  $n$  collections  $S_1, \dots, S_n$  of subsets of  $E$  where for each  $i \in [n]$ , the collection  $S_i \subset 2^E$  is a subset of the power-set of the elements. Given such an instance our task is to choose one subset  $s_i$  from each collection  $S_i$  such that their union  $\cup_{i \in [n]} s_i$  has maximum total weight, i.e.  $\sum_{e \in \cup_{i \in [n]} s_i} w_e$  is maximized.

**Covering Games.** Each *covering game* has a general covering problem as an underlying structure. Here, each of the  $n$  collections of subsets is controlled by a rational *player*, that is player  $i \in [n]$  has  $S_i$  as her strategy set. Denote  $S = S_1 \times \dots \times S_n$ . As for the general covering problem, each element  $e \in E$  has a weight  $w_e \in \mathbb{N}$ . For any subset of the elements  $E' \subseteq E$  denote  $W(E') = \sum_{e \in E'} w_e$ . Let  $W = W(E)$ . For each element  $e \in E$  denote by  $k_e = |\{i \in [n] : e \in s_i \text{ for some } s_i \in S_i\}|$  the *cardinality* of  $e$  which is the number of players that can possibly cover  $e$ . Let  $k = \max_{e \in E} k_e$  and  $k_{\min} = \min_{e \in E} k_e$ . A covering game is a *SAT-game* if  $|S_i| \leq 2$  for all player  $i \in [n]$ . In this case, elements correspond to clauses and players correspond to variables which can be set to true or false.

**Strategies and Strategy Profiles.** A *pure strategy* for player  $i$  is some specific strategy  $s_i \in S_i$ , while a *mixed strategy*  $P_i = (p(i, s_i))_{s_i \in S_i}$  is a probability distribution over  $S_i$ , where  $p(i, s_i)$  denotes the probability that player  $i$  chooses the pure strategy  $s_i$ .

A *pure strategy profile* is an  $n$ -tuple  $\mathbf{s} = (s_1, \dots, s_n)$  whereas a *mixed strategy profile*  $\mathbf{P} = (P_1, \dots, P_n)$  is represented by an  $n$ -tuple of mixed strategies. For a

mixed strategy profile  $\mathbf{P}$ , denote by  $p(\mathbf{s}) = \prod_{i \in [n]} p(i, s_i)$  the probability that the players choose the pure strategy profile  $\mathbf{s}$ .

**Utility Sharing Functions.** For each element  $e \in E$  there is a *payoff function*  $f_e$  that describes how much a player receives for covering  $e$ . In this paper we consider payoff functions that come from a common *utility sharing function*  $f$ . We want this function to depend only on local parameters. We consider two different kinds of utility sharing functions:

- A *cardinality dependent* utility sharing function depends on the number of players covering an element and the cardinality of the element, i.e.  $f : [k] \times [k] \mapsto \mathbb{N}$  and for all elements  $e \in E$  and  $j \in [k]$ ,  $f_e(j) = f(j, k_e) \cdot w_e$ .
- A *symmetric* utility sharing function depends only on the number of players covering an element, i.e.  $f : [k] \mapsto \mathbb{N}$  and for all elements  $e \in E$  and  $j \in [k]$ ,  $f_e(j) = f(j) \cdot w_e$ .

For both cases we assume that  $f$  is *non-increasing* in the number of players. Moreover, we assume that  $f$  does not overpay the players, i.e.  $j \cdot f(j) \leq 1$  for all  $j \in [k]$  in the symmetric case (and  $j \cdot f(j, l) \leq 1$  for all  $j \in [l], l \in [k]$  in the cardinality dependent case).

**Load and Player Utilities.** For a pure strategy profile  $\mathbf{s}$ , let  $\delta_e(\mathbf{s}) = |\{i \in [n] : e \in s_i\}|$  denote the *load* on element  $e \in E$ , i.e. the number of players covering  $e$ .

Fix a pure strategy profile  $\mathbf{s}$ . The *utility*  $u_i(\mathbf{s})$  of player  $i \in [n]$  is defined by the *payoff* from the elements she covers. Thus,  $u_i(\mathbf{s}) = \sum_{e \in s_i} f_e(\delta_e(\mathbf{s}))$ . For a mixed strategy profile  $\mathbf{P}$ , the *utility* of player  $i \in [n]$  is  $u_i(\mathbf{P}) = \sum_{\mathbf{s} \in \mathcal{S}} p(\mathbf{s}) \cdot u_i(\mathbf{s})$ .

**Social Utility.** Fix a pure strategy profile  $\mathbf{s}$ . Denote by  $E_{\mathbf{s}}$  the subset of elements that are covered by at least one player in  $\mathbf{s}$ , i.e.  $E_{\mathbf{s}} = \{e \in E : \delta_e(\mathbf{s}) > 0\}$ . The *social utility* in  $\mathbf{s}$  is the total weight  $W(E_{\mathbf{s}})$  of the covered elements. We abuse notation and denote this value also as  $W(\mathbf{s})$ . For a mixed strategy profile  $\mathbf{P}$  the social utility  $W(\mathbf{P}) = \sum_{\mathbf{s} \in \mathcal{S}} p(\mathbf{s}) \cdot W(E_{\mathbf{s}})$  is the expected total weight of the covered elements. Throughout denote by  $\mathbf{s}^*$  a pure strategy profile that maximizes the total weight of the covered elements, thus,  $\mathbf{s}^* = \arg \max_{\mathbf{s} \in \mathcal{S}} W(E_{\mathbf{s}})$ .

**Nash Equilibria and Potential Function.** A mixed strategy profile  $\mathbf{P}$  is a *Nash equilibrium* if and only if no player can increase her utility by unilaterally changing her strategy, that is,  $u_i(\mathbf{P}) \geq u_i(\mathbf{P}_{-i}, s_i)$  for all  $i \in [n]$  and  $s_i \in S_i$ . Depending on the type of strategy profile we distinguish between *pure* and *mixed* Nash equilibria. Given a pure strategy profile  $\mathbf{s}$ , a *selfish step* of player  $i \in [n]$  is a deviation to a strategy profile  $(\mathbf{s}_{-i}, s'_i)$  where  $u_i(\mathbf{s}_{-i}, s'_i) > u_i(\mathbf{s})$ , that is player  $i$  increases her utility.

For covering games, Rosenthal’s [21] *exact potential function*  $\Phi$  implies the existence of a pure Nash equilibrium. For every pure strategy profile  $\mathbf{s}$ , the potential  $\Phi(\mathbf{s})$  is defined by  $\Phi(\mathbf{s}) = \sum_{e \in E} \sum_{i=1}^{\delta_e(\mathbf{s})} f_e(i)$ . If a player performs a selfish step and increases her utility by  $\Delta$ , then  $\Phi(\mathbf{s})$  also increases by  $\Delta$ .

**Price of Anarchy.** Let  $\mathcal{G}(k)$  be the class of covering games where  $k_e \leq k$  for all  $e \in E$ . Fix a utility sharing function  $f$ . The *Price of Anarchy* for  $f$ , denoted by  $\text{PoA}_f$ , is the infimum, over all instances  $\Gamma \in \mathcal{G}(k)$  and Nash equilibria  $\mathbf{P}$ , of

the ratio  $\frac{W(\mathbf{P})}{W(\mathbf{s}^*)}$ . Thus,  $\text{PoA}_f(k) = \inf_{\Gamma \in \mathcal{G}(k), \mathbf{P}} \frac{W(\mathbf{P})}{W(\mathbf{s}^*)}$ . Similarly, define  $\text{PoA}_f$  by dropping the restriction on  $k$ .

### 3 Price of Anarchy Results

In this section, we study the price of anarchy for different utility sharing functions. We start with an upper bound that holds for all utility sharing functions.

**Theorem 1.** *Consider the class of covering games  $\mathcal{G}(k)$  with unweighted elements. Then,  $\text{PoA}_f(k) \leq 1 - \frac{1}{\binom{k-1}{(k-1)\Gamma} + \sum_{j \in [0, k-1]} \frac{1}{j^\Gamma}}$ .*

*This holds (a) for every cardinality dependent utility sharing function  $f$ , (b) even for SAT-games, if we restrict ourselves to symmetric utility sharing functions  $f$ .*

For SAT-games Giannakos et al. [11, Thm. 5] claim that using the cardinality dependent utility sharing function defined by  $f(1, l) = 1$  and  $f(j, l) = \frac{1}{2(l-1)}$  for  $j \geq 2$  achieves a price of anarchy of  $\frac{2}{3}$ . The following example shows that this does not hold:

*Example 1.* Given the maximum cardinality  $k$ , define a SAT-game with  $k$  players and  $k+1$  elements. We have  $w_e = 1$  for all  $e \in [k-1]$  and  $w_k = w_{k+1} = 2(k-1)$ . Each player  $i \in [k-1]$  can either cover element  $i$  or element  $k$ , while player  $k$  can choose between the elements  $k$  and  $k+1$ .

Let  $\mathbf{s}$  be the strategy profile where each players  $i \in [k]$  covers element  $i$ . It's not hard to see that  $\mathbf{s}$  is a Nash equilibrium with  $W(E_{\mathbf{s}}) = 3(k-1)$ . On the other hand there is a strategy profile  $\mathbf{s}^*$  (where only element 1 is not covered) with  $W(\mathbf{s}^*) = 5(k-1) - 1$ . For  $k \geq 4$ , this stands in conflict to the claim in [11, Thm. 5]; and for  $k \rightarrow \infty$ , we might only cover  $\frac{3}{5}$  of the optimum total weight. In fact, we believe that our upper bound in Theorem 1 holds also for SAT-games with cardinality dependent utility sharing functions.

We proceed by introducing a parameter  $\chi_f$  of the utility sharing function  $f$ . This parameter is a measure on how fast the utility sharing function decreases. We will use  $\chi_f$  in Theorem 2 to prove a general lower bound on the price of anarchy that depends on  $\chi_f$ .

**Definition 1.** *Given a cardinality dependent utility sharing function  $f$ , define  $\chi_f$  as the minimum value such that for all cardinalities  $l \in [k]$  we have*

$$j \cdot f(j, l) - f(\min\{j+1, l\}, l) \leq \chi_f \cdot f(1, l) \quad \text{for all } j \in [l].$$

**Theorem 2.** *Consider the class of covering games  $\mathcal{G}(k)$ . Let  $f$  be a cardinality dependent utility sharing function, where  $\alpha_{\min} = \min_{l \in [k]} f(1, l)$  and  $\alpha_{\max} = \max_{l \in [k]} f(1, l)$ . Then,  $\text{PoA}_f(k) \geq \frac{1}{\chi_f + 1} \cdot \frac{\alpha_{\min}}{\alpha_{\max}}$ .*

*Proof.* Let  $\mathbf{P}$  be an arbitrary mixed Nash equilibrium and  $\mathbf{s}^*$  be an optimum pure strategy profile. Since  $\mathbf{P}$  is a Nash equilibrium, it follows that  $u_i(\mathbf{P}) - u_i(\mathbf{P}_{-i}, \mathbf{s}_i^*) \geq 0$  for all player  $i \in [n]$ . So,

$$0 \leq \sum_{i \in [n]} u_i(\mathbf{P}) - \sum_{i \in [n]} u_i(\mathbf{P}_{-i}, \mathbf{s}_i^*) = \sum_{\mathbf{s} \in S} p(\mathbf{s}) \left( \sum_{i \in [n]} u_i(\mathbf{s}) - \sum_{i \in [n]} u_i(\mathbf{s}_{-i}, \mathbf{s}_i^*) \right) \quad (1)$$

By definition of player utility, for any pure strategy profile  $\mathbf{s}$ ,

$$\sum_{i \in [n]} u_i(\mathbf{s}) = \sum_{i \in [n]} \sum_{e \in s_i} f_e(\delta_e(\mathbf{s})) = \sum_{e \in E} \delta_e(\mathbf{s}) \cdot f_e(\delta_e(\mathbf{s})) = \sum_{j \in [k]} \sum_{\substack{e \in E_s, \\ \delta_e(\mathbf{s})=j}} j \cdot f_e(j). \quad (2)$$

Moreover,

$$\begin{aligned} \sum_{i \in [n]} u_i(\mathbf{s}_{-i}, \mathbf{s}_i^*) &= \sum_{i \in [n]} \sum_{e \in s_i^*} f_e(\delta_e(\mathbf{s}_{-i}, \mathbf{s}_i^*)) \geq \sum_{i \in [n]} \sum_{e \in s_i^*} f_e(\min\{k_e, \delta_e(\mathbf{s}) + 1\}) \\ &\geq \sum_{e \in E_{s^*}} f_e(\min\{k_e, \delta_e(\mathbf{s}) + 1\}) = \sum_{j=0}^k \sum_{\substack{e \in E_{s^*}, \\ \delta_e(\mathbf{s})=j}} f_e(\min\{k_e, j + 1\}) \end{aligned} \quad (3)$$

where the first inequality follows since  $f_e$  is a non-increasing function and the second inequality follows since  $\delta_e(\mathbf{s}^*) \geq 1$  for all  $e \in E_{s^*}$ .

So, for any pure strategy profile  $\mathbf{s}$ , (2) and (3) imply:

$$\begin{aligned} \sum_{i \in [n]} u_i(\mathbf{s}) - \sum_{i \in [n]} u_i(\mathbf{s}_{-i}, \mathbf{s}_i^*) &\leq \sum_{j \in [k]} \sum_{\substack{e \in E_s, \\ \delta_e(\mathbf{s})=j}} j \cdot f_e(j) - \sum_{j=0}^k \sum_{\substack{e \in E_{s^*}, \\ \delta_e(\mathbf{s})=j}} f_e(\min\{k_e, j + 1\}) \\ &= \sum_{j \in [k]} \sum_{\substack{e \in E_s \setminus E_{s^*}, \\ \delta_e(\mathbf{s})=j}} j f_e(j) + \sum_{j \in [k]} \sum_{\substack{e \in E_{s^*}, \\ \delta_e(\mathbf{s})=j}} [j f_e(j) - f_e(\min\{k_e, j + 1\})] - \sum_{e \in E_{s^*} \setminus E_s} f_e(1) \end{aligned}$$

By Definition 1 and the fact that  $f_e$  is non-increasing, we have  $j \cdot f_e(j) \leq (\chi_f + 1)f_e(1)$  for all  $e \in E$  and  $j \in [k_e]$ . Using this and Definition 1, we get:

$$\begin{aligned} &\sum_{i \in [n]} u_i(\mathbf{s}) - \sum_{i \in [n]} u_i(\mathbf{s}_{-i}, \mathbf{s}_i^*) \\ &\leq \sum_{j \in [k]} \sum_{\substack{e \in E_s \setminus E_{s^*}, \\ \delta_e(\mathbf{s})=j}} (\chi_f + 1) \cdot f_e(1) + \sum_{j \in [k]} \sum_{\substack{e \in E_{s^*}, \\ \delta_e(\mathbf{s})=j}} \chi_f \cdot f_e(1) - \sum_{e \in E_{s^*} \setminus E_s} f_e(1) \\ &= \sum_{e \in E_s \setminus E_{s^*}} (\chi_f + 1) \cdot f_e(1) + \sum_{e \in E_{s^*} \cap E_s} \chi_f \cdot f_e(1) - \sum_{e \in E_{s^*} \setminus E_s} f_e(1) \\ &= \sum_{e \in E_s \setminus E_{s^*}} (\chi_f + 1) \cdot f_e(1) + \sum_{e \in E_{s^*} \cap E_s} (\chi_f + 1) \cdot f_e(1) - \sum_{e \in E_{s^*}} f_e(1) \\ &\leq (\chi_f + 1) \cdot \alpha_{\max} \cdot W(E_s) - \alpha_{\min} \cdot W(E_{s^*}). \end{aligned}$$

With [\(III\)](#) we get  $0 \leq \sum_{s \in S} p(s) ((\chi_f + 1) \cdot \alpha_{\max} \cdot W(E_s) - \alpha_{\min} \cdot W(E_{s^*})) = (\chi_f + 1) \cdot \alpha_{\max} \cdot \sum_{s \in S} p(s) \cdot W(E_s) - \alpha_{\min} \cdot W(E_{s^*})$ . Rearranging terms yields  $\frac{W(\mathbf{P})}{W(s^*)} = \frac{\sum_{s \in S} p(s) \cdot W(E_s)}{W(E_{s^*})} \geq \frac{\alpha_{\min}}{(\chi_f + 1) \cdot \alpha_{\max}}$ . The theorem follows since  $\mathbf{P}$  is an arbitrary Nash equilibrium.  $\square$

In the following, we construct a utility sharing function such that the corresponding lower bound in [Theorem 2](#) is maximized. Observe, that  $\chi_f$  is independent of the values for  $\alpha_{\min}$  and  $\alpha_{\max}$ . So, without loss of generality, we can restrict our attention to a symmetric utility sharing function  $f$ , where  $\alpha_{\min} = \alpha_{\max}$ . Our task is to construct a symmetric utility sharing function that solves the following optimization problem:

$$\begin{aligned} & \text{minimize} && \chi && (4) \\ & \text{subject to} && j \cdot f(j) - f(j+1) \leq \chi \cdot f(1) \quad \forall j \in [k-1] \\ & && (k-1)f(k) \leq \chi \cdot f(1) \end{aligned}$$

Replacing " $\leq$ " with " $=$ " yields a homogeneous system of linear equations. The values for  $\chi$  and the utility sharing function  $f$  in the following theorem correspond to the solution of this system where  $f(1) = 1$ .

**Theorem 3.** *Given  $k$  we can construct a symmetric utility sharing function  $f$ , such that  $\text{PoA}_f(k) \geq 1 - \frac{1}{\frac{1}{(k-1)(k-1)!} + \sum_{i=0}^{k-1} \frac{1}{i!}}$ .*

*Proof.* Given  $k$ , let  $f$  be the symmetric utility sharing function defined by

$$f(j) = (j-1)! \frac{\frac{1}{(k-1)(k-1)!} + \sum_{i=j}^{k-1} \frac{1}{i!}}{\frac{1}{(k-1)(k-1)!} + \sum_{i=1}^{k-1} \frac{1}{i!}} \quad \text{for all } j \in [k]. \quad (5)$$

It is not hard to check that  $f$  is a valid utility sharing function, i.e.  $f$  is non-increasing and  $j \cdot f(j) \leq 1$  for all  $j$ . Moreover,  $f$  satisfies the constraints in [\(4\)](#) for  $\chi = \frac{1}{\frac{1}{(k-1)(k-1)!} + \sum_{i=1}^{k-1} \frac{1}{i!}}$ . Recall that  $\alpha_{\min} = \alpha_{\max}$  for symmetric utility sharing functions. The claim follows by applying [Theorem 2](#).  $\square$

In order to construct the utility sharing function  $f$  in [Theorem 3](#) we need to know the maximum cardinality  $k$  over all elements. However, since the value for  $\chi$  from the proof of [Theorem 3](#) is increasing with  $k$ , we can get the same lower bound if each element  $e \in E$  is only aware of her own cardinality  $k_e$ . For this case, the cardinality dependent utility sharing function is defined by replacing  $k$  with  $k_e$  in [\(5\)](#). This implies:

**Corollary 1.** *There exists a cardinality dependent utility sharing function  $f$ , such that  $\text{PoA}_f(k) \geq 1 - \frac{1}{\frac{1}{(k-1)(k-1)!} + \sum_{i=0}^{k-1} \frac{1}{i!}}$ .*

Observe that the lower bounds on the price of anarchy in [Theorem 3](#) and [Corollary 1](#) match exactly the upper bound in [Theorem 1](#).

There might also be cases where we want to use a utility sharing function that works for all  $k$ . For example, neither  $k$  is known a priori nor the elements can determine their own cardinality. For such cases, we get:

**Theorem 4.** *There exists a symmetric utility sharing function  $f$  with  $\text{PoA}_f(k) \geq 1 - \frac{1}{e}$ , which works for arbitrary  $k$ .*

*Proof.* This follows by applying Theorem 3 for  $k \rightarrow \infty$ . In this case,  $f$  reduces to  $f(j) = \frac{(j-1)!}{e-1} \left[ e - \sum_{i=0}^{j-1} \frac{1}{i!} \right]$  for all  $j \in \mathbb{N}$ , and  $\chi = \frac{1}{e-1}$ .  $\square$

We close this section with an alternative lower bound on the price of anarchy that depends on the maximum dimension  $d_{\max} = \max_{i \in [n]} |S_i|$  over all players and the minimum cardinality of an element  $k_{\min}$ . For certain cases (e.g. if  $k_{\min} > d_{\max}$ ) this is better than the bound in Theorem 4.

**Theorem 5.** *Consider the class of covering games  $\mathcal{G}$  where  $k_e \geq k_{\min}$  for all elements  $e \in E$  and  $d_{\max} = \max_{i \in [n]} |S_i|$ . Let  $f$  be a cardinality dependent utility sharing function with  $f(1, l) = 1$  for all cardinalities  $l \in [k]$ . Then,  $\text{PoA}_f(\mathcal{G}) \geq \frac{k_{\min}}{d_{\max} - 1 + k_{\min}}$ .*

## 4 Approximation Algorithm

In the previous section, we have shown results on the price of anarchy for covering games. In this section, we want to use those results to construct a distributed, local-search approximation algorithm for the covering problem.

The idea of the algorithm is simple:

- Choose an appropriate utility sharing function,
- start with an arbitrary strategy profile, and
- let the players unilaterally perform selfish steps until a pure Nash equilibrium is reached.

The approximation ratio is the price of anarchy for the chosen utility sharing function. Rosenthal’s potential function [21] can be used to bound the number of selfish step until a pure Nash equilibrium is reached. Unfortunately, the utility sharing functions in Theorem 3 and Theorem 4 do not provide a sub-exponential bound on the number of selfish steps, since the increase in the potential due to a single selfish step can be arbitrary small.

To overcome this, we will design a new symmetric utility sharing function  $f$ , where for each element  $e \in E$  the players receive strictly positive payoff only if at most a constant number  $k'$  of players cover this element, i.e.  $f(j) = 0$  for all  $j > k'$ .

We will show that the right choice of  $f$  yields a  $(1 - \frac{1}{e} - \varepsilon)$ -approximation algorithm, where  $\varepsilon = \varepsilon(k') = o(1)$ .

**Theorem 6.** *For every constant  $\varepsilon > 0$  there exists a local-search approximation algorithm with approximation ratio  $(1 - \frac{1}{e} - \varepsilon)$  that uses at most  $\mathcal{O}(\frac{1}{\varepsilon} \cdot \log \log \frac{1}{\varepsilon}) \cdot W$  selfish steps.*



*Proof.* Let  $k' \in \mathbb{N}$  be some positive integer (to be determined later) and construct  $f$  as a solution to the following optimization problem:

$$\begin{aligned} & \text{minimize} && \chi && (6) \\ & \text{subject to} && j \cdot f(j) - f(j + 1) \leq \chi \cdot f(1) \quad \forall j \in [k' - 1] \\ & && k' \cdot f(k') \leq \chi \cdot f(1) \end{aligned}$$

For the solution of the corresponding homogeneous system of linear equations with  $f(1) = 1$  we get:  $f(j) = (j - 1)! \frac{\sum_{i=j}^{k'} \frac{1}{i!}}{\sum_{i=1}^{k'} \frac{1}{i!}}$  for all  $j \in [k']$  and  $f(j) = 0$  for  $j > k'$ . It is not hard to check that  $f$  is a valid utility sharing function, i.e.  $f$  is non-increasing and  $j \cdot f(j) \leq 1$  for all  $j$ .

Observe, that  $f(1) = 1, k' \cdot f(k') = \frac{1}{\sum_{i=1}^{k'} \frac{1}{i!}}$  and  $f(j + 1) = j! \frac{\sum_{i=j+1}^{k'} \frac{1}{i!} - \frac{1}{j!}}{\sum_{i=1}^{k'} \frac{1}{i!}} = j \cdot f(j) - \frac{1}{\sum_{i=1}^{k'} \frac{1}{i!}}$ . Thus,  $f$  satisfies the constraints in (6) for  $\chi = \frac{1}{\sum_{i=1}^{k'} \frac{1}{i!}}$ . Applying Theorem 2 yields

$$\text{PoA}_f \geq 1 - \frac{1}{\sum_{i=0}^{k'} \frac{1}{i!}} = 1 - \frac{1}{e} - \varepsilon(k'),$$

where  $\varepsilon(k') = \frac{1}{\sum_{i=0}^{k'} \frac{1}{i!}} - \frac{1}{e} = \frac{\sum_{i=k'+1}^{\infty} \frac{1}{i!}}{e \cdot \sum_{i=0}^{k'} \frac{1}{i!}} = \frac{\sum_{i=k'+1}^{\infty} \frac{k'!}{i!}}{e \cdot \sum_{i=0}^{k'} \frac{k'!}{i!}} = \Theta(\frac{1}{(k'+1)!})$ . So for every constant  $\varepsilon$  we choose  $k' = \Theta\left(\frac{\log(\frac{1}{\varepsilon})}{\log \log(\frac{1}{\varepsilon})}\right)$ . We will now bound the maximum number of selfish steps. To do so, we will first show an upper bound on Rosenthal’s potential function  $\Phi(\mathbf{s})$ . Afterwards, we show a lower bound on the increase in  $\Phi$  due to a selfish step. For any pure strategy profile  $\mathbf{s}$  we have

$$\Phi(\mathbf{s}) = \sum_{e \in E} \sum_{i=1}^{\delta_e(\mathbf{s})} f_e(i) \leq \sum_{e \in E} \sum_{i=1}^{k'} f_e(i) \leq \sum_{e \in E} \sum_{i=1}^{k'} \frac{1}{i} \cdot w_e = H(k') \cdot W,$$

where  $H(k')$  is the harmonic number of order  $k'$ .

Recall, that  $w_e \in \mathbb{N}$  for all  $e \in E$ . Moreover, for all  $j \in [k']$ ,  $f(j)$  is an integer multiple of  $\frac{1}{\sum_{i=1}^{k'} \frac{k'!}{i!}}$ . To see this multiply the numerator and denominator by  $k'!$  and observe that both become integer. So if a player improves, then she improves by at least  $\frac{1}{\sum_{i=1}^{k'} \frac{k'!}{i!}}$ . Using the property of Rosenthal’s *exact* potential function (cf. Sec. 2), it follows that each selfish step increases  $\Phi$  by at least  $\frac{1}{\sum_{i=1}^{k'} \frac{k'!}{i!}} = \frac{1}{[(e-1)k'!]}$ . So the number of selfish steps is upper bounded by  $\lfloor (e - 1)k'! \rfloor \cdot H(k') \cdot W = \mathcal{O}(k'! \cdot \log k') \cdot W = \mathcal{O}(\frac{1}{\varepsilon} \cdot \log \log \frac{1}{\varepsilon}) \cdot W$ .  $\square$

Theorem 6 implies a polynomial time  $(1 - \frac{1}{e} - \varepsilon)$ -approximation algorithm for the case that the total weight  $W$  is polynomially bounded. This includes the important case of unweighted elements, where  $w_e = 1$  for all  $e \in E$ . In the following theorem we show that this restriction on  $W$  is necessary, since for arbitrary weights the problem of computing a pure Nash equilibrium is  $\mathcal{PLS}$ -complete (see [16] for an introduction to the complexity class  $\mathcal{PLS}$ ).

**Theorem 7.** *Consider the class of covering games with arbitrary weights. Then, for every symmetric utility sharing function with  $f(1) > f(2)$ , it is  $\mathcal{PLS}$ -complete to compute a pure Nash equilibrium. This holds even for SAT-games.*

For the general case, we present a (centralized) approximation algorithm, which is based on LP-rounding and generalizes an algorithm for MaxSat [12].

**Theorem 8.** *There exists a (centralized) polynomial-time  $(1 - \frac{1}{e})$ -approximation algorithm for the general covering problem.*

## 5 Conclusion

In this paper we use game theoretic concepts for the design of new local search approximation algorithms for a very general covering problem. Our approach is to design player payoff functions that minimize the price of anarchy and guarantee that any sequence of unilateral improvements by the players is of polynomial length. For the covering problem this yields essentially the best possible approximation ratio.

For future work, we propose to study how far such ideas can be utilized to get new local search approximation algorithms also for other interesting optimization problems. Certainly, our approach will not always yield the best possible approximation ratio. This gives rise to the new interesting concept of *selfish approximation ratio*, i.e. the best possible approximation ratio that can be achieved by selfish players.

## Acknowledgement

We thank Adrian Vetta and Thomas Sauerwald for many helpful discussions and comments on the topic.

## References

1. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact Price of Anarchy for Polynomial Congestion Games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 218–229. Springer, Heidelberg (2006)
2. Alimonti, P.: New Local Search Approximation Techniques for Maximum Generalized Satisfiability Problems. In: Bonuccelli, M.A., Crescenzi, P., Petreschi, R. (eds.) CIAC 1994. LNCS, vol. 778, pp. 40–53. Springer, Heidelberg (1994)
3. Azar, Y., Jain, K., Mirrokni, V.S. (Almost) Optimal Coordination Mechanisms for Unrelated Machine Scheduling. In: Proc. of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008), pp. 323–332 (2008)
4. Bilò, V.: On Satisfiability Games and the Power of Congestion Games. In: Kao, M.-Y., Li, X.-Y. (eds.) AAIM 2007. LNCS, vol. 4508, pp. 231–240. Springer, Heidelberg (2007)
5. Chen, H.-L., Roughgarden, T., Valiant, G.: Designing Networks with Good Equilibria. In: Proc. of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008), pp. 854–863 (2008)

6. Christodoulou, G., Koutsoupias, E.: The Price of Anarchy of Finite Congestion Games. In: Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005), pp. 67–73 (2005)
7. 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)
8. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The Complexity of Pure Nash Equilibria. In: Proc. of the 36th Annual ACM Symposium on Theory of Computing (STOC 2004), pp. 604–612 (2004)
9. Feige, U.: A Threshold of  $\ln n$  for Approximating Set Cover. *Journal of the ACM* 45, 314–318 (1998)
10. Gairing, M., Schoppmann, F.: Total Latency in Singleton Congestion Games. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 381–387. Springer, Heidelberg (2007)
11. Giannakos, A., Gourvès, L., Monnot, J., Paschos, V.T.: On the Performance of Congestion Games for Optimum Satisfiability Problems. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 220–231. Springer, Heidelberg (2007)
12. Goemans, M.X., Williamson, D.P.: New  $\frac{3}{4}$ -Approximation Algorithms for the Maximum Satisfiability Problem. *SIAM Journal on Discrete Mathematics* 7, 656–666 (1994)
13. Hochbaum, D.S.: *Approximation Algorithms for NP-Hard Problems*. PWS Publishing Company (1997)
14. Hochbaum, D.S., Pathria, A.: Analysis of the Greedy Approach in Problems of Maximum  $k$ -coverage. *Naval Research Logistics* 45(6), 615–627 (1998)
15. Immorlica, N., Li, L., Mirrokni, V.S., Schulz, A.S.: Coordination Mechanisms for Selfish Scheduling. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 55–69. Springer, Heidelberg (2005)
16. Johnson, D.S., Papadimitriou, C.H., Yannakakis, M.: How Easy is Local Search? *Journal of Computer and System Sciences* 37(1), 79–100 (1988)
17. Khanna, S., Motwani, R.: On Syntactic versus Computational Views of Approximability. *SIAM Journal on Computing* 28(1), 164–191 (1998)
18. Koutsoupias, E., Papadimitriou, C.H.: Worst-Case Equilibria. In: Meinel, C., Tiison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
19. Mavronicolas, M., Monien, B., Wagner, K.W.: Weighted Boolean Formula Games. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 469–481. Springer, Heidelberg (2007)
20. Nash, J.F.: Non-Cooperative Games. *Annals of Mathematics* 54(2), 286–295 (1951)
21. Rosenthal, R.W.: A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
22. Roughgarden, T.: Intrinsic Robustness of the Price of Anarchy. In: Proc. of the 41st Annual ACM Symposium on Theory of Computing (STOC 2009), pp. 513–522 (2009)
23. Vazirani, V.V.: *Approximation Algorithms*. Springer, Heidelberg (2003)
24. Vetta, A.: Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions. In: Proc. of the 43rd Annual Symposium on Foundations of Computer Science (FOCS 2002), pp. 416–425 (2002)

# Contract Auctions for Sponsored Search

Sharad Goel, Sébastien Lahaie, and Sergei Vassilvitskii

Yahoo! Research,  
111 West 40th Street, New York, New York, 10018

**Abstract.** In sponsored search auctions advertisers typically pay a fixed amount per click that their advertisements receive. In particular, the advertiser and the publisher enter into a contract (e.g., the publisher displays the ad; the advertiser pays the publisher 10 cents per click), and each party's subjective value for such a contract depends on their estimated click-through rates (CTR) for the ad. Starting from this motivating example, we define and analyze a class of contract auctions that generalize the classical second price auction. As an application, we introduce impression-plus-click pricing for sponsored search, in which advertisers pay a fixed amount per impression plus an additional amount if their ad is clicked. Of note, when the advertiser's estimated CTR is higher than the publisher's estimated CTR, both parties find negative click payments advantageous, where the advertiser pays the publisher a premium for the impression but the publisher then pays the advertiser per click.

## 1 Introduction

In the classical sealed-bid second-price auction, bidders report their value for the auctioned good, and the winner is the bidder with the highest reported value. Incentive compatibility is achieved by charging the winner the least amount for which they would have still won the auction (i.e., the winner pays the second highest bid). In contrast, consider a typical sponsored search auction where, for simplicity, we assume bidders compete for a single available impression: Advertisers report their value-per-click; the winner is the bidder from whom the publisher expects to receive the most revenue; and the winning bidder pays the least amount per click for which they would have still won the auction. While sponsored-search auctions are conceptually similar to traditional second price auctions, there is a key difference: Goods in traditional auctions are exchanged for deterministic payments, and in particular, the value of these payments is identical to the bidder and the auctioneer; in sponsored search auctions, impressions are exchanged for stochastic payments, and the value of such payments to the publisher and the advertiser depends on their respective estimated click-through rates (CTR). For example, if the advertiser's estimated CTR is higher than the publisher's, then the advertiser would expect to pay more than the publisher would expect to receive.

Starting from this motivating example of sponsored search, we define and develop a framework for contract auctions that generalize the second price auction.

We consider arbitrary agent valuations over a space of possible contracts; in particular, valuations may diverge for reasons other than mismatched probability estimates. As an application, we introduce impression-plus-click (IPC) sponsored search auctions, in which advertisers pay a fixed amount per impression and make an additional payment per click. Interestingly, when the advertiser’s estimated CTR is higher than the publisher’s estimate, both parties prefer negative click payments—or *paid per click* pricing: The advertiser pays the publisher a premium for the impression, and the publisher then pays the advertiser per click.

In the remainder of this introduction we review sponsored search auctions and related work. General contract auctions are developed in Section 2, and a dominant strategy incentive compatible mechanism is proposed. Impression-plus-click sponsored search auctions are introduced in Section 3. In Section 4 we analyze an impression-or-click auction, and consider connections to the hybrid auction model of Goel & Munagala [4]. We conclude in Section 5 by discussing potential offline applications of this work, including applications to insurance, book publication and executive compensation. Due to space constraints we omit some proofs from this version.

## 1.1 Background and Related Work

Sponsored search is the practice of auctioning off ad placement next to web search results; advertisers pay the search engine when their ads are clicked. These ad auctions are responsible for the majority of the revenue of today’s leading search engines [10]. Edelman et al. [3] and Varian [15] provide the standard model for sponsored search auctions and analyze its equilibrium properties (see also Lahaie et al. [11] for a survey of the literature in this area). We do not provide a description of this model here because our contract auction abstracts away from its details in order to cover pricing schemes beyond per-click or per-impression.

Harrenstein et al. [7] recently and independently developed the *qualitative Vickrey auction*, a mechanism similar to the general contract auction presented here. The primary differences between their work and ours concern subtleties in the bidding language, the tie-breaking rules, and the assumptions guaranteeing truthfulness. In this paper we detail our interpretation and results for contract auctions; our main contribution, however, is applying this framework to sponsored search, and in particular introducing impression-plus-click pricing.

Truthfulness under the standard model of sponsored search is well understood [1]. In mechanism design more generally, Holmstrom [8] characterizes truthful payment rules for type spaces that are smoothly path-connected (see also [12, 13]). In contrast, our truthfulness result for contract auctions does not assume any topology on the type space. Instead it is a consequence of the particular structure of the outcome space (the auctioneer may contract with only one agent) together with a novel consistency condition between the auctioneer and agents’ preferences.

Contract auctions generalize the single-item Vickrey auction [16], but are conceptually distinct from the well-known Vickrey-Clark-Groves (VCG) mechanism [2, 6]. An intuitive interpretation of the VCG mechanism is that it charges

each agent the externality that the agent imposes on others; thus, the mechanism only applies when utility is transferable between agents through payments. This is not possible when, for instance, the agents and auctioneer disagree on click-through rates, because there can be no agreement on how to quantify the externality. There are many reasons why disagreement might arise: clicks are low probability events whose distributions are hard to model [14], and the auctioneer and advertisers may disagree on which clicks were valid [5, 9].

The hybrid auction of Goel and Munagala [4] is a notable departure from the basic sponsored search model in that it attempts to reconcile differing publisher and advertiser click-through estimates. In a hybrid auction advertisers place per-click bids as well as per-impression bids, and the auctioneer then chooses one of the two pricing schemes. Goel and Munagala [4] show that, besides being truthful, their hybrid auction has many advantages over simple per-click keyword auctions. The auction allows advertisers to take into account their attitudes towards risk and may generate higher revenue, among other nice properties. The consistency condition given in this work distills the reason behind truthfulness in the hybrid auction, and our contract auction leads to variants and generalizations of the hybrid auction to multiple pricing schemes beyond CPC and CPM (e.g., CPA for any kind of action).

## 2 Contract Auctions

We define and develop an incentive compatible mechanism for contract auctions where agents have valuations over an arbitrary space of possible contracts. Suppose there are  $N$  agents  $A_1, \dots, A_N$  and finite sets  $C_1, \dots, C_N$  that denote the set of potential contracts each agent could enter into. Agents have valuation functions  $v_i : C_i \mapsto \mathbb{R}$  for their respective contracts, and the auctioneer's value for each contract is given by  $v_i^A : C_i \mapsto \mathbb{R}$ . *Contracts*, in this setting, are nothing more than abstract objects for which each party has a value. The auctioneer is to enter into a single contract, and our goal is to design a framework to facilitate this transaction.

The valuation functions are intended to represent purely subjective utilities, based, for example, on private beliefs or simply taste. In this sense, each agent values contracts in their own “currency,” which cannot directly be converted into values for other agents. We require that preferences be *consistent* in the following sense: Among contracts acceptable to a given bidder (i.e., those contracts for which the bidder has non-negative utility), the highest value contract to the auctioneer is one for which the bidder has zero utility. This statement is formalized by Definition [1](#).

**Definition 1.** *In the setting above, we say agent  $v_i$  and the auctioneer have consistent valuations if for each  $c_1 \in C_i$  with  $v_i(c_1) > 0$ , there exists  $c_2 \in C_i$  such that  $v_i(c_2) \geq 0$  and  $v_i^A(c_2) > v_i^A(c_1)$ .*

Consistency is equivalent to the following property:

$$\max_{\{c: v_i(c) \geq 0\}} v_i^A(c) > \max_{\{c: v_i(c) > 0\}} v_i^A(c).$$

We note that consistency is a weak restriction on the structure of valuations. In particular, if contracts include a “common currency” component, for which bidders and the auctioneer have an agreed upon value, then valuations are necessarily consistent.

Under this assumption of consistency, Mechanism  $\square$  defines a dominant strategy incentive-compatible mechanism for contract auctions. First, bidders report their valuation function to the auctioneer. In the applications we consider, this entails reporting a small set of parameters which defines the valuation function over the entire contract space. Next, among contracts for which agents have non-negative utility (i.e., “acceptable” or “individually-rational” contracts), the auctioneer identifies the contract for which it has maximum value; the winner of the auction is the bidder who submitted this maximum value acceptable contract. Finally, the auctioneer and the winner enter into the best contract from the winner’s perspective for which it would have still won the auction.

---

**Mechanism 1. A General Contract Auction**

---

- 1: Each agent  $A_1, \dots, A_N$  reports a valuation function  $\tilde{v}_i$ .
- 2: For  $1 \leq i \leq N$ , let  $S_i = \{c \in C_i \mid \tilde{v}_i(c) \geq 0\}$  be the set of contracts for which agent  $A_i$  claims to have non-negative valuation, and define

$$R_i = \max_{S_i} v_i^A(c) \tag{1}$$

to be the maximum value the auctioneer can achieve from each agent among these purportedly acceptable contracts.

- 3: Fix  $h$  so that  $R_{h(1)} \geq R_{h(2)} \geq \dots \geq R_{h(N)}$ , and let

$$S = \left\{ c \in C_{h(1)} \mid v_{h(1)}^A(c) \geq R_{h(2)} \right\}.$$

With agent  $A_{h(1)}$ , the auctioneer enters into any contract  $c^*$  such that

$$c^* \in \arg \max_S \tilde{v}_{h(1)}(c).$$


---

**Theorem 1.** *In the setting above, suppose agents have consistent valuations. Then Mechanism  $\square$  is dominant strategy incentive compatible.*

*Proof.* Fix an agent  $A_i$  and consider its strategy. Let  $R_{-i} = \max_{j \neq i} R_j$  where  $R_j$  is defined as in  $\square$ . If  $A_i$  were to win the auction, then it would necessarily enter into a contract among those in the set  $M_i = \{c \in C_i \mid v_i^A(c) \geq R_{-i}\}$ .

Suppose  $A_i$  has strictly positive valuation for some contract  $c_1 \in M_i$  (i.e.,  $\max_{M_i} v_i(c) > 0$ ). Then by the assumption of consistent valuations, there exists a contract  $c_2$  such that  $v_i(c_2) \geq 0$  and  $v_i^A(c_2) > v_i^A(c_1) \geq R_{-i}$ .

In particular, if  $A_i$  truthfully reports its valuation function, then we would have  $R_i \geq v_i^A(c_2) > R_{-i}$ , and hence  $A_i$  would win the auction. Furthermore, in this case the best  $A_i$  could do is to enter into a contract in the set  $\arg \max_{M_i} v_i(c)$ . Again, truthful reporting ensures that this optimal outcome occurs.

Now suppose  $\max_{M_i} v_i(c) \leq 0$ . In this case  $A_i$  has no possibility of positive gain, whether or not it wins the auction. However, by reporting truthfully, if  $A_i$  does win the auction the final contract would be selected from the set  $S_i = \{c \in C_i \mid v_i(c) \geq 0\}$ . That is, truthful reporting ensures that  $A_i$  achieves (the optimal) zero gain.  $\square$

We next show that the consistency condition plays a crucial role in achieving incentive compatibility by exhibiting an example with inconsistent valuations where truth telling is not a dominant strategy. Suppose the auctioneer has one item for sale and there are two agents  $A_i$  and  $A_j$ . Agent  $A_i$ , Irene, values the item at \$4 but only has \$2 to spend. There are three contracts she can enter into,  $c_1^i, c_2^i, c_3^i$ , intuitively buying the item for \$1, \$2, and \$3, resulting in utilities of 3, 2, and  $-1$ , the latter being negative since Irene has a limited budget of \$2. Agent  $A_j$ , Juliet, values the item at \$2 and has \$2 to spend. She can enter into three similar contracts,  $c_1^j, c_2^j$  and  $c_3^j$ , resulting in utilities of 1, 0 and  $-1$  respectively. From the auctioneer's point of view, his utility is the revenue,  $v^A(c_x^i) = v^A(c_x^j) = x$  for any  $x \in \{1, 2, 3\}$ . If the agents report their valuations truthfully, then  $R_i = R_j = 2$  and the auctioneer must break the tie. Unless the tie is broken deterministically in favor of Irene, she has an incentive to lie. Suppose she reports her valuation for  $c_3^i$  to be 1, pretending that she has enough money to afford the item. In that case  $R_i = 3$ , and  $R_j = 2$ , so Irene wins the item; but she can select any outcome so long as the auctioneer's utility is at least  $R_j = 2$ . She chooses  $c_2^i$ , which has a positive utility to her, but still makes \$2 for the auctioneer. Essentially, because the utilities of Irene and the auctioneer are not consistent, Irene can bluff to win the item.

*Remark 1.* In the above we have assumed the contract spaces  $C_i$  are finite. This restriction is imposed only to ensure the *maximum* operation is well-defined in Mechanism [□](#). We implicitly relax this condition in the following discussion, as it is clear the relevant maxima exist despite having infinite contract spaces.

Mechanism [□](#) generalizes the usual sealed-bid second-price auction. To see this, take  $C_i = \mathbb{R}$ , and let the contract  $p \in \mathbb{R}$  indicate agent  $A_i$ 's obligation to purchase the auctioned good at price  $p$ . If agent  $A_i$  values the good at  $w_i$ , then its value over contracts is given by  $v_i(p) = w_i - p$ , and in particular, its preferences over contracts is parametrized by  $w_i \in \mathbb{R}$ . The auctioneer has valuation  $v_i^A(p) = p$ . Now, letting  $\tilde{w}_i$  be  $A_i$ 's reported valuation, we have  $R_i = \tilde{w}_i$ . Furthermore,

$$S = \left\{ c \in C_{h(1)} \mid v_{h(1)}^A(c) \geq R_{h(2)} \right\} = [\tilde{w}_{h(2)}, \infty)$$

and so  $\arg \max_S \tilde{v}_{h(1)} = \tilde{w}_{h(2)}$ . That is, agent  $A_{h(1)}$  enters into the contract  $\tilde{w}_{h(2)}$ , agreeing to pay the second highest bid for the good.

### 3 The Impression-Plus-Click Pricing Model

We now consider a specific application of contract auctions for sponsored search: impression-plus-click pricing. For a given impression, define a contract  $(r_s, r_f) \in \mathbb{R}^2$  to require the advertiser pay  $r_s$  if a click occurs and  $r_f$  if no click occurs. This is a complete pricing scheme if the advertiser values only impressions and clicks. We note that so-called “brand advertisers” often have significant utility for simply displaying their ads, regardless of whether or not their ads are clicked. These contracts are equivalently parametrized by  $(r_m, r_c) \in \mathbb{R}^2$ , where the advertiser pays  $r_m$  per impression and an *additional*  $r_c$  per click. Using this latter,



additive, notation, an impression-plus-click (IPC) contract is represented as a point in the CPM-CPC price plane. A priori there are no restrictions on these contracts (e.g., one or both coordinates could be negative).

### 3.1 Contract Preferences

Suppose an advertiser  $A_i$  values an impression, regardless of whether it receives a click, at  $m_i \geq 0$ , values a click at  $w_i \geq 0$ , and estimates its CTR to be  $p_i > 0$ . Then, assuming risk neutrality, its value for the IPC contract  $(r_m, r_c)$  is

$$v_i(r_m, r_c) = (m_i + p_i w_i) - (r_m + p_i r_c).$$

Observe that the contract preferences of  $A_i$  are equivalent to those of an advertiser who values clicks at  $w_i + m_i/p_i$  and has no inherent value for impressions. Consequently, without loss of generality, we need only consider the case  $m_i = 0$ . We thus have the simplified expression:  $v_i(r_m, r_c) = p_i w_i - (r_m + p_i r_c)$ . The level curves of  $v_i$  are linear with slope  $-1/p_i$ :

$$\{(r_m, r_c) : v_i(r_m, r_c) = C\} = \{(r_m, K - r_m/p_i) : r_m \in \mathbb{R}\} \tag{2}$$

where  $K = w_i - C/p_i$ .

We suppose the advertiser requires limited liability in the following sense. For advertiser specific constants  $CPM_i > 0$  and  $CPC_i > 0$ , we assume the advertiser has strictly negative utility for any contract  $(r_m, r_c)$  with either  $r_m > CPM_i$  or  $r_c > CPC_i$ ; aside from this caveat, the advertiser is risk-neutral. In other words, advertisers effectively have a maximum amount they are willing to spend on clicks and impressions, but otherwise they are risk neutral.

The utility function of each advertiser  $A_i$  can be derived from four numbers: its value-per-click  $w_i$ , its estimated CTR  $p_i$ , and its price caps  $CPM_i$  and  $CPC_i$ . Equivalently,  $A_i$ 's utility function is determined by the two contracts

$$\{(r_m^i, CPC_i), (CPM_i, r_c^i)\}$$

where  $r_m^i = p_i(w_i - CPC_i)$  and  $r_c^i = w_i - CPM_i/p_i$ . These two IPC contracts lie on  $A_i$ 's zero-utility level line; that is,

$$v_i(r_m^i, CPC_i) = 0 \quad v_i(CPM_i, r_c^i) = 0.$$


Moreover, these contracts are the extreme points on this zero-utility line (i.e., they push up against the price caps). Observe that  $w_i$  is the  $y$ -intercept of the line through these two contract points, and  $p_i = (CPM_i - r_m^i) / (CPC_i - r_c^i)$  is the negative reciprocal of the slope of this line. Furthermore, the space of advertiser utility functions is parametrized by the set of contract pairs

$$U = \{(r_{m,1}, r_{c,1}), (r_{m,2}, r_{c,2}) \mid r_{m,1} \leq 0 < r_{m,2}, r_{c,2} \leq 0 < r_{c,1}\}. \tag{3}$$

From the (risk-neutral) publisher's perspective, the utility of a contract  $(r_m, r_c)$  entered into with advertiser  $A_i$  is

$$v_i^A(r_m, r_c) = r_m + p_i^A r_c$$

where  $p_i^A$  is the publisher's estimated CTR of the advertiser's ad. Figure 1(a) illustrates the contract preferences for an advertiser and a publisher.



(a) The solid red and black lines indicate the advertiser’s and publisher’s level curves in the CPM-CPC price plane. Here the publisher’s CTR estimate is lower than the advertiser’s.

(b) The star indicates a winning pure per-impression bid, the red line is the publisher’s  $R_*$  level line, and the two dots indicate the final contracts.

**Fig. 1.** Publisher and advertiser contract preferences

### 3.2 Designing the Impression-Plus-Click Auction

Given the advertiser and publisher preferences outlined in Section 3.1, we next apply Theorem 1 to design a dominant strategy incentive compatible IPC auction for sponsored search. We start with two preliminary lemmas.

**Lemma 1.** *Assume the setting and notation of Mechanism 1, and the advertiser and publisher preferences of Section 3.1. Then the agents have consistent valuations with the publisher. Furthermore, letting  $\{(\tilde{r}_m^i, \widetilde{CPC}_i), (\widetilde{CPM}_i, \tilde{r}_c^i)\}$  denote  $A_i$ ’s reported preferences, we have*

$$R_i = \begin{cases} \widetilde{CPM}_i + p_i^A \tilde{r}_c^i & \text{if } p_i^A \leq \tilde{p}_i \\ \tilde{r}_m^i + p_i^A \widetilde{CPC}_i & \text{if } p_i^A \geq \tilde{p}_i \end{cases}$$

where  $\tilde{p}_i = (\widetilde{CPM}_i - r_m^i) / (\widetilde{CPC}_i - r_c^i)$  is  $A_i$ ’s inferred (subjective) CTR.

*Proof.* The level curve  $L_0$  on which the advertiser has (true) zero utility is given by the line segment

$$L_0 = \{(r_m, r_c) \mid r_m + p_i r_c = p_i w_i, r_m \leq CPM_i, r_c \leq CPC_i\} \\ = \{(r_m, w_i - r_m/p_i) \mid p_i(w_i - CPC_i) \leq r_m \leq CPM_i\}$$

and the set  $S_i$  on which the advertiser has non-negative utility is given by the points below  $L_0$ :

$$S_i = \{(r_m, r_c) \mid \exists (r_m^*, r_c^*) \in L_0 \text{ such that } r_m \leq r_m^* \text{ and } r_c \leq r_c^*\}.$$

If  $(r_m, r_c) \in S_i \setminus L_0$  (i.e., if  $v_i(r_m, r_c) > 0$ ), then there exists  $(r_m^*, r_c^*) \in L_0$  such that either  $r_m^* > r_m$  or  $r_c^* > r_c$ . In either case,  $v_i^A(r_m^*, r_c^*) > v_i^A(r_m, r_c)$ , and so  $A_i$  and the publisher have consistent valuations.

To compute  $R_i$ , we first assume agent  $A_i$  truthfully reports its preferences. Consistent valuations implies that the publisher achieves its maximum value,

among contracts in  $S_i$ , on the set  $L_0$  where the advertiser has zero utility. For  $(r_m, r_c) \in L_0$ ,

$$v_i^A(r_m, r_c) = r_m + p_i^A r_c = r_m + (w_i - r_m/p_i)p_i^A = w_i p_i^A + r_m (1 - p_i^A/p_i).$$

Now note that (3.2) is an increasing function of  $r_m$  for  $p_i < p_i^A$ , and a decreasing function of  $r_m$  for  $p_i > p_i^A$ . Consequently, the maximum is achieved at the end-points of  $L_0$ . To extend to the case where  $A_i$  does not necessarily report truthfully, we need only replace  $A$ 's actual preferences with its reported preferences.  $\square$

**Lemma 2.** *Assume the setting and notation of Lemma 1. Fix  $1 \leq i \leq N$  and  $R_* \leq R_i$ . Then for  $S = \{(r_m, r_c) \in C_i \mid v_i^A(r_m, r_c) \geq R_*\}$  we have*

$$\arg \max_S \tilde{v}_i(r_m, r_c) = \begin{cases} \left( \widehat{CPM}_i, (R_* - \widehat{CPM}_i)/p_i^A \right) & \text{if } p_i^A < \tilde{p}_i \\ \left( R_* - p_i^A \widehat{CPC}_i, \widehat{CPC}_i \right) & \text{if } p_i^A > \tilde{p}_i \\ T & \text{if } p_i^A = \tilde{p}_i \end{cases}$$

where

$$T = \left\{ (r_m, (R_* - r_m)/p_i^A) \mid R_* - p_i^A \widehat{CPC}_i \leq r_m \leq \widehat{CPM}_i \right\}.$$

*Proof.* First note that since  $R_* \leq R_i$ ,  $\max_S \tilde{v}_i \geq 0$ . Now, the level curve  $L_A$  on which  $v_i^A(r_m, r_c) = R_*$  is given by

$$L_A = \{(r_m, r_c) \mid r_m + p_i^A r_c = R_*\} = \{(r_m, (R_* - r_m)/p_i^A) \mid r_m \in \mathbb{R}\}.$$

Furthermore,  $v_i^A(r_m, r_c) > R_*$  if and only if  $(r_m, r_c)$  lies above this line. That is,  $v_i^A(r_m, r_c) > R_*$  if and only if there exists a contract  $(r_m^*, r_c^*) \in L_A$  such that either  $r_m \geq r_m^*$  and  $r_c > r_c^*$ , or  $r_m > r_m^*$  and  $r_c \geq r_c^*$ . In either case,  $\tilde{v}_i(r_m^*, r_c^*) > \tilde{v}_i(r_m, r_c)$  and so  $\arg \max_S \tilde{v}_i \subseteq L_A$ . Since  $\max_S \tilde{v}_i \geq 0$ , we can further restrict ourselves to the set

$$\begin{aligned} T &= L_A \cap (-\infty, \widehat{CPM}_i] \times (-\infty, \widehat{CPC}_i] \\ &= \left\{ (r_m, (R_* - r_m)/p_i^A) \mid R_* - p_i^A \widehat{CPC}_i \leq r_m \leq \widehat{CPM}_i \right\}. \end{aligned}$$

For  $(r_m, r_c) \in T$ , and  $\tilde{w}_i$  indicating  $A_i$ 's inferred value per click, we have

$$\begin{aligned} \tilde{v}_i(r_m, r_c) &= \tilde{w}_i \tilde{p}_i - [r_m + \tilde{p}_i r_c] = \tilde{w}_i \tilde{p}_i - [r_m + (R_* - r_m)\tilde{p}_i/p_i^A] \\ &= \tilde{w}_i \tilde{p}_i - R_* \tilde{p}_i/p_i^A + r_m (\tilde{p}_i/p_i^A - 1). \end{aligned} \tag{4}$$

The result now follows by noting that (4) is increasing in  $r_m$  for  $p_i^A < \tilde{p}_i$ , decreasing for  $p_i^A > \tilde{p}_i$ , and constant for  $p_i^A = \tilde{p}_i$ .  $\square$

Together with Lemmas 1 and 2, the general contract auction of Mechanism 1 leads to the impression-plus-click auction described by Mechanism 2. First, each advertiser submits two contracts—ostensibly specifying its entire utility function. The publisher then computes its own utility for each of these  $2N$  contracts, and

the winner of the auction is the agent who submitted the contract with the highest value to the publisher. The “second-highest value” is the value of the best contract (again from the publisher’s perspective) among those submitted by the losing bidders. To determine the actual contract entered into, we consider two cases. If the highest value contract has higher CPM than the winner’s other bid, then the final contract is determined by decreasing the CPC on the highest value contract until the publisher’s value for that contract is equal to the second highest value. Analogously, if the highest value contract has lower CPM than the winner’s other bid, the final contract is determined by decreasing the CPM of the highest value contract.

---

**Mechanism 2.** An Impression-Plus-Click Auction

---

- 1: Advertisers  $A_1, \dots, A_N$  each report their valuation functions, encoded by the pair of extremal contracts as described in Section 3.1

$$\tilde{v}_i = \left\{ \left( r_{m,1}^i, r_{c,1}^i \right), \left( r_{m,2}^i, r_{c,2}^i \right) \right\},$$

where  $r_{m,1}^i \leq 0 < r_{m,2}^i$  and  $r_{c,2}^i \leq 0 < r_{c,1}^i$ .

- 2: For each report  $\tilde{v}_i$  define

$$R_i = \max \left( v_i^A \left( r_{m,1}^i, r_{c,1}^i \right), v_i^A \left( r_{m,2}^i, r_{c,2}^i \right) \right) = \max \left( r_{m,1}^i + r_{c,1}^i p_i^A, r_{m,2}^i + r_{c,2}^i p_i^A \right).$$

- 3: Fix  $h$  so that  $R_{h(1)} \geq R_{h(2)} \geq \dots \geq R_{h(N)}$ . The publisher enters into a contract with agent  $A_{h(1)}$ . The final contract  $c^*$  is determined as follows:

$$c^* = \begin{cases} \left( r_{m,2}^{h(1)}, \left( R_{h(2)} - r_{m,2}^{h(1)} \right) / p_{h(1)}^A \right) & \text{if } R_{h(1)} = v_{h(1)}^A \left( r_{m,2}^{h(1)}, r_{c,2}^{h(1)} \right) \\ \left( R_{h(2)} - p_{h(1)}^A r_{c,1}^{h(1)}, r_{c,1}^{h(1)} \right) & \text{otherwise} \end{cases}$$


---

**Theorem 2.** Consider the setting and notation of Mechanism 2 with the advertiser and publisher preferences of Section 3.7. Then

1.  $c^* \leq \left( r_{m,1}^{h(1)}, r_{c,1}^{h(1)} \right)$  or  $c^* \leq \left( r_{m,2}^{h(1)}, r_{c,2}^{h(1)} \right)$ , where the inequalities hold coordinate-wise.
2. The mechanism is dominant strategy incentive compatible. That is, it is a dominant strategy for each advertiser  $A_i$  to truthfully report

$$\left\{ \left( r_m^i, CPC_i^i \right), \left( CPM_i^i, r_c^i \right) \right\}.$$

## 4 The Impression-or-Click Pricing Model

With impression-plus-click pricing, advertisers pay publishers for each impression, and then pay an additional amount if their ad is clicked. The hybrid sponsored search auction of Goel & Munagala [4] can be thought of as *impression-or-click* (IOC) pricing. That is, the final selected contract is guaranteed to be either pure per-impression or pure per-click, but it is not known which it will

be until all bids have been submitted. The hybrid auction, as shown below, is equivalent to a special case of the general contract auction with the contract spaces restricted to the axes of the CPM-CPC plane:

$$C_i = \{(r_m, 0) \mid r_m \in \mathbb{R}\} \cup \{(0, r_c) \mid r_c \in \mathbb{R}\}. \tag{5}$$

Suppose both advertisers and publishers are risk neutral. As before, let  $p_i$  denote advertiser  $A_i$ 's subjective click-through rate, let  $p_i^A$  denote the publisher's estimated click-through rate for an impression awarded to  $A_i$ , and let  $w_i$  denote  $A_i$ 's value for a click. Then  $A_i$  has zero utility for the two contracts  $(p_i w_i, 0)$  and  $(0, w_i)$ . By the assumption of risk neutrality, these two contracts completely determine  $A_i$ 's preferences over all contracts. Hence,  $A_i$  can communicate its preferences by reporting the two numbers  $CPM_i = p_i w_i$  and  $CPC_i = w_i$ , corresponding to the maximum it is willing to pay for a per-impression and a per-click contract, respectively. The resulting IOC auction is outlined in Mechanism 3. Details of its derivation are straightforward and are omitted for space constraints.

---

**Mechanism 3.** An Impression-Or-Click Auction (Goel & Munagala)

---

- 1: Advertisers  $A_1, \dots, A_N$  each report their valuation functions, encoded by the constants  $\widetilde{CPM}_i, \widetilde{CPC}_i > 0$ .
- 2: For each report, define  $R_i = \max(CPM_i, p_i^A CPC_i)$ .
- 3: Fix  $h$  so that  $R_{h(1)} \geq R_{h(2)} \geq \dots \geq R_{h(N)}$ . Then the publisher enters into a contract with agent  $A_{h(1)}$ . The final contract  $c^*$  is determined as follows:

$$c^* = \begin{cases} (0, R_*/p_i^A) & \text{if } R_{h(1)} = p_{h(1)}^A CPC_{h(1)} \\ (R_*, 0) & \text{otherwise} \end{cases}$$


---

Although the hybrid and general contract auctions are equivalent when advertiser preferences are restricted to the CPM-CPC axis, they may lead to different outcomes when preferences are defined over the entire plane. Consider the IPC auction setting of Section 3, where we now assume that  $CPC_i = w_i$  and  $CPM_i = w_i p_i$ . That is, the most advertiser  $A_i$  is willing to pay per click or per impression is, respectively, its true per click value  $w_i$  and its true per impression value  $w_i p_i$ . In particular,  $A_i$  will not pay more than  $w_i$  per click even if it is compensated via negative per-impression payments. In this case, the two extremal contracts that define  $A_i$ 's utility function over the CPM-CPC plane are  $(CPM_i, 0)$  and  $(0, CPC_i)$ . With such a preference profile, we show that advertisers prefer the IPC auction over the IOC auction, and publishers are ambivalent between the two.

In both the IOC and IPC auctions, it is a dominant strategy to truthfully reveal ones' preferences: In the IOC auction advertisers report their maximum per-impression and per-click payments  $CPM_i$  and  $CPC_i$ ; in the IPC auction they report their pair of extremal contracts  $\{(CPM_i, 0), (0, CPC_i)\}$ . From the publisher's perspective, for each agent  $A_i$ ,  $R_i$  is the same in both auctions. Consequently, the winner of the auction is the same under either mechanism, and moreover, the expected (subjective) revenue  $R_*$  of the publisher is also the same. The publisher is thus ambivalent between the IOC and IPC auction designs.

From the advertisers’ view, however, the situation is quite different. Specifically, let  $c_{\text{IPC}}^*$  and  $c_{\text{IOC}}^*$  denote the final contract entered into by the winner  $A_{h(1)}$  under each mechanism. Then

$$v_{h(1)}(c_{\text{IPC}}^*) = \max_{Q_1} v_{h(1)} \quad v_{h(1)}(c_{\text{IOC}}^*) = \max_{Q_2} v_{h(2)}$$

where

$$Q_1 = \left\{ (r_m, r_c) \in \mathbb{R}^2 \mid v_{h(1)}^A(r_m, r_c) \geq R_* \right\}$$

$$Q_2 = \left\{ (r_m, r_c) \in \mathbb{R}^2 \mid v_{h(1)}^A(r_m, r_c) \geq R_*, \min(r_m, r_c) = 0 \right\}$$

That is, the IPC contract is optimized over the entire plane, whereas the IOC contract is optimized only over the axes. In particular,  $v_{h(1)}(c_{\text{IPC}}^*) \geq v_{h(1)}(c_{\text{IOC}}^*)$ . Since  $v_i^A(c_{\text{IPC}}^*) = v_i^A(c_{\text{IOC}}^*)$ , the line drawn between these two contracts has slope  $-1/p_i^A$  (as shown in Section 3.1). Furthermore, since  $v_i(c_{\text{IPC}}^*) = v_i(c_{\text{IOC}}^*)$  if and only if the line between the contracts has slope  $-1/p_i$ , we have  $v_{h(1)}(c_{\text{IPC}}^*) > v_{h(1)}(c_{\text{IOC}}^*)$  provided that  $p_i^A \neq p_i$ . Hence, in this setting, advertisers typically prefer the IPC over the IOC auction.

The distinction between the IPC and IOC settlement mechanisms is illustrated in Figure 1(b). When  $p_{h(1)}^A < p_{h(1)}$ , the publisher prefers (under both mechanisms) the winning advertiser’s pure per-impression bid  $\text{CPM}_{h(1)}$  over its pure per-click bid  $\text{CPC}_{h(1)}$ . In this case, the final IOC contract is a pure per-impression contract, where the per-impression payment is reduced from  $\text{CPM}_{h(1)}$  to an amount such that the ultimate value of the contract to the publisher is  $R_*$ . In contrast, the final IPC contract has the advertiser still paying  $\text{CPM}_{h(1)}$  per impression, but a “discount” is given to the advertiser via negative click payments (i.e., the publisher pays the advertiser for each click). This negative click payment is calculated so that the final value of the contract to the publisher is again  $R_*$ . The final contract in either auction lies on the  $R_*$  level curve of the publisher: In the IOC auction, the pure-impression contract  $\text{CPM}_{h(1)}$  is moved left along the CPM axis until hitting this level curve; in the IPC auction, the final contract is arrived at by moving the pure-impression contract down parallel to the CPC axis.

## 5 Discussion

General contract auctions facilitate transactions when parties have conflicting information, or when they simply have different inherent value for the specific terms of a contract. Such a situation is common in traditional business negotiations, and, at least implicitly, contracts in the offline world often balance the same tradeoffs encapsulated explicitly by impression-plus-click auctions. For example, with book publication, authors typically receive a one-time advance plus royalty fees (i.e., a percentage of total sales revenue). Thus, authors confident in the future success of their book should be willing to trade a smaller advance

for larger royalties. A similar tradeoff occurs with insurance premiums and deductibles: A driver who thinks he is unlikely to get into an accident should be willing to accept relatively high deductibles in exchange for relatively low premiums. Corporate executives face a similar situation when deciding between guaranteed salaries and performance-based bonuses. More generally, in instances where parties bargain between deterministic and stochastic payments, a design similar to the impression-plus-click auction may prove useful.

## References

- [1] Aggarwal, G., Goel, A., Motwani, R.: Truthful auctions for pricing search keywords. In: Proceedings of the 7th ACM Conference on Electronic Commerce, pp. 1–7 (2006)
- [2] Clarke, E.H.: Multipart pricing of public goods. *Public Choice* 11, 17–33 (1971)
- [3] 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)
- [4] Goel, A., Munagala, K.: Hybrid keyword search auctions. In: Proceedings of the 18th International World Wide Web Conference, pp. 221–230 (2009)
- [5] Goodman, J.: Pay per percentage of impressions: an advertising method that is highly robust to fraud. In: First Workshop on Sponsored Search Auctions (2005)
- [6] Groves, T.: Efficient collective choice when compensation is possible. *Review of Economic Studies* 46, 227–241 (1979)
- [7] Paul Harrenstein, B., de Weerd, M.M., Conitzer, V.: A qualitative Vickrey auction. In: Proceedings of the 10th ACM Conferences on Electronic Commerce (2009)
- [8] Holmstrom, B.: Groves schemes on restricted domains. *Econometrica* 47(5), 1137–1144 (1979)
- [9] 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)
- [10] Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: Proceedings of the 7th ACM conference on Electronic commerce, pp. 218–227 (2006)
- [11] Lahaie, S., Pennock, D.M., Saberi, A., Vohra, R.V.: Sponsored search auctions. In: *Algorithmic Game Theory*, ch. 28, pp. 699–716. Cambridge University Press, Cambridge (2007)
- [12] Milgrom, P., Segal, I.: Envelope theorems for arbitrary choice sets. *Econometrica* 70(2), 583–601 (2002)
- [13] Myerson, R.B.: Optimal auction design. *Mathematics of Operations Research* 6(1) (1981)
- [14] Richardson, M., Dominowska, E., Ragno, R.: Predicting clicks: estimating the click-through rate for new ads. In: Proceedings of the 16th International World Wide Web Conference, pp. 521–530 (2007)
- [15] Varian, H.R.: Position auctions. *International Journal of Industrial Organization* 25, 1163–1178 (2007)
- [16] Vickrey, W.: Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance* 16, 8–37 (1961)

# Bidding for Representative Allocations for Display Advertising\*

Arpita Ghosh, Preston McAfee, Kishore Papineni, and Sergei Vassilvitskii

Yahoo! Research

{arpita,mcafee,kpapi,sergei}@yahoo-inc.com

**Abstract.** Display advertising has traditionally been sold via guaranteed contracts – a guaranteed contract is a deal between a publisher and an advertiser to allocate a certain number of impressions over a certain period, for a pre-specified price per impression. However, as spot markets for display ads, such as the RightMedia Exchange, have grown in prominence, the selection of advertisements to show on a given page is increasingly being chosen based on price, using an auction. As the number of participants in the exchange grows, the price of an impressions becomes a signal of its value. This correlation between price and value means that a seller implementing the contract through bidding should offer the contract buyer a range of prices, and not just the cheapest impressions necessary to fulfill its demand.

Implementing a contract using a range of prices, is akin to creating a mutual fund of advertising impressions, and requires *randomized bidding*. We characterize what allocations can be implemented with randomized bidding, namely those where the desired share obtained at each price is a non-increasing function of price. In addition, we provide a full characterization of when a set of campaigns are compatible and how to implement them with randomized bidding strategies.

## 1 Introduction

Display advertising — showing graphical ads on regular web pages, as opposed to textual ads on search pages — is approximately a \$24 billion business. There are two ways in which an advertiser looking to reach a specific audience (for example, 10 million males in California in July 2009) can buy such ad placements. One is the traditional method, where the advertiser enters into an agreement, called a guaranteed contract, directly with the publishers (owners of the webpages). Here, the publisher guarantees to deliver a prespecified number (10 million) of impressions matching the targeting requirements (male, from California) of the contract in the specified time frame (July 2009). The second is to participate in a spot market for display ads, such as the RightMedia Exchange, where advertisers can buy impressions one pageview at a time: every time a user loads a

---

\* A full version of this paper appears in [6].



page with a spot for advertising, an auction is held where advertisers can bid for the opportunity to display a graphical ad to this user. Both the guaranteed and spot markets for display advertising now thrive side-by-side. There is demand for guaranteed contracts from advertisers who want to hedge against future uncertainty of supply. For example, an advertiser who must reach a certain audience during a critical period of time (e.g. around a forthcoming product launch, such as a movie release) may not want to risk the uncertainty of a spot market; a guaranteed contract insures the publisher as well against fluctuations in demand. At the same time, a spot market allows the advertisers to bid for specific opportunities, permitting very fine grained targeting based on user tracking. Currently, RightMedia runs over nine billion auctions for display ads everyday.

How should a *publisher* decide which of her supply of impressions to allocate to her guaranteed contracts, and which to sell on the spot market? One obvious solution is to fulfill the guaranteed demand first, and then sell the remaining inventory on the spot market. However, spot market prices are often quite different for two impressions that both satisfy the targeting requirements of a guaranteed contract, since *different impressions have different value*. For example, the impressions from two users with identical demographics can have different value, based on different search behavior reflecting purchase intent for one of the users, but not the other. Since advertisers on the spot market have access to more tracking information about each user<sup>1</sup>, the resulting bids may be quite different for these two users. Allocating impressions to guaranteed contracts first and selling the remainder on the spot market can therefore be highly suboptimal in terms of *revenue*, since two impressions that would fetch the same revenue from the guaranteed contract might fetch very different prices from the spot market<sup>2</sup>.

On the other hand, simply buying the cheapest impressions on the spot market to satisfy guaranteed demand is not a good solution in terms of *fairness* to the guaranteed contracts, and leads to increasing short term revenue at the cost of long term satisfaction. As discussed above, impressions in online advertising have a *common value component* because advertisers generally have different information about a given user. This information (e.g. browsing history on an advertiser site) is typically relevant to *all* of the bidders, even though only *one* bidder may possess this information. In such settings, *price is a signal of value*—in a model of valuations incorporating both common and private values, the price converges to the true value of the item in the limit as the number of bidders goes to infinity ([8, 11], see also [7] for discussion). On average, therefore, the price on the spot market is a good indicator of the value of the impression, and delivering

<sup>1</sup> For example, a car dealership advertiser may observe that a particular user has been to his webpage several times in the previous week, and may be willing to bid more to show a car advertisement to induce a purchase.

<sup>2</sup> Consider the following toy example: suppose there are two opportunities, the first of which would fetch 10 cents in the spot market, whereas the second would fetch only  $\epsilon$ ; both opportunities are equally suitable for the guaranteed contract which wants just one impression. Clearly, the first opportunity should be sold on the spot market, and the second should be allocated to the guaranteed contract.

cheapest impressions corresponds to delivering the lowest quality impressions to the guaranteed contract<sup>3</sup>.

A publisher with access to both sources of demand thus faces a trade-off between revenue and fairness when deciding which impressions to allocate to the guaranteed contract; this trade-off is further compounded by the fact that the publisher typically does not have access to all the information that determines the value of a particular impression. Indeed, publishers are often the least well informed participants about the value of running an ad in front of a user. For example, when a user visits a politics site, Amazon (as an advertiser) can see that the user recently searched Amazon for an ipod, and Target (as an advertiser) can see they searched target.com for coffee mugs, but the publisher only knows the user visited the politics site. Furthermore, the exact nature of this trade-off is unknown to the publisher in advance, since it depends on the spot market bids which are revealed only *after* the advertising opportunity is placed on the spot market.

**The publisher as a bidder.** To address the problem of unknown spot market demand (*i.e.*, the publisher would like to allocate the opportunity to a bidder on the spot market *if* the bid is “high enough”, else to a guaranteed contract), the publisher acts, in effect, as a *bidder* on behalf on the guaranteed contracts. That is, the publisher now plays two roles: that of a seller, by placing his opportunity on the spot market, and that of a bidding agent, bidding on behalf of his guaranteed contracts. If the publisher’s own bid turns out to be highest among all bids, the opportunity is won and is allocated to the guaranteed contract. Acting as a bidder allows the publisher to probe the spot market and decide whether it is more efficient to allocate the opportunity to an external bidder or to a guaranteed contract.

How should a publisher model the trade-off between fairness and revenue, and having decided on a trade-off, how should she place bids on the spot market? An ideal solution is (a) easy to implement, (b) allows for a trade-off between the quality of impressions delivered to the guaranteed contracts and short-term revenues, and (c) is robust to the exact tradeoff chosen. In this work we show precisely when such an ideal solution exists and how it can be implemented.

## 1.1 Our Contributions

In this paper, we provide an analytical framework to model the publisher’s problem of how to fulfill guaranteed advance contracts in a setting where there is an alternative spot market, and advertising opportunities have a common value component. We give a solution where the publisher bids on behalf of its guaranteed contracts in the spot market. The solution consists of two components: an *allocation*, specifying the fraction of impressions at each price allocated to a contract, and a *bidding strategy*, which specifies how to acquire this allocation by bidding in an auction.

---

<sup>3</sup> While allocating the cheapest inventory to the guaranteed contracts is indeed revenue maximizing in the short term, in the long term the publisher runs the risk of losing the guaranteed advertisers by serving them the least valuable impressions.

The quality, or value, of an opportunity is measured by its price<sup>4</sup>. A perfectly representative allocation is one which consists of the same proportion of impressions at every price—*i.e.*, a mix of high-quality and low quality impressions. The trade-off between revenue and fairness is modeled using a budget, or *average target spend* constraint, for each advertiser’s allocation: the publisher’s choice of target spend reflects her trade-off between short-term revenue and quality of impressions for that advertiser (this must, of course, be large enough to ensure that the promised number of impressions satisfying the targeting constraints can be delivered.) Given a target spend<sup>5</sup>, a maximally representative allocation is one which minimizes the distance to the perfectly representative allocation, subject to the budget constraint. We first show how to solve for a maximally representative allocation, and then show how to implement such an allocation by purchasing opportunities in an auction, using randomized bidding strategies.

**Organization.** We start out with the single contract case, where the publisher has just one existing guaranteed contract, in Section 2; this case is enough to illustrate the idea of maximally representative allocations and implementation via randomized bidding strategies. We move on to the more realistic case of multiple contracts in Section 3; we first prove a result about which allocations can be implemented in an auction in a decentralized fashion, and derive the corresponding decentralized bidding strategies, and comment on solution of the optimal allocation. Full details, along with experimental validations of these strategies appear in 6.

**Related Work.** The most relevant work is the literature on designing expressive auctions and clearing algorithms for online advertising [9, 2, 10]. This literature does not address our problem for the following reason. While it is true that guaranteed contracts have coarse targeting relative to what is possible on the spot market, most advertisers with guaranteed contracts choose not to use all the expressiveness offered to them. Furthermore, the expressiveness offered does not include attributes like relevant browsing history on an advertiser site, which could increase the value of an impression to an advertiser, simply because the publisher *does not have* this information about the advertising opportunity. Even with extremely expressive auctions, one might still want to adopt a mutual fund strategy to avoid the ‘insider trading’ problem. That is, if some bidders possess good information about convertibility, others will still want to randomize their bidding strategy since bidding a constant price means always losing on some good impressions. Thus, our problem cannot be addressed by the use of more expressive auctions as in [10] — the real problem is not lack of expressivity, but lack of information.

<sup>4</sup> We emphasize that the assumption being made is *not* about price being a signal of value, but rather that impressions do have a common value component – given that impressions have a common value, price reflecting value follows from the theorem of Milgrom [8]. This assumption is commonly observed in practice.

<sup>5</sup> We point out that we do not address the question of how to set target spends, or the related problem of how to price guaranteed contracts to begin with. *Given* a target spend, we propose a complete solution to the publisher’s problem.

Another area of research focuses on selecting the optimal set of guaranteed contracts. In this line of work, Feige et al. [5] study the computational problem of choosing the set of guaranteed contracts to maximize revenue. A similar problem is studied by in [3, 1]. We do not address the problem of how to select the set of guaranteed contracts, but rather take them as given and address the problem of how to fulfill these contracts in the presence of competing demand from a spot market.

## 2 Single Contract

We first consider the simplest case: there is a single advertiser who has a guaranteed contract with the publisher for delivering  $d$  impressions. There are a total of  $s \geq d$  advertising opportunities which satisfy the targeting requirements of the contract. The publisher can also sell these  $s$  opportunities via auction in a spot market to external bidders. The highest bid from the external bidders comes from a *distribution*  $F$ , with density  $f$ , which we refer to as the bid landscape. That is, for every unit of supply, the highest bid from all external bidders, which we refer to as the price, is drawn i.i.d from the distribution<sup>6</sup>  $f$ . We assume that the supply  $s$  and the bid landscape  $f$  are known to the publisher<sup>7</sup>. Recall that the publisher wants to decide how to allocate its inventory between the guaranteed contract and the external bidders in the spot market. Due to penalties as well as possible long term costs associated with underdelivering on guaranteed contracts, we assume that the publisher wants to deliver all  $d$  impressions promised to the guaranteed contract.

An *allocation*  $a(p)$  is defined as follows:  $a(p)/s$  is the proportion of opportunities at price  $p$  purchased on behalf of the guaranteed contract (the price is the highest (external) bid for an opportunity.) That is, of the  $sf(p)dp$  impressions available at price  $p$ , an allocation  $a(p)$  buys a fraction  $a(p)/s$  of these  $sf(p)dp$  impressions, *i.e.*,  $a(p)f(p)dp$  impressions. For example, a constant bid of  $p^*$  means that for  $p \leq p^*$ ,  $a(p) = 1$  with the advertiser always winning the auction, and for  $p > p^*$ ,  $a(p) = 0$  since the advertiser would never win.

Generally, we will describe our solution in terms of the allocation  $a(p)/s$ , which must integrate out to the total demand  $d$ : a solution where  $a(p)/s$  is larger for higher prices corresponds to a solution where the guaranteed contract is allocated more high-quality impressions. As another example,  $a(p)/s = d/s$  is a perfectly representative allocation, integrating out to a total of  $d$  impressions, and allocating the same fraction of impressions at every price point.

Not every allocation can be purchased by bidding in an auction, because of the inherent asymmetry in bidding— a bid  $b$  allows every price below  $b$  and rules out every price above; however, there is no way to rule out prices *below* a certain value. That is, we can choose to exclude high prices, but not low prices. Before describing our solution, we state what kinds of allocations  $a(p)/s$  can be purchased by bidding in an auction.

<sup>6</sup> Specifically, we do not consider adversarial bid sequences; we also do not model the effect of the publisher's own bids on others' bids.

<sup>7</sup> Publishers usually have access to data necessary to form estimates of these quantities.

**Proposition 1.** *A right-continuous allocation  $a(p)/s$  can be implemented (in expectation) by bidding in an auction if and only if  $a(p_1) \geq a(p_2)$  for  $p_1 \leq p_2$ .*

*Proof.* Given a right-continuous non-increasing allocation  $\frac{a(p)}{s}$  (that lies between 0 and 1), define  $H(p) := 1 - \frac{a(p)}{s}$ . Let  $p^* := \inf \{p : a(p) < s\}$ . Then,  $H$  is monotone non-decreasing and is right-continuous. Further,  $H(p^*) = 0$  and  $H(\infty) = 1$ . Thus,  $H$  is a cumulative distribution function. We place bids drawn from  $H$  (the probability of a strictly positive bid being  $a(0)/s$ ). Then the expected number of impressions won at price  $p$  is then exactly  $a(p)/s$ . Conversely, given that bids for the contract are drawn at random from a distribution  $H$ , the fraction of supply at price  $p$  that is won by the contract is simply  $1 - H(p)$ , the probability of its bid exceeding  $p$ . Since  $H$  is non-decreasing, the allocation (as a fraction of available supply at price  $p$ ) must be non-increasing in  $p$ .

Note that the distribution  $H$  used to implement the allocation is a different object from the bid landscape  $f$  against which the requisite allocation must be acquired— in fact, it is completely independent of  $f$ , and is specified only by the allocation  $a(p)/s$ . That is, *given an allocation*, the bidding strategy that implements the allocation in an auction is independent of the bid landscape  $f$  from which the competing bid is drawn.

### 2.1 Maximally Representative Allocations

Ideally the advertiser with the guaranteed contract would like the same proportion of impressions at every price  $p$ , *i.e.*,  $a(p)/s = d/s$  for all  $p$ . (We ignore the possibility that the advertiser would like a higher fraction of higher-priced impressions, since these cannot be implemented according to Proposition 1 above.) However, the publisher faces a trade-off between delivering high-quality impressions to the guaranteed contract and allocating them to bidders who value them highly on the spot market. We model this by introducing an average unit target spend  $t$ , which is the average price of impressions allocated to the contract. A smaller (bigger)  $t$  delivers more (less) cheap impressions. As we mentioned before,  $t$  is part of the input problem, and may depend, for instance, on the price paid by the advertiser for the contract.

Given a target spend, the *maximally representative* allocation is an allocation  $a(p)/s$  that is ‘closest’ (according to some distance measure) to the ideal allocation  $d/s$ , while respecting the target spend constraint. That is, it is the solution to the following optimization problem:

$$\begin{aligned}
 & \inf_{a(\cdot)} \int_p \mathbf{u} \left( \frac{a(p)}{s}, \frac{d}{s} \right) f(p) dp \\
 & \text{s.t.} \quad \int_p a(p) f(p) dp = d \\
 & \quad \int_p p a(p) f(p) dp \leq t d \\
 & \quad 0 \leq \frac{a(p)}{s} \leq 1.
 \end{aligned} \tag{1}$$

The objective,  $\mathbf{u}$ , is a measure of the deviation of the proposed fraction,  $a(p)/s$ , from the perfectly representative fraction,  $d/s$ . In what follows, we will consider the  $L_2$  measure

$$\mathbf{u} \left( \frac{a(p)}{s}, \frac{d}{s} \right) = \frac{s}{2} \left( \frac{a(p)}{s} - \frac{d}{s} \right)^2,$$

in [6] we also consider the Kullback-Leibler (KL) divergence. Why the choice of KL and  $L_2$  for “closeness”? Only Bregman divergences lead to a selection that is consistent, continuous, local, and transitive [4]. Further, in  $R^n$  only least squares is scale- and translation- invariant, and for probability distributions only KL divergence is *statistical* [4].

The first constraint in (II) is simply that we must meet the target demand  $d$ , buying  $a(p)/s$  of the  $sf(p)dp$  opportunities of price  $p$ . The second constraint is the *target spend* constraint: the total spend (the spend on an impression of price  $p$  is  $p$ ) must not exceed  $td$ , where  $t$  is a target spend parameter (averaged per unit). As we will shortly see, the value of  $t$  strongly affects the form of the solution. Finally, the last constraint simply says that the proportion of opportunities bought at price  $p$ ,  $a(p)/s$ , must never go negative or exceed 1.

**Optimality conditions.** Introduce Lagrange multipliers  $\lambda_1$  and  $\lambda_2$  for the first and second constraints, and  $\mu_1(p), \mu_2(p)$  for the two inequalities in the last constraint. The Lagrangian is

$$\begin{aligned} L = & \int \mathbf{u} \left( \frac{a(p)}{s}, \frac{d}{s} \right) f(p)dp + \lambda_1 \left( d - \int a(p)f(p)dp \right) + \lambda_2 \left( \int pa(p)f(p)dp - td \right) \\ & + \int \mu_1(p)(-a(p))f(p)dp + \int \mu_2(p)(a(p) - s)f(p)dp. \end{aligned}$$

By the Euler-Lagrange conditions for optimality, the optimal solution must satisfy

$$\mathbf{u}' \left( \frac{a(p)}{s}, \frac{d}{s} \right) = \lambda_1 - \lambda_2 p + \mu_1(p) - \mu_2(p),$$

where the multipliers  $\mu$  satisfy  $\mu_1(p), \mu_2(p) \geq 0$ , and each of these can be non-zero only if the corresponding constraint is tight.

These optimality conditions, together with Proposition II, give us the following:

**Proposition 2.** *The maximally representative allocation for a single contract can be implemented by bidding in an auction for any convex distance measure  $\mathbf{u}$ .*

The proof follows from the fact that  $\mathbf{u}'$  is increasing for convex  $\mathbf{u}$ .

**$L_2$  utility.** In this subsection, we derive the optimal allocation when  $\mathbf{u}$ , the distance measure, is the  $L_2$  distance, and show how to implement the optimal allocation using a randomized bidding strategy. In this case the bidding strategy turns out to be very simple: toss a coin to decide whether or not to bid, and, if bidding, draw the bid value from a uniform distribution. The coin tossing probability and the endpoints of the uniform distribution depend on the demand and target spend values.

First we give the following result about the continuity of the optimal allocation; this will be useful in deriving the values that parameterize the optimal allocation. See [6] for the proof.

**Proposition 3.** *The optimal allocation  $a(p)$  is continuous in  $p$ .*

Note that we do not assume *a priori* that  $a(\cdot)$  is continuous; the *optimal* allocation turns out to be continuous.

The optimality conditions, when  $\mathbf{u}$  is the  $L_2$  distance, are:

$$\frac{a(p)}{s} - \frac{d}{s} = \lambda_1 - \lambda_2 p + \mu_1(p) - \mu_2(p),$$

where the nonnegative multipliers  $\mu_1(p), \mu_2(p)$  can be non-zero only if the corresponding constraints are tight.

The solution to the optimization problem (II) then takes the following form: For  $0 \leq p \leq p_{\min}$ ,  $a(p)/s = 1$ ; for  $p_{\min} \leq p \leq p_{\max}$ ,  $a(p)/s$  is proportional to  $C - p$ , *i.e.*,  $a(p)/s = z(C - p)$ ; and for  $p \geq p_{\max}$ ,  $a(p)/s = 0$ .

To find the solution, we must find  $p_{\min}, p_{\max}, z$ , and  $C$ . Since  $a(p)/s$  is continuous at  $p_{\max}$ , we must have  $C = p_{\max}$ . By continuity at  $p_{\min}$ , if  $p_{\min} > 0$  then  $z(C - p_{\min}) = 1$ , so that  $z = \frac{1}{p_{\max} - p_{\min}}$ . Thus, the optimal allocation  $a(p)$  is always parametrized by two quantities, and has one of the following two forms:

1.  $a(p)/s = z(p_{\max} - p)$  for  $p \leq p_{\max}$  (and 0 for  $p \geq p_{\max}$ ).

When the solution is parametrized by  $z, p_{\max}$ , these values must satisfy

$$s \int_0^{p_{\max}} z(p_{\max} - p)f(p)dp = d \tag{2}$$

$$s \int_0^{p_{\max}} zp(p_{\max} - p)f(p)dp = td \tag{3}$$

Dividing (2) by (3) eliminates  $z$  to give an equation which is monotone in the variable  $p_{\max}$ , which can be solved, for instance, using binary search.

2.  $a(p)/s = 1$  for  $p \leq p_{\min}$ , and  $a(p)/s = \frac{p_{\max} - p}{p_{\max} - p_{\min}}$  for  $p \leq p_{\max}$  (and 0 thenceforth).

When the solution is parametrized by  $p_{\min}, p_{\max}$ , these values must satisfy

$$sF(p_{\min}) + \int_{p_{\min}}^{p_{\max}} s \frac{(p_{\max} - p)}{p_{\max} - p_{\min}} f(p)dp = d \tag{4}$$

$$\int_0^{p_{\min}} spf(p)dp + \int_{p_{\min}}^{p_{\max}} sp \frac{(p_{\max} - p)}{p_{\max} - p_{\min}} f(p)dp = td. \tag{5}$$

Note that the optimal allocation can be represented more compactly as

$$\frac{a(p)}{s} = \min\{1, z(p_{\max} - p)\}. \tag{6}$$

*Effect of varying target spend.* Varying the value of the target spend,  $t$ , while keeping the demand  $d$  fixed, leads to a tradeoff between representativeness and revenue from selling opportunities on the spot market, in the following way. The minimum possible target spend, while meeting the target demand (in expectation) is achieved by a solution where  $p_{\min} = p_{\max}$  and  $a(p)/s = 1$  for  $p$  less equal this value, and 0 for greater. The value of  $p_{\min}$  is chosen so that

$$\int_0^{p_{\min}} sf(p)dp = d \Rightarrow p_{\min} = F^{-1}\left(\frac{d}{s}\right).$$

This solution simply bids a flat value  $p_{\min}$ , and corresponds to giving the cheapest possible inventory to the advertiser, subject to meeting the demand constraint. This gives the minimum possible total spend for this value of demand, of

$$\underline{t}d = \int_0^{p_{\min}} spf(p)dp = sF(p_{\min})E[p|p \leq p_{\min}] = dE[p|p \leq p_{\min}]$$

(Note that the maximum possible total spend that is maximally representative while not overdelivering is  $R = \int pf(p)dp = dE[p] = d\bar{p}$ .)

As the value of  $t$  increases above  $\underline{t}$ ,  $p_{\min}$  decreases and  $p_{\max}$  increases, until we reach  $p_{\min} = 0$ , at which point we move into the regime of the other optimal form, with  $z = 1$ . As  $t$  is increased further,  $z$  decreases from 1, and  $p_{\max}$  increases, until at the other extreme when the spend constraint is essentially removed, the solution is  $\frac{a(p)}{s} = \frac{d}{s}$  for all  $p$ ; *i.e.*, a perfectly representative allocation across price. Thus the value of  $t$  provides a dial by which to move from the “cheapest” allocation to the perfectly representative allocation.

### 2.2 Randomized Bidding Strategies

The quantity  $a(p)/s$  is an optimal *allocation*, *i.e.*, a recommendation to the publisher as to how much inventory to allocate to a guaranteed contract at every price  $p$ . However, recall that the publisher needs to *acquire* this inventory on behalf of the guaranteed contract by bidding in the spot market. The following theorem shows how to do this when  $\mathbf{u}$  is the L2 distance.

**Theorem 1.** *The optimal allocation for the  $L_2$  distance measure can be implemented (in expectation) in an auction by the following random strategy: toss a coin to decide whether or not to bid, and if bidding, draw the bid from a uniform distribution.*

*Proof.* From (6) that the optimal allocation can be represented as

$$\frac{a(p)}{s} = \min\{1, z(p_{\max} - p)\}.$$

By Proposition 1, an allocation  $\frac{a(p)}{s} = \min\{1, z(p_{\max} - p)\}$  can be implemented by bidding in an auction using the following randomized bidding strategy: with probability  $\min\{zp_{\max}, 1\}$ , place a bid drawn uniformly at random from the range  $[\max\{p_{\max} - \frac{1}{z}, 0\}, p_{\max}]$ .

### 3 Multiple Contracts

We now study the more realistic case where the publisher needs to fulfill multiple guaranteed contracts with different advertisers. Specifically, suppose there are  $m$  advertisers, with demands  $d_j$ . As before, there are a total of  $s \geq \sum d_j$  advertising opportunities available to the publisher. 8 An allocation  $a_j(p)/s$  is the proportion

<sup>8</sup> In general, not all of these opportunities might be suitable for every contract; we do not consider this here for clarity of presentation. However the same ideas and methods can be applied in that case and the results are qualitatively similar.



of opportunities purchased on behalf of contract  $j$  at price  $p$ . Of course, the sum of these allocations cannot exceed 1 for any  $p$ , which corresponds to acquiring all the supply at that price.

As in the single contract case, we are first interested in what allocations  $a_j(p)$  are implementable by bidding in an auction. However, in addition to being implementable, we would like allocations that satisfy an additional practical requirement, explained below. Notice that the publisher, acting as a bidding agent, now needs to acquire opportunities to implement the allocations for *each of the guaranteed contracts*. When an opportunity comes along, therefore, the publisher needs to decide *which* of the contracts (if any) will receive that opportunity. There are two ways to do this: the publisher submits *one* bid on behalf of all the contracts; if this bid wins, the publisher then selects one amongst the contracts to receive the opportunity. Alternatively, the publisher can submit one bid for *each* contract; the winning bid then automatically decides which contract receives the opportunity. We refer to the former as a **centralized strategy** and the latter as a **decentralized strategy**.

There are situations where the publisher will need to choose the winning advertiser *prior to* seeing the price, that is, the highest bid from the spot market. For example, to reduce latency in placing an advertisement, the auction mechanism may require that the bids be accompanied by the advertisement (or its unique identifier). A decentralized strategy automatically fulfills this requirement, since the choice of winning contract does not depend upon knowing the price. In a centralized strategy, this requirement means that the relative fractions won at price  $p$ ,  $a_i(p)/a_j(p)$ , are *independent of the price  $p$* —when this happens, the choice of advertiser can be made (by choosing at random with probability proportional to  $a_j$ ) without knowing the price.

As before, we will be interested in implementing optimal (*i.e.*, maximally representative) allocations. We will, therefore, concentrate on characterizing allocations which can be implemented via a decentralized strategy. In the full version of the paper [6] we show how to compute the optimal allocation in the presence of multiple contracts.

### 3.1 Decentralization

In this section, we examine what allocations can be implemented via a decentralized strategy. Note that it is not sufficient to simply use a distribution  $H_j = 1 - \frac{a_j(p)}{a_j(0)}$  as in Proposition 1, since these contracts compete amongst each other as well. Specifically, using the distribution  $1 - \frac{a_j(p)}{a_j(0)}$  will lead to too few opportunities being purchased for contract  $j$ , since this distribution is designed to compete against  $f$  alone, rather than against  $f$  *as well as* the other contracts. We need to show how to choose distributions in such a way that lead to a fraction  $a_j(p)/s$  of opportunities being purchased for contract  $j$ , for every  $j = 1, \dots, m$ .

First, we argue that a decentralized strategy with given distributions  $H_j$  will lead to allocations that are non-increasing, as in the single contract case. A decentralized implementation uses distributions  $H_j$  to bid for impressions. Then, contract  $j$  wins an impression at price  $p$  with probability

$$a_j(p) = \int_p^\infty \left( \prod_{k \neq j} H_k(x) \right) h_j(x) dx,$$

since to win, the bid for contract  $j$  must be larger than  $p$  and larger than the bids placed by each of the remaining  $m - 1$  contracts. Since all the quantities in the integrand are nonnegative,  $a_j$  is non-increasing in  $p$ .

Now assume that  $a_j$  are differentiable a.e. and non-increasing. Let

$$\frac{A(p)}{s} := \sum_j \frac{a_j(p)}{s}$$

be the total fraction of opportunities at price  $p$  that the publisher needs to acquire. Clearly,  $a_j$  must satisfy  $A(p) \leq s, \forall p$ . Let  $p^* := \inf\{p : A(p) < s\}$ . Let

$$H_j(p) := \begin{cases} e^{\int_p^\infty a'_j(x)/(s-A(x)) dx} & p > p^* \\ 0 & \text{else} \end{cases} \tag{7}$$

Then,  $H_j(p) \geq 0$  and is continuous. Since  $a'_j(p)$  is non-increasing,  $H_j(p)$  is monotone non-decreasing. Further,  $H(\infty) = 1$  and  $H_j(p^*) = 0$ . Thus,  $H_j$  is a distribution function. We can verify that bidding according to  $H_j$  will result in the desired allocations (see [6] for details).

Thus, we have constructed distribution functions  $H_j(p)$  which implement the given non-increasing (and a.e. differentiable) allocations  $a_j(p)$ . If any  $a_j$  is increasing at any point, the set of campaigns cannot be decentralized. The following theorem generalizes Proposition 1:

**Theorem 2.** *A set of allocations  $a_j(p)$  can be implemented in an auction via a decentralized strategy iff each  $a_j(p)$  is non-increasing in  $p$ , and  $\sum_j a_j(p)/s \leq 1$ .*

Having determined which allocations can be implemented by bidding in an auction in a decentralized fashion, we turn to the question of finding suitable allocations to implement. As in the single contract case, we would like to implement allocations that are maximally representative, given the spend constraints.

As we show in [6], the optimal allocation is decentralizable in two cases:

1. The target spends are such that the solutions decouple. In this case the allocation for each contract is independent of the others; we solve for the parameters of each allocation as in Section 2.1.
2. The target spends are such that, for all  $j, k$ ,  $\frac{a_j(p)}{a_k(p)}$  is independent of  $p$ . In this case we need to solve for the common slope and  $p_{\min}$ , and the contract specific values  $p_{\max}^j$ , which together determine the allocation. This can be done using, for instance, Newton’s method.

When the target spends are such that the allocation is not decentralizable, the vector of target spends can be increased to reach a decentralizable allocation. One way is to scale up the target spends uniformly until they are large enough to admit a separable solution; this has the advantage of preserving the relative ratios of target spends. The minimum multiplier which renders the allocation decentralizable can be found numerically, using for instance binary search.

## 4 Conclusion

Moving guaranteed contracts into an exchange environment presents a variety of challenges for a publisher. Randomized bidding is a useful compromise between minimizing the cost and maximizing the quality of guaranteed contracts. It is akin to the mutual fund strategy common in the capital asset pricing model. We provide a readily computable solution for synchronizing an arbitrary number of guaranteed campaigns in an exchange environment. Moreover, the solution we detail appears stable with real data.

## References

- [1] Babaioff, M., Hartline, J., Kleinberg, R.: Selling banner ads: Online algorithms with buyback. In: 4th Workshop on Ad Auctions (2008)
- [2] Boutilier, C., Parkes, D., Sandholm, T., Walsh, W.: Expressive banner ad auctions and model-based online optimization for clearing. In: National Conference on Artificial Intelligence, AAAI (2008)
- [3] Constantin, F., Feldman, J., Muthukrishnan, S., Pal, M.: Online ad slotting with cancellations. In: 4th Workshop on Ad Auctions (2008)
- [4] Csizsar, I.: Why least squares and maximum entropy? an axiomatic approach to interference for linear inverse problems. *Annals of Statistics* 19(4), 2032–2066 (1991)
- [5] Feige, U., Immorlica, N., Mirrokni, V.S., Nazerzadeh, H.: A combinatorial allocation mechanism with penalties for banner advertising. In: Proceedings of ACM WWW, pp. 169–178 (2008)
- [6] Ghosh, A., McAfee, P., Papineni, K., Vassilvitskii, S.: Bidding for representative allocations for display advertising. CoRR, abs/0910-0880 (2009)
- [7] McAfee, R.P., McMillan, J.: Auctions and bidding. *Journal of Economic Literature* 25(2), 699–738 (1987)
- [8] Milgrom, P.R.: A convergence theorem for competitive bidding with differential information. *Econometrica* 47(3), 679–688 (1979)
- [9] Parkes, D., Sandholm, T.: Optimize-and-dispatch architecture for expressive ad auctions. In: 1st Workshop on Ad Auctions (2005)
- [10] Sandholm, T.: Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions 28(3), 45–58 (2007)
- [11] Wilson, R.: A bidding model of perfect competition. *Rev. Econ. Stud.* 44(3), 511–518 (1977)

# Social Networks and Stable Matchings in the Job Market<sup>\*</sup>

Esteban Arcaute<sup>1</sup> and Sergei Vassilvitskii<sup>2</sup>

<sup>1</sup> Institute for Computational and Mathematical Engineering, Stanford University  
arcaute@stanford.edu

<sup>2</sup> Yahoo! Research  
sergei@yahoo-inc.com

**Abstract.** In this paper we introduce and study a model that considers the job market as a two-sided matching market, and accounts for the importance of social contacts in finding a new job. We assume that workers learn *only* about positions in firms through social contacts. Given that information structure, we study both static properties of what we call *locally stable matchings*, a solution concept derived from stable matchings, and dynamic properties through a reinterpretation of Gale-Shapley’s algorithm as myopic best response dynamics.

We prove that, in general, the set of locally stable matching strictly contains that of stable matchings and it is in fact NP-complete to determine if they are identical. We also show that the lattice structure of stable matchings is in general absent. Finally, we focus on myopic best response dynamics inspired by the Gale-Shapley algorithm. We study the efficiency loss due to the informational constraints, providing both lower and upper bounds.

## 1 Introduction

When looking for a new job, the most often heard advice is to “ask your friends”. While in the modern world almost all of the companies have online job application forms, these are usually overloaded with submissions; and it is no secret that submitting a resume through someone on the inside greatly increases the chances of the application actually being looked at by a qualified person. This is the underlying premise behind the professional social networking site LinkedIn, which now boasts more than 40 million users. And, as pointed out by Jackson [10], has given a new meaning to the word ‘networking,’ with Merriam Webster’s Dictionary’s defining it as “the cultivation of productive relationships for employment or business.”

Sociologists have long studied this phenomenon, and have time and time again confirmed the role that social ties play in getting a new job. Granovetter’s seminal work [7, 8] headlines a long history of research into the importance of social contacts in labor markets. His results are striking, for example, 65 percent of

---

<sup>\*</sup> A full version of this work appears in [2].

managerial workers found their job through social contacts. Other studies (see, e.g., [15, 13, 9]) all echo the importance of social contacts in securing a new position.

While there are numerous reasons that social ties play such an important role, one may think that the employers themselves would prefer to evaluate *all* candidates for a position before making a hiring decision. This is in fact what happens in some segments of the job market. In the United States, the National Resident Match Program is a significant example of a centralized selection matching mechanism. Such centralized markets have been well studied in *two-sided matching theory*. Indeed, the NRMP is one of the most important practical applications of the celebrated *stable matching problem* in two-sided matching markets [16]. For an overview of two-sided matching markets, see [17].

However, the task of evaluating (and ranking) all possible candidates is often simply not feasible. Especially in today's economy, it is not rare to hear of hundreds of applicants for a position, obviously the vast majority cannot be interviewed, regardless of their qualifications. The recommendation by an employee thus carries extra weight in the decision process, precisely because it separates the specific application from the masses.

*Model.* In this work, we propose a new model that bridges the rigorous analysis of the two-sided matching theory with the observations made by social network analysis. Specifically, we develop a model of job markets where social contacts play a pivotal role; and then proceed to analyze it through the stable matching lens.

We integrate the usage of social contacts by allowing an applicant to apply *only* to jobs in firms employing her friends. Clearly this limitation depends on the underlying social graph. Intuitively, the equilibrium behavior in well connected social graphs should be closer to that in classical two sided matchings than in badly connected ones. But even in well connected social graphs this limitation leads to behaviors not observed in the traditional model. For example, a firm may lose all of its workers to the competition, and subsequently go out of business.

The model forces us to consider a setting where job applicants have only partial information on job opportunities. The main question we focus on in the paper is: how does the inclusion of such an informational constraint alter the model and predictions of traditional stable matching theory?

*Our Contributions*<sup>1</sup> In traditional two-sided matching theory, a matching where no worker-firm pair can find a profitable deviation is called *stable*. Analogously, we call our solution concept a *locally stable matching*, where the locality is qualified by the social network graph. We study structural properties of locally stable matchings by showing that, in general, the set of locally stable matchings does not form a distributive lattice, as is the case for global stable matchings. We also show that, in general, it is NP-complete to determine whether all locally stable matchings are also globally stable. Both of these results exploit a characteriza-

---

<sup>1</sup> The proofs of all our results are available online at [2].

tion of locally stable matchings in the special case of matching one worker per firm, for particular rankings over workers and firms.

We then turn our attention to dynamic analysis. We consider how a particular interpretation of the classic Gale-Shapley algorithm [6] performs under such informational constraints; we refer to our algorithm as the *local Gale-Shapley* algorithm. We first prove that, unlike the standard Gale-Shapley algorithm [17], the existence of informational constraints implies that the output of the algorithm is not independent of the order of proposals. Nevertheless, under weak stochastic conditions, we show that the local Gale-Shapley algorithm converges almost surely, assuming the same particular rankings over workers and firms as before.

Unlike the traditional Gale-Shapley algorithm, the algorithm in the limited information case is highly dependent on the initial conditions. To explore this further we define a minimal notion of efficiency, namely the number of firms still in business in the outcome matching, and quantify the efficiency loss under various initial conditions. Specifically, we show that if an adversary chooses an initial matching, he can ensure that some firms lose all of their workers; conversely there is a distribution on the preference lists used by the firms that guarantees that at least some firms remain in business, regardless of the actions of the adversary.

## Related Work

Our work touches on several threads of the literature. Most closely related is the work by Calvó-Armengol and Jackson [3, 4]. They consider how information dissemination through neighbors of workers on potential jobs can affect wage and employment dynamics in the job market. There are several key differences with our model. The most important one is that, in [3, 4], there is no competition for job openings between workers. Unemployment is the result of a random sampling process and not of strategic interactions between workers and firms. Also, all workers learn *directly* about potential job openings with some probability, and *indirectly* through their social contacts, whereas in our model a worker can only learn about potential job openings through her social contacts.

Also related is the work by Lee and Schwarz [11]. The authors consider the stable matching problem in the job market where a costly information acquisition step (interviewing) is necessary for both workers and firms to learn their preferences. Once interviewing is over, the standard Gale-Shapley algorithm is used to calculate the matching of workers to firms. Although the authors use stable matching as their solution concept, and only partial information on jobs and candidates is available, their assumptions imply that the information available to workers and firms is *unchanged* throughout the matching phase. In that sense, their work is related to the equilibrium analysis performed in our model, but is dramatically different when considering the evolution of the job market during the actual matching phase.

Finally, in [1, 12], the authors consider the problem of matching applicants to job positions. A matching is said to be *popular* if the number of happy applicants is as large as possible. This notion is related to the notion of efficiency used in our paper, namely that of maximizing the number of firms in business.

## 2 Definitions and Notation

Let  $W$  be a set of workers, and  $G = (W, E)$  be an undirected graph representing the social network among workers. Let  $F$  be the set of firms, each with  $k$  jobs, for some  $k > 0$ . We are interested in the case where there are as many workers as positions in all firms, i.e.,  $|W| = n = k|F|$ . Following standard notation, for a worker  $w \in W$ , let  $\Gamma(w)$  be the neighborhood of  $w$  in  $G$ .

An assignment of workers to firms can be described by a function mapping workers to jobs, or alternatively by a function mapping firms to workers. Following the definition from the two-sided matching literature, we define both functions simultaneously.

We assume that some companies are better to work for than others, and thus each worker  $w$  has a strict ranking  $\succ_w$  over firms such that, for firms  $f \neq f'$ ,  $w$  prefers being employed in  $f$  than in  $f'$  if and only if  $f \succ_w f'$ . Note however, that the ranking is blind to the individual positions within a firm: all of the  $k$  slots of a given firm are equivalent from the point of view of a worker.

Similarly, each firm  $f$  has a strict ranking  $\succ_f$  over workers. We assume that all workers strictly prefer being employed, and that all firms strictly prefer having all their positions filled. For any worker  $w$ , in a slight abuse of notation, we extend her ranking over firms to account for her being unemployed by setting  $f \succ_w w$  for all firms  $f$ ; in a similar way we extend the rankings of firms over workers.

### Definition 1 (Matching)

1. Case 1:  $k = 1$ . The function  $\mu : W \cup F \rightarrow W \cup F$  is a matching if the following conditions hold: (1) for all  $w \in W$ ,  $\mu(w) \in F \cup \{w\}$ ; (2) for all  $f \in F$ ,  $\mu(f) \in W \cup \{f\}$ ; and (3)  $\mu(w) = f$  if and only if  $\mu(f) = w$ .
2. Case 2:  $k > 1$ . The function  $\mu : W \cup F \rightarrow 2^W \cup F$  is a matching if the following conditions hold: (1) for all  $w \in W$ ,  $\mu(w) \in F \cup \{\{w\}\}$ ; (2) for all  $f \in F$ ,  $\mu(f) \in 2^W \cup \{f\}$ ; (3)  $\mu(w) = f$  if and only if  $w \in \mu(f)$ ; and (4)  $|\mu(f)| \leq k$ .

We say that a matching  $\mu$  is complete if:

$$\bigcup_{f \in F} \mu(f) = W.$$

Given a matching  $\mu$  and a firm  $f$ , let  $\min(\mu(f))$  be the *least preferred* worker employed by firm  $f$  (w.r.t. firm  $f$ 's ranking) if  $|\mu(f)| = k$ , and  $\min(\mu(f)) = f$  otherwise.

To study the notion of stable matchings, we adapt the usual concept of a *blocking pair*. Given the preferences of workers and firms, a matching  $\mu$ , a firm  $f$  and a worker  $w$ , we say that  $(w, f)$  is a *blocking pair* if and only if  $f \succ_w \mu(w)$  and  $w \succ_f \min(\mu(f))$ . In other words, worker  $w$  prefers firm  $f$  to her currently matched firm; and firm  $w$  prefers worker  $w$  to its least preferred current employee.

We now define a generalization of the standard notion of stable matching that accounts for the locality of information. Recall that a (global) matching is said

to be *stable* if there are no blocking pairs. However, in our paper we assume that *the workers can only discover possible firms by looking at their friends' places of employment*. This informally captures a significant mechanism of information transfer: although there may exist a firm  $f$  that would make  $(w, f)$  a blocking pair, if none of  $w$ 's friends work at  $f$ , then it becomes much less likely that  $w$  would learn of  $f$  on her own. We have the following definition.

**Definition 2 (Locally Stable Matching).** *Let  $G = (W, E)$  be the social network over the set of workers  $W$ . We say that a matching  $\mu$  is a locally stable matching with respect to  $G$  if, for all  $w \in W$  and  $f \in F$ ,  $(w, f)$  is a blocking pair if and only if  $\Gamma(w) \cap \mu(f) = \emptyset$  (i.e., no workers in  $w$ 's social neighborhood are employed by firm  $f$ ).*

Note that for a given worker  $w$ , the set of other workers she is competing against depends on both the social network  $G$  (i.e., her neighbors), and the current matching.

*Example 1 (Indirect Competition).* Assume  $k = 2$  and  $G$  is the path over  $W = \{w_1, w_2, w_3, w_4\}$ :  $w_1 - w_2 - w_3 - w_4$ . Consider worker  $w_4$ . If  $\mu(f_1) = \{w_3, w_4\}$  and  $\mu(f_2) = \{w_1, w_2\}$ , then  $w_4$  can only see positions in  $f_1$ . However, since  $w_2$  is adjacent to  $w_3$ ,  $w_2$  can see all position in  $f_1$ . Hence, if  $w_2 \succ_{f_1} w_3 \succ_{f_1} w_4$ ,  $w_2$  could get  $w_4$ 's position in  $f_1$ , leading to  $w_4$  being replaced by  $w_2$  even though  $w_2 \notin \Gamma(w_4)$ .

In the remainder of the paper, we characterize static properties of locally stable matchings, and then analyze dynamics similar to the Gale-Shapley algorithm.

### 3 Static Analysis

For  $k = 1$ , when the preferences of workers and firms are strict, it is known that the set of global stable matchings is a *distributive lattice*. In general, the distributive lattice structure of the set of global stable matchings is not present in the set of locally stable matchings. We first recall the Lattice Theorem (by Conway), and then show how, in general, it does not hold for locally stable matchings. The exposition of the Lattice Theorem is that found in [17] (Theorem 2.16).

Let  $\mu$  and  $\mu'$  be two matchings. Define the operation  $\vee_W$  over  $(\mu, \mu')$  as follows:  $\mu \vee_W \mu' : W \cup F \rightarrow W \cup F$  such that, for all  $w \in W$ ,  $\mu \vee_W \mu'(w) = \mu(w)$  if  $\mu(w) \succ_w \mu'(w)$ , and  $\mu \vee_W \mu'(w) = \mu'(w)$  otherwise. For all  $f \in F$ ,  $\mu \vee_W \mu'(f) = \mu'(f)$  if  $\mu(f) \succ_f \mu'(f)$ , and  $\mu \vee_W \mu'(f) = \mu(f)$  otherwise. We can similarly define  $\wedge_W$  by exchanging the roles of workers and firms.

**Theorem 1 (Lattice Theorem (Conway)).** *When all preferences are strict, if  $\mu$  and  $\mu'$  are stable matchings, then the functions  $\lambda = \mu \vee_W \mu'$  and  $\nu = \mu \wedge_W \mu'$  are both matchings. Furthermore, they are both stable.*



In general, given strict preferences<sup>2</sup> of workers and firms, Theorem 1 does not hold for the set of locally stable matchings. This is the content of the following example.

*Example 2 (Absence of Distributive Lattice).* In this example, we assume  $k = 1$ ,  $W = \{w_1, w_2, w_3\}$  and  $F = \{f_1, f_2, f_3\}$ . Further, let the preferences of all workers be  $f_1 \succ f_2 \succ f_3$ . Similarly, let the preferences of all firms be  $w_1 \succ w_2 \succ w_3$ . Finally, assume the graph  $G$  is the path with  $w_2$  and  $w_3$  at its endpoints.

Let  $\mu(w_i) = f_i$  (and  $\mu(f_i) = w_i$ ). It is clear that  $\mu$  is a 1-locally stable matching. Consider now  $\mu'$  be such that  $\mu'(w_1) = f_1$ ,  $\mu'(w_2) = f_3$  and  $\mu'(w_3) = f_2$  (and  $\mu'(f_1) = w_1$ ,  $\mu'(f_2) = w_3$  and  $\mu'(f_3) = w_2$ ). The only blocking pair here is  $(w_2, f_2)$ , but  $f_2 = \mu'(w_3)$  and  $w_3 \notin \Gamma(w_2)$ . Hence  $\mu'$  is a 1-locally stable matching.

We now construct  $\lambda = \mu \vee_W \mu'$ . For all  $i$ ,  $\lambda(w_i) = f_i$ . Now  $\lambda(f_1) = w_1$  but  $\lambda(f_2) = \lambda(f_3) = w_3$ . Hence  $\lambda$  is not a matching.

*Assumption.* In the remainder of the paper we focus on a specific family of preferences over workers and firms. Uniqueness of global stable matching is a desirable property in matching markets as it allows for sharp predictions of the outcome at equilibrium. Clark [5] studies thoroughly the question and identifies a set of sufficient conditions on the preferences, called *aligned preferences*, for the global stable matching to be unique. The study of aligned preferences have recently received attention in the economics literature [14, 18].

In this paper we consider a subset of aligned preferences, where all workers share the same ranking over firms, and firms share the same ranking over workers. This assumption is made for technical reasons - we believe our results extend to the case of general aligned preferences.

**Assumption 1.** *There exist a labeling of the nodes in  $W = \{w_1, \dots, w_n\}$  such that all firms rank workers as follows:  $w_i \succ w_j$  if and only if  $i < j$ . Similarly, we assume there exists a labeling of the firms  $F = \{f_1, \dots, f_{n_f}\}$  such that all workers rank the firms as follows:  $f_i \succ f_j$  if and only if  $i < j$ .*

We first show that, for  $k = 1$ , the set of locally stable matchings is equivalent to the set of topological orderings over the partial order induced by  $G$  and the labeling of the workers.

**Theorem 2 (Characterization of Locally Stable Matchings).** *Assume  $k = 1$ , and let  $G(W, E)$  be the social network over the set of workers. Let  $D(W, E')$  be a directed graph over  $W$  such that  $(w_i, w_j) \in E'$  if and only if  $i < j$  and  $(w_i, w_j) \in E$ . Let  $\mu$  be a complete matching of workers to firms. Construct the following ordering  $\phi_\mu$  over  $W$  induced by  $\mu$ : the  $i^{\text{th}}$  node in the ordering is the node  $w$  such that  $\mu(w) = f_i$ , i.e.  $\phi_\mu(w) = i$ .*

*The matching  $\mu$  is a 1-locally stable matching if and only if  $\phi_\mu$  is a topological ordering on  $D$ .*

<sup>2</sup> The absence of the distributive lattice has been previously observed when the preferences are not strict, see Roth [16].

There are several important corollaries to the characterization from Theorem 2. First, the set complete locally stable matchings can be exponentially large. Thus, by introducing informational constraints, the uniqueness property of global stable matchings under aligned preferences is, in general, lost under locally stable matchings.

**Corollary 3 (Number of Locally Stable Matchings).** *Assume  $k = 1$  and the social network  $G(W, E)$  is the star centered at worker  $w_1$ . Then there are  $(n - 1)!$  distinct locally stable matchings.*

It is interesting to ask whether there are specific properties of the social network  $G$  that guarantee the existence of a labeling under which there is a unique complete locally stable matching. As shown in the next corollary, it is NP-complete to answer positively such question.

**Corollary 4.** *Let  $(k, G(W, E))$  be given. It is NP-complete to test if there is a labeling  $\{w_1, w_2, \dots, w_n\}$  of the workers such that, if all firms rank the workers according to that labeling, the complete locally stable matching is unique.*

For general  $k > 1$ , we only have a set of sufficient conditions for complete locally stable matchings to be unique. See 2 for more details.

## 4 Algorithmic Questions

We are thus interested in decentralized algorithms that can find a locally stable matching. In this section we propose a decentralized version of Gale-Shapley's algorithm. Assumption 1 is again enforced in this section. We first prove that our algorithm converges. Unlike the case without informational constraints, our algorithm does not always select the same locally stable matching.

Recall that the Gale-Shapley algorithm is initialized by an empty matching 17. Since the empty matching is a locally stable matching, our algorithm requires to be initialized by a non-empty matching. We thus explore our algorithm's performance under adversarial initial complete matchings. We use the number of firms with no employees as a proxy for efficiency 3. We characterize the potential efficiency loss by providing upper and lower bounds on the number of firms with no employees.

### 4.1 Local Gale-Shapley Algorithm

One can interpret the Gale-Shapley algorithm from two-sided matching theory as a constrained version of *myopic best response dynamics* in the following way.

The dynamics proceed in rounds, which we index by  $q \in \mathbb{N}$ . Let  $\mu^{(q)}$  be the matching at the beginning of round  $q$ . Let  $w^{(q)} \in W$  be the *active* worker, where  $w^{(q)}$  is sampled uniformly at random from  $W$ , and independently from

---

<sup>3</sup> Our bounds naturally translate into unemployment rate, a common indicator of the efficiency of the job market.

previous rounds. We call such sampling process the *activation process*. Such activation process can be thought of as follows: assume all workers decide to explore employment opportunities according to a random clock with an exponential distribution with a given mean (the same mean for all workers). When the clock of  $w_i$  “sets off”,  $w_i$  becomes active and looks for a better job. It is easy to see that the sequence of active nodes has the same distribution as taking independent uniform samples from  $W$ .

In myopic best response dynamics,  $w^{(q)}$  would consider its current firm  $\mu^{(q)}(w^{(q)})$  and compare it to the best firm  $f$  it could be employed by given  $\mu^{(q)}$  (i.e. the best firm where the worst employee was worse than  $w$  given the matching  $\mu^{(q)}$ ). If its current firm was better, it would pass. Else it would quit its job and get employed by  $f$  (leading to a worker being fired, or an empty position being filled).

Gale-Shapley’s algorithm is a constrained version of the above dynamics as it requires the active worker to consider the best firm it has not considered before (in other words it requires the active worker to remember what firms he has already failed to get a position at).

We consider a local and decentralized version of the myopic best-response dynamics proposed above. We call it “local Gale-Shapley” algorithm. Instead of restricting the strategy space of the active worker using “memory” as in Gale-Shapley’s algorithm, we restrict it using the graph  $G(W, E)$  in the following way:  $w^{(q)}$  compares its current firm in  $\mu^{(q)}$  to the best firm that employs one of its neighbors in  $G$  it could be employed by given  $\mu^{(q)}$ . An alternative way to describe the process is that the active node  $w^{(q)}$  applies for a job at all the firms employing its neighbors that she strictly prefers to her current employer, and selects the best offer she gets (that offer might eventually be to stay at her current job).

More formally, the algorithm proceeds in rounds indexed by  $q \in \mathbb{N}$ . During round  $q \geq 0$ :

- the active worker  $w^{(q)}$  is sampled, independently from previous rounds, uniformly at random from  $W$ .
- Next,  $w^{(q)}$  applies to all firms she strictly prefers to  $\mu^{(q)}(w^{(q)})$ , her current employer.
- The active worker receives some offers:
  - if at least one offer is received,  $w^{(q)}$  quits her current employer and joins the best firm that sent an offer;
  - if no offers are received,  $w^{(q)}$  stays at her current job.

It is important to note that, unlike the Gale-Shapley algorithm, this variant of best-response dynamics can lead to a firm losing all its employees as demonstrated below.

*Example 3 (Firm with no Employees).* Let  $n = 4$  and  $k = 2$ . Thus there are four workers and two firms. Assume that  $G = K_4$ . Consider the following initial matching:

$$\mu^{(0)}(f_1) = \{w_3, w_4\} \text{ and } \mu^{(0)}(f_2) = \{w_1, w_2\}$$

in other words, the best company has the worst workers. Then if we activate workers  $w_1$  and  $w_2$  before activating  $w_3$  or  $w_4$ , both  $w_1$  and  $w_2$  would quit  $f_2$  and work for  $f_1$ , getting both  $w_3$  and  $w_4$  fired. In that setting,  $f_2$  has no employees, and thus the process ends.

It is also important to understand the need of the activation process. Recall that the matching found by the Gale-Shapley algorithm is independent on the order of activation of the workers [17]. When considering locally stable matchings, this is no longer the case even if the underlying graph is the complete graph. Let us reconsider Example 3.

*Example 4.* Now consider the resulting matching when the activation sequence is as follows:  $\{w_1, w_4, w_2, w_3\}$ . First,  $w_1$  leaves  $f_2$  and gets a position at  $f_1$ . This makes  $w_4 = \min(\mu^{(0)}(f_1))$  unemployed. Next, since we activate  $w_4$ , she gets the free position from  $f_2$ . Next  $w_2$  leaves  $f_2$  and gets a position at  $f_1$ , which results in  $w_3$  loosing her job. Finally,  $w_3$  gets the free position at  $f_2$ . Thus the resulting matching is now

$$\mu(f_1) = \{w_1, w_2\}, \text{ and } \mu(f_2) = \{w_3, w_4\}$$

which is a locally stable matching different from that obtained with the activation sequence in Example 3.

An important question is whether this local decentralized version of best response dynamics converges as it is not immediately clear it can't cycle. This is the content of our first result.

**Theorem 5 (Convergence of Local Gale-Shapley Algorithm).** *Given the social network  $G(W, E)$ , for any initial matching  $\mu^{(0)}$ , the local Gale-Shapley algorithm started at  $\mu^{(0)}$  converges almost surely to a locally stable matching.*

## 4.2 Worst Case Efficiency

In this subsection we consider the following question. Given that firms can go out of business when running the local Gale-Shapley algorithm, can we measure the quality of matchings selected by the algorithm. We explore the previous question assuming a given initial complete matching  $\mu^{(0)}$ .

We consider the following setting. An adversary observes  $G(W, E)$  (but not the ranking over workers used by firms) and produces a probability distribution  $\mathcal{P}_M$  over initial matchings. The ranking of workers (possibly taken from a distribution) is then revealed, a sample from  $\mathcal{P}_M$  is taken to produce  $\mu^{(0)}$ ; and the local Gale-Shapley algorithm run.

To compare the efficiency of different final matchings we simply look at the total number of firms losing all of their employees and subsequently going out of business. One can easily imagine more intricate notions of efficiency, our point here is that even in this austere model, the power of the adversary is non-trivial.

**The power of the adversary.** We first show that even without knowing the relative rankings of the individual workers, the adversary is powerful enough to force some firms to go out of business.

**Theorem 6 (Lower Bound on Firms).** *Let  $G(W, E)$  be given. Let  $\Delta$  be its maximum degree, and  $M$  a maximum matching in  $G$ . Then there exist a probability distribution  $\mathcal{P}_M$  over complete assignment matchings such that*

$$\mathbb{E}[N_{\text{job}}] \geq \left\lfloor \frac{|M|}{k(2\Delta)} \right\rfloor \frac{1}{2^k k! (2\Delta - 1)^k}$$

where  $N_{\text{job}}$  is the number of firms going out of business; and the expectation is taken both over the distribution  $\mathcal{P}_M$  and over the activation process.

Further, one can find  $\mathcal{P}_M$  in time polynomial in  $n$ .

An important observation is that not only does the adversary force some firms to go out of business, but he controls the identities of these firms. Thus, if we measure efficiency by the identity of the firms of the positions filled in a matching, Theorem 6 provides a lower bound on the efficiency loss of the local Gale-Shapley algorithm (under adversarial initial conditions).

**The power of the social planner.** Given the lower bound from Theorem 6 on the expected number of firms going out of business, we can ask the following question: can similar guarantees be proven if a social planner had full control over the ranking used by firms? More precisely, given  $G(W, E)$ , if the ranking over workers used by firms was a sample from a random variable, can the social planner guarantee, in expectation, a minimal number of firms that will not go out of business *regardless of the power given to the adversary*? The following theorem answers positively that question.

**Theorem 7 (Upper Bound on Firms).** *Let  $G(W, E)$  be given. There exist a probability distribution over the ranking used by firms such that*

$$\mathbb{E}[N_{\text{job}}] \leq n_f - \left\lceil \frac{|I|}{k} \right\rceil$$

where  $I$  is a maximum independent set of  $G$  ( $n_f$  is the number of firms and  $k$  the number of positions at each firm)

Note that, just as in Theorem 6 we were able to identify the firms forced out of business (the top firms) but not the unemployed workers, in Theorem 7 we are able to identify the workers that are going to be employed (the top employees), but not which firms will remain in business.

**Discussion.** We have now shown that neither the adversary, nor the social planner have all the power — we can reinterpret the results above as a game between these two players. The game proceeds as follows, the adversary picks the initial assignment matching (possibly random), and the social planner chooses

the ordering on the workers (possibly random). Once they both pick an action we run the local Gale-Shapley algorithm.

Theorem 6 then states that, even if the social planner knows the probability distribution selected by the firm adversary, there is a probability distribution over initial assignments that the firm adversary can use such that, in expectation, at least some number of firms go out of business.

Theorem 7 states the converse: Even if the firm adversary knows the probability distribution selected by the social planner, there is a deterministic ordering of the workers such that at least some number of workers will never lose their job.

We note that by looking at the number of firms going out of business, we have used a very minimal notion of efficiency. It is not hard to imagine more complex notions which may take into account the relative rankings of the firms going out of business or workers remaining unemployed. We further note that for dense graphs, where the size of the independent set, and the independent matchings are quite small, our bounds are quite loose. Our main contribution here is not the precise bound on  $N_{fob}$ , although that remains an interesting open question, but rather the fact that the adversary has non-trivial power, and the initial matching plays a pivotal role in determining the final outcome.

## 5 Conclusions

In this work we have introduced a new model for incorporating social network ties into classical stable matching theory. Specifically, we show that restricting the firms willing to consider a worker only to those employing his friends has a profound impact on the system. We defined the notion of locally stable matchings and showed that while a simple variation of the Gale-Shapley mechanism converges to a stable solution, this solution may be far from efficient; and, unlike in traditional Gale-Shapley, the initial matching plays a large role in the final outcome. In fact, if the adversary controls the initial matching, he can force some firms to be left with *no* workers in the final solution.

The model we propose is ripe for extensions and further analysis. To give an example, we have assumed that as employees leave the firm, it may find itself with empty slots that it cannot fill (and go out of business). However, this is precisely the time when it can start looking actively for workers, by advertising online, recruiting through headhunters, etc. This has the effect of it becoming visible to the unemployed workers in the system. Understanding the dynamics and inefficiencies of final matchings under this scenario is one interesting open question.

## Acknowledgments

We would first like to thank Ramesh Johari for many discussions on the model studied and its static analysis. We also want to thank David Liben-Nowell for the discussions that lead to the proofs of Theorems 7 and 8. Finally, we would like to thank Ravi Kumar and Matt Jackson for many useful discussions.

This research was supported by the National Science Foundation and by the Defense Advanced Research Projects Agency under the ITMANET program.

## References

- [1] Abraham, D.J., Irving, R.W., Kavitha, T., Mehlhorn, K.: Popular matchings. *SIAM Journal of Computing* 37(4), 1030–1045 (2007)
- [2] Arcaute, E., Vassilvitskii, S.: Social networks and stable matchings in the job market. Working paper, arXiv.org (October 2009), <http://www.arxiv.org>
- [3] Calvó-Armengol, A., Jackson, M.O.: The effects of social networks on employment and inequality. *American Economic Review* 94(3), 426–454 (2004)
- [4] Calvó-Armengol, A., Jackson, M.O.: Networks in labor markets: Wage and employment dynamics and inequality. *Journal of Economic Theory* 132(1), 27–46 (2007)
- [5] Clark, S.: The uniqueness of stable matchings. *Contributions to Theoretical Economics* 6(1) (2006)
- [6] Gale, D., Shapley, L.: College admissions and the stability of marriage. *The American Mathematical Monthly* 69(1), 9–15 (1962)
- [7] Granovetter, M.: The strength of weak ties. *American Journal of Sociology* 78(6), 1360–1380 (1973)
- [8] Granovetter, M.: *Getting a Job: A Study of Contacts and Careers*. University Of Chicago Press, Chicago (1995)
- [9] Ioannides, Y.M., Loury, L.D.: Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature* 42(4), 1056–1093 (2004)
- [10] Jackson, M.O.: *Social and Economic Networks*. Princeton University Press, Princeton (2008)
- [11] Lee, R., Schwarz, M.: Interviewing in two-sided matching markets. In: 2008 American Economic Association Annual Meeting (2008)
- [12] Mahdian, M.: Random popular matchings. In: *Proceedings 7th ACM Conference on Electronic Commerce (EC 2006)*, Ann Arbor, Michigan, USA, June 11–15, 2006, pp. 238–242 (2006)
- [13] Montgomery, J.: Social networks and labor markets outcomes. *American Economic Review* 81(5), 1408–1418 (1991)
- [14] Niederle, M., Yariv, L.: Decentralized matching with aligned preferences. NBER Working Papers 14840, National Bureau of Economic Research, Inc. (April 2009)
- [15] Rees, A., Shultz, G.P.: *Workers in an Urban Labor Market*. University of Chicago Press, Chicago (1970)
- [16] Roth, A.E.: The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy* 92(6), 991–1016 (1984)
- [17] Roth, A.E., Sotomayor, M.A.O.: *Two-Sided Matching*. Cambridge University Press, Cambridge (1990)
- [18] Sorensen, M.: How smart is smart money? a two-sided matching model of venture capital. *The Journal of Finance* 62(6), 2725–2762 (2007)

# Maximizing the Minimum Load: The Cost of Selfishness

Leah Epstein<sup>1</sup>, Elena Kleiman<sup>1</sup>, and Rob van Stee<sup>2,\*</sup>

<sup>1</sup> Department of Mathematics, University of Haifa, 31905 Haifa, Israel  
lea@math.haifa.ac.il, elena.kleiman@gmail.com

<sup>2</sup> Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany  
vanstee@mpi-inf.mpg.de

**Abstract.** We consider a scheduling problem where each job is controlled by a selfish agent, who is only interested in minimizing its own cost, which is defined as the total load on the machine that its job is assigned to. We consider the objective of maximizing the minimum load (cover) over the machines. Unlike the regular makespan minimization problem, which was extensively studied in a game theoretic context, this problem has not been considered in this setting before.

We study the price of anarchy (POA) and the price of stability (POS). We show that on related machines, both these values are unbounded. We then focus on identical machines. We show that the POS is 1, and we derive tight bounds on the POA for  $m \leq 6$  and nearly tight bounds for general  $m$ . In particular, we show that the POA is at least 1.691 for larger  $m$  and at most 1.7. Hence, surprisingly, the POA is less than the POA for the makespan problem, which is 2. To achieve the upper bound of 1.7, we make an unusual use of weighting functions. Finally, in contrast we show that the mixed POA grows exponentially with  $m$  for this problem, although it is only  $\Theta(\log m / \log \log m)$  for the makespan.

## 1 Introduction

Classical optimization problems, and network optimization problems in particular, are often modeled as non-cooperative strategic games. Many solution concepts are used to study the behavior of selfish agents in non-cooperative games. Probably the best known concept is that of the Nash equilibrium. This is a state which is stable in the sense that no agent can gain from unilaterally switching strategies. Following recent interest of computer scientists in game theory [17,13,19], we study Nash equilibria for a scheduling problem where the goal is maximizing the minimum load.

This goal function is motivated by issues of Quality of Service and fair resource allocation. It is useful for describing systems where the complete system relies on keeping all the machines productive for as long as possible, as the entire system fails in case even one of the machines ceases to be active. From the networking aspect, this problem has applications to basic problems in network optimization

---

\* Research supported by the German Research Foundation (DFG).



such as fair bandwidth allocation. Consider pairs of terminal nodes that wish to communicate; we would like to allocate bandwidth to the connections in a way that no link unnecessarily suffers from starvation, and all links get a fair amount of resources. Another motivation is efficient routing of traffic. Consider parallel links between pairs of terminal nodes. Requests for shifting flow are assigned to the links. We are interested in having the loads of the links balanced, in the sense that each link should be assigned a reasonable amount of flow, compared to the other links. Yet another incentive to consider this goal function is congestion control by fair queuing. Consider a router that can serve  $m$  shifting requests at a time. The data pieces of various sizes, need to be shifted, are arranged in  $m$  queues (each queue may have a different data rate), each pays a price which equals the delay that it causes in the waiting line. Our goal function ensures that no piece gets a "preferred treatment" and that they all get at least some amount of delay.

The problem of maximizing the minimum load, seeing jobs as selfish agents, can be modeled as a routing problem. In this setting, machines are associated with parallel links between a source and a destination. The links have bounded capacities, and a set of users request to send a certain amount of unsplitable flow between the two nodes. Requests are to be assigned to links and consume bandwidth which depends on their sizes. The cost charged from a user for using a link equals to the total amount of the utilized bandwidth of that link. Thus, the selfish users prefer to route their traffic on a link with small load. This scenario is similar to the model proposed by Koutsoupias and Papadimitriou [13], but our model has a different social goal function. To demonstrate the non-triviality of the problem, see Figure 1.

The novelty of our study compared to other work in the area is that the social goal is very different from the private goals of the players.

In our scheduling model, the *coordination ratio*, or *price of anarchy* (POA) [18] is the worst case ratio between the social value (i.e., minimum delay of any machine, or cover) of an optimal schedule, denoted by OPT, and the value of any Nash equilibrium. If both these values are 0 then we define the POA to be 1. The *price of stability* (POS) [1] is the worst case ratio between the social value of an optimal solution, and the value of the *best* Nash equilibrium. Similarly, if both these values are 0 then we define the POS to be 1.

In addition, we study the *mixed* POA (MPOA), where we consider mixed Nash equilibria that result from mixed strategies, where the player's choices are not deterministic and are regulated by probability distributions on a set  $M$  of pure strategies. A mixed Nash equilibrium is characterized by the property that there is no incentive for any job to deviate from its probability distribution (a deviation is any modification of its probability vector over machines), while probability distributions of other players remain unchanged. The existence of such an equilibrium over mixed strategies for non-cooperative games was shown by Nash in his famous work [16]. The values MPOA and MPOS are defined similarly to the pure ones, but mixed Nash equilibria are being considered instead of pure ones. Clearly, any pure NE is also a mixed NE.

*Our results and related work.* The non-selfish version of the problem has been well studied (known by different names such as "machine covering" and "Santa Claus problem") in the computer science literature (see e.g. [4,2,6]). Various game-theoretic aspects of max-min fairness in resource allocation games were considered before this paper (e.g. in [20]), but unlike the makespan minimization problem POA and POS of which were extensively studied (see [13,3,15]), these measures were not previously considered for the uncoordinated machine covering problem in the setting of selfish jobs. A different model, where machines are selfish rather than jobs with the same social goal was studied recently in [7,5].


For identical machines, we show that the POS is equal to 1. As our main result, we study the pure POA and show close bounds on the overall value of the POA ( $\text{POA} = \sup_m \text{POA}(m)$ , where  $\text{POA}(m)$  is the POA on  $m$  machines), i.e., that it is at least 1.691 and at most 1.7. This in contrast with the makespan minimization problem, where it is known that the POA for  $m$  identical machines is  $\frac{2m}{m+1}$ , giving an overall bound of 2 [10,21]. This is rather unusual, as the cover maximization problem is typically harder than the makespan minimization problem, thus it could be expected that the POA for the covering problem would be higher.

For the analysis of our upper bound we use the weighting function technique, which is uncommon in scheduling problems. Moreover, we use not only the weight function but also its inverse function in our analysis. Surprisingly, these lower and upper bounds are approximation ratios of well known algorithms for Bin-Packing (Harmonic [14] and First-Fit [12], respectively). We furthermore prove that the POA is monotonically non-decreasing as a function of  $m$ . For small numbers of machines, in the full version we provide the exact values of POA: we find  $\text{POA}(2) = \text{POA}(3) = 3/2$  and  $\text{POA}(4) = 13/8 = 1.625$ . We show  $\text{POA}(m) \geq \frac{5}{3}$  for  $m > 5$ . As for the MPOA, we show that its value is very large as a function of  $m$ , and  $\text{MPOA}(2) = 2$ .

In contrast to these results, we can show that for uniformly related machines even the POS is unbounded already for two machines with a speed ratio *larger than 2*, and the POA is unbounded for a speed ratio of *at least 2*. The same property holds for  $m$  machines (where the speed ratio is defined to be the maximum speed ratio between any pair of machines). These results are very different from the situation for the makespan minimization social goal. For that problem, the POS is 1 for any speed combination. Chumaj and Vöcking [3] showed that the overall POA is  $\Theta(\frac{\log m}{\log \log m})$  (see also [9]).

## 2 The Model

In this section, we define the more general model of scheduling on related machines, which we will consider first. A set of  $n$  jobs  $J = \{1, 2, \dots, n\}$  is to be assigned to a set of  $m$  machines  $M = \{M_1, \dots, M_m\}$ , where machine  $M_i$  has a speed  $s_i$ . If  $s_i = 1$  for  $i = 1, \dots, m$ , the machines are called identical. This is an important and widely studied special case of uniformly related machines. The size of job  $1 \leq k \leq n$  is denoted by  $p_k$ . An assignment or schedule is a function  $A : J \rightarrow M$ . The load of machine  $M_i$ , which is also called the delay of this



**Fig. 1.** An example of two packings with different social values. This example demonstrates the non-triviality of the problem. There are three jobs of size 0.8, three jobs of size 0.4 and two jobs of size 0.1. The three machines are identical. The assignment on the right hand side is not a Nash equilibrium, since a job of size 0.1 would reduce its delay from 1.4 to 1.3 by migrating to another machine. The social value of this assignment is 1.2. The assignment on the left hand side is a Nash equilibrium, but its social value is only 1.

machine, is  $L_i = \sum_{k:\mathcal{A}(k)=M_i} \frac{p_k}{s_i}$ . The value, or the *social value* of a schedule is the minimum delay of any machine, also known as the *cover*. We denote it by  $\text{COVER}(\mathcal{A})$ . This problem is a dual to the makespan scheduling problem.

The non-cooperative machine covering game  $MC$  is characterized by a tuple  $MC = \langle N, (\mathcal{M}_k)_{k \in N}, (c_k)_{k \in N} \rangle$ , where  $N$  is the set of atomic players. Each player  $k \in N$  controls a single job of size  $p_k > 0$  and selects the machine to which it will be assigned. We associate each player with the job it wishes to run, that is,  $N = J$ . The set of strategies  $\mathcal{M}_k$  for each job  $k \in N$  is the set  $M$  of all machines. i.e.  $\mathcal{M}_k = M$ . Each job must be assigned to one machine only. Preemption is not allowed. The outcome of the game is an assignment  $\mathcal{A} = (\mathcal{A}_k)_{k \in N} \in \times_{k \in N} M_k$  of jobs to the machines, where  $\mathcal{A}_k$  for each  $1 \leq k \leq n$  is the index of the machine that job  $k$  chooses to run on. Let  $\mathcal{S}$  denote the set of all possible assignments. The cost function of job  $k \in N$  is denoted by  $c_k : \mathcal{S} \rightarrow \mathbb{R}$ . The cost  $c_k^i$  charged from job  $k$  for running on machine  $M_i$  in a given assignment  $\mathcal{A}$  is defined to be the load observed by machine  $i$  in this assignment, that is  $c_k(i, \mathcal{A}_{-k}) = L_i(\mathcal{A})$ , when  $\mathcal{A}_{-k} \in \mathcal{S}_{-k}$ ; here  $\mathcal{S}_{-k} = \times_{j \in N \setminus \{k\}} \mathcal{S}_j$  denotes the actions of all players except for player  $k$ .

The goal of the selfish jobs is to run on a machine with a load which is as small as possible. At an assignment that is a (pure) Nash equilibrium or NE assignment for short, there exists no machine  $M_{i'}$  for which  $L_{i'}(\mathcal{A}) + \frac{p_k}{s_{i'}} < L_i(\mathcal{A})$  for some job  $k$  which is assigned to machine  $M_i$  (see Figure 1(a) for an example). For this selfish goal of players, a pure Nash equilibrium (with deterministic agent choices) always exists [11, 8]. We can also consider mixed strategies, where players use probability distributions. Let  $t_k^i$  denote the probability that job  $k \in N$  chooses to run on machine  $M_i$ . A strategy profile is a vector  $p = (t_k^i)_{k \in N, i \in M}$  that specifies the probabilities for all jobs and all machines. Every strategy profile  $p$

induces a random schedule. The *expected load*  $\mathbb{E}(L_i)$  of machine  $M_i$  in setting of mixed strategies is  $\mathbb{E}(L_i) = \frac{1}{s_i} \sum_{k \in N} p_k t_k^i$ . The *expected cost* of job  $k$  if assigned on machine  $M_i$  (or its *expected delay* when it is allocated to machine  $M_i$ ) is  $\mathbb{E}(c_k^i) = \frac{p_k}{s_i} + \sum_{j \neq k} p_j t_j^i / s_i = \mathbb{E}(L_i) + (1 - t_k^i) \frac{p_k}{s_i}$ . The probabilities  $(t_k^i)_{k \in N, i \in M}$  give rise to a (*mixed*) Nash equilibrium if and only if any job  $k$  will assign non-zero probabilities only to machines  $M_i$  that minimize  $c_k^i$ , that is,  $t_k^i > 0$  implies  $c_k^i \leq c_k^j$  for any  $j \in M$ . The social value of a strategy profile  $p$  is the *expected minimum load* over all machines, i.e.  $\mathbb{E}(\min_{i \in M} L_i)$ .

We omit some of the proofs due to space constraints.

### 3 Related Machines

In the setting of related machines, we show that for large enough speed ratios the POA and the POS are unbounded already for two machines.

Denote the speeds of the  $m$  machines by  $s_1, s_2, \dots, s_m$ , where  $s_i \leq s_{i+1}$  and let  $s = \frac{s_m}{s_1}$ . Without loss of generality, we assume that the fastest machine has speed  $s \geq 1$ , and the slowest machine has speed 1.

**Theorem 1.** *On related machines with speed ratio  $s$ ,  $\text{POA}(s) = \infty$  for  $s \geq 2$  and  $\text{POS}(s) = \infty$  for  $s > 2$ .*

*Proof.* We only consider the case  $s > 2$  here. Consider an instance that contains  $m$  identical sized jobs of size  $s$ . Clearly,  $\text{COVER}(\text{OPT}) = 1$  for this input.

For  $s > 2$ , we show that any assignment where each job is assigned to a different machine is not a Nash equilibrium. In fact, in such an assignment, any job assigned to the first machine sees a load of  $s$ , while if it moves to the  $m$ -th machine, its load becomes  $\frac{2s}{s} = 2 < s$ . Thus, any NE assignment has a cost of 0 and the claim follows.  $\square$

### 4 The POS for $m$ Identical Machines

**Theorem 2.** *On identical machines,  $\text{POS} = 1$  for any  $m$ .*

*Proof.* We show that for every instance of the machine covering game, among the optimal assignments there exists an optimal assignment which is also an NE. Our proof technique is based upon the technique which was used in [8,11] to prove that in job scheduling games where the selfish goal of the players is run on the least loaded machine (like in our machine covering game), any sequence of improvement steps converges to an NE.

We first define a complete order relation on the assignments, and then show that an optimal assignment which is the “highest” among all optimal assignments with respect to this order is always an NE.

**Definition 1.** *A vector  $(l_1, l_2, \dots, l_m)$  is larger than  $(l'_1, l'_2, \dots, l'_m)$  with respect to the inverted lexicographic order, if for some  $i$ ,  $l_i > l'_i$  and  $l_k = l'_k$  for all  $k > i$ . An assignment  $s$  is called larger than  $s'$  according to the inverted lexicographic*

order if the vector of machine loads  $L(s) = (L_1(s), L_2(s), \dots, L_m(s))$ , sorted in non-increasing order, is larger in the inverted lexicographic order than the vector  $L(s') = (L_1(s'), L_2(s'), \dots, L_m(s'))$ , sorted in non-increasing order. We denote this relation by  $s \succ_{L^{-1}} s'$ .

**Lemma 1.** *For any instance of the machine covering game, a maximal optimal packing w.r.t. the inverted lexicographic order is an NE.*

Since for any set of  $n$  jobs there are finitely many possible assignments, among the assignments that are optimal with respect to our social goal there exists at least one which is maximal w.r.t. the total order  $\succ_{L^{-1}}$ , and according to Lemma 1 this assignment is an NE. As no NE assignment can have a strictly greater social value than the optimal one, we conclude that  $\text{POS} = 1$ .  $\square$


## 5 The Price of Anarchy

Figure 1 clearly demonstrates that not every NE schedule is optimal. We next measure the extent of deterioration in the quality of NE schedules due to the effect of selfish and uncoordinated behavior of the players (jobs), in the worst case. As mentioned before, the measure metrics we use are the POA and the POS. In Figure 2, we give some lower bounds on the POA for small  $m$  (without proof).

Consider a pure NE assignment of jobs to machines, denoted by  $\mathcal{A}$ , for an instance of the machine covering game. We assume that the social value of  $\mathcal{A}$ , that is, the load of the least loaded machine in  $\mathcal{A}$ , is 1. Otherwise, we can simply scale all sizes of jobs in the instances which we consider so that  $\text{COVER}(\mathcal{A}) = 1$ .

We denote a machine which is loaded by 1 in  $\mathcal{A}$  by  $P$ . All other machines are called *tall* machines. We would like to estimate the load of  $P$  in the optimal assignment. Let  $C = \text{COVER}(\text{OPT})$ . Obviously,  $C \geq 1$ , and the total sum of jobs sizes, denoted by  $W$ , satisfies  $W \geq mC$ . First, we introduce some assumptions on  $\mathcal{A}$ . Note that the modifications needed to be applied such that this instance will satisfy these assumptions do not increase  $\text{COVER}(\mathcal{A})$ , do not violate the conditions for NE and do not decrease  $\text{COVER}(\text{OPT})$ .

1. Machine  $P$  contains only tiny jobs, that is, jobs of infinitesimal size.  
 Since no machine has a smaller load, replacing the jobs on this machine by tiny jobs keeps the schedule as an NE. The value  $\text{COVER}(\text{OPT})$  may only increase.
2. For a tall machine in  $\mathcal{A}$  which has two jobs, both jobs have a size of 1.  
 If one of them is larger, then the second job would want to move to  $P$ , so this case cannot occur. If some such job is smaller, its size can be increased up to 1 without affecting the NE.
3. This assignment is minimal with respect to the number of machines (among assignments for which  $\text{COVER}(\text{OPT}) \geq C$ ). In particular, no machine in  $\mathcal{A}$  has a single job.  
 Else, if some machine has a single job, remove this machine and the job from  $\mathcal{A}$ , and the machine with it from the optimal assignment  $\text{OPT}$ . Assign



**Fig. 2.** Lower bounds for the price of anarchy for small  $m$ . Machines are on the horizontal axis, jobs are on the vertical axis. The squares in the first few figures represent jobs of size 1.

any remaining jobs that ran on this machine in OPT arbitrarily among the remaining machines. This gives a new assignment with  $\text{COVER}(\text{OPT}) \geq C$ , and less machines.

4. Given jobs of sizes  $p_1 \leq p_2 \leq \dots \leq p_t$  assigned to a machine  $Q$  in  $\mathcal{A}$ , then  $p_2 + \dots + p_t = 1$ . In fact,  $p_2 + \dots + p_t > 1$  is impossible since this would mean that the job of size  $p_1$  has an incentive to move. If the sum is less, enlarge the size  $p_t$  to  $1 - p_{t-1} - \dots - p_2$ . This does not affect the NE conditions, and keeps the property  $\text{COVER}(\text{OPT}) \geq C$ .
5. Consider a machine  $Q \neq P$  in  $\mathcal{A}$  which has  $t \geq 3$  jobs assigned to it. Let  $a, b$  denote the sizes of the smallest and largest jobs on it, respectively. Then,  $b < 2a$ .

Otherwise, if we have an assignment where  $b \geq 2a$ , replace  $b$  with two jobs of size  $\frac{b}{2}$ . This modification preserves the NE, as the new jobs do not have an incentive to move; Let  $T$  denote the total size of jobs on machine  $Q$ . As  $a$  does not want to move to  $P$ ,  $T \leq 1 + a$  holds. As in this case  $a \leq \frac{b}{2}$ , we have  $T \leq 1 + \frac{b}{2}$ , whereas  $1 + \frac{b}{2}$  would have been the load of  $P$  if the job  $\frac{b}{2}$  moved there.

**Lemma 2.** *No job has a size larger than 1.*

*Proof.* There is no machine in  $\mathcal{A}$  with a single job. □

**Lemma 3.** *There is no job of size in  $[\frac{2}{3}, 1)$  assigned to a machine  $Q$  ( $Q \neq P$ ) in  $\mathcal{A}$ .*

*Proof.* If there is such a job, then it has at least two jobs assigned together with it, each of size greater than  $\frac{1}{3}$  (due to assumption 5), which contradicts assumption 4. □

We define a weight function  $w(x)$  on sizes of jobs.

$$w(x) = \begin{cases} \frac{1}{2} & , \text{ for } x = 1 \\ \frac{x}{2-x} & , \text{ for } x \in (\frac{1}{2}, \frac{2}{3}) \\ \frac{x}{x+1} & , \text{ for } x \in (0, \frac{1}{2}] \end{cases}$$

The motivation for the weight function is to define the weight of a job  $j$  to be at least the fraction of its size out of the total size of jobs assigned to the same machine in  $\mathcal{A}$ . In fact, for a job  $j$  of size  $x$ , the total size of jobs assigned in an NE to the same machine as  $j$  is no larger than  $1 + x$ . Moreover if  $x > \frac{2}{3}$ , then by our assumptions, it is possible to prove that the total size of these jobs is at most  $2 - x$ .

Its inverse function  $f(y)$  is

$$f(y) = \begin{cases} 1 & , \text{ for } y = \frac{1}{2} \\ \frac{2y}{y+1} & , \text{ for } y \in (\frac{1}{3}, \frac{1}{2}) \\ \frac{y}{1-y} & , \text{ for } y \in (0, \frac{1}{3}] \end{cases}$$

Note that  $f(y)$  is continuous at  $\frac{1}{3}$  but not at  $\frac{1}{2}$ . Both functions are monotonically increasing.

We now state several lemmas, which follow from the properties of this weight function defined above.

**Lemma 4.** *The total weight (by  $w$ ) of jobs on  $P$  in  $\mathcal{A}$  is less than 1.*

*Proof.* Follows from the fact that  $P$  has a load 1 and for any  $x$ ,  $w(x) < x$ . □

**Lemma 5.** *The total weight (by  $w$ ) of jobs assigned to a machine  $Q \neq P$  in  $\mathcal{A}$  is at most 1.*

*Proof.* By Assumption 3', each machine  $Q$  in  $\mathcal{A}$  has at least two jobs. The claim clearly holds for a machine with two jobs, as by Assumption 2 both these jobs have size 1, and by the definition of  $w$  have a total weight of  $\frac{1}{2} + \frac{1}{2} = 1$ . For a machine which has at least three jobs assigned to it, there are two cases. If there are no jobs with size in  $(\frac{1}{2}, \frac{2}{3})$  (i.e., all jobs are of size in  $(0, \frac{1}{2}]$ ), then let  $T$  be the total size of jobs on  $Q$ . For each job of size  $x_i$ , as the job does not want to move to machine  $P$ ,  $x_i + 1 \geq T$  holds. Combining this with the definition of  $w$  for jobs of size in  $(0, \frac{1}{2}]$ , we get that  $w(x_i) = \frac{x_i}{x_i+1} \leq \frac{x_i}{T}$ . Summing this up over all the jobs on  $Q$  proves the claim.

Else, for a job  $x_i \in (\frac{1}{2}, \frac{2}{3})$  (there can be only one such job assigned to  $Q$ , by Assumption 4), we have  $T \leq 2 - x_i$ . Otherwise, let  $a$  be the size of the smallest job on  $Q$ . Then, by Assumption 4,  $T = 1 + a$ . As  $T > 2 - x_i$ , we get  $x_i + a > 1$ . Therefore, as there is an additional job of size of at least  $a$  assigned to  $Q$ , we get that the total size of all jobs except for the smallest job is more than 1, contradicting Assumption 4.

Combining this with the definition of  $w$  for jobs of size in  $(\frac{1}{2}, \frac{2}{3})$ , we get that  $w(x_i) = \frac{x_i}{2-x_i} \leq \frac{x_i}{T}$  for this job too. □

**Lemma 6.** *There is a machine in the optimal assignment with a total weight strictly smaller than 1.*

*Proof.* The total weight of all jobs is less than  $m$  by Lemmas 4 and 5. □

**Lemma 7.** *The total size of any set of jobs with total weight below 1 is at most 1.7.*

*Proof (Sketch).* To prove the claim, we use the property that there is at most one item of weight 1. If it exists, then there is at most one item of weight larger than  $\frac{1}{3}$ . If such an item of weight 1 does not exist, then the ratio between size and weight is at most 1.5. We find the supremum possible weight in each one of three cases.

Consider a set of jobs  $I$  of a total weight strictly below 1. Note that for any  $x$ ,  $w(x) \leq \frac{1}{2}$ . If there is no job of weight  $\frac{1}{2}$  in  $I$ , then since  $\frac{2}{y+1} \leq \frac{3}{2}$  for  $y \geq \frac{1}{3}$  and  $\frac{1}{1-y} \leq \frac{3}{2}$  for  $y \leq \frac{1}{3}$ , the size of any job is at most  $\frac{3}{2}$  times its weight, and thus the total size of the jobs in  $I$  does not exceed  $\frac{3}{2}$ .

Otherwise, there is exactly one job of weight  $\frac{1}{2}$ , and its size is 1. We therefore need to show that the total size of any set of jobs  $I'$  which has total weight below  $\frac{1}{2}$  is at most 0.7. There is at most one job of weight in  $(\frac{1}{3}, \frac{1}{2})$  in  $I'$ . If there is one such job we show that without loss of generality, there is at most one other job in  $I'$ , and its weight is in  $(0, \frac{1}{3}]$ . Else, we show that there are at most two jobs, and their weights are in  $(0, \frac{1}{3}]$ .

Note that  $f_1(y) = \frac{y}{1-y}$  is a convex function, thus, for any pair of jobs of weights  $\alpha, \beta \in (0, \frac{1}{3}]$ ,  $f_1(0) + f_1(\alpha + \beta) \geq f_1(\alpha) + f_1(\beta)$  holds. As we can define  $f_1(0) = 0$ , it turns into  $f_1(\alpha + \beta) \geq f_1(\alpha) + f_1(\beta)$ .

Thus, any two jobs of total weight of at most  $\frac{1}{3}$  can be combined into a single job while as a result, their total size cannot decrease. This is due to the convexity of  $f_1$  and the fact that it is monotonically increasing. The replacement may only increase the weight, and respectively, the size. If there exists a job of weight larger than  $\frac{1}{3}$ , then the total weight of jobs of weight at most  $\frac{1}{3}$  is at most  $\frac{1}{6}$ , so they can all be combined into a single job. Moreover, among any three jobs of a total weight of at most  $\frac{1}{2}$ , there exists a pair of jobs of total weight no larger than  $\frac{1}{3}$ , which can be combined as described above, so if there is no job of weight larger than  $\frac{1}{3}$ , still jobs can be combined until at most two jobs remain. Thus, there are only two cases to consider.

*Case 1.* There is one job of weight in  $(0, \frac{1}{3}]$  and one job of weight in  $(\frac{1}{3}, \frac{1}{2})$ . Since the inverse function  $f$  is monotonically increasing as a function of  $y$  and their weight does not exceed  $\frac{1}{2}$ , we can assume that their total weight is  $\frac{1}{2} - \gamma$  for a negligible value of  $\gamma > 0$  (by increasing the weight the job of the smaller weight, which may only increase the total size). Letting  $d < \frac{1}{6}$  denote the weight of the smaller job (since if  $d \geq \frac{1}{6}$  then  $\frac{1}{2} - \gamma - d < \frac{1}{3}$ ), we get a size of at most

$$\frac{d}{1-d} + \frac{2(\frac{1}{2} - \gamma - d)}{\frac{1}{2} - \gamma - d + 1} < \frac{d}{1-d} + \frac{2-4d}{3-2d}$$

(by letting  $\gamma \rightarrow 0$ ). This function is increasing (as a function of  $d$ ) so its greatest value is for  $d \rightarrow \frac{1}{6}$  and it is 0.7. As the inverse function  $f$  is monotonically increasing, this case also encompass the case where there is only one job of weight in  $(\frac{1}{3}, \frac{1}{2})$ , and no jobs of weight in  $(\frac{1}{3}, \frac{1}{2})$ .

*Case 2.* There are at most two jobs, where each job has a weight in  $(0, \frac{1}{3}]$ . If there is at most one job, then its size is at most  $\frac{1}{2}$ . We therefore focus on the



case of exactly two such jobs. Recall that the total size of the two jobs is larger than  $\frac{1}{3}$  (since they cannot be combined). The total weight of these jobs is less than  $\frac{1}{2}$ , so we can assume that their total weight is  $1/2 - \gamma$  for a negligible value of  $\gamma > 0$  (increasing the respective size). Let the weight of the large of these jobs be  $d > \frac{1}{6}$ . We get (by letting  $\gamma \rightarrow 0$ ) a total size of at most

$$\frac{d}{1-d} + \frac{\frac{1}{2} - \gamma - d}{\frac{1}{2} + d + \gamma} < \frac{d}{1-d} + \frac{1-2d}{2d+1},$$

where  $\frac{1}{6} < d \leq \frac{1}{3}$ . This function is monotonically decreasing in  $(\frac{1}{6}, \frac{1}{4}]$  and increasing in  $(\frac{1}{4}, \frac{1}{3}]$ , and its values at the endpoints  $\frac{1}{6}$  and  $\frac{1}{3}$  are both 0.7.  $\square$

**Theorem 3.** *For covering identical machines, the POA is at most 1.7.*

*Proof.* This follows from Lemmas 6 and 7.  $\square$

**Theorem 4.** *For covering identical machines, the POA is at least 1.691.*

*Proof.* We first define a sequence  $t_i$  of positive integers, which is often used in the literature for analysis and proving of lower bounds for online bin packing algorithms. Let  $t_1 = 1$  and  $t_{i+1} = t_i(t_i + 1)$  for  $i \geq 1$ .

Let  $m = t_k$  for an integer  $k$ . Consider the following assignment  $\mathcal{A}$ , that has  $\frac{m}{t_i+1}$  machines with  $t_i + 1$  jobs of size  $\frac{1}{t_i}$ , (for  $1 \leq i < k$ ) and one machine (i.e.,  $\frac{m}{t_k}$  machines) with  $t_k$  jobs of size  $\frac{1}{t_k}$ . We assume that the machines are sorted in a non-increasing order w.r.t. their load. We define the load class  $i$ ,  $1 \leq i \leq k$  as the subset of  $\frac{m}{t_i+1}$  machines with the same load  $L_i = \frac{t_i+1}{t_i} = 1 + \frac{1}{t_i} > 1$  in this assignment. As  $t_i$  is an increasing sequence of integers, it follows that  $L_i$  is monotonically non-increasing as a function of  $i$ . Since  $L_k = 1$ , the social value of this assignment is 1. We now verify that it is a Nash equilibrium. As  $L_i > 1$  for any  $1 \leq i < k$ , no job will benefit from leaving the machine of class  $L_k$ . It is enough to show that any job assigned to a machine of a class  $L_i$  ( $1 \leq i < k - 1$ ) would not benefit from moving to the machine of class  $k$ . Since machine  $i$  has  $t_i + 1$  jobs of size  $\frac{1}{t_i}$ , the migration of such a job to the machine of load 1 would again result in a load of  $1 + \frac{1}{t_i}$ , thus the job would not benefit from the migration.

In the socially optimal assignment, each machine has a set of jobs of distinct sizes,  $1, \frac{1}{2}, \frac{1}{6}, \dots, \frac{1}{t_{k-1}}, \frac{1}{t_k}$ . The social value of this assignment is  $\sum_{i=1}^k \frac{1}{t_i}$ . Thus, the POA equals  $\sum_{i=1}^k \frac{1}{t_i}$ . For  $k \rightarrow \infty$ , this value tends to  $h_\infty = \sum_{i=1}^\infty \frac{1}{t_i} = 1.69103\dots$ , the well-known worst-case ratio of the Harmonic algorithm for bin packing. As the POA is monotonically non-decreasing as a function of the number of the machines, we conclude that this is a lower bound for any number of machines larger than  $t_k$ . In particular, the overall POA for identical machines is at least 1.69103.  $\square$

*Conjecture 1.* For covering identical machines,  $\text{POA} = \sum_{i=1}^\infty \frac{1}{t_i} = 1.691\dots$

## 6 Mixed Equilibria

In the setting of mixed strategies we consider the case of identical machines, similarly to [13]. In that work, it was shown that the mixed POA for two machines is  $\frac{3}{2}$ . In this section we prove that the mixed POA for two machines is equal to 2.

We start by showing that for  $m$  identical machines, the mixed POA can be exponentially large as a function of  $m$ , unlike the makespan minimization problem, where the mixed POA is  $\Theta(\frac{\log m}{\log \log m})$  [13].

**Theorem 5.** *The mixed POA for  $m$  identical machines is at least  $\frac{m^m}{m!}$ .*

*Proof.* Consider the following instance  $G \in MC$  of the Machine Covering game.  $N = \{1, 2, \dots\}$  such that  $p_1 = \dots = p_n = 1$ , and let  $M_j = \{M_1, \dots, M_m\}$ , for  $j = 1, \dots, n$ . Each of the jobs  $p_i$ ,  $i = 1, \dots, n$  chooses each machine with probability  $t_i^j = 1/m$ . Each job sees the same expected load for each machine, and thus has no incentive to change its probability distribution vector. We get a schedule having a non-zero cover, where each job chooses to run on a different machine, with a probability of  $\frac{m!}{m^m}$ . So, for the mixed Nash equilibrium the expected minimum load is  $\frac{m!}{m^m}$ . But the coordinated optimal solution achieved by deterministically allocating each job to its own machine has a social value  $\text{COVER}(\text{OPT}(G)) = 1$ , and so it follows that the mixed  $\text{POA}(G) = \frac{m^m}{m!}$ . We conclude that the mixed  $\text{POA} \geq \frac{m^m}{m!}$ .

**Theorem 6.** *The mixed POA for two identical machines is exactly 2.*

## 7 Summary and Conclusion

In this paper we have studied a non-cooperative variant of the machine covering problem for identical and related machines, where the selfish agents are the jobs. We considered both pure and mixed strategies of the agents. We provided various results for the POA and the POS that are the prevalent measures of the quality of the equilibria reached with uncoordinated selfish agents.

For the pure POA for  $m$  identical machines, we provided nearly tight lower and upper bounds of 1.691 and 1.7, respectively. An obvious challenge would be bridging this gap. As stated in the paper, we believe that the actual bound is the lower bound we gave.

## References

1. 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 the 45th Symposium on Foundations of Computer Science (FOCS 2004), pp. 295–304 (2004)
2. Bansal, N., Sviridenko, M.: The santa claus problem. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC), pp. 31–40 (2006)

3. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms* 3(1) (2007)
4. Deuermeier, B.L., Friesen, D.K., Langston, M.A.: Scheduling to maximize the minimum processor finish time in a multiprocessor system. *SIAM Journal on Discrete Mathematics* 3(2), 190–196 (1982)
5. Dhangwatnotai, P., Dobzinski, S., Dughmi, S., Roughgarden, T.: Truthful approximation schemes for single-parameter agents. In: 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 15–24 (2008)
6. Ebenlendr, T., Noga, J., Sgall, J., Woeginger, G.J.: A note on semi-online machine covering. In: Erlebach, T., Persinao, G. (eds.) WAOA 2005. LNCS, vol. 3879, pp. 110–118. Springer, Heidelberg (2006)
7. Epstein, L., van Stee, R.: Maximizing the minimum load for selfish agents. In: Laber, E.S., Bornstein, C., Nogueira, L.T., Faria, L. (eds.) LATIN 2008. LNCS, vol. 4957, pp. 264–275. Springer, Heidelberg (2008)
8. Even-Dar, E., Kesselman, A., Mansour, Y.: Convergence time to nash equilibrium in load balancing. *ACM Trans. Algorithms* 3(3), 32 (2007)
9. Feldmann, R., Gairing, M., Lüking, 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)
10. Finn, G., Horowitz, E.: A linear time approximation algorithm for multiprocessor scheduling. *BIT Numerical Mathematics* 19(3), 312–320 (1979)
11. Fotakis, D., 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)
12. Johnson, D.S., Demers, A., Ullman, J.D., Garey, M.R., Graham, R.L.: Worst-case performance bounds for simple one-dimensional packing algorithms. *SIAM Journal on Computing* 3(4), 299–325 (1974)
13. 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)
14. Lee, C.C., Lee, D.T.: A simple online bin packing algorithm. *Journal of the ACM* 32(3), 562–572 (1985)
15. Mavronicolas, M., Spirakis, P.G.: The price of selfish routing. In: Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001), pp. 510–519 (2001)
16. Nash, J.: Non-cooperative games. *Annals of Mathematics* 54(2), 286–295 (1951)
17. Nisan, N., Ronen, A.: Algorithmic mechanism design. *Games and Economic Behavior* 35, 166–196 (2001)
18. Roughgarden, T.: *Selfish routing and the price of anarchy*. MIT Press, Cambridge (2005)
19. Roughgarden, T., Tardos, É.: How bad is selfish routing? *Journal of the ACM* 49(2), 236–259 (2002)
20. Sahasrabudhe, A., Kar, K.: Bandwidth allocation games under budget and access constraints. In: CISS, pp. 761–769 (2008)
21. Schuurman, P., Vredeveld, T.: Performance guarantees of local search for multiprocessor scheduling. *INFORMS Journal on Computing* 19(1), 52–63 (2007)

# Competitive Repeated Allocation without Payments<sup>\*</sup>

Mingyu Guo<sup>1</sup>, Vincent Conitzer<sup>1</sup>, and Daniel M. Reeves<sup>2</sup>

<sup>1</sup> Duke University, Department of Computer Science, Durham, NC, USA  
{mingyu, conitzer}@cs.duke.edu

<sup>2</sup> Yahoo! Research, 111 W. 40th St. 17th Floor, New York, NY, USA  
dreeves@yahoo-inc.com

**Abstract.** We study the problem of allocating a single item repeatedly among multiple competing agents, in an environment where monetary transfers are not possible. We design (Bayes-Nash) incentive compatible mechanisms that do not rely on payments, with the goal of maximizing expected social welfare. We first focus on the case of two agents. We introduce an artificial payment system, which enables us to construct *repeated* allocation mechanisms *without payments* based on *one-shot* allocation mechanisms *with payments*. Under certain restrictions on the discount factor, we propose several repeated allocation mechanisms based on artificial payments. For the simple model in which the agents' valuations are either high or low, the mechanism we propose is 0.94-competitive against the optimal allocation mechanism with payments. For the general case of any prior distribution, the mechanism we propose is 0.85-competitive. We generalize the mechanism to cases of three or more agents. For any number of agents, the mechanism we obtain is at least 0.75-competitive. The obtained competitive ratios imply that for repeated allocation, artificial payments may be used to replace real monetary payments, without incurring too much loss in social welfare.

## 1 Introduction

An important class of problems at the intersection of computer science and economics deals with allocating resources among multiple competing agents. For example, an operating system allocates CPU time slots to different applications. The resources in this example are the CPU time slots and the agents are the applications. Another example scenario, closer to daily life, is “who gets the TV remote control.” Here the resource is the remote control and the agents are the members of the household. In both scenarios the resources are allocated repeatedly among the agents, and monetary transfers are infeasible (or at least inconvenient). In this paper, we investigate problems like the above. That is, we study how to allocate resources in a repeated setting, without relying on payments. Our objective is to maximize social welfare, i.e., allocative efficiency.

The problem of allocating resources among multiple competing agents when monetary transfers are possible has been studied extensively in both the one-shot mechanism design setting [9,6,20,16,19,15] and the repeated setting [11,7,10,5]. A question that has recently been drawing the attention of computer scientists is how to design mechanisms without payments to achieve competitive performance against mechanisms with

---

<sup>\*</sup> 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).

payments [2][13].<sup>1</sup> This paper falls into this category. We consider mechanisms without payments in repeated settings. A paper that lays out many of the foundations for repeated games is due to Abreu *et al.* [2], in which the authors investigate the problem of finding pure-strategy sequential equilibria of repeated games with imperfect monitoring. Their key contribution is the state-based approach for solving repeated games, where in equilibrium, the game is always in a *state* which specifies the players' long-run utilities, and on which the current period's payoffs are based. There are many papers that rely on the same or a similar state-based approach [22][18][17][8].

The following papers are more related to our work: Fudenberg *et al.* [14] give a folk theorem for repeated games with imperfect public information. Both [14] and our paper are built on the (dynamic programming style) *self-generating* technique in [2] (it is called *self-decomposable* in [14]). However, [14] considers self-generation based on a certain supporting hyperplane, which is guaranteed to exist only when the discount factor goes to 1.<sup>2</sup> Therefore, their technique does not apply to our problem because we are dealing with non-limit discount factors.<sup>3</sup> Another difference between [14] and our paper is that we are designing specific mechanisms in this paper, instead of trying to prove the existence of a certain class of mechanisms. With non-limit discount factors, it is generally difficult to precisely characterize the set of feasible utility vectors (optimal frontier) for the agents. Several papers have already proposed different ways of approximation (for cases of non-limit discount factors). Athey *et al.* [4] study approximation by requiring that the payoffs of the agents must be symmetric. In what, from a technical perspective, appears to be the paper closest to the work in this paper, Athey and Bagwell [3] investigate collusion in a repeated game by approximating the optimal frontier by a line segment (the same technique also appears in the work of Abdulkadiroğlu and Bagwell [1]). One of their main results is that if the discount factor reaches a certain threshold (still strictly less than 1), then the approximation comes at no cost. That is, the optimal (first-best) performance can be obtained. However, their technique only works for finite type spaces, as it builds on uneven tie-breaking.

The main contribution of this paper can be summarized as follows. First, we introduce a new technique for approximating the optimal frontier for repeated allocation problem. Our technique works for non-limit discount factors and is not restricted to symmetric payoffs or finite type spaces. The technique we propose is presented in the form of an artificial payment system, which corresponds to approximating the optimal frontier by triangles. The artificial payment system enables us to construct repeated allocation mechanisms without payments based on one-shot allocation mechanisms with payments. We analytically characterize several repeated allocation mechanisms that do

<sup>1</sup> In the previous work, as well as in this paper, the first-best result can be achieved by mechanisms with payments.

<sup>2</sup> In [14], it is shown that any feasible and individually rational equilibrium payoff vector  $v$  can be achieved in a perfect public equilibrium (self-generated based on certain supporting hyperplanes), as long as the discount factor reaches a threshold  $\underline{\beta}$ . However, the threshold  $\underline{\beta}$  depends on  $v$ . If we consider all possible values of  $v$ , then we essentially require that the discount factor/threshold approach 1, since any discount factor that is strictly less than 1 does not work (for some  $v$ ).

<sup>3</sup> In this paper, we also require that the discount factor reaches a threshold, but here the threshold is a constant that works for all possible priors.

not rely on payments, and prove that they are competitive against the optimal mechanism with payments.

This paper also contributes to the line of research on designing competitive mechanisms without payments. The proposed artificial payment system provides a link between mechanisms with payments and mechanisms without payments. By proposing specific competitive mechanisms that do not rely on payments, our paper also provides an answer to the question: *Are monetary payments necessary for designing good mechanisms?* Our results imply that in repeated settings, artificial payments are “good enough” for designing allocation mechanisms with high social welfare. Conversely, it is easy to see that for one-shot settings, artificial payments are completely useless in the problem we study (single-item allocation).

The idea of designing mechanisms without payments to achieve competitive performance against mechanisms with payments was explicitly framed by Procaccia and Tennenholtz [21], in their paper titled *Approximate Mechanism Design Without Money*. That paper carries out a case study on locating a public facility for agents with single-peaked valuations. (The general idea of approximate mechanism design without payments dates back further, at least to work by Dekel *et al.* [13] in a machine learning framework.) To our knowledge, along this line of research, we are the first to study allocation of private goods. Unlike the models studied in the above two papers [13,21], where agents may have consensus agreement, when we are considering the allocation of private goods, the agents are fundamentally in conflict. Nevertheless, it turns out that even here, some positive results can be obtained if the allocation is carried out repeatedly. Thus, we believe that our results provide additional insights to this line of research.

## 2 Model Description

We study the problem of allocating a single item repeatedly between two (and later in the paper, more than two) competing agents. Before each allocation period, the agents learn their (private) valuations for having the item in that period (but not for any future periods). These preferences are independent and identically distributed, across agents as well as periods, according to a distribution  $F$ . We assume that these valuations are non-negative and have finite expectations.  $F$  does not change over time. There are infinitely many periods, and agents’ valuations are discounted according to a discount factor  $\beta$ . Our objective is to design a mechanism that maximizes expected social welfare under the following constraints (we allow randomized mechanisms):

- *(Bayes-Nash) Incentive Compatibility*: Truthful reporting is a Bayes-Nash equilibrium.
- *No Payments*: No monetary transfers are ever made.

In the one-shot mechanism design setting, incentive compatibility is usually achieved through payments. This ensures that agents have no incentive to overbid, because they may have to make large payments. In the repeated allocation setting, there are other ways to achieve incentive compatibility: for example, if an agent strongly prefers to obtain the item in the current period, the mechanism can ensure that she is less likely

to obtain it in future periods. In a sense, this is an artificial form of payment. Such payments introduce some new issues that do not always occur with monetary payments, including that each agent effectively has a limited budget (corresponding to a limited amount of future utility that can be given up); and if one agent makes a payment to another agent by sacrificing some amount of future utility, the corresponding increase in the latter agent's utility may be different from the decrease in the former agent's utility.

### 3 State-Based Approach

Throughout the paper, we adopt the state-based approach introduced in Abreu *et al.* [2]. In their paper, the authors investigated the problem of finding pure-strategy sequential equilibria of repeated games with imperfect monitoring. Their problem can be rephrased as follows: Given a game, what are the possible pure-strategy sequential equilibria? Even though in our paper we are considering a different problem (we are *designing* the game), the underlying ideas still apply. In their paper, states correspond to possible equilibria, while in our paper, states correspond to feasible mechanisms. In this section, we review a list of basic results and observations on the state-based approach, specifically in the context of repeated allocation.

Let  $M$  be an incentive compatible mechanism without payments for a particular (fixed) repeated allocation problem, defined by a particular type distribution and a discount factor. If, under  $M$ , the expected long-term utilities of agents 1 and 2 (at the beginning) are  $x$  and  $y$  respectively, then we denote mechanism  $M$  by state  $(x, y)$ . All mechanisms that can be denoted by  $(x, y)$  are considered equivalent. If we are about to apply mechanism  $M$ , then we say the agents are in state  $(x, y)$ . In the first period, based on the agents' reported values, the mechanism specifies both how to allocate the item in this period, and what to do in the future periods. The rule for the future is itself a mechanism. Hence, a mechanism specifies how to allocate the item within the first period, as well as the state (mechanism) that the agents will be in in the second period. We have that  $(x, y) = E_{v_1, v_2}[(r_1(v_1, v_2), r_2(v_1, v_2)) + \beta(s_1(v_1, v_2), s_2(v_1, v_2))]$ , where  $v_1, v_2$  are the first-period valuations,  $r_1, r_2$  are the immediate rewards obtained from the first-period *allocation rule*, and  $(s_1, s_2)$  gives the second-period state, representing the *transition rule*.

State  $(x, y)$  is called a *feasible* state if there is a feasible mechanism (that is, an incentive compatible mechanism without payments) corresponding to it. We denote the set of feasible states by  $S^*$ . Let  $e$  be an agent's expected valuation for the item in a single period.  $E = \frac{e}{1-\beta}$  is the maximal expected long-term utility an agent can receive (corresponding to the case where she receives the item in every period). Let  $O$  be the set of states  $\{(x, y) | 0 \leq x \leq E, 0 \leq y \leq E\}$ . We have that  $S^* \subseteq O - \{(E, E)\} \subsetneq O$ .

$S^*$  is convex, for the following reason. If  $(x_1, y_1)$  and  $(x_2, y_2)$  are both feasible, then  $(\frac{x_1+x_2}{2}, \frac{y_1+y_2}{2})$  is also feasible (it corresponds to the randomized mechanism where we flip a coin to decide which of the two mechanisms to apply).  $S^*$  is symmetric with respect to the diagonal  $y = x$ : if  $(x, y)$  is feasible, then so is  $(y, x)$  (by switching the roles of the two agents).

The approximate shape of  $S^*$  is illustrated in Figure 1. There are three noticeable extreme states:  $(0, 0)$  (nobody ever gets anything),  $(E, 0)$  (agent 1 always gets the item),



and  $(0, E)$  (agent 2 always gets the item).  $S^*$  is confined by the x-axis (from  $(0, 0)$  to  $(E, 0)$ ), the y-axis (from  $(0, 0)$  to  $(0, E)$ ), and, most importantly, the bold curve, which corresponds to the optimal frontier. The square specified by the dotted lines represents  $O$ .




Fig. 1. The shape of  $S^*$




Fig. 2. Bow shape approximated by triangle

Our objective is to find the state  $(x^*, y^*) \in S^*$  that maximizes  $x^* + y^*$  (expected social welfare). By convexity and symmetry, it does not hurt to consider only cases where  $x^* = y^*$ .

We now define a notion of when one set of states is *generated* by another. Recall that a mechanism specifies how to allocate the item within the first period, as well as which state the agents transition to for the second period. Let  $S$  be any set of states with  $S \subseteq O$ . Let us assume that, in the second period, exactly the states in  $S$  are feasible. That is, we assume that, if and only if  $(x, y) \in S$ , starting at the second period, there exists a feasible mechanism under which the expected utilities of agent 1 and 2 are  $x$  and  $y$ , respectively. Based on this assumption, we can construct incentive compatible mechanisms starting at the first period, by specifying an *allocation rule* for the first period, as well as a *transition rule* that specifies the states in  $S$  to which the agents will transition for the beginning of the second period. Now, we only need to make sure that the first period is incentive compatible. That is, the allocation rule in the first period, combined with the rule for selecting the state at the start of the second period, must incentivize the agents to report their true valuations in the first period. We say the set of resulting feasible states for the first period is *generated* by  $S$ , and is denoted by  $Gen(S)$ .

The following claim provides a general guideline for designing feasible mechanisms.

**Claim 1.** *For any  $S \subseteq O$ , if  $S \subseteq Gen(S)$ , then  $S \subseteq S^*$ . That is, if  $S$  is self-generating, then all the states in  $S$  are feasible.*

We now consider starting with the square  $O$  that contains  $S^*$  and iteratively generating sets. Let  $O^0 = O$  and  $O^{i+1} = Gen(O^i)$  for all  $i$ . The following claim, together with Claim 1, provide a general approach for computing  $S^*$ .



**Claim 2.** *The  $O^i$  form a sequence of (weakly) decreasing sets that converges to  $S^*$  if it converges at all. That is,  $S^* = \text{Gen}(S^*)$ .  $S^* \subseteq O^i$  for all  $i$ .  $O^{i+1} \subseteq O^i$  for all  $i$ . If  $O^i = O^{i+1}$ , then  $O^i = S^*$ .*

The above guideline leads to a numerical solution technique for finite valuation spaces. With a properly chosen numerical discretization scheme, we are able to compute an underestimation of  $O^i$  for all  $i$ , by solving a series of linear programs. The underestimations of the  $O^i$  always converge to an underestimation of  $S^*$  (a subset of  $S^*$ ). That is, we end up with a set of feasible mechanisms. We are also able to show that as the discretization step size goes to 0, the obtained feasible set approaches  $S^*$ . That is, the numerical solution technique produces an optimal mechanism in the limit as the discretization becomes finer. Details of the numerical solution technique are omitted due to space constraint.

One drawback of the numerical approach is that the obtained mechanism does not have an elegant form. This makes it harder to analyze. From the agents’ perspective, it is difficult to comprehend what the mechanism is trying to do, which may lead to irrational behavior. Another drawback of the numerical approach is that it only applies to cases of finite valuation spaces. For the rest of the paper, we take a more analytical approach. We aim to design mechanisms that can be more simply and elegantly described, work for any valuation space, and are (hopefully) close to optimality.

At the end of Section 4.2, we will compare the performances of the mechanisms obtained numerically and the mechanisms obtained by the analytical approach.

## 4 Competitive Analytical Mechanism

In this section, we propose the idea of an artificial payment system. Based on this, we propose several mechanisms that can be elegantly described, and we can prove that these mechanisms are close to optimality.

### 4.1 Artificial Payment System

Let us recall the approximate shape of  $S^*$  (Figure 2). The area covered by  $S^*$  consists of two parts. The lower left part is a triangle whose vertices are  $(0, 0)$ ,  $(E, 0)$ , and  $(0, E)$ . These three states are always feasible, and so are their convex combinations. The upper right part is a bow shape confined by the straight line and the bow curve from  $(0, E)$  to  $(E, 0)$ . To solve for  $S^*$ , we are essentially solving for the largest bow shape satisfying that the union of the bow shape and the lower-left triangle is self-generating. Here, we consider an easier problem. Instead of solving for the largest bow shape, we solve for the largest triangle (whose vertices are  $(0, E)$ ,  $(E, 0)$ , and  $(x^*, x^*)$ ) so that the union of the two triangles is self-generating (illustrated in Figure 2). That is, we want to find the largest value of  $x^*$  that satisfies that the set of convex combinations of  $(0, 0)$ ,  $(E, 0)$ ,  $(0, E)$ , and  $(x^*, x^*)$  is self-generating.

The triangle approximation corresponds to an *artificial payment system*. Let  $(x^*, x^*)$  be any feasible state satisfying  $x^* \geq \frac{E}{2}$ . Such a feasible state always exists (e.g.,  $(\frac{E}{2}, \frac{E}{2})$ ). We can implement an artificial payment system based on  $(x^*, x^*)$ ,  $(E, 0)$ ,

and  $(0, E)$ , as follows. At the beginning of a period, the agents are told that the default option is that they move to state  $(x^*, x^*)$  at the beginning of the next period. However, if agent 1 wishes to pay  $v_1$  ( $v_1 \leq \beta x^*$ ) units of artificial currency to agent 2 (and agent 2 is not paying), then the agents will move to  $(x^* - \frac{v_1}{\beta}, x^* + \frac{E-x^*}{x^*} \frac{v_1}{\beta})$ . That is, the future state is moved  $\frac{v_1}{\beta}$  units to the left along the straight line connecting  $(0, E)$  and  $(x^*, x^*)$ . (This corresponds to going to each of these two states with a certain probability.) By paying  $v_1$  units of artificial currency, agent 1's expected utility is decreased by  $v_1$  (the expected utility is decreased by  $\frac{v_1}{\beta}$  at the start of the next period). When agent 1 pays  $v_1$  units of artificial currency, agent 2 receives only  $\frac{E-x^*}{x^*} v_1$  (also as a result of future utility). In effect, a fraction of the payment is lost in transmission. Similarly, if agent 2 wishes to pay  $v_2$  ( $v_2 \leq \beta x^*$ ) units of artificial currency to agent 1 (and agent 1 is not paying), then the agents will move to  $(x^* + \frac{E-x^*}{x^*} \frac{v_2}{\beta}, x^* - \frac{v_2}{\beta})$ . That is, the future state is moved  $\frac{v_2}{\beta}$  units towards the bottom along the straight line connecting  $(x^*, x^*)$  and  $(E, 0)$ . If both agents wish to pay, then the agents will move to  $(x^* - \frac{v_1}{\beta} + \frac{E-x^*}{x^*} \frac{v_2}{\beta}, x^* - \frac{v_2}{\beta} + \frac{E-x^*}{x^*} \frac{v_1}{\beta})$ , which is a convex combination of  $(0, 0)$ ,  $(0, E)$ ,  $(E, 0)$ , and  $(x^*, x^*)$ .

Effectively, both agents have a *budget* of  $\beta x^*$ , and when an agent pays the other agent, there is a *gift tax* with rate  $1 - \frac{E-x^*}{x^*}$ .

Based on the above artificial payment system, our approach is to design repeated allocation mechanisms without payments, based on one-shot allocation mechanisms with payments. In order for this to work, the one-shot allocation mechanisms need to take the gift tax into account, and an agent's payment should never exceed the budget limit.

The budget constraint is difficult from a mechanism design perspective. We circumvent this based on the following observation. An agent's budget is at least  $\beta \frac{E}{2} = \frac{e\beta}{2-2\beta}$ , which goes to infinity as  $\beta$  goes to 1. As a result, for sufficiently large discount factors, the budget constraint will not be binding. For the remainder of this paper, we ignore the budget limit when we design the mechanisms. Then, for each obtained mechanism, we specify how large the discount factor has to be for the mechanism to be well defined (that is, the budget constraint is not violated). This allows us to work around the budget constraint. The drawback is obvious: our proposed mechanisms only work for discount factors reaching a (constant) threshold (though it is not as restrictive as studying the limit case as  $\beta \rightarrow 1$ ).

## 4.2 High/Low Types

We start with the simple model in which the agents' valuations are either  $H$  (high) with probability  $p$  or  $L$  (low) with probability  $1 - p$ . Without loss of generality, we assume that  $L = 1$ . We will construct a repeated allocation mechanism without payments based on the following *pay-only* one-shot allocation mechanism:

*Allocation:* If the reported types are the same, we determine the winner by flipping a (fair) coin. If one agent's reported value is high and the other agent's reported value is low, then we allocate the item to the agent reporting high.

*Payment:* An agent pays 0 if its reported type is low. An agent pays  $\frac{1}{2}$  if its reported type is high (whether she wins or not); this payment does not go to the other agent.

**Claim 3.** *The above pay-only mechanism is (Bayes-Nash) incentive compatible.*

Now we return to repeated allocation settings. Suppose  $(x^*, x^*)$  is a feasible state. That is, we have an artificial payment system with gift tax rate  $1 - \frac{E-x^*}{x^*}$ . We apply the above one-shot mechanism, with the modifications that when an agent pays  $\frac{1}{2}$ , it is paying artificial currency instead of real currency, and the other agent receives  $\frac{1}{2} \frac{E-x^*}{x^*}$ . We note that the amount an agent receives is only based on the other agent's reported value. Therefore, the above modifications do not affect the incentives.

Under the modified mechanism, an agent's expected utility equals  $\frac{T}{2} - P + P \frac{E-x^*}{x^*} + \beta x^*$ . In the above expression,  $T = 2p(1-p)H + p^2H + (1-p)^2$  is the expected value of the higher reported value.  $\frac{T}{2}$  is then the ex ante expected utility received by an agent as a result of the allocation.  $P = \frac{E}{2}$  is the expected amount of artificial payment an agent pays.  $P \frac{E-x^*}{x^*}$  is the expected amount of artificial payment an agent receives.  $\beta x^*$  is the expected future utility by default (if no payments are made).

If both agents report low, then, at the beginning of the next period, the agents go to  $(x^*, x^*)$  by default. If agent 1 reports high and agent 2 reports low, then the agents go to  $(x^* - \frac{1}{2\beta}, x^* + \frac{E-x^*}{2\beta x^*})$ , which is a convex combination of  $(x^*, x^*)$  and  $(0, E)$ . If agent 1 reports low and agent 2 reports high, then the agents go to  $(x^* + \frac{E-x^*}{2\beta x^*}, x^* - \frac{1}{2\beta})$ , which is a convex combination of  $(x^*, x^*)$  and  $(E, 0)$ . If both agents report high, then the agents go to  $(x^* - \frac{1}{2\beta} + \frac{E-x^*}{2\beta x^*}, x^* - \frac{1}{2\beta} + \frac{E-x^*}{2\beta x^*})$ , which is a convex combination of  $(x^*, x^*)$  and  $(0, 0)$ . Let  $S$  be the set of all convex combinations of  $(0, 0)$ ,  $(E, 0)$ ,  $(0, E)$ , and  $(x^*, x^*)$ . The future states given by the above mechanism are always in  $S$ . If an agent's expected utility under this mechanism is greater than or equal to  $x^*$ , then  $S$  is self-generating. That is,  $(x^*, x^*)$  is feasible as long as  $x^*$  satisfies  $x^* \leq \frac{T}{2} - P + P \frac{E-x^*}{x^*} + \beta x^*$ .

We rewrite it as  $ax^{*2} + bx^* + c \leq 0$ , where  $a = 1 - \beta$ ,  $b = 2P - \frac{T}{2}$ , and  $c = -EP$ . The largest  $x^*$  satisfying the above inequality is simply the larger solution of  $ax^{*2} + bx^* + c = 0$ , which is  $\frac{\frac{T}{2} - 2P + \sqrt{(2P - \frac{T}{2})^2 + 4(1-\beta)EP}}{2(1-\beta)}$ .

This leads to a feasible mechanism  $M^*$  (corresponding to state  $(x^*, x^*)$ ). The expected social welfare under  $M^*$  is  $2x^*$ , where  $x^*$  equals the above solution.

We have not considered the budget limit. For the above  $M^*$  to be well-defined (satisfying the budget constraint), we need  $\beta x^* \geq \frac{1}{2}$ . Since  $x^* \geq \frac{E}{2} = \frac{e}{2-2\beta} \geq \frac{1}{2-2\beta}$ , we only need to make sure that  $\frac{\beta}{2-2\beta} \geq \frac{1}{2}$ . Therefore, if  $\beta \geq \frac{1}{2}$ , then  $M^*$  is well-defined. For specific priors,  $M^*$  could be well-defined even for smaller  $\beta$ .

Next, we show that (whenever  $M^*$  is well-defined)  $M^*$  is very close to optimality. Consider the *first-best allocation mechanism*: the mechanism that always successfully identifies the agent with the higher valuation and allocates the item to this agent (for free). This mechanism is not incentive compatible, and hence not feasible. The expected social welfare achieved by the first-best allocation mechanism is  $\frac{T}{1-\beta}$ , which is an upper bound on the expected social welfare that can be achieved by any mechanism with (or without) payments (it is a strict upper bound, as the dAGVA mechanism [12] is efficient, incentive compatible, and budget balanced).

**Definition 1.** When the agents’ valuations are either high or low, the prior distribution over the agents’ valuations is completely characterized by the values of  $H$  and  $p$ . Let  $W$  be the expected social welfare under a feasible mechanism  $M$ . Let  $W^F$  be the expected social welfare under the first-best allocation mechanism. If  $W \geq \alpha W^F$  for all  $H$  and  $p$ , then we say  $M$  is  $\alpha$ -competitive. We call  $\alpha$  a competitive ratio of  $M$ .

**Claim 4.** Whenever  $M^*$  is well-defined for all  $H$  and  $p$ , (e.g.,  $\beta \geq \frac{1}{2}$ ),  $M^*$  is 0.94-competitive.

As a comparison, the lottery mechanism that always chooses the winner by flipping a fair coin has competitive ratio (exactly) 0.5 (if  $H$  is much larger than  $L$  and unlikely to occur).

In the following table, for different values of  $H$ ,  $p$ , and  $\beta$ , we compare  $M^*$  to the near-optimal *feasible* mechanism obtained with the numerical solution technique. The table elements are the expected social welfare under  $M^*$ , the near-optimal feasible mechanism, the first-best allocation mechanism, and the lottery mechanism.

	$M^*$	Optimal	First-best	Lottery
$H = 2, p = 0.2, \beta = 0.5$	2.6457	2.6725	2.7200	2.4000
$H = 4, p = 0.4, \beta = 0.5$	5.5162	5.7765	5.8400	4.4000
$H = 16, p = 0.8, \beta = 0.5$	30.3421	30.8000	30.8000	26.0000
$H = 2, p = 0.2, \beta = 0.8$	6.6143	6.7966	6.8000	6.0000
$H = 2, p = 0.8, \beta = 0.8$	9.4329	9.8000	9.8000	9.0000
$H = 16, p = 0.8, \beta = 0.8$	75.8552	77.0000	77.0000	65.0000

### 4.3 General Valuation Space

In this section, we generalize the earlier approach to general valuation spaces. We let  $f$  denote the probability density function of the prior distribution. (A discrete prior distribution can always be smoothed to a continuous distribution that is arbitrarily close.)

We will construct a repeated allocation mechanism without payments based on the following *pay-only* one-shot allocation mechanism:

*Allocation:* The agent with the higher reported value wins the item.

*Payment:* An agent pays  $\int_0^v t f(t) dt$  if it reports  $v$ .

This mechanism is actually a dAGVA mechanism [12], which is known to be (Bayes-Nash) incentive compatible.

The process is similar to that in the previous section. Due to space constraints, we omit the details. At the end, we obtain a feasible mechanism  $M^*$ . The expected social welfare under  $M^*$  is  $2x^*$ , where  $x^*$  equals  $\frac{\frac{T}{2} - 2P + \sqrt{(2P - \frac{T}{2})^2 + 4(1-\beta)EP}}{2(1-\beta)}$ . Here,  $T =$

<sup>4</sup> “The” dAGVA mechanism often refers to a specific mechanism in a class of Bayes-Nash incentive compatible mechanisms, namely one that satisfies budget balance. In this paper, we will use “dAGVA mechanisms” to refer to the entire class, including ones that are not budget-balanced. Specifically, we will only use dAGVA mechanisms in which payments are always nonnegative.

$\int_0^\infty \int_0^\infty \max\{t, v\} f(t) f(v) dt dv$  is the expected value of the higher valuation.  $P = \int_0^\infty \int_0^v t f(t) dt f(v) dv$  is the expected amount an agent pays.

For the above  $M^*$  to be well-defined, we need the budget  $\beta x^*$  to be greater than or equal to  $\int_0^\infty t f(t) dt = e$  (the largest possible amount an agent pays). Since  $x^* \geq \frac{E}{2} = \frac{e}{2-2\beta}$ , we only need to make sure  $\frac{\beta e}{2-2\beta} \geq e$ . Therefore, if  $\beta \geq \frac{2}{3}$ , then  $M^*$  is well-defined. For specific priors,  $M^*$  may be well-defined for smaller  $\beta$ .

Next, we show that (whenever  $M^*$  is well-defined)  $M^*$  is competitive against the first-best allocation mechanism for all prior distribution  $f$ . Naturally, the competitive ratio is slightly worse than the one obtained previously for high/low valuations. We first generalize the definition of competitiveness appropriately.

**Definition 2.** Let  $W$  be the expected social welfare under a feasible mechanism  $M$ . Let  $W^F$  be the expected social welfare under the first-best allocation mechanism. If  $W \geq \alpha W^F$  for all prior distributions, then we say that  $M$  is  $\alpha$ -competitive. We call  $\alpha$  a competitive ratio of  $M$ .

**Claim 5.** Whenever  $M^*$  is well-defined for all prior distributions (e.g.,  $\beta \geq \frac{2}{3}$ ),  $M^*$  is 0.85-competitive.

### 5 Three or More Agents

We have focused on allocation problems with two agents. In this section, we generalize our analytical approach to cases of three or more agents.

Let  $n$  be the number of agents. We will continue with the state-based approach. That is, a mechanism (state) is denoted by a vector of  $n$  nonnegative real values. For example, if under mechanism  $M$ , agent  $i$ 's long-term expected utility equals  $x_i$ , then mechanism  $M$  is denoted by  $(x_1, x_2, \dots, x_n)$ . If we are about to apply mechanism  $M$ , then we say the agents are in state  $(x_1, x_2, \dots, x_n)$ .

For any  $n$ , it is easy to see that the set of feasible states is convex and symmetric with respect to permutations of the agents. A state is called *fair* if all its elements are equal. For example,  $(1, 1, 1)$  is a fair state ( $n = 3$ ). When there is no ambiguity about the number of agents, the fair state  $(x, x, \dots, x)$  is denoted simply by  $x$ .

An artificial payment system can be constructed in a way that is similar to the case of two agents. Let  $\mu_{n-1}$  be any feasible fair state for the case of  $n - 1$  agents. Then, the following  $n$  states are also feasible for the case of  $n$  agents:

$$(0, \underbrace{\mu_{n-1}, \dots, \mu_{n-1}}_{n-1}), (\mu_{n-1}, 0, \underbrace{\mu_{n-1}, \dots, \mu_{n-1}}_{n-2}), \dots, (\underbrace{\mu_{n-1}, \dots, \mu_{n-1}}_{n-1}, 0).$$

We denote the above  $n$  states by  $s_i$  for  $i = 1, 2, \dots, n$ . Let  $\hat{S}$  be the set of all feasible states with at least one element that equals 0.  $\hat{S}$  is self-generating. Suppose we have a fair state  $\mu_n$  for the case of  $n$  agents. Let  $S$  be the smallest convex set containing  $\mu_n$  and all the states in  $\hat{S}$ . The  $s_i$  are in both  $\hat{S}$  and  $S$ . An artificial payment system can be implemented as follows (for the case of  $n$  agents): The agents will go to state  $\mu_n$  by default. If for all  $i$ , agent  $i$  chooses to pay  $v_i$  units of artificial currency, then we move

to a new state whose  $i^{\text{th}}$  element equals  $\mu_n - \frac{v_i}{\beta} + \gamma \sum_{j \neq i} \frac{v_j}{\beta}$ . Here  $\gamma = \frac{\mu_{n-1} - \mu_n}{\mu_n}$ . The new state  $M$  is in  $S$ . (The reason is the following. If only agent  $i$  is paying, and it is paying  $nv_i$  instead of  $v_i$ , then the new state  $M_i$  is  $(\underbrace{\mu_n + \gamma \frac{nv_i}{\beta}, \dots, \mu_n + \gamma \frac{nv_i}{\beta}}_{i-1}, \underbrace{\mu_n + \gamma \frac{nv_i}{\beta}, \dots, \mu_n + \gamma \frac{nv_i}{\beta}}_{n-i})$ , which is a convex combination of  $\mu_n$  and  $s_i$ . The

average of the  $M_i$  over all  $i$  is just  $M$ . Thus  $M$  is a convex combination of  $\mu_n$  and the  $s_i$ , which implies  $M \in S$ .)

With the above artificial payment system, by allocating the item to the agent with the highest reported value and charging the agents dAGVA payments, we get an incentive compatible mechanism. We denote agent  $i$ 's reported value by  $v_i$  for all  $i$ . The dAGVA payment for agent  $i$  equals  $E_{v_{-i}}(I(v_i \geq \max\{v_{-i}\}) \max\{v_{-i}\})$ , where  $I$  is the characteristic function (which evaluates to 1 on *true* and to 0 otherwise) and  $v_{-i}$  is the set of reported values from agents other than  $i$ .

We still use  $P$  to denote the expected amount of payment from an agent. We use  $T$  to denote the expected value of the highest reported value. The expected utility for an agent is then  $\frac{T}{n} - P + (n - 1) \frac{\mu_{n-1} - \mu_n}{\mu_n} P + \beta \mu_n$ .

To show  $S$  is self-generating, we only need to show  $\mu_n$  is in  $Gen(S)$ . That is,  $\mu_n$  is a feasible fair state as long as  $\mu_n$  satisfies the following inequality:  $\mu_n \leq \frac{T}{n} - P + (n - 1) \frac{\mu_{n-1} - \mu_n}{\mu_n} P + \beta \mu_n$ .

The largest solution of  $\mu_n$  equals  $\frac{\frac{T}{n} - nP + \sqrt{(nP - \frac{T}{n})^2 + 4(1-\beta)(n-1)\mu_{n-1}P}}{2(1-\beta)}$ .

The above expression increases when the value of  $\mu_{n-1}$  increases. The highest value for  $\mu_1$  is  $E$  (when there is only one agent, we can simply give the item to the agent for free). A natural way of solving for a good fair state  $\mu_n$  is to start with  $\mu_1 = E$ , then apply the above technique to solve for  $\mu_2$ , then  $\mu_3$ , etc.

Next, we present a claim that is similar to Claim 5.

**Claim 6.** *Let  $n$  be the number of agents. Let  $M_n^*$  be the mechanism obtained by the technique proposed in this section. Whenever  $\beta \geq \frac{n^2}{n^2 + 4}$ ,  $M_n^*$  is well defined for all priors, and is  $\alpha_n$ -competitive, where  $\alpha_1 = 1$ , and for  $n > 1$ ,*

$$\alpha_n = \min_{\{1 \leq u \leq \frac{n}{n-1}\}} n \frac{\frac{n}{n} - n + nu - u + \sqrt{(n - nu + u - \frac{n}{n})^2 + 4\alpha_{n-1} \frac{n - nu + u}{n}}}{2u}$$

For all  $i$ ,  $\alpha_i \geq \frac{3}{4}$  holds.

As a comparison, the lottery mechanism that always chooses the winner uniformly at random has competitive ratio (exactly)  $\frac{1}{n}$ , which goes to 0 as  $n$  goes to infinity.

<sup>5</sup> It should be noted that when one agent pays 1, then every other agent receives  $\gamma$ . In a sense,  $\gamma$  already incorporates the fact that the payment must be divided among multiple agents.

<sup>6</sup> The above argument assumes that the available budget is at least  $n$  times the maximum amount an agent pays.

## References

1. Abdulkadiroğlu, A., Bagwell, K.: Trust, reciprocity and favors in cooperative relationships. Working paper (2005)
2. Abreu, D., Pearce, D., Stacchetti, E.: Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58(5), 1041–1063 (1990)
3. Athey, S., Bagwell, K.: Optimal collusion with private information. Working papers 99-17, Massachusetts Institute of Technology (MIT), Department of Economics (September 1999)
4. Athey, S., Bagwell, K., Sanchirico, C.: Collusion and price rigidity. *Review of Economic Studies* 71(2), 317–349 (2004)
5. Athey, S., Segal, I.: An efficient dynamic mechanism. Technical report, Working Paper (2007)
6. Bailey, M.J.: The demand revealing process: to distribute the surplus. *Public Choice* 91, 107–126 (1997)
7. Bergemann, D., Välimäki, J.: Efficient dynamic auctions. Cowles Foundation Discussion Papers 1584, Cowles Foundation, Yale University (October 2006)
8. Casella, A.: Storable votes. *Games and Economic Behavior* 51, 391–419 (2005)
9. Cavallo, R.: Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In: Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Hakodate, Japan, pp. 882–889 (2006)
10. Cavallo, R.: Efficiency and redistribution in dynamic mechanism design. In: Proceedings of the ACM Conference on Electronic Commerce (EC), Chicago, IL, USA, pp. 220–229 (2008)
11. Cavallo, R., Parkes, D.C., Singh, S.: Optimal coordinated planning amongst self-interested agents with private state. In: Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI), Cambridge, MA, USA (2006)
12. d’Aspremont, C., Gérard-Varet, L.-A.: Incentives and incomplete information. *Journal of Public Economics* 11, 25–45 (1979)
13. Dekel, O., Fischer, F., Procaccia, A.D.: Incentive compatible regression learning. In: Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), Philadelphia, PA, USA, pp. 884–893. Society for Industrial and Applied Mathematics (2008)
14. Fudenberg, D., Levine, D.I., Maskin, E.: The folk theorem with imperfect public information. *Econometrica* 62(5), 997–1039 (1994)
15. 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, pp. 210–219 (2008)
16. Guo, M., Conitzer, V.: Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior* 67(1), 69–98 (2009); Special Section of Games and Economic Behavior Dedicated to the 8th ACM Conference on Electronic Commerce
17. Hopenhayn, H.A., Hauser, C.: Trading favors: Optimal exchange and forgiveness. 2004 Meeting Papers 125, Society for Economic Dynamics (2004)
18. Möbius, M.M.: Trading favors (May 2001) (manuscript)
19. Moulin, H.: Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory* 144(1), 96–119 (2009)
20. Porter, R., Shoham, Y., Tennenholtz, M.: Fair imposition. *Journal of Economic Theory* 118, 209–228 (2004); Early version appeared in IJCAI 2001
21. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. In: Proceedings of the ACM Conference on Electronic Commerce (EC), Stanford, CA, USA, pp. 177–186 (2009)
22. Sannikov, Y.: Games with imperfectly observable actions in continuous time. *Econometrica* 75(5), 1285–1329 (2007)



# Pseudonyms in Cost-Sharing Games<sup>\*</sup>

Paolo Penna<sup>1</sup>, Florian Schoppmann<sup>2</sup>, Riccardo Silvestri<sup>3</sup>, and Peter Widmayer<sup>4</sup>

<sup>1</sup> Dipartimento di Informatica ed Applicazioni, Università di Salerno,  
via S. Allende 2, Baronissi (SA), Italy  
penna@dia.unisa.it

<sup>2</sup> Department of Computer Science, Stanford University, Stanford, CA, 94305, USA  
fschopp@stanford.edu

<sup>3</sup> Dipartimento di Informatica, Università di Roma “La Sapienza”,  
via Salaria 113, 00198 Roma, Italy  
silvestri@di.uniroma1.it

<sup>4</sup> Institut für Theoretische Informatik, ETH Zürich, CAB H 15 Universitätstrasse 6,  
8092 Zürich, Switzerland  
widmayer@inf.ethz.ch

**Abstract.** This work initiates the study of cost-sharing mechanisms that, in addition to the usual incentive compatibility conditions, make it disadvantageous for the users to employ pseudonyms. We show that this is possible only if all serviced users pay the *same* price, which implies that such mechanisms do not exist even for certain subadditive cost functions. In practice, a user can increase her utility by lying in one way (misreport her willingness to pay) or another (misreport her identity). We prove also results for approximately budget-balanced mechanisms. Finally, we consider mechanisms that rely on some kind of “reputation” associated to the pseudonyms and show that they are provably better.

## 1 Introduction

Incentives play a crucial role in distributed systems of almost any sort. Typically users want to get resources (a service) without contributing or contributing very little. For instance, file-sharing users in peer-to-peer systems are only interested in downloading data, even though uploading is essential for the system to survive (but maybe costly for some of the users). Some successful systems, like BitTorrent [15, 6], have already incorporated incentive compatibility considerations in their design (users can download data only if they upload some content to others).

One can regard such systems as *cost-sharing mechanisms* in which the overall cost must be recovered from the users in a reasonable and *incentive compatible* manner; the mechanism determines which users get the service and at what price. The study of these mechanisms is a very important topic in economics and in cooperative game theory, in which individuals (users) can coordinate their strategies (i.e., they can form coalitions and collude). Cost-sharing mechanisms should guarantee several essential properties:

---

<sup>\*</sup> Work supported by the European Project IST-15964 “Algorithmic Principles for Building Efficient Overlay Computers” (AEOLUS), by the European Project COST 295 (DYNAMO), and by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD). Part of this work has been done while the first author was visiting ETH Zürich. Research done while the second author was at the University of Paderborn, Germany.



Users recover the overall cost and are not charged more than necessary (*budget-balance*), even coalitions of users cannot benefit from misreporting their willingness to pay (*group-strategyproofness*), no user is excluded a priori (*consumer sovereignty*), no user is charged more than her willingness to pay (*voluntary participation*), and no user receives money (*no positive transfer*). To get a feel of the difficulties of obtaining such mechanisms, consider the following:

*Example 1 (identical prices)*. There are three players, and the overall cost depends on how many users get the service: Servicing three users costs 2 while servicing one or two users costs only 1. Consider a mechanism that iteratively drops all users who cannot afford an equal fraction of the total cost (three users pay  $2/3$  each, two users pay  $1/2$  each, and one user pays 1). Unfortunately, this mechanism is *not even strategyproof*<sup>1</sup>: For valuations  $(0.6, 0.6, 1)$  the first two users are dropped, so their utility is zero (no service and nothing to pay). If the first user misreports her valuation to 1 she gets the service for price 0.5 (now only the second user is dropped), and her utility (valuation minus price to pay =  $0.6 - 0.5 = 0.1$ ) is strictly better than before.

As Example 1 shows, a groupstrategyproof and budget-balanced mechanism has to charge the serviced users *different* prices in general. In this work, we consider a new form of manipulation that, to the best of our knowledge, has not been considered in the cost-sharing literature before: In most of the Internet applications a user can easily create several virtual identities [9, 11, 32, 10, 5]. The typical scenario is that there is a universe of possible names (e.g., all strings of up to 40 characters), and only a subset of these correspond to actual users. Each user can replace her name with some *pseudonym*, that is, another name in the universe of names which does not correspond to any other user (e.g., an email account that has not been taken by anybody else). In this new scenario, each user can manipulate the mechanism in *two* ways: by misreporting her name and/or by misreporting her willingness to pay for the service. Ideally, we would like resistance against both kinds of manipulations.

*Example 2 (manipulation via pseudonyms)*. Consider a different mechanism for Example 1: Order players alphabetically by their name. The first two players bidding at least  $1/2$  only have to pay  $1/2$ . Otherwise, the price is always 1. This mechanism is group-strategyproof [4]; however, it can be manipulated via pseudonyms. For instance, in the situation

<i>Names</i>	<i>Alice Bob Cindy</i>
<i>Valuations</i>	0.6 0.6 1
<i>Prices</i>	0.5 0.5 1

there is an obvious incentive for Cindy to use a pseudonym “Adam”:

<i>Names</i>	<i>Alice Bob Adam</i>
<i>Valuations</i>	0.6 0.6 1
<i>Prices</i>	0.5 1 0.5

Now Bob is dropped and the remaining two users get the service for a price of 0.5 and thus Cindy’s utility is better off.

<sup>1</sup> A mechanism is *strategyproof* if no single player can benefit from misreporting her willingness to pay (i.e., group-strategyproof for maximum coalition size 1).

A natural idea to discourage pseudonyms is to *randomize* the mechanism by ordering the players randomly. In this way users can be motivated to use their own name, provided they trust the fact that the mechanism picks a truly random shuffle of the (reported) names *and* they cannot guess the randomness before they bid. Unfortunately, the notion of group-strategyproofness is problematic for randomized mechanisms [12] (see also the full version of this work [26]). Moreover, randomness is generally considered a scarce resource and derandomization of mechanisms is especially difficult because we can run the mechanism (“the auction”) only once [1].

## 1.1 Our Contribution and Related Work

In our model, the overall cost depends only on the number of users that get the service and the mechanism is deterministic. Call a mechanism *renameproof* if, in the scenario above, no user has an incentive to change her name (i.e., create a pseudonym and obtain the service for a better price). We first prove that for *all* budget-balanced groupstrategyproof mechanisms the following equivalence holds:

$$\text{renameproof} \iff \text{identical prices} \quad (1)$$

By “identical prices” we mean that all serviced users pay the same fraction of the total cost (identical prices imply renameproofness and our contribution is to prove that the converse holds). In some cases, this result gives a *characterization* of such mechanisms, while in other cases it implies that these mechanisms *do not exist!* This is because, even in our restricted scenario, budget-balance and groupstrategyproofness can be achieved only with mechanisms that use *different* prices [4].

We actually prove a more general version of Equivalence (1) which applies also to *approximately* budget-balanced ones. In proving this, we exhibit an intriguing connection with a problem of hypergraph coloring, a linear algebra result by Gottlieb [13] and, ultimately, with the Ramsey Theorem (the impossibility of coloring large hypergraphs with a constant number of colors without creating monochromatic components).

The notion of renameproof mechanisms leads naturally to a weaker condition which we call *reputationproof* because it captures the idea that names have some reputation associated, and a user cannot create a pseudonym with a better reputation than that of her true name [5]. In some sense, we can regard reputation as a way to “derandomize” some of the mechanisms that are renameproof in expectation. We show that reputation does help because there exist cost functions that admit budget-balanced, groupstrategyproof, and reputationproof mechanisms, while no budget-balanced mechanism can be simultaneously groupstrategyproof and renameproof (all budget-balanced mechanisms can be manipulated in one way or another). We find it interesting that certain new constructions of mechanisms [4], which were originally introduced to overcome the limitations of identical prices, are reputationproof (though not renameproof).

## 1.2 Connections with Prior Work

The notion of renameproof mechanism is weaker than that of *falsenameproof* mechanism [32] where users can submit *multiple* bids to the mechanism, each one under a

different name (pseudonym)<sup>4</sup>. In the context of combinatorial auctions, where budget-balance is ignored, groupstrategyproofness and falsenemeproofness are independent notions [32]; several constructions of mechanisms that are both strategyproof and false-nameproof are known (see e.g. [32,31]).

It is common opinion that manipulations via pseudonyms arise because a single user can gain a large influence on the game by “voting” (bidding) many times [9,5,11,10,21,32]. Indeed, most of the research focuses on solutions that make it impossible (or very difficult) to vote more than once [7,30,25]. While this is enough in combinatorial auctions [7], manipulations are still possible by declaring false names in cost-sharing games where the mechanism must guarantee budget-balance.

Cost-sharing mechanisms are usually studied in more general settings where costs do not depend only on the number of users. The main technique for obtaining groupstrategyproof mechanisms is due to Moulin [23]. These mechanisms are budget-balanced for arbitrary submodular<sup>5</sup> costs [23]. When costs are not submodular, Moulin mechanisms can achieve only *approximate* budget-balance [16,14,20]. An alternative method called *two-price* mechanisms has recently been presented in [4]. The authors showed that, for some problems where the cost depends only on the number of serviced users, groupstrategyproof mechanisms using identical prices cannot be budget-balanced, while their two-price mechanisms are both groupstrategyproof and budget-balanced. This holds for subadditive costs that depend only on the number of users [4]. In [18] it is shown that the so-called *sequential* mechanisms are groupstrategyproof and budget-balanced for supermodular costs. Only partial characterizations of general budget-balanced and groupstrategyproof mechanisms are known [14,27,28,17,18]. Acyclic mechanisms [22] can achieve budget-balance for all non-decreasing cost-functions by considering a weaker version of groupstrategyproofness (see also [3]).

Finally, [8] studies the *public excludable good* problem and gives sufficient conditions for which every mechanism must use identical prices. Since this problem is a special case of those studied here, we obtain an alternative axiomatic characterization of identical prices.

*Road map.* The formal model and the definition of renameproof mechanism are given in Section 2. Equivalence (I) and its relaxation to approximately budget-balanced mechanisms are proved in Section 3. Reputationproof mechanisms are defined and analyzed in Section 4. Sketched or omitted proofs are given in full detail in [26].

## 2 The Formal Model

We consider names as integers taken from a universe  $\mathcal{N} = \{1, \dots, n\}$ . A user can register herself under one or more names and two users cannot share the same name. A user who first registered under name  $i$  can thus create a number of additional names which we call her *pseudonyms*, while the first name she created is her *true name*. The

<sup>2</sup> For instance, a falsenameproof mechanism for combinatorial auctions guarantees that a user interested in a bundle of items cannot obtain the bundle for a lower price by submitting several bids. Each bid is made under a different identity and for some of the objects for sale.

<sup>3</sup> A cost function is submodular if  $C(A \cup B) \leq C(A) + C(B) - C(A \cap B)$ .

set of all users (the true names) is private knowledge, that is, the mechanism does not know if a name  $i$  is a true name or a pseudonym, and it cannot distinguish if two names have been created by the same user. We assume further that a user cannot use two names simultaneously and thus can only make a *single bid*. For instance, the system may be able to detect that two names correspond to the same IP address and thus to the same user. Moreover, every negative result proved under the single bid assumption also holds when users can make multiple bids.

A *mechanism* for a cost-sharing game is a pair  $(S, P)$  defined as follows. The input to the mechanism is an  $n$ -dimensional *bid vector*  $v$  where the  $i^{\text{th}}$  coordinate  $v_i$  is either  $\perp$ , indicating that no user submitted her bid using the name  $i$ , or it is equal to the bid submitted under the name  $i$ . The mechanism outputs the names of the winners and their prices: the user who submitted her bid using a name  $i \in S(v)$  receives the service at the price  $P(v, i)$ ; all other users do not get served and do not pay. When user  $i$  bids using the name  $j$  and her valuation for the service is  $v_i^*$ , she derives a *utility* equal to

$$S(v, j) \cdot v_i^* - P(v, j),$$

where  $S(v, j)$  is equal to 1 if  $j \in S(v)$ , and 0 otherwise. We consider only mechanisms  $(S, P)$  that satisfy the following standard requirements: *voluntary participation* meaning that  $P(v, j) = 0$  if  $j \notin S(v)$  and  $P(v, j) \leq v_j$  otherwise; *no positive transfer* meaning that prices  $P(v, j)$  are always nonnegative; *consumer sovereignty*, meaning that every user can get the service if bidding sufficiently high, regardless of the bids of the other users.

There is a *symmetric cost function*  $C$  whose value depends only on the number of users that get the service, that is,  $C(S) = C(T)$  whenever  $S$  and  $T$  are two subsets of names of the same size. Following [4], we speak of a *symmetric cost-sharing game*. A mechanism  $(S, P)$  is  $\alpha$ -*budget-balanced* if for every bid vector  $v$ , it holds that

$$C(S(v)) \leq \sum_{i \in S(v)} P(v, i) \leq \alpha \cdot C(S(v))$$

and it is *budget-balanced* if this condition holds for  $\alpha = 1$ . We sometimes write  $C_s$  in place of  $C(S)$ , where  $s = |S|$ .

We require that, in the scenario in which pseudonyms are not used, users cannot improve their utilities by misreporting their valuations. A mechanism  $(S, P)$  is *group-strategyproof* if no group of users can raise the utility of some of its members without lowering the utility of some other member. Mechanism  $(S, P)$  is *strategyproof* if this condition is required to hold for  $G$  of size one only.

The next desideratum is that, when a user reports truthfully her valuation, she cannot improve her utility when bidding with a pseudonym. For any bid vector  $b$ , we let  $U(b)$  be the set of names that have been used to submit the bids, that is, those  $i$ 's such that  $b_i \neq \perp$ . We also let  $b_{i \rightarrow j}$  be the vector obtained from  $b$  by exchanging  $b_i$  with  $b_j$ , where  $i \in U(b)$  and  $j \notin U(b)$ .

**Definition 1 (renameproof mechanism).** A *cost-sharing mechanism*  $(S, P)$  is *renameproof* if the following holds. For any bid vector  $v$  and for any  $i$  and  $j$  such that  $i \in U(v)$  and  $j \notin U(v)$

$$S(v, i) \cdot v_i - P(v, i) \geq S(v_{i \rightarrow j}, j) \cdot v_i - P(v_{i \rightarrow j}, j).$$

We think of  $v$  as the bid vector in which user  $i$  reports truthfully her valuation and her name. Thus, the above definition says that no user can improve her utility by using a pseudonym in place of her name, no matter if the other users report truthfully their names and valuations. Symmetry implies that, in a renameproof mechanism, the utility of an agent must be constant over all names. Note that this does not imply that the price must be the same for all serviced users (see Example 4 below). We decide to present the definition in this form since this will naturally lead to a weaker condition used by mechanisms based on “reputation” in Section 4.

### 3 Renameproof Mechanisms and (Non-) Identical Prices

A general approach to design (approximately) budget-balanced mechanisms is to define a suitable *cost-sharing method* for the cost function of the problem. That is, a function  $\xi$  which (approximately) divides the overall cost among the serviced users:  $\xi(S, i)$  is the price associated to user  $i$  when the subset  $S$  is served, and

$$1 \leq \sum_{i \in S} \xi(S, i) / C(S) \leq \alpha,$$

for  $\alpha \geq 1$  and for every subset  $S$  of users. Such a function is an  $\alpha$ -*budget-balanced* cost-sharing method (or simply budget-balanced if  $\alpha = 1$ ). The main technique to construct groupstrategyproof cost-sharing mechanisms is due to Moulin [23]:

*Example 3 (Moulin mechanism  $M_\xi$ ).* Initially we set  $S$  as all users in  $U(b)$ . At each round we remove all users in  $S$  whose bid is less than the price  $\xi(S, i)$  offered by the mechanism. We iterate this step until all users in the current set  $S$  accept the offered price, or no user is left. We service the final set  $S$  obtained in this way and charge each user  $i \in S$  an amount  $\xi(S, i)$ .

We first observe that renameproofness by itself is not a problem since the *average-cost* mechanism [24] which charges all users the same price is budget-balanced and renameproof. However, this mechanism is (group) strategyproof only when the average cost does not increase with the number of users (see Example 1). Therefore, it is natural to ask if there exist other mechanisms that are renameproof. To answer this question, we consider a general class of mechanisms that define a unique cost-sharing method and that are known to include *all* groupstrategyproof mechanisms [23]. Such mechanisms are termed *separable*:

**Definition 2 (separable mechanism).** A cost-sharing mechanism  $(S, P)$  is separable if it induces a unique cost-sharing method  $\xi = \xi_{(S, P)}$ . That is, there exists a cost-sharing method  $\xi = \xi_{(S, P)}$  such that  $P(v, i) = \xi(S(v), i)$  for all bids  $v$  and for all  $i$ .

Our first observation is that, in order to be renameproof, the price assigned to a user should be independent from her name:

**Definition 3.** A cost-sharing method  $\xi$  is name independent if the price associated to a user does not depend on her name. That is, for every  $R \subset \mathcal{N}$  and for any two  $i, j \notin R$ , it holds that  $\xi(R \cup \{i\}, i) = \xi(R \cup \{j\}, j)$ .

**Theorem 1.** *Every separable renameproof mechanism induces a name independent cost-sharing method.*

Because of the previous result, we will focus on name independent cost-sharing methods. The main idea is to regard each name-independent cost-sharing method as a “fractional coloring” of the complete hypergraph on  $n$  nodes.

**Definition 4.** *An  $\alpha$ -balanced  $(n, s)$ -coloring is a function  $x$  assigning a nonnegative weight  $x_R$  to every  $(s - 1)$ -subset  $R \subset \mathcal{N}$  such that, for every  $s$ -subset  $S \subseteq \mathcal{N}$ , it holds that*

$$1 \leq \sum_{R \in \text{subsets}(S)} x_R \leq \alpha \tag{2}$$


where  $\text{subsets}(S)$  denotes the family of all  $(s - 1)$ -subsets of  $S$ .

The connection between a name-independent  $\xi$  and  $(n, s)$ -colorings can be made explicit by considering

$$x_R^{(\xi, s)} := \xi(R \cup \{i\}, i) / C_s \tag{3}$$

and by observing that  $x^{(\xi, s)}$  must be  $\alpha$ -balanced if  $\xi$  is  $\alpha$ -budget-balanced. With this equivalent representation, we can argue about the existence of cost-sharing mechanisms using non-identical prices by studying “non-uniform”  $(n, s)$ -colorings.

*Example 4 (three users).* When servicing three users, Definition 4 boils down to the problem of assigning weights to the edges of the complete graph over  $n$  nodes so that every triangle has total weight in between 1 and  $\alpha$ . For  $n = 4$ , one possible way is as follows:



where on the right we show the corresponding prices (cost-sharing scheme). Indeed, this example actually yields a budget-balanced renameproof groupstrategyproof mechanism for  $n = 4$  (see [26] for the details).

Despite the above example, our main result below says that one cannot avoid identical prices as soon as the number of possible names is not very small:

**Theorem 2.** *Any  $\alpha$ -budget-balanced renameproof separable mechanism must charge each serviced user a price which is at least  $\frac{C_s}{s} (2^{s-1} (1 - \alpha) + \alpha)$  and at most  $\frac{C_s}{s} (2^{s-1} (\alpha - 1) + 1)$ , unless the number  $s$  of serviced users is more than half the number of names (i.e., unless  $s > n/2$ ). In particular, if the mechanism is budget-balanced (i.e.,  $\alpha = 1$ ) then all serviced users are charged the same price  $C_s/s$ .*

<sup>4</sup> An  $r$ -subset is a subset of cardinality  $r$ .

*Proof (Sketch).* We prove bounds on  $\alpha$ -balanced  $(n, s)$ -colorings. Let us consider the following incidence matrix  $A$  with  $\binom{n}{s}$  rows and  $\binom{n}{s-1}$  columns (originally defined in [13]). For any  $s$ -subset (row)  $S$  and any  $(s-1)$ -subset (column)  $R$ , define:

$$A_{S,R} = \begin{cases} 1 & \text{if } R \subset S, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

By definition  $A_S \cdot x = \sum_{R \in \text{subsets}(S)} x_R$  for any  $s$ -subset  $S$ . Thus the set of all  $\alpha$ -balanced  $(n, s)$ -colorings is given by the following polytope:

$$\mathbf{polytope} := \{x \in \mathbb{R}^q \mid \mathbf{1} \leq A \cdot x \leq \alpha\}$$

where  $q = \binom{n}{s-1}$  and  $\alpha = (\alpha, \dots, \alpha)$  for every real  $\alpha$ . Intuitively, each  $x_R$  corresponds to the price  $\xi(S, i)/C_s$  for  $R = S \setminus \{i\}$ , which is independent of  $i$ .

It is possible to prove that, for every  $x$  in **polytope** and for every  $(s-1)$ -subset  $R \subset \mathcal{N}$ , there exists  $\bar{x}$  in **polytope** such that  $x_R = \bar{x}_R$  and all the values  $\bar{x}_Q$  depend only on the size of  $R \cap Q$ . That is,  $\bar{x}_Q = \bar{x}_{Q'}$  for all  $(s-1)$ -subsets  $Q$  and  $Q'$  such that  $|Q \cap R| = |Q' \cap R|$ . Therefore that the minimum/maximum value of any component of  $x$  in **polytope** can be found by looking only at vectors  $\bar{x}$  as above. Let  $p_k$  denote the unique value such that  $\bar{x}_Q = p_k$  for all  $Q$  such that  $|R \cap Q| = k$ , for  $k = 0, \dots, s-1$ . (This is well-defined because  $2s \leq n$ .) Now let  $b := A \cdot \bar{x}$ , i.e.,  $b_k = k \cdot p_{k-1} + (s-k) \cdot p_k$ . We can express  $p_{s-1} = \bar{x}_R = x_R$  as a function that grows monotonically (increasing or decreasing) in any component  $b_k$  of  $b$ . Plugging in that  $1 \leq b_k \leq \alpha$  for every  $b_k$ , we obtain that  $\frac{2^{s-1}(1-\alpha)+\alpha}{s} \leq x_R \leq \frac{2^{s-1}(\alpha-1)+1}{s}$ , for any  $x$  in **polytope** and for any  $(s-1)$ -subset  $R$ . Due to Theorem 1 and the relation in (3), this implies the theorem.  $\square$

Notice that every groupstrategyproof mechanism must be separable [23]. Since we have just shown that (a) every budget-balanced renameproof separable mechanism is an average-cost mechanism, but since we also know (recall Example 1) that (b) average-cost mechanisms are not groupstrategyproof if costs are not submodular, we obtain the following characterization and impossibility result:

**Corollary 1.** *If the number of possible names is at least twice the number of users, then the following holds. If costs are submodular, then the Moulin mechanism charging all players equally is the only renameproof, budget-balanced and groupstrategyproof mechanism (up to welfare-equivalence).<sup>5</sup> When costs are not submodular, there is no mechanism that is at the same time renameproof, groupstrategyproof and budget-balanced.*

We note that the bounds of Theorem 2 become rather weak as  $s$  grows (and  $\alpha$  is fixed). Although the corresponding bounds on  $\alpha$ -balanced  $(n, s)$ -colorings are tight (see the full proof in [26]), we next show a different kind of bounds that essentially rule out constructions of mechanisms based on different prices.

**Theorem 3.** *For any  $\alpha$ , for any  $s$ , and for any  $\delta$ , there exists  $N = N(\alpha, s, \delta)$  such that the following holds for all  $n \geq N$ . For every  $\alpha$ -budget-balanced name-independent cost-sharing method  $\xi$  there exists a subset  $S$  of  $s$  users such that their prices satisfy  $|\xi(S, i) - \xi(S, i')| \leq \delta$  for every  $i$  and  $i'$  in  $S$ .*

<sup>5</sup> Two mechanisms are welfare-equivalent if they produce the same utilities.



*Proof (Sketch).* Round every  $x_R^{(\xi,s)}$  to the closest power of  $\delta' := \delta/C_s$ . This gives an  $(n, s)$ -coloring which takes at most  $c = \lceil \log_{\delta'} \alpha \rceil$  different values. The Ramsey Theorem (see e.g. [19]) says that, for sufficiently large  $n$ , there is an  $s$ -subset  $S$  such that any two  $(s-1)$ -subsets  $R, R' \subset S$  get the same value. In particular, for  $R = S \setminus \{i\}$  and  $R' = S \setminus \{i'\}$  this implies that  $|x_R^{(\xi,s)} - x_{R'}^{(\xi,s)}| \leq \delta'$ . The theorem follows from  $\delta' = \delta/C_s$  and from [3].  $\square$

**Two applications.** We present two applications of our results. The first one is the *public excludable good* problem, which plays a central role in the cost-sharing literature and arises as a special case of many optimization problems [8]. It corresponds to the simple cost function  $C_0 = 0$  and  $C_s = 1$  for all  $s \geq 1$ . Corollary 1 implies that the Shapley value mechanism [1] is the only deterministic one that meets budget-balance, groupstrategyproofness, and renameproofness. This gives an alternative axiomatic characterization of the Shapley value to the one in [8, Theorem 2].

The second application concerns certain *bin packing* or *scheduling* problems, in which users (items or jobs) are homogeneous [2, 4]. For instance, costs are of the form  $C_s = \lceil s/m \rceil$ , where  $m$  can be interpreted as the capacity of the bins (bins are identical and the cost is the number of bins needed) or the number of machines (the cost is the makespan of the computing facility offering the service). For these problems, Bleischwitz et al [4] proved that budget-balanced groupstrategyproof mechanisms must use *non-identical* prices. Corollary 1 implies that deterministic budget-balanced mechanisms which are both groupstrategyproof and renameproof *do not exist* (these two notions are *mutually exclusive*). Theorem 3 essentially rules out constructions with non-identical prices, and identical prices can only achieve 2-approximate budget-balance.

### 4 Reputation Helps

In this section we consider the scenario in which the mechanism has some additional information about the users. For instance, there might be *time stamps* for the times at which new identities have been created associated to users. Similarly, we might employ *reputation functions* that make it impossible for a user to obtain a better ranking [5]. An abstraction of these two scenarios is to consider the case in which each user can replace her name only with a “larger” one.

**Definition 5 (reputationproof mechanism).** A cost-sharing mechanism is reputationproof if it satisfies the condition of Definition 1 limited to  $j > i$ .

We show a *sufficient* condition for obtaining mechanisms that are both groupstrategyproof and reputationproof. Here the cost-sharing method can be represented by an  $n \times n$  matrix  $\xi = \{\xi_i^s\}$ , where each  $\xi^s$  is a vector of  $s$  prices. This is an example of the kind of the cost-sharing schemes we use here:

$$\xi = \begin{pmatrix} 1/2 & 1/2 & 1/2 & 1/2 \\ 1/2 & 1/2 & 1 & \\ 1/2 & 1/2 & & \\ 1 & & & \end{pmatrix} \tag{5}$$

---

<sup>6</sup> The Shapley value mechanism services the largest set  $S$  such that all users in  $S$  bid at least  $1/|S|$ .



where the scheme is regarded as a triangular matrix in which we leave empty the (irrelevant) values  $\xi_i^s$  with  $i > s$ . Users are ranked according to their (reported) names and the highest price is offered to the user with largest (reported) name.

**Definition 6.** Let  $\xi = \{\xi_j^s\}$  be a cost-sharing scheme containing only two values low and high  $>$  low, and say that  $\xi^s$  offers low prices if it is equal to (low, . . . , low). Such a cost-sharing scheme  $\xi$  is a same-two-price scheme if the following holds for every  $s$ . If  $\xi^s$  does not offer low prices, then  $\xi^{s-1}$  offers low prices and  $\xi^s = (\text{low}, \dots, \text{low}, \text{high})$ .

The main idea in [4] is to drop some “indifferent” users who bid the low price:

**Definition 7.** Let  $\xi$  be a same-two-price scheme and call indifferent those users bidding exactly the low price. The corresponding same-two-price mechanism is as follows:

1. Drop all users who do not even bid the low price;
2. If the resulting subset has size  $s$  such that  $\xi^s$  offers low prices, then service this set and charge all these users the low price;
3. Otherwise do the following:
  - (a) If there are indifferent users, then drop the one with the last name and charge all others the low price;
  - (b) Otherwise, drop the last user if she is not willing to pay the high price and charge all others the same price;
  - (c) If no user has been dropped in Steps (3a-3b) then charge the low price to all but the last user who is charged the high price.

**Theorem 4.** Every same-two-price mechanism is both groupstrategyproof and reputationproof.

*Proof (Sketch).* The same-two-price mechanism is a special case of two-price mechanisms in [4] and thus it is groupstrategyproof. We prove that it is also reputationproof as follows. If  $v$  is such that  $i$  is offered the high price (Step 3b) then the same happens to  $j$  for  $v_{i \rightarrow j}$ , for any  $j > i$  with  $j \notin U(v)$ . In all other cases,  $i$  obtains the highest possible utility because she gets served if bidding strictly more than the low price (an invariant of the mechanism is that the utility of every user bidding at most the low price is 0).  $\square$

*Example 5 (two-machine scheduling).* The cost of scheduling  $s$  identical jobs on two identical machines is the makespan, that is,  $C_s = \lceil s/2 \rceil$ . No budget-balanced mechanism which is groupstrategyproof can be renameproof (see Section 3). In contrast, there exists a budget-balanced groupstrategyproof mechanism which is also reputationproof: The same-two-price mechanism corresponding to the cost-sharing methods of the form (5).

Unfortunately, not all two-price mechanisms from [4] are reputationproof. The argument is similar to the discussion in [26] about Moulin mechanisms. Concerning the notion of reputationproof mechanism, we note that BitThief [21] is a sophisticated client that free rides on BitTorrent using a weakness of the protocol: newcomers are allowed to download data “for free” because they are supposed to have nothing to upload yet. BitTorrent resembles a mechanism that is *not* reputationproof because it offers a better price to users that have “no reputation”.

## 5 Concluding Remarks and Open Questions

Despite the fact that in symmetric cost-sharing games all users play the same role, in most of the cases one must employ *non-identical* prices in order to get budget-balanced and groupstrategyproof mechanisms [4]. In sharp contrast, we have shown that in order to make it disadvantageous for users to use pseudonyms, one must use *identical* prices, thus implying that groupstrategyproofness and renameproofness can be achieved only *separately*. The results apply also to *randomized* mechanisms in which using the “true” name is a dominant strategy for all coin tosses. The notion of reputationproof mechanism captures in a natural way the use of “reputation” to overcome these difficulties.

It would be very interesting to characterize the class of “priority-based” mechanisms (sequential [18], acyclic [22], two-price [4]) that are also reputationproof. Another important issue is to consider multiple bids (false-nameproofness [32, 31]) and consider non-symmetric (“combinatorial”) cost-sharing games [16, 14, 29].

An interesting issue concerns the “social cost” of pseudonyms in cost-sharing games. For instance, we might consider the economic *efficiency loss* [29] caused by the use of pseudonyms (namely, the efficiency of arbitrary mechanisms versus those that use identical prices) and thus quantify how much “reputation” helps.

*Acknowledgements.* We wish to thank Roger Wattenhofer for several useful discussions, and to Carmine Ventre for comments on an earlier version of this work. We are grateful to an anonymous referee for many insightful comments. and, in particular, for suggesting the use of negative bids to simplify some of the proofs.

## References

1. Aggarwal, G., Fiat, A., Goldberg, A.V., Hartline, J.D., Immorlica, N., Sudan, M.: Derandomization of auctions. In: Proc. of 37th STOC (2005)
2. Bleischwitz, Y., Monien, B.: Fair cost-sharing methods for scheduling jobs on parallel machines. In: Calamoneri, T., Finocchi, I., Italiano, G.F. (eds.) CIAC 2006. LNCS, vol. 3998, pp. 175–186. Springer, Heidelberg (2006)
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: Kučera, L., Kučera, A. (eds.) MFCS 2007. LNCS, vol. 4708, pp. 657–668. Springer, Heidelberg (2007)
5. Cheng, A., Friedman, E.: Sybilproof reputation mechanisms. In: Proc. of the 3rd Workshop on Economics of Peer-to-Peer Systems (P2PECON), pp. 128–132 (2005)
6. Cohen, B.: Incentives Build Robustness in BitTorrent. In: Proc. of the 1st Workshop on Economics of Peer-to-Peer Systems, P2PECON (2003)
7. Conitzer, V.: Limited verification of identities to induce false-name-proofness. In: Proc. of the 11th TARK, pp. 102–111 (2007)
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. Douceur, J.R.: The sybil attack. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 251–260. Springer, Heidelberg (2002)

10. Feldman, M., Papadimitriou, C., Chuang, J., Stoica, I.: Free-riding and whitewashing in peer-to-peer systems. In: Proc. of ACM SIGCOMM workshop PINS, pp. 228–236 (2004)
11. Friedman, E., Resnick, P.: The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy* 10(2), 173–199 (2001)
12. Goldberg, A., Hartline, J.: Collusion-resistant mechanisms for single-parameter agents. In: Proc. of the 16th SODA, pp. 620–629 (2005)
13. Gottlieb, D.H.: A certain class of incidence matrices. In: Proc. Amer. Math. Soc., pp. 1233–1237 (1966)
14. Immorlica, N., Mahdian, M., Mirrokni, V.S.: Limitations of cross-monotonic cost sharing schemes. In: Proc. of the 16th SODA, pp. 602–611 (2005)
15. Izal, M., Urvoy-Keller, G., Biersack, E., Felber, P., Hamra, A., Garcés-Erice, L.: Dissecting BitTorrent: Five Months in a Torrent’s Lifetime. In: Barakat, C., Pratt, I. (eds.) PAM 2004. LNCS, vol. 3015, pp. 1–11. Springer, Heidelberg (2004)
16. Jain, K., Vazirani, V.V.: Applications of approximation algorithms to cooperative games. In: Proc. of the 33rd STOC, pp. 364–372 (2001)
17. Juarez, R.: Collusion-proof cost sharing. Technical report, Mimeo University of Hawaii (2007)
18. Juarez, R.: Group strategyproof cost sharing: the role of indifferences. Technical report, Mimeo University of Hawaii (2008)
19. Jukna, S.: *Extremal Combinatorics With Applications in Computer Science*. Springer, Heidelberg (2001)
20. Könemann, J., Leonardi, S., Schäfer, G., van Zwam, S.H.M.: A group-strategyproof cost sharing mechanism for the steiner forest game. *SIAM J. Comput.* 37(5), 1319–1341 (2008)
21. Locher, T., Moor, P., Schmid, S., Wattenhofer, R.: Free Riding in BitTorrent is Cheap. In: Proc. of the 5th HotNets (2006)
22. Mehta, A., Roughgarden, T., Sundararajan, M.: Beyond Moulin Mechanisms. In: Proc. of ACM EC, pp. 1–10 (2007)
23. Moulin, H.: Incremental cost sharing: characterization by coalition strategy-proofness. *Social Choice and Welfare* 16(2), 279–320 (1999)
24. Moulin, H., Shenker, S.: Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory* 18(3), 511–533 (2001)
25. Newsome, J., Shi, E., Song, D., Perrig, A.: The sybil attack in sensor networks: analysis & defenses. In: Proc. of the 3rd IPSN, pp. 259–268 (2004)
26. Penna, P., Schoppmann, F., Silvestri, R., Widmayer, P.: Pseudonyms in cost-sharing games. Technical report, AEOLUS (2009), <http://aeolus.ceid.upatras.gr/>
27. 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)
28. 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)
29. Roughgarden, T., Sundararajan, M.: New trade-offs in cost-sharing mechanisms. In: Proc. of the 38th STOC, pp. 79–88 (2006)
30. Wagman, L., Conitzer, V.: Optimal false-name-proof voting rules with costly voting. In: Proc. of 23rd AAAI, pp. 190–195 (2008)
31. Yokoo, M.: Characterization of Strategy/False-name Proof Combinatorial Auction Protocols: Price-oriented, Rationing-free Protocol. In: Proc. of the 18th IJCAI, pp. 733–742 (2003)
32. 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)

# Computing Optimal Contracts in Series-Parallel Heterogeneous Combinatorial Agencies\*

Yuval Emek<sup>1</sup> and Michal Feldman<sup>2</sup>

<sup>1</sup> School of Electrical Engineering, Tel Aviv University, Tel Aviv, 69978 Israel  
yuvale@eng.tau.ac.il

<sup>2</sup> School of Business Administration and the Center for the Study of Rationality,  
The Hebrew University of Jerusalem, Jerusalem, 91904 Israel  
michal.feldman@huji.ac.il

**Abstract.** We study an economic setting in which a principal motivates a team of strategic agents to exert costly effort toward the success of a joint *project*. The action taken by each agent is *hidden* and affects the (binary) outcome of the agent's individual *task* in a stochastic manner. A Boolean function, called *technology*, maps the individual tasks' outcomes to the outcome of the whole project. The principal induces a Nash equilibrium on the agents' actions through payments that are conditioned on the project's outcome (rather than the agents' actual actions) and the main challenge is that of determining the Nash equilibrium that maximizes the principal's net utility, referred to as the *optimal contract*.

Babaioff, Feldman and Nisan suggest and study a basic *combinatorial agency* model for this setting, and provide a full analysis of the AND technology. Here, we concentrate mainly on OR technologies and on *series-parallel* (SP) technologies, which are constructed inductively from their building blocks — the AND and OR technologies. We provide a complete analysis of the computational complexity of the optimal contract problem in OR technologies, which resolves an open question and disproves a conjecture raised by Babaioff et al. In particular, we show that while the AND case admits a polynomial time algorithm, computing the optimal contract in an OR technology is NP-hard. On the positive side, we devise an FPTAS for the OR case and establish a scheme that given any SP technology, provides a  $(1 + \epsilon)$ -approximation for all but an  $\hat{\epsilon}$ -fraction of the relevant instances (for which a failure message is output) in time polynomial in the size of the technology and in the reciprocals of  $\epsilon$  and  $\hat{\epsilon}$ .

## 1 Introduction

We consider the setting in which a principal motivates a team of rational agents to exert costly effort towards the success of a joint project, where their actions

---

\* This work was partially supported by the Israel Science Foundation (grant number 1219/09) and by the Leon Recanati Fund of the Jerusalem school of business administration.

are hidden from her. The outcome (usually, success or failure of the project) is stochastically determined by the set of actions taken by the agents and is visible to all. As agents' actions are invisible, their compensation depends on the outcome and the principal's challenge is to design contracts (conditional payments to the agents) as to maximize her net utility, given the payoff that she obtains from a successful outcome.

The problem of hidden-action in production teams has been extensively studied in the economics literature [7,9,15,8,16]. More recently, the problem has been examined from a computational perspective [5,11,2]. This line of research complements the field of Algorithmic Mechanism Design (AMD) [11,10,13,4,12] that received much attention in the last decade. While AMD studies the design of mechanisms in scenarios characterized by private information held by the individual agents, our focus is on the complementary problem, that of hidden-action taken by the individual agents. In [1], the authors concentrated on the case of homogeneous users, i.e., agents with identical capabilities and costs. The current work extends the original work to the more complex (yet realistic) case, that of heterogeneous agents.

For example, consider an executive board that assigns stock options to the company's employees in attempt to motivate them to excel so that the value of the company increases. While the exact contribution of each individual may be difficult to measure, the stock's market price is visible to all, hence it serves as the groundwork in determining future payments to the staff. Given the significance of each employee (position, rank, etc.), what is the optimal incentive (in terms of stock options) he should get? What is the complexity of computing the optimal incentives in the above examples? This is the type of questions that motivate us in this work.

*The model.* We use the model presented in [1] (which is an extension of the model devised in [17]). In this model, a principal employs a set<sup>1</sup>  $N$  of agents in a joint *project*. Each agent  $i$  takes an action  $a_i \in \{0, 1\}$ , which is known only to him, and succeeds or fails in his own *task* probabilistically and independently. The individual outcome of agent  $i$  is denoted by  $x_i \in \{0, 1\}$ . If the agent shirks ( $a_i = 0$ ), he succeeds in his individual task ( $x_i = 1$ ) with probability  $0 < \gamma_i < 1$  and incurs no cost. If, however, he decides to exert effort ( $a_i = 1$ ), he succeeds with probability  $0 < \delta_i < 1$ , where  $\delta_i > \gamma_i$ , but incurs some positive real *cost*  $c_i > 0$ .

A key component of the model is the way in which the individual outcomes determine the outcome of the whole project. We assume a monotone Boolean function  $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$  that determines whether the project succeeds as a function of the individual outcomes of the  $n$  agents' tasks (and is not determined by any set of  $n - 1$  agents). Two fundamental examples of such Boolean functions are AND and OR. The AND function is the logical conjunction of  $x_i$  ( $\varphi(x_1, \dots, x_n) = \bigwedge_{i \in N} x_i$ ), representing the case in which the project

---

<sup>1</sup> Unless stated otherwise, we assume that  $N = [n]$ , where  $[n]$  denotes the set  $\{1, \dots, n\}$ .

succeeds only if *all* agents succeed in their tasks. In this case, we say that the agents *complement* each other. The OR function represents the other extreme, in which the project succeeds if *at least one* of the agents succeeds in his task. This function is the logical disjunction of  $x_i$  ( $\varphi(x_1, \dots, x_n) = \bigvee_{i \in N} x_i$ ), and we say that the agents *substitute* each other.

A more general class of monotone Boolean functions is that of *series-parallel* (SP) functions. This class is defined inductively as follows. The un-argument identity function is considered SP. Consider some two SP functions  $\varphi_l : \{0, 1\}^{n_l} \rightarrow \{0, 1\}$  and  $\varphi_r : \{0, 1\}^{n_r} \rightarrow \{0, 1\}$ . The Boolean functions  $\varphi_l \wedge \varphi_r : \{0, 1\}^{n_l+n_r} \rightarrow \{0, 1\}$ , defined as the logical conjunction of  $\varphi_l$  and  $\varphi_r$ , and  $\varphi_l \vee \varphi_r : \{0, 1\}^{n_l+n_r} \rightarrow \{0, 1\}$ , defined as the logical disjunction of  $\varphi_l$  and  $\varphi_r$ , are also considered SP. We refer to the former (respectively, the latter) as a *series composition* (resp., a *parallel composition*) of  $\varphi_l$  and  $\varphi_r$ , hence the name series-parallel. Since series and parallel compositions are associative, it follows that the class of SP Boolean functions is indeed a generalization of both AND and OR Boolean functions.

Given the action profile  $a = (a_1, \dots, a_n) \in \{0, 1\}^n$  and a monotone Boolean function  $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$ , the *effectiveness* of the action profile  $a$ , denoted by  $f(a)$ , is the probability that the whole project succeeds under  $a$  and  $\varphi$  according to the distribution specified above. That is, the effectiveness  $f(a)$  is defined as the probability that  $\varphi(x_1, \dots, x_n) = 1$ , where  $x_i \in \{0, 1\}$  is determined probabilistically (and independently) by  $a_i$ : if  $a_i = 0$ , then  $x_i = 1$  with probability  $\gamma_i$ ; if  $a_i = 1$ , then  $x_i = 1$  with probability  $\delta_i$ . The monotonicity of  $\varphi$  and the assumption that  $\delta_i > \gamma_i$  for every  $i \in N$  imply the monotonicity of the effectiveness function  $f$ , i.e., if we denote by  $a_{-i} \in \{0, 1\}^{n-1}$  the vector of actions taken by all agents excluding agent  $i$  (namely,  $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ ), then the effectiveness function must satisfy  $f(1, a_{-i}) > f(0, a_{-i})$  for every  $i \in N$  and  $a_{-i} \in \{0, 1\}^{n-1}$ . Note that it is inherent to our model (in fact, the model of [II](#)) that the effectiveness  $f(a)$  consists of a “probabilistic component” that determines the individual outcomes  $x_1, \dots, x_n$  and a “deterministic component” that maps these individual outcomes to success or failure of the whole project.

The agents’ success probabilities, the costs of exerting effort, and the monotone Boolean function that determines the final outcome define the *technology*, formally defined as the five-tuple  $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$ , where  $N$  is a (finite) set of agents;  $\gamma_i$  (respectively,  $\delta_i$ ) is the probability that  $x_i = 1$  when agent  $i$  shirks (resp., when agent  $i$  exerts effort), where  $\delta_i > \gamma_i$ ;  $c_i$  is the cost incurred on agent  $i$  for exerting effort; and  $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$  is the monotone Boolean function that maps the individual outcomes  $x_1, \dots, x_n$  to the outcome of the whole project. We sometimes abuse notation and refer to the Boolean function  $\varphi$  as the technology. It is important to emphasize that the technology is assumed to be known by the principal and the agents.

Since exerting effort entails some positive cost, an agent will not exert effort unless induced to do so by appropriately designed incentives. The principal can motivate the agents by offering them individual *payments*. However, due to the non-visibility of the agents’ actions, the individual payments cannot be

directly contingent on the actions of the agents, but rather only on the success of the whole project. The *conditional payment* to agent  $i$  is thus given by a real value  $p_i \geq 0$  that is granted to agent  $i$  by the principal if the project succeeds (otherwise, the agent receives 0 payment<sup>2</sup>).

The expected *utility* of agent  $i$  under the profile of actions  $a = (a_1, \dots, a_n)$  and the conditional payment  $p_i$  is  $p_i \cdot f(a)$  if  $a_i = 0$ ; and  $p_i \cdot f(a) - c_i$  if  $a_i = 1$ . Given a real *payoff*  $v > 0$  that the principal obtains from a successful outcome of the project, the principal wishes to design the payments  $p_i$  as to maximize her own expected *utility* defined as  $U_a(v) = f(a) \cdot (v - \sum_{i \in N} p_i)$ , where the action profile  $a$  is assumed to be at Nash-equilibrium with respect to the payments  $p_i$  (i.e., no agent can improve his utility by a unilateral deviation). As multiple Nash equilibria may (and actually do) exist, we focus on the one that maximizes the utility of the principal. This is as if we let the principal choose the desired Nash equilibrium, and “suggest” it to the agents. The following observation is established in [1].

**Observation 1.** *The best conditional payments (from the principal’s point of view) that induce the action profile  $a \in \{0, 1\}^n$  as a Nash equilibrium are  $p_i = 0$  for agent  $i$  who shirks ( $a_i = 0$ ), and  $p_i = \frac{c_i}{\Delta_i(a_{-i})}$  for agent  $i$  who exerts effort ( $a_i = 1$ ), where  $\Delta_i(a_{-i}) = f(1, a_{-i}) - f(0, a_{-i})$ . (Note that the monotonicity of the effectiveness function guarantees that  $\Delta_i(a_{-i})$  is always positive.)*

The last observation implies that once the principal chooses the action profile  $a \in \{0, 1\}^n$ , her (maximum) expected utility is determined to be  $U_a(v) = f(a) \cdot (v - p(a))$ , where  $p(a)$  is the total *payment* (in case of a successful outcome of the project), given by  $p(a) = \sum_{i|a_i=1} \frac{c_i}{\Delta_i(a_{-i})}$ . Therefore the principal’s goal is merely to choose a subset  $S \subseteq N$  of agents that exert effort (the rest of the agents shirk) so that her expected utility is maximized. The agent subset  $S$  is referred to as a *contract* and we say that the principal *contracts with agent  $i$*  if  $i \in S$ . We sometimes abuse notation and denote  $f(S)$ ,  $p(S)$  and  $U_S(v)$  instead of  $f(a)$ ,  $p(a)$  and  $U_a(v)$ , respectively, where  $a_i = 1$  if  $i \in S$  and  $a_i = 0$  if  $i \notin S$ . Given the principal’s payoff  $v > 0$ , a contract  $T \subseteq N$  is said to be *optimal* if  $U_T(v) \geq U_S(v)$  for every contract  $S \subseteq N$ .

While finding the optimal set of payments that induces a particular set of agents to exert effort is a straightforward task (and can be efficiently computed), finding an optimal contract for a given payoff  $v > 0$  is the main challenge addressed in this paper. Given a technology  $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$ , we refer to the collection of contracts that can be obtained as an optimal contract for some payoff as the *orbit* of  $t$  (ties between different contracts are broken according to a lexicographic order<sup>3</sup>). Once the contract  $S \subseteq N$  is chosen, the expected utility of the principal  $U_S(v) = f(S)(v - p(S))$  becomes a linear function of the payoff  $v$ . Therefore each contract  $S$  corresponds to some line in  $\mathbb{R}^2$ .

<sup>2</sup> We impose the *limited liability* constraint, implying that the principal can pay the agents but not fine them. Thus, all payments must be non-negative.

<sup>3</sup> This implies that there are no two contracts with the same effectiveness in the orbit.



It follows that computing the orbit of  $t$  is equivalent to identifying the (positive) top envelope of the line collection  $\{U_S(\cdot) \mid S \subseteq N\}$  in  $\mathbb{R}^2$ .

It is easy to see that for sufficiently low payoffs, no agent will ever be contracted while for sufficiently high payoffs, all agents will always be contracted. Therefore the *trivial contracts*  $\emptyset$  and  $N$  are always in the orbit. Let  $v^* = \inf\{v > 0 \mid N \text{ is optimal for } v\}$ . Clearly, the trivial contract  $N$  is optimal for every  $v > v^*$  and the infinite interval  $(v^*, \infty)$  does not exhibit any transitions in the orbit. We refer to the payoffs in the interval  $(0, v^*]$  as the *relevant* payoffs.

*Our results.* Multi-agent projects may exhibit delicate combinatorial structures of dependencies between the agents' actions, which can be represented by a wide range of monotone Boolean functions. In the two extremes of this range reside two simple and natural functions, namely AND and OR, which correspond to the respective cases of pure complementarities and pure substitutabilities. However, real-life technologies are usually composed of various components that exhibit different combinations of complementarities and substitutabilities. The class of SP Boolean functions represents exactly those technologies that can be inductively constructed from AND and OR components.

SP Boolean functions are of great interest to computer science. For instance, they play an important role in combinatorial games due to their equivalence to game trees (and-or trees). In addition, many of the graph-theoretic problems that are computationally hard in general have been shown to admit efficient solutions when applied to *series-parallel* graphs, which are the graph theoretic equivalent of SP functions. Perhaps the best example in our context is the *network reliability* problem [14], which reduces to the optimal contract problem in network technologies [1]. While the network reliability problem is  $\#P$ -complete on general networks, it admits an efficient algorithm when applied to series-parallel networks.

Obviously, a first step in the analysis of SP technologies is the analysis of their building blocks, namely, the AND and OR technologies. The AND case was fully analyzed in [1]. In particular, it was (implicitly) shown that the optimal contract of any AND technology can be computed in polynomial time. In contrast, the OR case was left unresolved to the most part. Specifically, it was left as an open question whether the optimal contract problem on OR technologies can be solved in polynomial time.

We provide a complete analysis of the computational complexity of the optimal contract problem on OR technologies. Our first theorem addresses the hardness of this variant.

**Theorem 1.** *The problem of computing the optimal contract in OR technologies is NP-hard*<sup>4</sup>.

A formal proof of this theorem appears in the full version; an overview is provided in Section 2. Note that aside from establishing the computational hardness of

<sup>4</sup> The problem remains NP-hard even for the special case in which  $c_i = 1$  and  $\delta_i = 1 - \gamma_i$  for every  $i \in N$ .



the problem, our analysis implies the existence of OR technologies which admit exponential-size orbits, thus refuting a conjecture raised in [1].

On the positive side, in Section 3.1 we devise a scheme for SP technologies which serves as the key ingredient in establishing the following approximations. For OR technologies (a special case of SP technologies), we prove Theorem 2 in Section 3.2.

**Theorem 2.** *The problem of computing the optimal contract in OR technologies admits a fully polynomial-time approximation scheme (FPTAS).*

General SP technologies are considerably more involved and the approximability of the optimal contract problem on such technologies remains an open question. However, an interesting insight into this question is provided by a scheme that approximates all but a small fraction of the relevant payoffs. Due to lack of space, the proof of the following theorem is deferred to the full version.

**Theorem 3.** *Given an SP technology  $t$  and two real parameters  $0 < \epsilon, \hat{\epsilon} \leq 1$ , there exists a scheme that on input payoff  $v > 0$ , either returns a  $(1 + \epsilon)$ -approximate solution for  $v$  or outputs a failure message, in time  $\text{poly}(|t|, 1/\epsilon, 1/\hat{\epsilon})$ . Assuming that  $F \subseteq \mathbb{R}_{>0}$  is the set of reals on which the scheme outputs a failure message, it is guaranteed that  $\int_0^\infty 1_F(v)dv \leq \hat{\epsilon}v^*$ , where  $1_F$  is the characteristic function of  $F$ .*

It may be the case that the hardness of the optimal contract problem on SP technologies is somehow “concentrated” exactly in those payoffs which cannot be reached by the scheme of Theorem 3. However, if an instance of the problem is chosen uniformly at random out of the “relevant instances”, then with high probability our scheme provides a good approximation for this instance. (Recall that the trivial contract  $N$  is optimal for any non-relevant payoff.) In fact, the payoffs  $v$  on which the scheme of Theorem 3 outputs a failure message belong to a small (polynomial) number of sub-intervals of  $(0, v^*]$ ; by making the parameter  $\hat{\epsilon}$  smaller, we decrease the guaranteed bound on the size of each such sub-interval.

It is interesting to contrast the aforementioned results with the *observable-action* case, where the agents’ actions are not hidden and may be contracted on, which admits a polynomial time algorithm for SP technologies [3].

Finally, we obtain a positive result regarding the general case. Consider an arbitrary technology  $t$  and let  $\mathcal{S}$  be a collection of contracts. Given some real  $\alpha > 1$ , we say that  $\mathcal{S}$  is an  $\alpha$ -approximation of  $t$ ’s orbit if for every payoff  $v$ , there exists a contract  $S \in \mathcal{S}$  such that  $U_S(v) \geq \frac{U_T(v)}{\alpha}$ , where  $T$  is optimal for  $v$ . Due to lack of space, the proof of the following theorem is deferred to the full version.

**Theorem 4.** *For every technology  $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$  and for any  $\epsilon > 0$ , the orbit of  $t$  admits a  $(1 + \epsilon)$ -approximation of size  $\text{poly}(|t|, 1/\epsilon)$ .*

Unfortunately, in the case of arbitrary technologies (as opposed to OR technologies) we do not know how to construct the approximating collection efficiently.

## 2 NP-Hardness of OR Technologies

We present a polynomial time Turing reduction from X3SAT (Problem LO4 in [6]) to the problem of computing an optimal contract for an OR technology. Recall that a 3-CNF formula  $\phi$  is solvable under X3SAT if there exists a truth assignment for the variables of  $\phi$  that assigns true to exactly one literal in every clause. The X3SAT problem is known to be NP-hard even if the literals in  $\phi$  are all positive. Given a 3-CNF formula  $\phi$  with  $m$  clauses and  $n$  variables in which all literals are positive, we construct an OR technology  $t = \langle N, \{\gamma_j\}_{j=1}^{n+5}, \{\delta_j\}_{j=1}^{n+5}, \{c_j\}_{j=1}^{n+5}, \varphi \rangle$  such that (1) the agent set  $N$  contains  $n + 5$  agents; (2) the cost incurred on agent  $j$  for exerting effort is  $c_j = 1$  for every  $j \in N$ ; and (3)  $\gamma_j = 1 - \delta_j$  for every  $j \in N$ . The construction is designed to guarantee that by performing  $O(n)$  queries, each reveals the optimal contract for some carefully chosen payoff, we can decide whether  $\phi$  is solvable under X3SAT.

Let  $\mathcal{W} = \{0, 1, 2, 3\}^{m+2} \times \{0, 1\}^2$ . Each agent  $j \in N$  is assigned with a vector  $\mathbf{u}^j = (u_0^j, \dots, u_{m+3}^j) \in \mathcal{W}$ . The first  $n$  agents correspond to the  $n$  variables of the 3-CNF formula  $\phi$  and affect coordinates  $1, \dots, m$  in a manner that reflects the appearance of their corresponding variables in the  $m$  clauses. The additional 5 agents affect coordinates  $0, m + 1, m + 2, m + 3$  and are provided for the sake of analysis. We extend the assignment of vectors to sets of agents (a.k.a. contracts) in a natural way: given a contract  $S \subseteq N$ , we define the vector  $\mathbf{u}^S = \sum_{j \in S} \mathbf{u}^j$ . (Note that different contracts may be assigned with the same vector.) The assignment of vectors to contracts guarantees that the formula  $\phi$  is solvable under X3SAT if and only if there exists a contract  $S$  with vector  $\mathbf{u}^S = (1, \dots, 1)$ .

The parameters of  $\{\gamma_j\}_{j=1}^{n+5}$  and  $\{\delta_j\}_{j=1}^{n+5}$  are defined as follows. Consider the vector  $\mathbf{x} = (x_0, \dots, x_{m+3})$  in  $\mathcal{W}$ . Let  $\sigma(\mathbf{x}) = \sum_{i=0}^{m+1} x_i 4^i$  and fix  $\mu = 4^{5(m+2)}$ . The evaluation of  $\mathbf{x}$  is defined to be  $\tau(\mathbf{x}) = \left(1 + \frac{1}{\mu}\right)^{\sigma(\mathbf{x})} \cdot \mu^{2x_{m+2}} \cdot \mu^{5x_{m+3}}$ . Let  $\epsilon = \mu^{-\kappa}$ , where  $\kappa$  is a sufficiently large constant. We would have wanted to fix  $\gamma_j = 1 - \delta_j = \tau(\mathbf{u}^j) \cdot \epsilon$  for every  $j \in N$ . Unfortunately, the standard binary representation of  $\tau(\mathbf{u}^j)$  may be much larger than the binary representation of  $\phi$  for some  $j$ , and in particular, exponential in  $m$ . To overcome this obstacle, we use a carefully chosen estimation of  $\tau(\mathbf{u}^j)$ , so that on the one hand, the desired properties of the evaluation function are preserved, and on the other hand, the binary representation of  $\gamma_j$  (and  $\delta_j$ ) is polynomial in  $m$ . In particular, the choice of  $\{\gamma_j\}_{j=1}^{n+5}$  and  $\{\delta_j\}_{j=1}^{n+5}$  guarantees that for every two contracts  $S, T \subseteq N$ ,  $f(S) > f(T)$  if and only if  $|S| > |T|$  or  $|S| = |T|$  and  $\mathbf{u}^S$  is lexicographically smaller than  $\mathbf{u}^T$ .

We argue that if some contracts  $S$  with  $\mathbf{u}^S = (1, \dots, 1)$  exist, then at least one of them is in the orbit. This is done as follows. A vector  $\mathbf{x} = (x_0, \dots, x_{m+3})$  is said to be *protected* if  $x_{m+2} = x_{m+3} = 1$ . The key lemma of our proof asserts that any contract assigned with a protected vector  $\mathbf{x}$  cannot be dominated by any two contracts assigned with different vectors. Following some standard geometric arguments, we conclude that the contracts assigned with  $\mathbf{x}$  cannot be dominated by any set of (other) contracts. More formally, for every  $0 \leq k \leq n+5$ , we denote

$\Psi_k(\mathbf{x}) = \{S \subseteq N \mid \mathbf{u}^S = \mathbf{x} \text{ and } |S| = k\}$ , and show that for any protected vector  $\mathbf{x}$ , if  $\Psi_k(\mathbf{x})$  is not empty, then at least one contract in  $\Psi_k(\mathbf{x})$  is in the orbit. In particular, assuming that  $\mathbf{x} = (1, \dots, 1)$ , if  $\Psi_k(\mathbf{x}) \neq \emptyset$ , then there exist a contract  $S \in \Psi_k(\mathbf{x})$  and a payoff  $v_k^*$  such that  $S$  is optimal for  $v_k^*$ .

Computing the payoff  $v_k^*$  for every  $1 \leq k \leq n + 5$  remains our ultimate challenge. To achieve this goal, we define two additional vectors  $\mathbf{w} = (2, 1, 1, \dots, 1) \in \mathcal{W}$  and  $\mathbf{y} = (0, 1, 1, \dots, 1) \in \mathcal{W}$ . The choice of the additional vectors guarantees that if  $\Psi_k(\mathbf{x})$  is not empty, then neither are  $\Psi_k(\mathbf{y})$  and  $\Psi_k(\mathbf{w})$ . Suppose that  $\Psi_k(\mathbf{x}) \neq \emptyset$  and fix  $\lambda_k^{w,x} = \max\{v[S, T] \mid S \in \Psi_k(\mathbf{w}) \text{ and } T \in \Psi_k(\mathbf{x})\}$  and  $\lambda_k^{x,y} = \min\{v[S, T] \mid S \in \Psi_k(\mathbf{x}) \text{ and } T \in \Psi_k(\mathbf{y})\}$ , where  $v[S, T]$  is the intersection payoff of  $S$  and  $T$ , i.e.,  $U_S(v[S, T]) = U_T(v[S, T])$ . We show that the optimal contract for every  $\lambda_k^{w,x} < v < \lambda_k^{x,y}$  must be in  $\Psi_k(\mathbf{x})$ . The analysis is completed by identifying some payoff  $\lambda_k^{w,x} < v_k^* < \lambda_k^{x,y}$  such that the binary representation of  $v_k^*$  is polynomial in  $m$ .

The decision whether the formula  $\phi$  is solvable under X3SAT is now carried out as follows. For  $k = 1, \dots, n + 5$ , we query on the optimal contract  $S_k$  for the payoff  $v_k^*$ . If  $\mathbf{u}^{S_k}$  is of the form  $(1, \dots, 1)$  for some  $k$ , then  $\phi$  must be solvable. Otherwise, there does not exist any such contract and  $\phi$  is not solvable.

### 3 Approximations

#### 3.1 A Polynomial Time Scheme for SP Technologies

Consider some technology  $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$  and fix  $\Delta = \min\{\prod_{i \in N} \gamma_i, \prod_{i \in N} (1 - \delta_i)\}$ . In the full version we show that  $f(S) \in [\Delta, 1 - \Delta]$  for every contract  $S \subseteq N$ .

Our scheme is executed by an algorithm, referred to as **Algorithm Calibrate**. Consider an SP technology  $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$  input to **Algorithm Calibrate** and let  $0 < \rho \leq 1$  be the *performance parameter* of the algorithm. **Algorithm Calibrate** generates a collection  $\mathcal{C}$  of contracts in time  $O\left(\frac{n^3 \log^2(1/\Delta)}{\rho^2}\right)$ . (Note that the binary representation of  $\{\gamma_i\}_{i=1}^n$  and  $\{\delta_i\}_{i=1}^n$  requires  $\Omega(\log(1/\Delta))$  bits.) We will soon prove that for every contract  $T \subseteq N$ , there exists a contract  $S \in \mathcal{C}$  such that  $f(S) \geq \frac{f(T)}{(1+\rho)}$ , and  $p(S) \leq (1 + \rho)p(T)$ .

Let  $\eta = \frac{\rho \ln 2}{2n-1}$ , and let  $r = \max\{k \in \mathbb{Z}_{\geq 0} \mid \Delta(1 + \eta)^k < \frac{1}{2}\}$ . Since  $r < \log_{1+\eta}\left(\frac{1}{2\Delta}\right) = \log \frac{1}{2\Delta} \cdot \log_{1+\eta}(2)$ , and since  $\log_{1+\eta}(2) \leq \frac{1}{\eta}$ , we conclude that  $r < \frac{1}{\eta} \log \frac{1}{\Delta}$ . We partition the interval  $[\Delta, 1 - \Delta]$  into  $2r + 3$  smaller intervals  $[\Delta, \Delta(1 + \eta)], [\Delta(1 + \eta), \Delta(1 + \eta)^2], \dots, [\Delta(1 + \eta)^{r-1}, \Delta(1 + \eta)^r], [\Delta(1 + \eta)^r, \frac{1}{2}], [\frac{1}{2}, \frac{1}{2}], (\frac{1}{2}, 1 - \Delta(1 + \eta)^r], (1 - \Delta(1 + \eta)^r, 1 - \Delta(1 + \eta)^{r-1}], \dots, (1 - \Delta(1 + \eta)^2, 1 - \Delta(1 + \eta)], (1 - \Delta(1 + \eta), 1 - \Delta]$ . The collection of these smaller intervals is called the *scale*. The *precision* of the scale is defined as  $1 + \eta$ . We say that contract  $S$  is *calibrated* to interval  $\mathcal{I}$  in the scale if  $f(S) \in \mathcal{I}$ .

**Observation 2.** *Let  $S, S' \in N$  be some two contracts. The scale is designed to ensure that if  $S$  and  $S'$  are calibrated to the same interval, then  $\frac{f(S')}{1+\eta} \leq f(S) \leq (1 + \eta)f(S')$  and  $\frac{1-f(S')}{1+\eta} \leq 1 - f(S) \leq (1 + \eta)(1 - f(S'))$ .*

Throughout the execution, Algorithm **Calibrate** maintains a collection  $\mathcal{C}$  of contracts. The algorithm guarantees that no two contracts in  $\mathcal{C}$  are calibrated to the same interval, thus  $|\mathcal{C}| \leq 2r + 3$  at any given moment.

Every SP function  $\varphi$  is constructed inductively from two simpler SP functions by either a series composition or by a parallel composition. Therefore the function  $\varphi$  can be represented by a full binary tree  $\mathcal{T}$ , referred to as the *composition tree* of  $\varphi$ . The leaves of  $\mathcal{T}$  represents the identity functions of  $\varphi$ 's arguments. An internal node is said to be an  $\wedge$ -node (respectively, an  $\vee$ -node) if it represents a series (resp., parallel) composition of the functions represented by its children.

Consider the SP technology  $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$  and let  $\mathcal{T}$  be the tree that represents the Boolean function  $\varphi$ . Let  $x$  be some node in  $\mathcal{T}$  and consider the subtree  $\mathcal{T}_x$  of  $\mathcal{T}$  rooted at  $x$ . The subtree  $\mathcal{T}_x$  corresponds to some (SP) subtechnology  $t_x$  of  $t$ . Let  $N_x$  denote the set of agents in  $t_x$  (corresponding to the leaves of  $\mathcal{T}_x$ ) and let  $m_x$  denote the number of nodes in  $\mathcal{T}_x$  (as  $\mathcal{T}$  is a full binary tree, we have  $m_x = 2|N_x| - 1$ ). Given some contract  $S \subseteq N_x$ , we denote the effectiveness and payment of  $S$  under  $t_x$  by  $f_x(S)$  and  $p_x(S)$ , respectively.

Suppose that  $x$  is an internal node in  $\mathcal{T}$  with left child  $l$  and right child  $r$ . Let  $S = L \cup R$  be some contract in  $t_x$ , where  $L \subseteq N_l$  and  $R \subseteq N_r$ . Clearly, if  $x$  is an  $\wedge$ -node, then  $f_x(S) = f_l(L) \cdot f_r(R)$ ; and if  $x$  is an  $\vee$ -node, then  $f_x(S) = 1 - (1 - f_l(L))(1 - f_r(R))$ . It is simple to verify that if  $x$  is an  $\wedge$ -node, then  $p_x(S) = \frac{p_l(L)}{f_r(R)} + \frac{p_r(R)}{f_l(L)}$ ; and if  $x$  is an  $\vee$ -node, then  $p_x(S) = \frac{p_l(L)}{1 - f_r(R)} + \frac{p_r(R)}{1 - f_l(L)}$ .

Algorithm **Calibrate** traverses the composition tree  $\mathcal{T}$  in a postorder fashion. Consider some leaf  $x$  in  $\mathcal{T}$  that corresponds to agent  $i \in N$ . The algorithm calibrates the contracts  $\emptyset$  and  $\{i\}$  to a (fresh) scale according to their effectiveness under the technology  $t_x$ , that is,  $f_x(\emptyset) = \gamma_i$  and  $f_x(\{i\}) = \delta_i$ . If both  $\emptyset$  and  $\{i\}$  are calibrated to the same interval  $\mathcal{I}$ , then  $\{i\}$  is removed from the scale. The resulting contract(s) in the scale constitutes the collection  $\mathcal{C}_x$ .

Now, consider some internal node  $x$  in  $\mathcal{T}$  with left child  $l$  and right child  $r$  and suppose that the algorithm has already constructed the collections  $\mathcal{C}_l$  and  $\mathcal{C}_r$  for the technologies  $t_l$  and  $t_r$ , respectively. The collection  $\mathcal{C}_x$  for the technology  $t_x$  is constructed as follows. Let  $\mathcal{S} = \{L \cup R \mid L \in \mathcal{C}_l \text{ and } R \in \mathcal{C}_r\}$ . (Note that  $\mathcal{S}$  contains  $|\mathcal{C}_l| \cdot |\mathcal{C}_r| = O(r^2)$  contracts of the technology  $t_x$ .) The contracts in  $\mathcal{S}$  are calibrated to a (fresh) scale according to the effectiveness function  $f_x(\cdot)$ . Consequently, there may exist some interval in the new scale to which two (or more) contracts are calibrated (a conflict).

Let  $\mathcal{I}$  be an interval in the scale and suppose that  $S_1, \dots, S_k \in \mathcal{S}$  were all calibrated to  $\mathcal{I}$  ( $k > 1$ ), that is,  $f_x(S_i) \in \mathcal{I}$  for every  $1 \leq i \leq k$ . Assume without loss of generality that  $S_k$  admits a minimum payment under  $t_x$ , i.e.,  $p_x(S_k) \leq p_x(S_i)$  for every  $1 \leq i < k$ . The algorithm then resolves the conflict by removing the contracts  $S_1, \dots, S_{k-1}$  from the scale so that  $S_k$  remains the only contract calibrated to  $\mathcal{I}$ . In that case we say that the contracts  $S_1, \dots, S_{k-1}$  were *compensated* by the contract  $S_k$ . The contracts that remain in the scale constitutes the collection  $\mathcal{C}_x$ . Thus the new collection  $\mathcal{C}_x$  contains at most one contract for every interval and we may proceed with the next stage of the algorithm.

At the end of this postorder process, when Algorithm **Calibrate** reaches the root  $z$  of  $\mathcal{T}$ , it returns the collection  $\mathcal{C} = \mathcal{C}_z$ .

We turn to the analysis of Algorithm **Calibrate**. The running time of the algorithm is determined by the number of nodes in  $\mathcal{T}$  (which is  $2n - 1$ ) and by the size of the collection  $\mathcal{C}_x$  for every node  $x$  in the tree. The latter cannot exceed the number of intervals in the scale which is  $O\left(\frac{1}{\eta} \log \frac{1}{\Delta}\right)$ . In order to analyze the performance guarantee of the algorithm, we first define the following notion. Given two contracts  $S, S' \subseteq N$  and some real  $\alpha > 1$ , we say that  $S$  is an  $\alpha$ -estimation of  $S'$  under the technology  $t$  if the following three conditions hold: (1)  $\frac{f(S')}{\alpha} \leq f(S) \leq \alpha f(S')$ ; (2)  $\frac{1-f(S')}{\alpha} \leq 1 - f(S) \leq \alpha(1 - f(S'))$ ; and (3)  $p(S) \leq \alpha p(S')$ . We say that a collection  $\mathcal{S}$  of contracts is an  $\alpha$ -estimation of the technology  $t$  if for every contract  $S' \subseteq N$  there exists a contract  $S \in \mathcal{S}$  such that  $S$  is an  $\alpha$ -estimation of  $S'$  under  $t$ . The following Lemma is established by induction on the height of the composition tree  $\mathcal{T}$ .

**Lemma 1.** *The collection  $\mathcal{C}_x$  is a  $(1 + \eta)^{m_x}$ -estimation of the technology  $t_x$  for every node  $x$  in the composition tree  $\mathcal{T}$ .*

Lemma 1 implies that  $\mathcal{C}$  serves as a  $(1 + \eta)^{2n-1}$ -estimation of  $t$ . By the definition of  $\eta = \frac{\rho \ln 2}{2n-1}$ , we have  $(1 + \eta)^{2n-1} \leq e^{\rho \ln 2} = 2^\rho \leq 1 + \rho$ , which establishes the following corollary.

**Corollary 1.** *Given an SP technology  $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$  and a performance parameter  $0 < \rho \leq 1$ , it is guaranteed that Algorithm **Calibrate** generates a collection  $\mathcal{C} \subseteq 2^N$  that serves as a  $(1 + \rho)$ -estimation of  $t$  in time  $O\left(\frac{n^3 \log^2(1/\Delta)}{\rho^2}\right)$ .*

### 3.2 An FPTAS for OR Technologies

In this section we prove several properties of OR technologies which will be used to present an FPTAS. We first establish the sub-modularity of OR technologies. We say that a function  $h : 2^N \rightarrow \mathbb{R}$  is *strictly sub-modular* if  $h(S) + h(T) \geq h(S \cup T) + h(S \cap T)$  for every  $S, T \subseteq N$ , where equality holds (if and) only if  $S \subseteq T$  or  $T \subseteq S$ .

**Lemma 2.** *The effectiveness function of every OR technology is strictly sub-modular.*

Consider an arbitrary OR technology  $t = \langle N, \{\gamma_i\}_{i=1}^n, \{\delta_i\}_{i=1}^n, \{c_i\}_{i=1}^n, \varphi \rangle$ . Let  $T \subseteq N$  be some contract,  $|T| \geq 2$ , and consider the partition  $T = R_1 \cup R_2$ ,  $R_1 \cap R_2 = \emptyset$ , such that  $|R_1|, |R_2| \geq 1$ . A direct consequence of Lemma 2 is that  $f(R_1) + f(R_2) > f(T)$ . Another consequence is that  $p_j(R_i) < p_j(T)$  for every  $i = 1, 2$  and every agent  $j \in R_i$ , thus  $p(R_1) + p(R_2) < p(T)$ . These consequences of Lemma 2 are employed to establish the following key property.

**Lemma 3.** *Let  $v > 0$  be some payoff and let  $T$  be an optimal contract for  $v$  under the OR technology  $t$ . If  $v < (1 + \hat{\sigma})p(T)$  for some positive real  $\hat{\sigma} \leq 1/n$ , then there exists some agent  $j \in T$  such that  $f(\{j\}) > (1 - \hat{\sigma})f(T)$ .*

We are now ready to establish an FPTAS for the optimal contract problem on OR technologies. Let  $\epsilon > 0$  be the performance parameter of the FPTAS. (Recall that for every  $\epsilon > 0$ , the FPTAS returns a solution which is at most  $1 + \epsilon$  times worse than the optimal solution in time  $\text{poly}(|t|, 1/\epsilon)$ .) Subsequently, we assume that  $\epsilon \leq 1/n$  at the price of incurring an extra additive  $\text{poly}(|t|)$  term on the running time.

Fix  $\sigma = \epsilon$  and  $\hat{\sigma} = \frac{\epsilon}{1+\epsilon}$ , and let  $\mathcal{C}$  be the collection generated by Algorithm **Calibrate** when invoked on  $t$  with performance parameter  $\rho = \frac{\sigma\hat{\sigma}}{1+2\hat{\sigma}}$ . The FPTAS will consider the contracts in  $\mathcal{C} \cup \{\{j\} \mid j \in N\}$ , namely, the contracts in  $\mathcal{C}$  and all the singleton contracts. Consider an arbitrary payoff  $v > 0$  and let  $T \subseteq N$  be an optimal contract for  $v$ . In order to establish Theorem **2**, we have to prove that there exists a contract  $S \in \mathcal{C} \cup \{\{j\} \mid j \in N\}$  such that  $U_T(v)/U_S(v) \leq 1 + \epsilon$ .

Assume first that  $v < (1 + \hat{\sigma})p(T)$ . Since  $\hat{\sigma} < \sigma \leq 1/n$ , we may apply Lemma **3** and conclude that there exists some agent  $j \in N$  such that  $f(\{j\}) > (1 - \hat{\sigma})f(T)$ . By Lemma **2**, we have  $p(\{j\}) \leq p(T)$ , hence  $\frac{U_T(v)}{U_{\{j\}}(v)} = \frac{f(T)(v-p(T))}{f(\{j\})(v-p(\{j\}))} \leq \frac{f(T)}{f(\{j\})} < \frac{1}{1-\hat{\sigma}}$ . The assertion follows by the choice of  $\hat{\sigma} = \frac{\epsilon}{1+\epsilon}$ .

Now, assume that  $v \geq (1 + \hat{\sigma})p(T)$ . Let  $S$  be the contract in  $\mathcal{C}$  that serves as a  $(1 + \rho)$ -estimation of  $T$ . Since  $f(S) \geq f(T)/(1 + \rho)$  and  $p(S) \leq (1 + \rho)p(T)$ , we have  $\frac{U_T(v)}{U_S(v)} = \frac{f(T)(v-p(T))}{f(S)(v-p(S))} \leq (1 + \rho) \frac{v-p(T)}{v-(1+\rho)p(T)} \leq (1 + \rho) \frac{(1+\hat{\sigma})p(T)-p(T)}{(1+\hat{\sigma})p(T)-(1+\rho)p(T)} = \frac{(1+\rho)\hat{\sigma}}{\hat{\sigma}-\rho}$ . The requirement  $\frac{(1+\rho)\hat{\sigma}}{\hat{\sigma}-\rho} \leq 1 + \epsilon = 1 + \sigma$  is guaranteed by the choice of the performance parameter  $\rho = \frac{\sigma\hat{\sigma}}{1+2\hat{\sigma}}$  as  $\frac{(1+\rho)\hat{\sigma}}{\hat{\sigma}-\rho} \leq 1 + \sigma \iff \hat{\sigma} + \rho\hat{\sigma} \leq \hat{\sigma} + \sigma\hat{\sigma} - \rho - \rho\hat{\sigma} \iff \rho(1 + 2\hat{\sigma}) \leq \sigma\hat{\sigma}$ .

## 4 Conclusions

The hidden action problem lies at the heart of economic theory and has been recently studied from an algorithmic perspective. In this article we continue the study initiated by Babaioff et al **[1]** of the computational complexity of optimal team incentives under hidden actions. Our contribution focuses on OR technologies and on the more general family of *series-parallel* (SP) technologies. In particular, we establish the NP-hardness of the problem of computing an optimal contract in an OR technology (an open problem in **[1]**). We also show that there exist OR technologies with exponentially large orbits (disproving a conjecture of **[1]**).

On the positive side, we devise an FPTAS for OR technologies. For SP technologies, we establish a scheme that provides a  $(1 + \epsilon)$ -approximation for all but an  $\hat{\epsilon}$ -fraction of the relevant instances in time polynomial in the size of the technology and in the reciprocals of  $\epsilon$  and  $\hat{\epsilon}$ . The existence of an approximation scheme for SP technologies remains an open question. Put together, this article makes a significant step in understanding the combinatorial agency setting, which is an example for the important interaction between game theory, economic theory and computer science.

## Acknowledgments

We are indebted to Noam Nisan for his help on various parts of the paper.

## References

1. Babaioff, M., Feldman, M., Nisan, N.: Combinatorial Agency. In: ACM Conference on Electronic Commerce (2006)
2. Babaioff, M., Feldman, M., Nisan, N.: Mixed Strategies in Combinatorial Agency. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 353–364. Springer, Heidelberg (2006)
3. Emek, Y., Heitner, I.: Combinatorial agency: the observable-action model (2006) (manuscript)
4. Feigenbaum, J., Shenker, S.: Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. In: Proceedings of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (2002)
5. Feldman, M., Chuang, J., Stoica, I., Shenker, S.: Hidden-action in multi-hop routing. In: ACM Conference on Electronic Commerce, pp. 117–126 (2005)
6. Garey, M.R., Jhonson, D.S.: Computers and intractability: a guide to the theory of NP-completeness. W. H. Freeman, New York (1979)
7. Holmstrom, B.: Moral Hazard in Teams. *Bell Journal of Economics* 13, 324–340 (1982)
8. Legeros, P., Matthews, S.: Efficient and Nearly-Efficient Partnerships. *Review of Economic Studies* 60(3), 599–611 (1993)
9. Mookherjee, D.: Optimal Incentive Schemes with Many Agents. *Review of Economic Studies* 51(3), 433–446 (1984)
10. Nisan, N.: Algorithms for selfish agents. In: Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, pp. 1–15 (1999)
11. Nisan, N., Ronen, A.: Algorithmic Mechanism Design. In: Proceedings of the 31st Symposium on Theory of Computing (1999)
12. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
13. Papadimitriou, C.: Algorithms, Games, and the Internet. In: Proceedings of the 33rd Symposium on Theory of Computing, pp. 749–753 (2001)
14. Provan, J.S., Ball, M.O.: The Complexity of Counting Cuts and of Computing the Probability that a Graph is Connected. *SIAM Journal on Computing* 12(4), 777–788 (1983)
15. Rasmusen, E.: Moral Hazard in Risk-Averse Teams. *Rand Journal of Economics* 18, 428–435 (1987)
16. Strausz, R.: Moral Hazard in Sequential Teams. Departmental Working Paper. Free University of Berlin (1996)
17. Winter, E.: Incentives and Discrimination. *American Economic Review* 94, 764–773 (2004)



# Nash Equilibria for Voronoi Games on Transitive Graphs\*

Rainer Feldmann<sup>1</sup>, Marios Mavronicolas<sup>2</sup>, and Burkhard Monien<sup>1</sup>

<sup>1</sup> Faculty of Computer Science, Electrical Engineering and Mathematics,  
University of Paderborn, 33102 Paderborn, Germany

{obelix, bm}@upb.de

<sup>2</sup> Department of Computer Science, University of Cyprus, Nicosia CY-1678, Cyprus  
mavronic@cs.ucy.ac.cy

**Abstract.** In a *Voronoi game*, each of  $\kappa \geq 2$  players chooses a vertex in a graph  $G = \langle V(G), E(G) \rangle$ . The *utility* of a player measures her *Voronoi cell*: the set of vertices that are closest to her chosen vertex than to that of another player. In a *Nash equilibrium*, unilateral deviation of a player to another vertex is not profitable. We focus on various, symmetry-possessing classes of *transitive graphs*: the *vertex-transitive* and *generously vertex-transitive graphs*, and the more restricted class of *friendly graphs* we introduce; the latter encompasses as special cases the popular *d-dimensional bipartite torus*  $T_d = T_d(2p_1, \dots, 2p_d)$  with even *sides*  $2p_1, \dots, 2p_d$  and *dimension*  $d \geq 2$ , and a subclass of the *Johnson graphs*.

*Would transitivity enable bypassing the explicit enumeration of Voronoi cells?* To argue in favor, we resort to a technique using *automorphisms*, which suffices alone for generously vertex-transitive graphs with  $\kappa = 2$ .

To go beyond the case  $\kappa = 2$ , we show the *Two-Guards Theorem for Friendly Graphs*: whenever two of the three players are located at an *antipodal* pair of vertices in a friendly graph  $G$ , the third player receives a utility of  $\frac{|V(G)|}{4} + \frac{|\Omega|}{12}$ , where  $\Omega$  is the intersection of the three Voronoi cells. If the friendly graph  $G$  is *bipartite* and has odd *diameter*, the utility of the third player is fixed to  $\frac{|V(G)|}{4}$ ; this allows discarding the third player when establishing that such a triple of locations is a Nash equilibrium. Combined with appropriate automorphisms *and without explicit enumeration*, the *Two-Guards Theorem* implies the existence of a Nash equilibrium for *any* friendly graph  $G$  with  $\kappa = 4$ , with colocation of players allowed; if colocation is forbidden, existence still holds under the additional assumption that  $G$  is bipartite and has odd diameter.

For the case  $\kappa = 3$ , we have been unable to bypass the *explicit enumeration of Voronoi cells*. Combined with appropriate automorphisms and *explicit enumeration*, the *Two-Guards Theorem* implies the existence of a Nash equilibrium for (i) the 2-dimensional torus  $T_2$  with odd diameter  $\sum_{j \in [2]} p_j$  and  $\kappa = 3$ , and (ii) the hypercube  $H_d$  with odd  $d$  and  $\kappa = 3$ .

In conclusion, *transitivity does not seem sufficient for bypassing explicit enumeration*: far-reaching challenges in combinatorial enumeration are in sight, even for values of  $\kappa$  as small as 3.

---

\* Supported by the IST Program of the European Union under contract number 15964 (AEOLUS).



## 1 Introduction

Recently, there has been a vast amount of research on *non-cooperative games on networks*, inspired from diverse application domains from computer and communication networks, such as resource allocation, routing, scheduling and facility location. In this work, we shall extend the study of the *pure Nash equilibria* [13,14] associated with a particular game inspired from facility location and called the *Voronoi game*; it was introduced in [3] and further studied in [12].

The Voronoi game [3,12] is reminiscent of the classical *Hotelling games* [9], with a number of *vendors* in some continuous *metric space*— see, e.g., [10]. Each vendor comes with goods for sale; simultaneously with others, she must choose a *location* for her *facility*; her objective is to maximize the region of points that are closest to her than to any other vendor, called her *Voronoi cell*. In a *Nash equilibrium* [13,14], no vendor can increase her profit by switching to a different point. The Voronoi game  $\langle G, [\kappa] \rangle$  is the discrete analog where an undirected graph  $G = \langle V(G), E(G) \rangle$  is used instead of a metric space. There are  $\kappa$  *players*, each choosing a vertex; a player's *utility* measures her *Voronoi cell*: the set of vertices closest to her than to another player. Closest to more than one player, a *boundary vertex* contributes uniformly to the utilities of its closest players; so, it needs to be taken into account. (The (zero-measure) boundary points in a Hotelling game and can be discarded.)

In a *Nash equilibrium* [13,14], no player can unilaterally increase her utility by switching to another vertex. So, *existence of Nash equilibria is contingent upon enumeration properties of Voronoi cells*. Hence, *explicit enumeration of Voronoi cells manifests as a combinatorial bottleneck to identifying Nash equilibria, even if the locations of the players are given. Is this bottleneck inherent?*

It is  $\mathcal{NP}$ -complete to decide the existence of a Nash equilibrium for an *arbitrary* Voronoi game  $\langle G, [\kappa] \rangle$  [3, Section 4]. (For a *constant*  $\kappa$ , the decision problem is in  $\mathcal{P}$  through exhaustive search.) A simple counterexample was presented of a (*not* vertex-transitive) graph with no Nash equilibrium for  $\kappa = 2$  [3, Section 4]. Subsequent work [12] provided combinatorial characterizations of Nash equilibria for *rings*, determining the ring size allowing for a Nash equilibrium.

Zhao *et al.* [15] proposed recently the *isolation game* on an arbitrary *metric space* as a generalization of the Voronoi game. Several results on the Nash equilibria associated with isolation games were shown in [15], and later in [2].

Here is our motivation in two sentences: *How easily would transitivity enable by-passing the explicit enumeration of Voronoi cells? What are the broadest classes of (transitive) graphs for which transitivity would so succeed for a given number of players?* Hereby, we embark on the broad class of *vertex-transitive* graphs, which "look" the same from each vertex. However, we shall focus on restricted classes of vertex-transitive graphs. To start with, in a *generously vertex-transitive graph*, an arbitrary pair of vertices can be swapped (cf. [5, Section 12.1] or [8, Section 4.3]). Although there are examples of vertex-transitive graphs that are *not* generously vertex-transitive (e.g., the *cube-connected-cycles* [11, Section 3.2.1]), the class of generously vertex-transitive graphs includes a sufficiently rich subclass: a *friendly graph* is a generously vertex-transitive graph where, in addition,

every vertex is on some shortest path between an *antipodal* pair of vertices. Our prime example of a friendly graph is the *d-dimensional, bipartite torus*  $T_d$ , which encompasses the *d-dimensional hypercube*  $H_d$  as a special case (Lemma 1). Yet, we identify a special subclass of the *Johnson graphs* [6, Section 1.6] as another example of a friendly graph (Lemma 2).

In this endeavor, we seek to exploit the algebraic and combinatorial structure of friendly graphs in devising techniques to *compare the cardinalities of Voronoi cells without explicitly enumerating them*; such techniques will allow *establishing the existence of Nash equilibria by bypassing explicit enumeration*. This idea is naturally inspired from the technique of *bijective proofs* in Combinatorics (see, for example, [4, Section 2] and references therein), which shows that two (finite) sets have the same cardinality by providing a bijection between them. In particular, we shall resort to *automorphisms* of friendly graphs.

*Such resorting suffices to settle the case of generously vertex-transitive graphs with  $\kappa = 2$* . Specifically, we prove that *every* location for the two players yields a Nash equilibrium for a generously vertex-transitive graph (Proposition 3). Unfortunately, this simple idea may not extend beyond generously vertex-transitive graphs in a general way: we prove that some particular vertex-transitive but not generously vertex-transitive graph, namely the *cube-connected-cycles*, has no Nash equilibrium for  $\kappa = 2$ . This fact follows immediately from a general *necessary condition* we establish for *any* vertex-transitive graph to admit a Nash equilibrium: There is a pair of vertices to locate the two players so that they receive different utilities (Proposition 4). This counterexample extends the earlier one of Dürr and Thang [3, Section 4].

*We have been unable to go beyond the case  $\kappa = 2$  without assuming some additional structure on the graph G*. Towards this end, we establish the (perhaps surprising) *Two-Guards Theorem for Friendly Graphs* concerning the case  $\kappa = 3$  (Theorem 6): If two of the players are located at an antipodal pair of vertices in a friendly graph  $G$ , the third player receives a utility of  $\frac{|V(G)|}{4} + \frac{|\Omega|}{12}$ , where  $\Omega$  denotes the intersection of the three Voronoi cells. For a *bipartite* friendly graph with *odd* diameter, the *Two-Guards Theorem for Friendly Graphs* has an interesting extension: independently of her location, the third player receives a *fixed* utility of  $\frac{|V(G)|}{4}$  (Corollary 7). So, a corresponding paradigm emerges for establishing the existence of a Nash equilibrium: *locate two of players at an antipodal pair and prove that none of them can unilaterally improve*.

Through this paradigm, we have been able to *bypass explicit enumeration for the case  $\kappa = 4$* . Assuming that *colocation of players is allowed*, we establish, through a simple proof, the existence of a Nash equilibrium for (i) an *arbitrary* friendly graph with  $\kappa = 4$  (Theorem 8). However, *forbidding colocation* has required (still for  $\kappa = 4$ ) the additional assumption that (ii) the friendly graph is bipartite and has odd diameter (Theorem 9); the proof is more challenging and uses suitable automorphisms. The key idea for the proofs of both results has been that when one of the four players deviates, *there still remain two players located at an antipodal pair of vertices*; in turn, this allows applying the *Two-Guards Theorem for Friendly Graphs* and its extension.

For the case  $\kappa = 3$ , we have developed techniques for the explicit enumeration of Voronoi cells, which make the most technically challenging part of this work. These enumeration techniques have enabled settling the existence of a Nash equilibrium in the following special cases: (iii) The 2-dimensional torus  $T_2$  with odd diameter  $\sum_{j \in [2]} p_j$  (Theorem 10), and (iv) the hypercube  $H_d$  with odd  $d$  (Theorem 7); for the proof of (iv), we have derived explicit combinatorial formulas (as nested sums of binomial coefficients) for the utilities of three players located arbitrarily in the hypercube  $H_d$ . Although the hypercube  $H_d$  is a special case of the torus  $T_d$ , these two existence results are *incomparable*: (iv) applies to the hypercube  $H_d$  with odd  $d$ , while (iii) applies to the torus  $T_d$  with  $d = 2$ .

To complement the existence results for  $\kappa = 3$  in (iii) and (iv), we have carried out an extensive set of experiments. Their results provide strong evidence that there is *no* Nash equilibrium for the cases of (v) the 2-dimensional torus  $T_2$  with *even* diameter  $\sum_{j \in [2]} p_j$ , and (vi) the hypercube  $H_d$  with *even*  $d$ ; so, they suggest that the assumptions made for (iii) and (iv) are *essential*.

## 2 Vertex-Transitive Graphs

We shall consider a simple, connected and undirected **graph**  $G = \langle V(G), E(G) \rangle$ . A **path** in  $G$  is a sequence  $v_0, v_1, \dots, v_\ell$  of vertices such that for each index  $i \in [\ell]$ ,  $\{v_{i-1}, v_i\} \in E(G)$ ; the **length** of the path is the number  $\ell$  of its edges. A **cycle** is a path  $v_0, v_1, \dots, v_\ell$  with  $v_\ell = v_0$ . For a pair of vertices  $u, v \in V(G)$ , the **distance** between  $u$  and  $v$ , denoted as  $\text{dist}_G(u, v)$  (or just  $\text{dist}(u, v)$ ) is the length of the shortest path between  $u$  and  $v$ . The **diameter** of  $G$  is given by  $\text{diam}(G) = \max_{u, v \in V(G)} \text{dist}(u, v)$ . Say that the pair of vertices  $u, v \in V(G)$  is **antipodal** if  $\text{dist}(u, v) = \text{diam}(G)$ ; so,  $u$  (resp.,  $v$ ) is an **antipode** to  $v$  (resp.,  $u$ ). For a set of vertices  $V' \subseteq V(G)$ , denote  $\Omega(V') = \{u \in V(G) \mid \text{the distance } \text{dist}(u, v) \text{ is the same for all vertices } v \in V'(G)\}$ . Note that for a *bi-partite* graph  $G$ , if the set  $V'$  contains two vertices at odd distance from each other, then  $\Omega(V') = \emptyset$ . Two graphs  $G = (V, E)$  and  $G' = (V', E')$  are **isomorphic** if there is a bijection  $\varphi : V \rightarrow V'$  such that for each pair of vertices  $u, v \in V$ ,  $\{u, v\} \in E$  if and only if  $\{\varphi(u), \varphi(v)\} \in E'$ ;  $\varphi$  is an **isomorphism** from  $G$  to  $G'$ . An **automorphism** of  $G$  is an isomorphism from  $G$  to itself. Note that for an automorphism  $\varphi$ , for each pair of vertices  $u, v \in V(G)$ ,  $\text{dist}_G(u, v) = \text{dist}_{G'}(\varphi(u), \varphi(v))$ .

We continue with some notions of *transitivity*; for more details, we refer the reader to [1, Chapters 15 & 16], [5, Section 12.1], [6, Chapter 3], [7, Section 6.1], or [8, Section 4.3]. The graph  $G$  is **vertex-transitive** if for each pair of vertices  $u, v \in V$ , there is an automorphism  $\phi$  of  $G$  such that  $\phi(u) = v$ ; roughly speaking, a vertex-transitive graph "looks" the same from each vertex. The graph  $G$  is **generously vertex-transitive** if for each pair of vertices  $u, v \in V$ , there is an automorphism  $\phi$  of  $G$  such that  $\phi(u) = v$  and  $\phi(v) = u$ ; so, each pair of vertices can be swapped. To the best of our knowledge, the following definition is new. A graph  $G$  is **friendly** if the following two conditions hold:

- (F.1)  $G$  is generously vertex-transitive.
- (F.2) For any pair of antipodal vertices  $\alpha, \beta \in V(G)$ , and for any arbitrary vertex  $\gamma \in V(G)$ ,  $\gamma$  is on a shortest path between  $\alpha$  and  $\beta$ .

So, each vertex in a friendly graph has a *unique* antipode. (In fact, vertex-transitivity (rather than Condition (F.1)) and Condition (F.2) suffice for this.) Fix an arbitrary integer  $d \geq 2$ , called the *dimension*, and a sequence of integers  $p_1, \dots, p_d \geq 1$ , called the *sides*. The  *$d$ -dimensional bipartite torus*  $T_d = T_d[2p_1, \dots, 2p_d]$  is the graph  $T_d$  with  $V(T_d) = \{0, 1, \dots, 2p_1 - 1\} \times \dots \times \{0, 1, \dots, 2p_d - 1\}$  and  $E(T_d)$  consists of those edges  $\{\alpha, \beta\}$  such that  $\alpha$  and  $\beta$  differ in exactly one component  $j \in [d]$  and  $|\alpha_j - \beta_j| \equiv 1 \pmod{2p_j}$ ; the *dimension* of the edge  $\{\alpha, \beta\} \in E(T_d)$  is the dimension  $j \in [d]$  in which  $\alpha$  and  $\beta$  differ. Note that the bipartite graph  $T_d$  is the cartesian product of  $d$  even cycles, where the cycle in dimension  $j \in [d]$  has length  $2p_j$ . We shall often abuse notation to call each integer  $j \in [d]$  a *dimension* of the graph  $T_d$ .

Fix a pair of vertices  $\alpha, \beta \in V(T_d)$ . Then,  $\text{dist}(\alpha, \beta) = \sum_{j \in [d]} \text{dist}_j(\alpha_j, \beta_j)$ , where for each dimension  $j \in [d]$ ,  $\text{dist}_j(\alpha_j, \beta_j)$  is the distance between the components  $\alpha_j$  and  $\beta_j$  on the cycle of length  $2p_j$  in dimension  $j$ . Note that a pair of vertices  $\alpha = \langle \alpha_1, \dots, \alpha_d \rangle$  and  $\bar{\alpha} = \langle (\alpha_1 + p_1) \bmod 2p_1, \dots, (\alpha_d + p_d) \bmod 2p_d \rangle$  is antipodal in the torus  $T_d$ . Clearly,  $\text{diam}(T_d) = \sum_{j \in [d]} p_j$ . Since an even cycle fulfils Condition (F.2) and  $T_d$  is the cartesian product of even cycles, so does  $T_d$ . Induced by an arbitrary pair of vertices  $\alpha, \beta \in V(T_d)$  is the automorphism  $\Psi : V(T_d) \rightarrow V(T_d)$  where: for each vertex  $\chi$ ,  $\Psi(\chi) = \langle \psi_1(\chi_1), \dots, \psi_d(\chi_d) \rangle$ , where for each dimension  $j \in [d]$ ,  $\psi_j(\chi_j) = (\alpha_j + \beta_j - \chi_j) \bmod (2p_j)$ ; clearly,  $\Psi(\alpha) = \beta$  and  $\Psi(\beta) = \alpha$ , and (F.1) follows. Hence, we obtain:

**Lemma 1.** *The  $d$ -dimensional bipartite torus  $T_d = T_d[2p_1, \dots, 2p_d]$  is friendly.*

Note that the *non-bipartite* torus is *not* friendly. As a special case, the  *$d$ -dimensional hypercube*  $H_d$  is the  $d$ -dimensional torus  $T_d[2, \dots, 2]$ ; so, each vertex is a binary vector  $\alpha \in \{0, 1\}^d$ , and the distance between two vertices is the usual *Hamming distance* between the two binary vectors. So, the diameter of  $H_d$  equals the dimension  $d$ . The  *$d$ -dimensional cube-connected-cycles*  $CCC_d$  is constructed from the  $d$ -dimensional hypercube  $H_d$  (cf. [11, Section 3.2.1]). Note that  $CCC_d$  fails Condition (F.2) in the definition of friendly graphs. We will later conclude that  $CCC_d$  fails also Condition (F.1). Since it is vertex-transitive, generously vertex-transitive graphs are a *strict* subset of vertex-transitive graphs.

Let  $\nu$ ,  $k$  and  $\ell$  be fixed positive integers with  $\nu \geq k \geq \ell$ ; let  $\mathcal{U}$  be a fixed *ground set* of size  $\nu$ . Define the graph  $J(\nu, k, \ell)$  as follows (cf. [6, Section 1.6]). The vertices of  $J(\nu, k, \ell)$  are the subsets of  $\mathcal{U}$  with size  $k$ ; two subsets are adjacent if their intersection has size  $\ell$ . If  $\varphi$  is a permutation of  $\mathcal{U}$  and  $S \subseteq \mathcal{U}$ , then define  $\varphi(S) = \{\varphi(s) \mid s \in S\}$ . Clearly, each permutation of  $\mathcal{U}$  determines a permutation of the subsets of  $\mathcal{U}$ , and in particular a permutation of the subsets with size  $k$ . If  $S, T \subseteq \mathcal{U}$ , then  $|S \cap T| = |\varphi(S) \cap \varphi(T)|$ . So,  $\varphi$  is an automorphism of  $J(\nu, k, \ell)$ . For  $\nu \geq 2k$ , the graph  $J(\nu, k, k - 1)$  is known as a *Johnson graph*. We prove:

**Lemma 2.** *The graph  $J(\nu, k, \ell)$  is generously vertex-transitive for all  $\nu \geq k \geq \ell$ . It is friendly if  $\nu = 2k$  and  $\ell = k - 1$ .*

### 3 Voronoi Games

Fix any integer  $\kappa \geq 2$ ; denote  $[\kappa] = \{1, \dots, \kappa\}$ . The *Voronoi game*  $\langle G, [\kappa] \rangle$  is the strategic game  $\langle [\kappa], \{S_i\}_{i \in [\kappa]}, \{U_i\}_{i \in [\kappa]} \rangle$ , where for each *player*  $i \in [\kappa]$ , (i)  $S_i = V(G)$  and (ii) for each *profile*  $\mathbf{s} \in S_1 \times \dots \times S_\kappa$ , the *utility* of player  $i$  in the profile  $\mathbf{s}$  is given by  $U_i(\mathbf{s}) = \sum_{v \in \text{Vor}_i(\mathbf{s})} \frac{1}{\mu_v(\mathbf{s})}$ , where the *Voronoi cell* of player  $i \in [\kappa]$  in the profile  $\mathbf{s}$  is the set

$$\text{Vor}_i(\mathbf{s}) = \{v \in V(G) \mid \text{dist}(s_i, v) \leq \text{dist}(s_{i'}, v) \text{ for each player } i' \in [\kappa]\},$$

and the *multiplicity* of vertex  $v \in V(G)$  in the profile  $\mathbf{s}$  is the integer  $\mu_v(\mathbf{s}) = |\{i' \in [\kappa] \mid v \in \text{Vor}_{i'}(\mathbf{s})\}|$ . Clearly, the Voronoi game  $\langle G, [\kappa] \rangle$  is *constant-sum*.

For a profile  $\mathbf{s}$  and a player  $i \in [\kappa]$ ,  $\mathbf{s}_{-i} \oplus v$  denotes the profile obtained by replacing vertex  $s_i$  in  $\mathbf{s}$  with vertex  $v$ . Say that  $\mathbf{s}$  is a *Nash equilibrium* [13,14] if for each player  $i \in [\kappa]$ , for each vertex  $v \in V(G)$ ,  $U_i(\mathbf{s}) \geq U_i(\mathbf{s}_{-i} \oplus v)$ .

The *support* of the profile  $\mathbf{s}$  is the set  $\text{support}(\mathbf{s}) = \{s_i \mid i \in [\kappa]\}$ , the set of vertices chosen by the players. Given a profile  $\mathbf{s}$ , an automorphism  $\phi$  of  $G$  maps each strategy  $s_i$  with  $i \in [\kappa]$  to the strategy  $\phi(s_i)$ ; so,  $\phi$  induces an *image profile*  $\phi(\mathbf{s}) = \langle \phi(s_1), \dots, \phi(s_\kappa) \rangle$ . Say that profiles  $\mathbf{s}$  and  $\mathbf{t}$  are *equivalent* if there is an automorphism  $\phi$  of  $G$  such that  $\mathbf{t} = \phi(\mathbf{s})$ . We observe that for a pair of equivalent profiles  $\mathbf{s}$  and  $\mathbf{t}$ , and for each player  $i \in [\kappa]$ ,  $U_i(\mathbf{s}) = U_i(\mathbf{t})$ .

Given a profile  $\mathbf{s}$ , an automorphism  $\phi$  of  $G$  induces an *image support*  $\phi(\text{support}(\mathbf{s})) = \text{support}(\phi(\mathbf{s}))$ . A pair of players  $i, i' \in [\kappa]$  is *symmetric* for the profile  $\mathbf{s}$  if there is an automorphism  $\phi$  of  $G$  such that (i)  $\phi(\text{support}(\mathbf{s})) = \text{support}(\mathbf{s})$ , and (ii)  $\phi(s_i) = s_{i'}$ . Say that  $\mathbf{s}$  is *colocational* if there is a pair of distinct players  $i, i' \in [3]$  such that  $s_i = s_{i'}$ ; say that  $\mathbf{s}$  is *balanced* if for each pair of vertices  $u, v \in V(G)$ ,  $|\{i \in [\kappa] \mid s_i = u\}| = |\{i \in [\kappa] \mid s_i = v\}|$ . (Note that a non-colocational profile is balanced.) We observe that for a symmetric pair of players  $i, i' \in [\kappa]$  for the balanced profile  $\mathbf{s}$ ,  $U_i(\mathbf{s}) = U_{i'}(\mathbf{s})$ . The profile  $\mathbf{s}$  is *symmetric* if each pair of players  $i, i' \in [\kappa]$  is symmetric for  $\mathbf{s}$ ; then, clearly, for any pair of players  $i, i' \in [\kappa]$ ,  $U_i(\mathbf{s}) = U_{i'}(\mathbf{s})$ . The profile  $\mathbf{s}$  is *antipodal* if its support includes an antipodal pair of vertices.

### 4 Two Players

For the case  $\kappa = 2$ , we show:

**Proposition 3.** *Assume that  $G$  is generously vertex-transitive and  $\kappa = 2$ , and fix an arbitrary profile  $\mathbf{s}$ . Then,  $\mathbf{s}$  is a Nash equilibrium with  $U_1(\mathbf{s}) = U_2(\mathbf{s}) = \frac{|V|}{2}$ .*

*Proof.* Since  $G$  is generously vertex-transitive, it follows that  $\mathbf{s}$  is symmetric. Hence,  $U_1(\mathbf{s}) = U_2(\mathbf{s}) = \frac{|V|}{2}$ . Fix now any player  $i \in [2]$  and a vertex  $u \in V$ . Since  $G$  is generously vertex-transitive, it follows that  $\mathbf{s}_{-i} \oplus u$  is symmetric.

Hence,  $U_i(\mathbf{s}_{-i} \oplus u) = U_{[2] \setminus \{i\}}(\mathbf{s}_{-i} \oplus u) = \frac{|V|}{2}$ . So,  $U_i(\mathbf{s}_{-i} \oplus u) = U_1(\mathbf{s})$ . Since  $i$  was chosen arbitrarily, it follows that  $\mathbf{s}$  is a Nash equilibrium.  $\square$

Compare Proposition 3 to a corresponding result for Hotelling games on a (finite) line segment with two players: there is only one Nash equilibrium where both players are located in the middle of the line segment and receive the same utility 9. This result confirms to the *Principle of Minimum Differentiation* 9 (for Hotelling games): in a Nash equilibrium, players must be indifferent. Since all vertices are indifferent in a vertex-transitive graph, Lemma 3 confirms to the analog of the principle for Voronoi games. We next show:

**Proposition 4.** *Assume that  $G$  is vertex-transitive and  $\kappa = 2$ . Assume that there are vertices  $\alpha$  and  $\beta$  such that  $U_1(\langle \alpha, \beta \rangle) \neq U_2(\langle \alpha, \beta \rangle)$ . Then, the Voronoi game  $\langle G, [2] \rangle$  has no Nash equilibrium.*

*Proof.* Assume, without loss of generality, that  $U_1(\langle \alpha, \beta \rangle) < U_2(\langle \alpha, \beta \rangle)$ . Consider an arbitrary profile  $\langle \gamma, \delta \rangle$ ; we shall prove that  $\langle \gamma, \delta \rangle$  is *not* a Nash equilibrium.

1. Assume first that  $U_1(\langle \gamma, \delta \rangle) \neq U_2(\langle \gamma, \delta \rangle)$ . Without loss of generality, take that  $U_1(\langle \gamma, \delta \rangle) > U_2(\langle \gamma, \delta \rangle)$ . So,  $U_2(\langle \gamma, \delta \rangle) < \frac{V(G)}{2}$ . But,  $U_2(\langle \gamma, \gamma \rangle) = \frac{V(G)}{2}$ , and player 2 improves by switching to  $\gamma$ .
2. Assume now that  $U_1(\langle \gamma, \delta \rangle) = U_2(\langle \gamma, \delta \rangle)$ ; so,  $U_2(\langle \gamma, \delta \rangle) = \frac{V(G)}{2}$ . Since  $G$  is vertex-transitive, there is an automorphism  $\psi$  of  $G$  with  $\psi(\alpha) = \gamma$ . Then,  $U_2(\langle \gamma, \psi(\beta) \rangle) = U_2(\langle \psi(\alpha), \psi(\beta) \rangle) = U_2(\langle \alpha, \beta \rangle) > U_1(\langle \alpha, \beta \rangle) = U_1(\langle \psi(\alpha), \psi(\beta) \rangle) = U_1(\langle \gamma, \psi(\beta) \rangle)$ . So,  $U_2(\langle \gamma, \psi(\beta) \rangle) > \frac{V(G)}{2}$ , and player 2 improves by switching to  $\psi(\beta)$ .

Hence, the profile  $\langle \gamma, \delta \rangle$  is *not* a Nash equilibrium, as needed.  $\square$

We now use Proposition 4 to show that the Voronoi game  $\langle CCC_3, [2] \rangle$  has no Nash equilibrium. By Proposition 3, the cube-connected cycles  $CCC_d$  is *not* generously vertex-transitive (in general). So, an impossibility result in Algebraic Graph Theory is concluded from an impossibility result about Nash equilibria.

## 5 Two-Guards Theorems

For a profile  $\langle \alpha, \beta, \gamma \rangle$ . For each index  $\ell \in \{0, 1, 2\}$ , define the sets

$$\begin{aligned} \mathcal{A}_\ell(\langle \alpha, \beta, \gamma \rangle) &= \{ \delta \in V(G) \mid \text{dist}_G(\delta, \gamma) \sim_\ell \text{dist}_G(\delta, \alpha) \}, \\ \mathcal{B}_\ell(\langle \alpha, \beta, \gamma \rangle) &= \{ \delta \in V(G) \mid \text{dist}_G(\delta, \gamma) \sim_\ell \text{dist}_G(\delta, \beta) \}, \end{aligned}$$

where  $\sim_0$  is  $<$ ,  $\sim_1$  is  $=$ , and  $\sim_2$  is  $>$ . Clearly,  $\mathcal{A}_0, \mathcal{A}_1$  and  $\mathcal{A}_2$  (resp.  $\mathcal{B}_0, \mathcal{B}_1$  and  $\mathcal{B}_2$ ) partition  $V(G)$ . So,  $\mathcal{A}_0$  (resp.,  $\mathcal{B}_0$ ) contains all vertices that are closer to  $\gamma$  than to  $\alpha$  (resp., than to  $\beta$ );  $\mathcal{A}_1$  (resp.,  $\mathcal{B}_1$ ) contains all vertices that are equally close to each of  $\alpha$  and  $\gamma$  (resp., to each of  $\beta$  and  $\gamma$ );  $\mathcal{A}_2$  (resp.,  $\mathcal{B}_2$ ) contains all

vertices that are closer to  $\alpha$  (resp., to  $\beta$ ) than to  $\gamma$ . For each index  $\ell \in \{0, 1, 2\}$ , we shall use the shorter notations  $\mathcal{A}_\ell$  and  $\mathcal{B}_\ell$  for  $\mathcal{A}_\ell(\langle\alpha, \beta, \gamma\rangle)$  and  $\mathcal{B}_\ell(\langle\alpha, \beta, \gamma\rangle)$ , respectively, when the profile  $\langle\alpha, \beta, \gamma\rangle$  is clear from context. The sets  $\mathcal{A}_\ell$  and  $\mathcal{B}_\ell$ , with  $\ell \in \{0, 1, 2\}$  determine the utility of player 3 in the profile  $\langle\alpha, \beta, \gamma\rangle$  as

$$U_3(\langle\alpha, \beta, \gamma\rangle) = |\mathcal{A}_0 \cap \mathcal{B}_0| + \frac{1}{2} |\mathcal{A}_0 \cap \mathcal{B}_1| + \frac{1}{2} |\mathcal{A}_1 \cap \mathcal{B}_0| + \frac{1}{3} |\mathcal{A}_1 \cap \mathcal{B}_1|.$$

We first prove:

**Lemma 5.** *For any antipodal pair of vertices  $\alpha$  and  $\beta$ , and for any arbitrary vertex  $\gamma$  in a friendly graph  $G$ , consider an automorphism  $\Phi$  of  $G$  such that  $\Phi(\beta) = \gamma$  and  $\Phi(\gamma) = \beta$ . Then, for each vertex  $\chi \in V(G)$ :*

- (C.1) *For each index  $\ell \in \{0, 1, 2\}$ ,  $\chi \in \mathcal{A}_\ell$  if and only if  $\Phi(\chi) \in \mathcal{A}_\ell$ .*
- (C.2)  *$\chi \in \mathcal{B}_0$  if and only if  $\Phi(\chi) \in \mathcal{B}_2$  (and  $\chi \in \mathcal{B}_2$  if and only if  $\Phi(\chi) \in \mathcal{B}_0$ ).*
- (C.3)  *$\chi \in \mathcal{B}_1$  if and only if  $\Phi(\chi) \in \mathcal{B}_1$*

The fact that  $\Phi$  is an automorphism suffices for Conditions (C.2) and (C.3); the assumptions that (i) the pair  $\alpha, \beta$  is antipodal, and (ii)  $G$  is friendly are only needed for Condition (C.1). We now show:

**Theorem 6.** *Fix an antipodal pair of vertices  $\alpha$  and  $\beta$ , and an arbitrary vertex  $\gamma$  in a friendly graph  $G$ . Then,  $U_3(\langle\alpha, \beta, \gamma\rangle) = \frac{1}{4} |V(G)| + \frac{1}{12} |\Omega(\{\alpha, \beta, \gamma\})|$ .*

*Proof.* By Lemma 5 (Conditions (C.1) and (C.2)), it follows that for each index  $\ell \in \{0, 1\}$ , for each vertex  $\chi \in V(G)$ ,  $\chi \in \mathcal{A}_\ell \cap \mathcal{B}_0$  if and only if  $\Phi(\chi) \in \mathcal{A}_\ell \cap \mathcal{B}_2$ . Since the function  $\Phi$  is a bijection, the restriction  $\Phi : \mathcal{A}_\ell \cap \mathcal{B}_0 \rightarrow \mathcal{A}_\ell \cap \mathcal{B}_2$  is a bijection. Hence,  $|\mathcal{A}_\ell \cap \mathcal{B}_0| = |\mathcal{A}_\ell \cap \mathcal{B}_2|$ . It follows that for each index  $\ell \in \{0, 1\}$ ,  $|\mathcal{A}_\ell| = |\mathcal{A}_\ell \cap \mathcal{B}_0| + |\mathcal{A}_\ell \cap \mathcal{B}_1| + |\mathcal{A}_\ell \cap \mathcal{B}_2| = 2|\mathcal{A}_\ell \cap \mathcal{B}_0| + |\mathcal{A}_\ell \cap \mathcal{B}_1|$ . Hence,

$$U_3(\langle\alpha, \beta, \gamma\rangle) = \frac{1}{2} \left( |\mathcal{A}_0| + \frac{1}{2} |\mathcal{A}_1| \right) + \frac{1}{12} |\mathcal{A}_1 \cap \mathcal{B}_1|.$$

Consider the Voronoi game  $\langle G, [2] \rangle$ , with players 1 and 3. Then,  $U_3(\langle\alpha, \gamma\rangle) = |\mathcal{A}_0| + \frac{1}{2} |\mathcal{A}_1|$ . By Lemma 3,  $U_3(\langle\alpha, \gamma\rangle) = \frac{1}{2} |V(G)|$ . Hence,  $|\mathcal{A}_0| + \frac{1}{2} |\mathcal{A}_1| = \frac{1}{2} |V(G)|$ , so that  $U_3(\langle\alpha, \beta, \gamma\rangle) = \frac{1}{4} |V(G)| + \frac{1}{12} |\Omega(\{\alpha, \beta, \gamma\})|$ , as needed.  $\square$

An immediate implication of Theorem 6 for a bipartite friendly graph  $G$  with odd diameter will now follow. Since  $\text{dist}_G(\alpha, \beta)$  is odd for an arbitrary antipodal pair  $\alpha$  and  $\beta$ ,  $\Omega(\{\alpha, \beta, \gamma\}) = \emptyset$  for an arbitrary vertex  $\gamma$ ; hence, we obtain:

**Corollary 7.** *Fix an antipodal pair of vertices  $\alpha$  and  $\beta$ , and an arbitrary vertex  $\gamma$  in a bipartite friendly graph  $G$  with odd diameter. Then,  $U_3(\langle\alpha, \beta, \gamma\rangle) = \frac{1}{4} |V(G)|$ .*



## 6 Four Players

**With Colocation.** We show:

**Theorem 8.** *Consider a friendly graph  $G$ . Then, the Voronoi game  $\langle G, [4] \rangle$  has a Nash equilibrium. Specifically, for any arbitrary antipodal pair  $\alpha, \beta$ , the profile  $\mathbf{s} = \langle \alpha, \alpha, \beta, \beta \rangle$  is a Nash equilibrium.*

*Proof.* Consider a pair of players  $i \in \{1, 2\}$  and  $i' \in \{3, 4\}$ ; so,  $s_i = \alpha$  and  $s_{i'} = \beta$ . Since  $G$  is doubly vertex-transitive, there is an automorphism  $\phi$  of  $G$  such that  $\phi(s_i) = s_{i'}$  and  $\phi(s_{i'}) = s_i$ . Note that  $\phi(\text{support}(\mathbf{s})) = \text{support}(\mathbf{s})$ . Hence, the pair of players  $i, i'$  is symmetric for  $\mathbf{s}$ . Since  $\mathbf{s}$  is balanced,  $U_i(\mathbf{s}) = U_{i'}(\mathbf{s})$ . Hence,  $U_i(\langle \alpha, \alpha, \beta, \beta \rangle) = \frac{|V(G)|}{4}$  for each player  $i \in [4]$ . We now prove that no player  $i \in [4]$  improves by switching to vertex  $\alpha'$ . Without loss of generality, fix  $i = 1$ . Consider the Voronoi game  $\langle G, [3] \rangle$  with players 1, 2 and 3. By Theorem 7,  $U_1(\langle \alpha', \alpha, \beta \rangle) = \frac{|V(G)|}{4} + \frac{|\Omega(\{\alpha', \alpha, \beta\})|}{12}$ . The utility of player 1 decreases from  $\langle \alpha', \alpha, \beta \rangle$  to  $\langle \alpha', \alpha, \beta, \beta \rangle$  at least due to the fact that the vertices in  $\Omega(\{\alpha', \alpha, \beta\})$  will be shared with player 4 (additionally to the players 1, 2 and 3); this partial decrease is  $\frac{|\Omega(\{\alpha', \alpha, \beta\})|}{3} - \frac{|\Omega(\{\alpha', \alpha, \beta\})|}{4} = \frac{|\Omega(\{\alpha', \alpha, \beta\})|}{12}$ . So,  $U_1(\langle \alpha', \alpha, \beta, \beta \rangle) \leq U_1(\langle \alpha', \alpha, \beta \rangle) - \frac{|\Omega(\{\alpha', \alpha, \beta\})|}{12} = \frac{|V(G)|}{4}$ , as needed.  $\square$

**Without Colocation.** We show:

**Theorem 9.** *Consider a bipartite friendly graph  $G$  with odd diameter. Then, the Voronoi game  $\langle G, [4] \rangle$  has a Nash equilibrium without colocation. Specifically, for any arbitrary pair of two distinct antipodal pairs  $\alpha, \beta$  and  $\gamma, \delta$ , respectively, the profile  $\langle \alpha, \beta, \gamma, \delta \rangle$  is a Nash equilibrium.*

*Proof.* Consider first the bijection  $\psi : V(G) \rightarrow V(G)$  which maps each vertex to its unique antipode. (Since  $G$  is friendly, such a bijection exists.) So,  $\psi(\alpha) = \beta$  and  $\psi(\gamma) = \delta$ ; also,  $\psi^2 = \text{id}$ . It is simple to verify that  $\psi$  is an automorphism of  $G$ :

Since  $G$  is generously vertex-transitive, there is an automorphism  $\varphi$  of  $G$  such that  $\varphi(\alpha) = \gamma$  and  $\varphi(\gamma) = \alpha$ . Since  $\varphi$  preserves distances, it follows that  $\varphi(\beta) = \delta$  and  $\varphi(\delta) = \beta$ . Note that each pair of players is symmetric for the profile  $\langle \alpha, \beta, \gamma, \delta \rangle$  due to some automorphism from  $\psi, \varphi, \varphi\psi$  and  $\psi\varphi$ ; hence, the profile is symmetric. So, it follows that for each player  $i \in [4]$ ,  $U_i(\langle \alpha, \beta, \gamma, \delta \rangle) = \frac{1}{4} |V(G)|$ .

To prove that the symmetric profile  $\langle \alpha, \beta, \gamma, \delta \rangle$  is a Nash equilibrium, we only have to prove that one of the players cannot improve by switching. So, assume that player 3 switches to vertex  $\hat{\gamma}$ . Consider the Voronoi game  $\langle T_d, [3] \rangle$  with players 1, 2 and 3. Since the pair  $\alpha$  and  $\beta$  is antipodal, Corollary 7 implies that  $U_3(\langle \alpha, \beta, \hat{\gamma} \rangle) = \frac{1}{4} |V(G)|$ . Clearly,  $U_3(\langle \alpha, \beta, \hat{\gamma}, \delta \rangle) \leq U_3(\langle \alpha, \beta, \hat{\gamma} \rangle)$ . It follows that  $U_3(\langle \alpha, \beta, \hat{\gamma}, \delta \rangle) \leq \frac{1}{4} |V(G)|$ . So,  $U_3(\langle \alpha, \beta, \hat{\gamma}, \delta \rangle) \leq U_3(\langle \alpha, \beta, \gamma, \delta \rangle)$ , as needed.  $\square$



## 7 Three Players

A profile  $\langle \alpha, \beta, \gamma \rangle$  is *linear* if  $\text{dist}(\alpha, \beta) + \text{dist}(\beta, \gamma) = \text{dist}(\alpha, \gamma)$ ; then,  $\beta$  is called the *middle vertex* and player 2 is called the *middle player*.

**Tori with Odd Diameter.** We show:

**Theorem 10.** *Consider the 2-dimensional torus  $T_2$  with odd diameter  $\sum_{j \in [2]} p_j$ . Then, the Voronoi game  $\langle T_2, [3] \rangle$  has a Nash equilibrium.*

*Proof.* Assume, without loss of generality, that  $p_1 > p_2$ . Set  $\alpha = \langle 0, 0 \rangle$ ,  $\beta = \langle 1, 0 \rangle$  and  $\gamma = \langle p_1, p_2 \rangle$ . We will prove that the profile  $\langle \alpha, \beta, \gamma \rangle$  is a Nash equilibrium. Note that  $\alpha$  and  $\gamma$  are an antipodal pair of vertices. By Lemma 1 and Corollary 7,  $U_2(\langle \alpha, \delta, \gamma \rangle) = \frac{1}{4} \prod_{j \in [d]} (2p_j)$  for any vertex  $\delta \in V(T_d)$ ; thus, we only have to prove that neither player 1 nor 3 can improve her utility by switching. We prove:

**Lemma 11 (Player 3 Cannot Improve).** (1)  $U_3(\langle \alpha, \beta, \gamma \rangle) = 2p_1p_2 - p_2$ . (2) For each vertex  $\hat{\gamma} \in V(T_2)$ ,  $U_3(\langle \alpha, \beta, \hat{\gamma} \rangle) \leq 2p_1p_2 - p_2$ .

**Lemma 12 (Player 1 Cannot Improve).** (1)  $U_1(\langle \alpha, \beta, \gamma \rangle) = p_1p_2 + p_2$ . (2) For each vertex  $\hat{\alpha} \in V(T_2)$ ,  $U_1(\langle \hat{\alpha}, \beta, \gamma \rangle) \leq p_1p_2 + p_2$ .

For Lemma 12, the proof of (1) uses Lemma 11; the proof of (2) uses ideas from the proof of Theorem 6. □

**Hypercubes.** We finally consider the Voronoi game  $\langle H_d, [3] \rangle$ . Consider a profile  $\mathbf{s} = \langle \alpha_1, \dots, \alpha_\kappa \rangle$  for the Voronoi game  $\langle H_d, [\kappa] \rangle$ . Say that dimension  $j \in [d]$  is *irrelevant* for the profile  $\mathbf{s}$  if bit  $j$  is the same in all binary words  $\alpha_i$ , with  $i \in [\kappa]$ . Denote as  $\text{irr}(\mathbf{s})$  the number of irrelevant dimensions for  $\mathbf{s}$ ; clearly,  $0 \leq \text{irr}(\mathbf{s}) \leq d$ . The profile  $\mathbf{s}$  is *irreducible* if it has no irrelevant dimension. Clearly, an antipodal profile is irreducible. We continue with two observations.

**Lemma 13.** *Consider an antipodal pair of vertices  $\alpha$  and  $\beta$  for the hypercube  $H_d$ . Then, for any vertex  $\gamma \in V(H_d)$ , the profile  $\langle \alpha, \beta, \gamma \rangle$  is linear.*

**Lemma 14.** *Fix an irreducible profile  $\langle \alpha, \beta, \gamma \rangle$  for the Voronoi game  $\langle H_d, [3] \rangle$ . Then, (i)  $\text{dist}(\alpha, \beta) + \text{dist}(\beta, \gamma) + \text{dist}(\alpha, \gamma) = 2d$ , and (ii) there is an equivalent profile  $\langle 0^d, 1^{p+q}0^r, 1^p0^q1^r \rangle$ , for some suitable triple of integers  $p, q, r \in \mathbb{N}$ .*

We first determine the utility of an arbitrary player in an irreducible profile. For each index  $i \in \{0, 1\}$ , define the combinatorial function  $M_i : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$  with

$$M_i(x, t) = \begin{cases} \sum_{j=\frac{x-1}{2}-t}^{\frac{x-1}{2}} \binom{x}{j} + \frac{i}{2} \binom{x}{\frac{x-1}{2}-t}, & \text{if } x \text{ is odd} \\ \sum_{j=\frac{x}{2}+1-t}^{\frac{x}{2}-1} \binom{x}{j} + \frac{1-i}{2} \binom{x}{\frac{x}{2}-t} + \frac{1}{2} \binom{x}{\frac{x}{2}}, & \text{if } x \text{ is even} \end{cases}$$

We now determine the utility of an arbitrary player in an irreducible profile.

**Theorem 15.** Fix integers  $x, y, z \in \mathbb{N}$  with  $x + y + z = d$ , and consider the irreducible profile  $\mathbf{s} = \langle \alpha, \beta, \gamma \rangle$  with  $\alpha = 0^d$ ,  $\beta = 1^x 0^y 1^z$  and  $\gamma = 1^{x+y} 0^z$ . Then,

$$U_2(\mathbf{s}) = \frac{1}{4} 2^d + \begin{cases} 2 \sum_{t \in [\frac{z}{2}]} \binom{z}{\frac{z}{2}-t} M_0(x, t) M_0(y, t), & (x \text{ or } y \text{ is odd}), z \text{ is even} \\ \frac{1}{12} \binom{x}{\frac{x}{2}} \binom{y}{\frac{y}{2}} \binom{z}{\frac{z}{2}} + \frac{1}{6} \sum_{t \in [\frac{z}{2}]} \binom{x}{\frac{x}{2}-t} \binom{y}{\frac{y}{2}-t} \binom{z}{\frac{z}{2}} \\ + 2 \sum_{t \in [\frac{z}{2}]} \binom{z}{\frac{z}{2}-t} M_0(x, t) M_0(y, t), & x, y \text{ and } z \text{ are even} \\ 2 \sum_{t=0}^{\frac{z-1}{2}} \binom{z-1}{\frac{z-1}{2}-t} M_1(x, t) M_1(y, t), & (x \text{ or } y \text{ is even}), z \text{ is odd} \\ \frac{1}{6} \sum_{t=0}^{z-1} \binom{x}{\frac{x}{2}-t} \binom{y}{\frac{y}{2}-t} \binom{z-1}{\frac{z-1}{2}-t} \\ + 2 \sum_{t=0}^{\frac{z-1}{2}} \binom{z-1}{\frac{z-1}{2}-t} M_1(x, t) M_1(y, t) & x, y \text{ and } z \text{ are odd} \end{cases}$$

Theorem 15 immediately implies:

**Corollary 16.** Consider the Voronoi game  $\langle H_d, [3] \rangle$ . Fix an antipodal profile  $\langle \alpha, \beta, \gamma \rangle$ , with  $\text{dist}(\alpha, \gamma) = d$ ,  $\text{dist}(\alpha, \beta) = p$  and  $\text{dist}(\beta, \gamma) = q$ . Then,

$$U_2(\langle \alpha, \beta, \gamma \rangle) = \frac{1}{4} 2^d + \begin{cases} 0 & p \text{ or } q \text{ is odd} \\ \frac{1}{12} \binom{p}{\frac{p}{2}} \binom{q}{\frac{q}{2}}, & p \text{ and } q \text{ are even} \end{cases} .$$

We now show:

**Theorem 17.** For any odd integer  $d$ , the Voronoi game  $\langle H_d, [3] \rangle$  has a Nash equilibrium (specifically, any antipodal profile  $\langle \alpha, \beta, \gamma \rangle$  with  $\text{dist}_{H_d}(\alpha, \beta) = 1$ ).

*Proof.* Fix such an antipodal profile  $\langle \alpha, \beta, \gamma \rangle$ . By Lemma 13,  $\langle \alpha, \beta, \gamma \rangle$  is linear; so,  $\text{dist}(\beta, \gamma) = d - 1$ . Since  $d$  is odd, Corollary 7 implies that for any vertex  $\chi \in V(H_d)$ ,  $U_2(\langle \alpha, \chi, \gamma \rangle) = \frac{1}{4} 2^d$ ; so, player 2 cannot improve her utility  $U_2(\alpha, \beta, \gamma)$  by switching. So, in order to prove that the profile  $\langle \alpha, \beta, \gamma \rangle$  is a Nash equilibrium, we only need to consider players 1 and 3. By Lemma 14 (and its proof), there is an equivalent profile  $\langle \alpha, \beta, \gamma \rangle$  with  $\alpha = 0^d$ ,  $\beta = 1^{p+q} 0^r$  and  $\gamma = 1^p 0^q 1^r$ , where  $p+q = \text{dist}_{H_d}(\alpha, \beta) = 1$ ,  $p+r = \text{dist}_{H_d}(\alpha, \gamma) = d$ , and  $q+r = \text{dist}_{H_d}(\alpha, \gamma) = d-1$ . It follows that  $p = 1$ ,  $q = 0$  and  $r = d - 1$ , so that  $\alpha = 0^d$ ,  $\beta = 10^{d-1}$  and  $\gamma = 1^d$ . We use an appropriate automorphism and Theorem 15 to get that  $U_1(\langle 0^d, 10^{d-1}, 1^d \rangle) = \frac{1}{4} 2^d + \frac{1}{2} \binom{d-1}{\frac{d-1}{2}}$ . By the constant-sum property and the inequality  $\binom{d-1}{\frac{d-1}{2}} \leq 2^{d-1}$ , we get that  $U_3(\langle 0^d, 10^{d-1}, 1^d \rangle) \geq \frac{1}{4} 2^d$ . We now use Corollary 16 to prove that  $\langle 0^d, 10^{d-1}, 1^d \rangle$  is a Nash equilibrium.  $\square$

## 8 Open Problems

The full power of *Two-Guards*-like theorems is yet to be realized. Are there similar theorems when either  $G$  comes from some broader class encompassing the friendly graphs, or  $\kappa > 3$ ? More concretely, it is very interesting to generalize

Theorem 15 and find combinatorial formulas for the three players' utilities when  $G = T_d$ ; this may enable generalizing Theorem 17 to the torus  $T_d$  (with odd  $d$ ). It is also interesting to study the *uniqueness* of Nash equilibria; in particular, we know no *non-antipodal* Nash equilibrium on some friendly graph  $G$  with  $\kappa \geq 3$ .

Beyond Theorems 8 and 9, nothing is known for the case  $\kappa \geq 4$ . (For example, we do not know if these results can be extended to broader classes encompassing friendly graphs.) We invite the reader to prove or disprove the following conjectures: (1) The game  $\langle H_d, [\kappa] \rangle$  with odd diameter  $d$  has a Nash equilibrium, whatever  $\kappa$  is. (2) The game  $\langle H_d, [\kappa] \rangle$  with even  $\kappa$  has a Nash equilibrium, whatever  $d$  is. (3) The game  $\langle H_d, [\kappa] \rangle$  with even  $d$  and odd  $\kappa$  has *no* Nash equilibrium.

## References

1. Biggs, N.: Algebraic Graph Theory, 2nd edn. Cambridge Mathematical Library (1993)
2. Biló, V., Flammini, M., Monaco, G., Moscardelli, L.: On the Performance of Nash Equilibria in Isolation Games. In: Proceedings of the 15th Annual International Conference on Computing and Combinatorics. LNCS, vol. 5609, pp. 17–26. Springer, Heidelberg (2001)
3. Dürr, C., Thang, N.K.: Nash Equilibria in Voronoi Games on Graphs. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) ESA 2007. LNCS, vol. 4698, pp. 17–28. Springer, Heidelberg (2007)
4. Gessel, I.M., Stanley, R.P.: "Algebraic Enumeration". In: Graham, R.L., Grötschel, M., Lovász, L. (eds.) Handbook of Combinatorics. The MIT Press & North-Holland (1995)
5. Godsil, C.D.: Algebraic Combinatorics. Chapman and Hall, Boca Raton (1993)
6. Godsil, C., Royle, G.: Algebraic Graph Theory. Graduate Texts in Mathematics. Springer, Heidelberg (2001)
7. Gross, J.L., Yellen, J. (eds.): Handbook of Graph Theory. CRC Press, Boca Raton (2004)
8. Hahn, G., Sabidussi, G. (eds.): Graph Symmetry — Algebraic Methods and Applications. Springer, Heidelberg (1997)
9. Hotelling, H.: Stability in Competition. The Economic Journal 39(153), 41–57 (1929)
10. Kilkenny, M., Thisse, J.: Economics of Location: A Selective Survey. Computers and Operations Research 26(14), 1369–1394 (1999)
11. Leighton, F.T.: Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes. Morgan Kaufmann, San Francisco (1991)
12. Mavronicolas, M., Monien, B., Papadopoulou, V.G., Schoppmann, F.: Voronoi Games on Cycle Graphs. In: Ochmański, E., Tyszkiewicz, J. (eds.) MFCS 2008. LNCS, vol. 5162, pp. 503–514. Springer, Heidelberg (2008)
13. Nash, J.F.: Equilibrium Points in  $n$ -Person Games. Proceedings of the National Academy of Sciences of the United States of America 36, 48–49 (1950)
14. Nash, J.F.: Non-Cooperative Games. Annals of Mathematics 54, 286–295 (1951)
15. Zhao, Y., Chen, W., Teng, S.-H.: The Isolation Game: A Game of Distances. In: Hong, S.-H., Nagamochi, H., Fukunaga, T. (eds.) ISAAC 2008. LNCS, vol. 5369, pp. 148–158. Springer, Heidelberg (2008)

# Selfish Scheduling with Setup Times

Laurent Gourvès<sup>1,\*</sup>, Jérôme Monnot<sup>1,\*</sup>, and Orestis A. Telelis<sup>2,\*\*</sup>

<sup>1</sup> LAMSADE, CNRS FRE 3234, Université de Paris-Dauphine, 75775 Paris, France

{laurent.gourves,monnot}@lamsade.dauphine.fr

<sup>2</sup> Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

telelis@gmail.com

**Abstract.** We study multiprocessor scheduling games with setup times on identical machines. Given a set of scheduling policies (coordination mechanism) on the machines, each out of  $n$  players chooses a machine to assign his owned job to, so as to minimize his individual completion time. Each job has a processing length and is of a certain type. Same-type jobs incur a setup overhead to the machine they are assigned to. We study the Price of Anarchy with respect to the makespan of stable assignments, that are pure Nash or strong equilibria for the underlying strategic game. We study in detail the performance of a well established preemptive scheduling mechanism. In an effort to improve over its performance, we introduce a class of mechanisms with certain properties, for which we examine existence of pure Nash and strong equilibria. We identify their performance limitations, and analyze an optimum mechanism out of this class. Finally, we point out several interesting open problems.

## 1 Introduction

We study multiprocessor job scheduling games with setup times, where each out of  $n$  players assigns his owned job for execution to one out of  $m$  machines. Jobs are of certain *types*. On any machine, jobs of a given type may be executed *only* after a type-dependent preprocessing (performed once for all same-type jobs) called *setup*. Each machine schedules its assigned jobs according to a *scheduling policy* (algorithm). Given the deployed scheduling policies, players assign their jobs selfishly, to minimize their individual completion times. We examine the impact of selfish behavior on the *overall* (social) cost of *stable* assignments and how this can be alleviated, by deployment of appropriate scheduling policies on the machines. Stable assignments are pure Nash equilibria (PNE) of an underlying strategic game, or strong equilibria (SE); the latter extend PNE by being resilient to coalitional deviations [1]. The overall cost of stable assignments is measured by the latest completion time among all players, known as *makespan*. System performance degradation due to selfish behavior is measured by the

---

\* This work is supported by French National Agency (ANR), project COCA JC09 461625.

\*\* This work was carried out during the tenure of an ERCIM “Alain Benssouan” Fellowship Programme.

*Price of Anarchy*, the worst-case ratio of the makespan of the most expensive equilibrium, relative to the optimum achievable makespan [2].

Our work is motivated by concerns in performance optimization of large-scale distributed systems (computational grids, P2P file sharing systems etc.). In these systems autonomous entities (end-users, communities, enterprises) compete strategically for resources (distributed storage, processing power, bandwidth) to increase their individual profit. Setup overheads in these systems may well dominate the net processing load of tasks assigned by users; consider e.g. loading application environments, booting operating systems, establishing QoS for network connections. Scheduling with setup times also provides for modeling another very natural situation; the case where many autonomous users may benefit from the output of an identical process. In this case users may only care for the output of the setup (which takes non-negligible time) but their jobs may have virtually zero load for the machine. As an example, consider the setup corresponding to computation and writing of output to a remote file which users simply read. Then the machine may also have to deal *obviously* of users' presence; only one of them may actually declare his presence by requesting execution of the setup, whereas the rest simply benefit from the request.

The vast number of resources and users in modern distributed environments renders centralized deployment of global resource management policies expensive and inefficient; a central authority may only specify local operational rules per resource, to coordinate users towards globally efficient system utilization. The deployment of such local rules was formalized in [3], under the notion of *coordination mechanisms* and demonstrated for scheduling and network congestion games. In scheduling, local policies are scheduling algorithms deployed on the machines; their set is referred to as a coordination mechanism. The policies may be preemptive or non-preemptive, deterministic or randomized; a policy decides the order of execution of assigned jobs on each machine and may also introduce delays in a systematic manner. A coordination mechanism induces a strategic game, by affecting the players' completion times. The purpose of designing coordination mechanisms is to induce a strategic game that has stable assignments (in our case, PNE or SE outcomes) with low *PoA*, that can be found efficiently.

A series of works concerned the study of strategic games induced by well established and novel scheduling policies applied on basic scheduling settings [2, 4, 3, 5, 6, 7, 8, 9, 10]. The *PoA* of strong equilibria (SPoA) was first studied for selfish scheduling games in [11] under the preemptive mechanism introduced in [2]. We give a brief account of these works below, in section 2. Our focus is on *strongly local scheduling policies* under which, the completion time of a job  $j$  on a machine  $i$  is solely dependent on the parameters (with respect to  $i$ ) of jobs assigned to  $i$ . It is called simply *local* if the completion time of  $j$  depends on parameters of jobs assigned to  $i$  across all machines. We investigate in detail the performance of strongly local mechanisms that can handle the challenges outlined above.

**Contribution.** We analyze the performance of deterministic strongly local coordination mechanisms for selfish scheduling with setup times, on identical

machines. We give a detailed analysis of a preemptive scheduling mechanism, referred to as **Makespan** (section 4), which was introduced in [2] and studied subsequently in [7,11,12]. We show existence of SE for the game induced by **Makespan**. If  $k$  denotes the number of different types of jobs, the *PoA* of **Makespan** is  $m$ , when  $m \leq k$  and  $k + 1 - \epsilon$ , for  $1 > \epsilon \geq 1/m$  when  $m > k$ . We prove  $SPoA = 3/2$  for  $m = 2$  and 2 for  $m \geq 3$ . In section 5 we study a class of deterministic strongly local mechanisms, referred to as *type ordering* mechanisms, that can schedule jobs obliviously of the number of players with zero processing lengths. We prove that any deterministic type ordering mechanism induces a strategic game that has PNE, for any number of machines, and SE for 2 machines. We prove a lower bound of  $\frac{m+1}{2}$  for the *PoA* of type ordering mechanisms, and argue that other intuitive solutions are no more powerful. In section 6 we analyze the performance of an optimal type ordering mechanism. It achieves a *PoA* of  $\frac{m+1}{2}$ , when  $m \leq k$ , and  $\frac{k+3}{2} - \epsilon$  ( $\epsilon = \frac{k}{m}$ , when  $m$  is even and  $\epsilon = \frac{k-1}{m-1}$  otherwise) when  $m > k$ . We conclude with challenging open problems in section 7.

## 2 Related Work

Performance of coordination mechanisms with respect to Nash equilibria in multi-processor scheduling *games* has been the subject of several recent works [2, 4,3,7,6,5,11,12,8,9]. The preemptive mechanism known as **Makespan** was introduced and studied in [2] in the scheduling setting of *uniformly related* machines; each machine  $i$  has speed  $v_i$  and each job  $j$  has processing length  $\ell_j$ , so that the time needed by  $i$  to execute  $j$  is  $\ell_j/v_i$ . The **Makespan** mechanism schedules jobs in parallel on each machine, so that they all have the same completion time. **Makespan** was shown to have  $PoA = \Theta(\frac{\log m}{\log \log m})$  on uniformly related machines in [4,5] (see also [13]). In case of *identical* machines, all speeds are equal and the *PoA* of PNE is known to be  $\frac{2m}{m+1}$  by the works of [14,15]. This holds also for the (*S*)*PoA* of strong equilibria, which were shown to exist in any machine model [11]. [12] studied the *SPoA* as a function of the number of different speeds.

Scheduling games in *unrelated* machines were additionally studied in [7,8,9]. In the *unrelated* machines model a job's processing time depends solely on the machine it is assigned to. In [7], bounds on the *PoA* of well known deterministic and randomized scheduling policies were studied for unrelated machines. A wide class of non-preemptive strongly local scheduling policies was shown in [8] to have  $PoA \geq \frac{m}{2}$ . This class contains well known policies such as “longest” and “shortest job first”. The authors designed a simply *local* mechanism that induces PNE with  $PoA = O(\log^2 m)$ . A local mechanism with PNE having  $PoA = O(\log m)$  was given recently in [9]. Deterministic and randomized *non-clairvoyant* mechanisms for scheduling on unrelated machines were studied in [10].

Preemptive multi-processor scheduling *with setup times* [16] (see also [17][SS6]) requires a minimum makespan preemptive schedule, such that the setup is executed by a machine between execution of two job portions of different type. The best known approximation algorithm has a performance guarantee of  $\frac{4}{3}$  [18] (see [19] for a previous  $\frac{3}{2}$  factor). For equal setup times a PTAS is given in [18] and an FPTAS for 2 machines in [20]. See [21] for a slightly different version.

### 3 Definitions

We consider  $m$  identical machines, indexed by  $i \in \mathcal{M} = \{1, \dots, m\}$  and  $n$  jobs  $j \in \mathcal{J} = \{1, \dots, n\}$ , each owned by a self-interested player. We use interchangeably the terms *job* and *player*. Every job  $j$  has a *type*  $t_j \in \mathcal{U}$ , where  $\mathcal{U}$  is the universe of all possible types<sup>1</sup>. The subset of  $\mathcal{U}$  corresponding to the set of jobs  $\mathcal{J}$  is denoted by  $\mathcal{T} = \{t_j | j \in \mathcal{J}\}$  and define  $k = |\mathcal{T}|$ . We refer to any specific type by  $\theta$ . Each job  $j \in \mathcal{J}$  and each type  $\theta \in \mathcal{U}$  are respectively associated to processing length  $\ell_j \geq 0$  and setup time  $w(\theta) \geq 0$ . If  $w(\theta) = 0$ , then  $\ell_j > 0$  for all  $j$  with  $t_j = \theta$ . Otherwise, we allow  $\ell_j = 0$ . Agent  $j \in \mathcal{J}$  chooses a strategy  $s_j \in \mathcal{M}$  (a machine to assign his job to). A strategy profile (assignment) is denoted by  $s = (s_1, \dots, s_n)$ ;  $s_{-j}$  refers to  $s$ , without the strategy of  $j$ .

The cost of a player  $j$  under assignment  $s$  is  $c_j(s)$ , the completion time of his job.  $c_j(s)$  depends on the scheduling policy deployed on machine  $s_j$ . The completion time of machine  $i \in \mathcal{M}$  under  $s$  is  $C_i(s) = \max_{j:s_j=i} c_j(s)$ . The social cost function is the makespan  $C(s) = \max_i C_i(s) = \max_j c_j(s)$ . We use  $L_i(s) = \sum_{j:s_j=i} \ell_j$  for the total job processing load on machine  $i$  under assignment  $s$ , excluding setup times.  $L_{i,\theta}(s) = \sum_{t_j=\theta, s_j=i} \ell_j$  is the total processing length of type  $\theta$  assigned on machine  $i$ .  $T_i(s)$  denotes the subset of types that have jobs on machine  $i$  under  $s$ . A scheduling policy is a scheduling algorithm. The set of scheduling policies deployed on all machines is a coordination mechanism.

**Definition 1** [1, 11]. *A strategy profile  $s$  is a strong equilibrium if for every subset of players  $J \subseteq \mathcal{J}$  and every assignment  $s' = (s_{-J}, s'_J)$ , where  $s'_j \neq s_j$  for all  $j \in J$ , there is at least one player  $j_0 \in J$  with  $c_{j_0}(s) \leq c_{j_0}(s')$ .*

The makespan of a socially optimum assignment  $s^*$  can be lower bounded as:

$$\begin{aligned}
 \text{(a)} \quad & mC(s^*) \geq \sum_{\theta \in \mathcal{T}} w(\theta) + \sum_{j \in \mathcal{J}} \ell_j & (1) \\
 \text{(b)} \quad & C(s^*) \geq w(t_j) + \ell_j & \text{for any } j \in \mathcal{J} \\
 \text{(c)} \quad & (k-1)C(s^*) \geq \sum_{\xi \in \mathcal{T} \setminus \{\theta\}} w(\xi) & \text{for any } \theta \in \mathcal{T}
 \end{aligned}$$

The only restriction that we impose on the scheduling policies is that the setup of any type  $\theta$  on any machine  $i$  is executed before execution of type  $\theta$  jobs on  $i$ .

### 4 On the Makespan Mechanism

We study an adaptation of the preemptive mechanism introduced in [2] and referred to as **Makespan** [7, 10]. In any assignment  $s$  under **Makespan** it is  $c_j(s) = C_{s_j}(s)$  for every  $j \in \mathcal{J}$ ; completion time of  $j$  equals completion time of the machine that  $j$  is assigned to. **Makespan** schedules jobs on a machine in parallel by usage of time multiplexing. They are broken into small pieces that are executed in a round-robin fashion; each job is assigned a fraction of the machine's time proportionally to its processing length.

<sup>1</sup> E.g. the set of application environments installed on each machine.



**Theorem 1.** *Strong Equilibria exist in the scheduling game with setup times, under the Makespan mechanism.*

See full version for the proof; it generalizes a related proof of *theorem 3.1* in [11]. Thus, theorem 1 is valid even for uniformly related or unrelated machines.

**Theorem 2.** *The PoA of the Makespan mechanism for the scheduling game with setup times is  $m$  when  $m \leq k$ , at most  $k + 1 - k/m$  when  $m > k$  and at least  $\frac{k+1}{1+\epsilon}$  for  $m \geq 3k - 2$  and  $\epsilon = \frac{2k-1}{m-k+1}$ .*

*Proof. Case  $m \leq k$ :* The most expensive PNE  $s$  has makespan  $C(s) \leq \sum_{\theta} w(\theta) + \sum_j \ell_j$  (all jobs on one machine). By (1a), it is  $PoA \leq m$ . For the lower bound take  $k = m$  types of jobs, each type  $\theta$  having  $w(\theta) = 1$  and containing  $m$  jobs of zero processing length. A job of each type is assigned to each machine in  $s$ . Then  $C(s) = m$  and  $s$  is clearly a PNE. In the social optimum  $s^*$  each type is assigned to a dedicated machine, thus  $C(s^*) = 1$ .

*Case  $m > k$ :* Assume that for the most expensive PNE  $s$  it is  $C(s) = C_1(s)$ . Let  $x$  be a job of type  $t_x = \theta$  executed on machine 1.  $x$  cannot decrease its cost  $c_x(s) = C_1(s)$  by switching to any machine  $i \neq 1$ . Then  $C_1(s) \leq C_i(s) + \ell_x$  if  $\theta \in T_i(s)$  or  $C_1(s) \leq C_i(s) + w(\theta) + \ell_x$  otherwise. We sum up the inequalities over all machines, assuming that  $\theta$  does not appear on  $\alpha$  machines, and add  $C_1(s)$  to both sides to obtain  $mC_1(s) \leq \sum_{i=1}^m C_i(s) + \alpha w(\theta) + (m - 1)\ell_x \leq m \sum_{\xi \in \mathcal{T}} w(\xi) + (m - 1)\ell_x + \sum_{j \in \mathcal{J}} \ell_j$ . Divide by  $m$  and rewrite it as:

$$C_1(s) \leq \frac{m-1}{m} \left( \sum_{\xi \in \mathcal{T} \setminus \{\theta\}} w(\xi) + w(\theta) + \ell_x \right) + \frac{1}{m} \left( \sum_{\xi \in \mathcal{T}} w(\xi) + \sum_{j \in \mathcal{J}} \ell_j \right)$$

Using (1a,b,c),  $C(s) \leq \frac{m-1}{m} ((k - 1)C(s^*) + C(s^*)) + C(s^*) = (k + 1 - \frac{k}{m})C(s^*)$ .

For the lower bound take  $k$  types,  $m \geq 3k - 2$ , and let  $w(1) = 0$  and  $w(\theta) = 1$  for  $\theta \in \{2, \dots, k\}$ . There are  $k + 1$  jobs of type 1 and length 1 and  $\frac{m-1}{\epsilon}$  jobs of type 1 and length  $\epsilon = \frac{2k-1}{m-k+1}$ . Types  $\theta \in \{2, \dots, k\}$  have  $m - 1$  jobs each, of processing length 0. A PNE  $s$  is as follows.  $k + 1$  jobs of type 1 and length 1 are assigned to machine 1. One job from each type  $\theta \geq 2$  is assigned to each machine  $i = 2, \dots, m$ .  $\frac{1}{\epsilon}$  jobs of type 1 and length  $\epsilon$  are also assigned to each machine  $i \geq 2$ . Thus  $C_i(s) = k$  for  $i \geq 2$  and  $C_1(s) = k + 1$ . No job may decrease its completion time (equal to the makespan of the machine it is assigned to) by switching machine. In the optimum assignment  $s^*$  assign two jobs of type 1 - with lengths 1 and  $\epsilon$  - to each machine  $i = 1 \dots k + 1$ . Every machine  $i = k + 2 \dots 2k$ , has  $m - 1$  jobs of type  $i - k$ , each of length 0. Every machine  $i = 2k + 1 \dots m$ , has  $1/\epsilon + 1$  jobs of type 1, of length  $\epsilon$ . The makespan of  $s^*$  is  $1 + \epsilon$ .  $\square$

**Theorem 3.** *The Price of Anarchy of strong equilibria under Makespan for the scheduling game with setup times is 2 for  $m \geq 3$ , and  $\frac{3}{2}$  for  $m = 2$  machines.*

*Proof.* We give the proof for the case  $m \geq 3$  (the proof for  $m = 2$  is deferred to the full version). Let  $s$  be a SE,  $s^*$  the socially optimum assignment, and  $C(s) = C_1(s)$ . If  $C_1(s) \leq C(s^*)$  we get  $SPoA = 1$ . If  $C_1(s) > C(s^*)$ , there is



machine  $i \neq 1$  with  $C_i(s) \leq C(s^*)$ , because otherwise  $s$  would not be a SE; all jobs would reduce their completion time by switching from  $s$  to  $s^*$ . For any job  $x$  with  $s_x = 1$ , it is  $c_x(s) \leq c_x(s_{-x}, i)$ . Thus  $C_1(s) = c_x(s) \leq C_i(s) + w(t_x) + \ell_x$ . Thus  $C(s) = c_x(s) \leq 2C(s^*)$ , because  $C_i(s) \leq C(s^*)$  and **(10b)**. For the lower bound, take 3 machines and 4 jobs, with  $t_1 = t_2 = \theta_1$  and  $t_3 = t_4 = \theta_2$ . Set  $w(\theta_1) = \epsilon$ ,  $\ell_1 = \ell_2 = 1$  and  $w(\theta_2) = 1$ ,  $\ell_3 = \ell_4 = \epsilon$ . An assignment where jobs 1, 2 play machine 1 and jobs 3, 4 play machines 2, 3 respectively is a strong equilibrium of makespan  $2 + \epsilon$ . In the social optimum jobs 3, 4 are assigned to the same machine and 1 and 2 on dedicated machines; the makespan becomes then  $1 + 2\epsilon$ . Thus  $SPoA \geq \frac{2+\epsilon}{1+2\epsilon} \rightarrow 2$ , as  $\epsilon \rightarrow 0$ .  $\square$

## 5 Type Ordering Mechanisms

We describe a class of (deterministic) *type ordering* mechanisms, for *batch scheduling* of same-type jobs. Each machine  $i$  groups jobs of the same type  $\theta$ , into a *batch* of type  $\theta$ . A type batch is executed as a whole; the setup is executed first, followed by *preemptive* execution of all jobs in the batch, in a **Makespan** fashion. Jobs within the same batch have equal completion times and are scheduled preemptively in parallel. Type batches are executed serially by each machine.

Policies in type ordering mechanisms satisfy a version of the property of *Independence of Irrelevant Alternatives* (IIA) **(8)**. Under the IIA property, for any set of jobs  $J_i \subseteq \mathcal{J}$  assigned to machine  $i \in \mathcal{M}$  and for any pair of types  $\theta, \theta' \in \mathcal{U}$  with jobs in  $J_i$  if the  $\theta$ -type batch has smaller completion time than the  $\theta'$ -type batch, then the  $\theta$  batch has a smaller completion time than the  $\theta'$  batch in any set  $J_i \cup \{j\}$ ,  $j \in \mathcal{J} \setminus J_i$ . Presence of  $j$  does not affect the relative order of execution of  $\theta$  and  $\theta'$  batches. The IIA property was used in **(8)** for proving a lower bound on the *PoA* of a class of job ordering mechanisms in the context of unrelated machines scheduling. Type ordering policies do not introduce delays in the execution of batches, but only decide the relative order of their execution, based on a batch's type index and setup time. They do not use the number of jobs within each batch; otherwise the IIA property may not be satisfied. Job lengths are used only for **Makespan**-wise scheduling within batches. Hence type ordering mechanisms function obliviously of “hidden” players with zero job lengths.

We prove next existence of PNE for any number of machines, and of SE for  $m = 2$  under type ordering mechanisms. An algorithm for finding PNE follows. Let  $o(i)$  be the ordering of types on machine  $i$ , and  $O = \{o(i) | i \in \mathcal{M}\}$  be the set of all orderings of the mechanism. By  $\prec_o$  denote the precedence relation of types, prescribed by  $o \in O$ . Let  $M_o$  be the set of machines that schedule according to  $o \in O$ . Initialize  $o \in O$  arbitrarily, and **repeat** until all jobs are assigned:

1. Find the earliest type  $\theta$  according to  $\prec_o$ , with at least one unassigned job.
2. Let  $j$  be the largest length unassigned job with  $t_j = \theta$ .
3. Pick  $i \in \mathcal{M}$  minimizing completion time of  $j$  **(9)** (break ties in favor of  $i \in M_o$ ).
4. **If**  $i \in M_o$  set  $s_j = i$  **else** switch ordering  $o$  to  $o(i)$ .

<sup>2</sup>  $j$  incurs processing load  $w(t_j) + \ell_j$  if a  $t_j$ -type job is not already assigned to  $i$ .

**Theorem 4.** *The scheduling game with setup times has pure Nash equilibria, under type ordering mechanisms.*

*Proof.* The algorithm terminates in polynomial time; once a job is assigned, it is never considered again and within every  $O(m + n)$  iterations some job is always assigned. For any type  $\theta$ , denote by  $\hat{s}_\theta$  the partial assignment up to the time after the last job of type  $\theta$  has been assigned. We show by contradiction that no job  $j$  has incentive to deviate under an assignment  $s$  returned by the algorithm.

Assume that  $j$  does have incentive to deviate from  $s_j$ , and let  $s'$  be the resulting assignment after deviation of  $j$ . At the time corresponding to the partial assignment  $\hat{s}_{t_j}$ , there is no type  $\theta \neq t_j$  and machine  $i$  such that  $\theta \in T_i(\hat{s}_{t_j})$  and  $t_j \prec_{o(i)} \theta$ . If it was the case, the first job of type  $\theta \neq t_j$  assigned to  $i$  would have been chosen before jobs of type  $t_j$  were exhausted, which contradicts step 1. of the algorithm. Thus, batches of type  $t_j$  are scheduled - under  $\hat{s}_{t_j}$  - last on all machines with  $t_j \in T_i(\hat{s}_{t_j})$ . Furthermore, if  $j$  wishes to deviate to a machine  $i \neq s_j$ , then  $c_j(s) = c_j(\hat{s}_{t_j}) > C_i(\hat{s}_{t_j}) + \ell_j = c_j(s')$ , if  $t_j \in T_i(\hat{s}_{t_j})$ , and  $c_j(s) = c_j(\hat{s}_{t_j}) > C_i(\hat{s}_{t_j}) + w(t_j) + \ell_j = c_j(s')$ , if  $t_j \notin T_i(\hat{s}_{t_j})$ . Let  $j'$  be the last job of type  $t_j$  assigned to machine  $s_j$  (it may be  $j' = j$ ). Because  $\ell_{j'} \leq \ell_j$ , it is also  $c_{j'}(\hat{s}_{t_j}) = c_j(\hat{s}_{t_j}) > C_i(\hat{s}_{t_j}) + \ell_{j'}$  or  $c_{j'}(\hat{s}_{t_j}) = c_j(\hat{s}_{t_j}) > C_i(\hat{s}_{t_j}) + w(t_{j'}) + \ell_{j'}$  accordingly. By the time  $j'$  was assigned, the completion time of  $i$  was at most  $C_i(\hat{s}_{t_j})$ . This contradicts step 3. of the algorithm with respect to  $j'$ .  $\square$

Note that the previous result can be easily extended to hold in the case of *uniformly related* machines (to appear in the full version - see also [13]).

**Theorem 5.** *For the scheduling game with setup times under type ordering mechanisms, any pure Nash equilibrium is strong, when  $m = 2$ .*

*Proof.* Assume that  $s$  is PNE, but not SE, and let  $J \subseteq \mathcal{J}$  be a coalition of jobs that have incentive to deviate jointly. Define  $J_1 = \{j \in J | s_j = 1\}$ ,  $J_2 = \{j \in J | s_j = 2\}$ ; since  $s$  is PNE,  $J_1, J_2 \neq \emptyset$ . Let  $\theta_i$  be the earliest type according to  $\prec_{o(i)}$  with jobs in  $J_i$  and denote by  $J'_i$  type  $\theta_i$  jobs in  $J_i$ . Take two jobs  $j_1 \in J'_1$ ,  $j_2 \in J'_2$ , and let  $s'$  be the resulting assignment after deviation.

**CASE 1:**  $\theta_1 \neq \theta_2$ . Since  $s$  is a PNE, it must be  $\theta_1 \prec_{o(1)} \theta_2$  and  $\theta_2 \prec_{o(2)} \theta_1$  because, if e.g.  $\theta_2 \prec_{o(1)} \theta_1$ ,  $j_2$  would have incentive to deviate unilaterally to machine 1, since it wishes to deviate jointly with coalition  $J$ . Hence  $c_{j_2}(s') \geq c_{j_1}(s) - \sum_{j \in J'_1} \ell_j + \sum_{j \in J'_2} \ell_j + w(\theta_2)$  if  $J'_1$  does not contain the entire batch of type  $\theta_1$  and  $c_{j_2}(s') \geq c_{j_1}(s) - \sum_{j \in J'_1} \ell_j - w(\theta_1) + \sum_{j \in J'_2} \ell_j + w(\theta_2)$  otherwise. So, in the worst case, we get  $c_{j_2}(s') \geq c_{j_1}(s) - \sum_{j \in J'_1} \ell_j - w(\theta_1) + \sum_{j \in J'_2} \ell_j + w(\theta_2)$ . Similarly,  $c_{j_1}(s') \geq c_{j_2}(s) - \sum_{j \in J'_2} \ell_j - w(\theta_2) + \sum_{j \in J'_1} \ell_j + w(\theta_1)$ . Summing up these two inequalities, we obtain  $c_{j_2}(s') + c_{j_1}(s') \geq c_{j_2}(s) + c_{j_1}(s)$  which is impossible since it must be  $c_{j_2}(s') < c_{j_2}(s)$  and  $c_{j_1}(s') < c_{j_1}(s)$ .

**CASE 2:**  $\theta_1 = \theta_2$ . Then, in the worst case we obtain  $c_{j_2}(s') = c_{j_1}(s) - \sum_{j \in J'_1} \ell_j + \sum_{j \in J'_2} \ell_j$  and  $c_{j_1}(s') = c_{j_2}(s) - \sum_{j \in J'_2} \ell_j + \sum_{j \in J'_1} \ell_j$ . The rest of the proof is similar to the previous case.  $\square$

For  $m \geq 3$ , a PNE under type ordering mechanisms is not generally SE (see full version). The following result identifies performance limitations of type ordering mechanisms, due to lack of a priori knowledge of  $\mathcal{T} \subseteq \mathcal{U}$ .

**Theorem 6.** *The Price of Anarchy of the scheduling game with setup times is at least  $\frac{m+1}{2}$  for every deterministic type ordering mechanism.*

*Proof.* For any deterministic type ordering mechanism, assume there is a subset  $\mathcal{T} \subseteq \mathcal{U}$  of  $k = 2m - 1$  types, say  $\mathcal{T} = \{1, \dots, 2m - 1\}$ , such that: all types of  $\mathcal{T}$  are scheduled in order of ascending index in  $a$  machines and in order of descending index in  $d = m - a$  machines. Then, there is a family of instances with  $PoA \geq \frac{m+1}{2}$ . Next we prove existence of  $\mathcal{T}$ . Set  $w(\theta) = 1$  for all  $\theta \in \mathcal{U}$ . When  $a = m$  or  $d = m$ , take an instance of  $m$  zero length jobs for each type  $\theta \in \{1, \dots, m\}$ . Placing one job of each type on every machine yields a PNE with makespan  $m$ . An assignment of makespan 1 has all same-type jobs assigned to a dedicated machine, thus  $PoA \geq m$ . When  $a \geq 1$  and  $d \geq 1$ , the instance has:

- $a$  jobs of zero length for each type  $\theta \in \{1, \dots, m - 1\}$
- $d$  jobs of zero length for each type  $\theta \in \{m + 1, \dots, 2m - 1\}$
- $m - 1$  jobs of of zero length and type  $m$
- one job of length 1 and type  $m$
- no jobs for  $\theta \in \mathcal{U} \setminus \mathcal{T}$

Assign one job of type  $\theta \in \{1, \dots, m - 1\}$  on each of the  $a$  ascending type index machines, and one job of type  $\theta \in \{m + 1, \dots, 2m - 1\}$  on each of the  $d$  descending type index machines. Put one job of type  $m$  on every machine. This is a PNE of makespan  $m + 1$ . Placing all jobs of type  $\theta \in \{i, 2m - i\}$  on machine  $i$  yields makespan 2. Thus it is  $PoA \geq \frac{m+1}{2}$ .

We show existence of  $\mathcal{T}$  for sufficiently large universe  $\mathcal{U}$ . We use the fact that any sequence of  $n$  different real numbers has a monotone (not necessarily contiguous) subsequence of  $\sqrt{n}$  terms (a corollary of Theorem 4.4, page 39 in [22]). By renaming types in  $\mathcal{U}$  we can assume w.l.o.g. that  $\mathcal{U}$  is ordered monotonically (index-wise) on machine 1, and set  $T_1 = \mathcal{U}$ . Then, there is  $T_2 \subseteq T_1$  such that  $|T_2| \geq \sqrt{|T_1|}$  and all the types of  $T_2$  are ordered monotonically according to index, on machines 1 and 2. After  $m - 1$  applications of the corollary, we obtain a set  $T_m \subseteq T_{m-1} \subseteq \dots \subseteq T_1 = \mathcal{U}$  with  $|T_m| \geq |\mathcal{U}|^{2^{1-m}}$  and all its types are scheduled monotonically to their index on every machine. We set  $\mathcal{T} = T_m$ , and take a universe  $\mathcal{U}$  of types with  $|\mathcal{U}| = (2m - 1)^{2^{m-1}}$ , to ensure existence of  $\mathcal{T}$  with  $k = |\mathcal{T}| = 2m - 1$  types.  $\square$

Let us note that “longest” or “shortest batch first” policies are no more powerful than type ordering mechanisms; they reduce to them for zero length jobs.

## 6 An Optimal Type Ordering Mechanism

We analyze the  $PoA$  of a type ordering mechanism termed AD (for *Ascending-Descending*), that schedules type batches by ascending type index on half of the machines, and by descending type index on the rest. If  $m$  is odd one of the policies is applied to one machine more. First we prove the following lemma.

**Lemma 1.** *Let  $\mathcal{T}' \subseteq \mathcal{T}$  include types with non-zero setup times. If two jobs of the same type in  $\mathcal{T}'$  play an ascending and a descending index machine respectively under the AD mechanism, their type batches are scheduled last on the respective machines.*

*Proof.* We show the result by contradiction. Let jobs  $x_1, x_2$  with  $t_{x_1} = t_{x_2} = \theta$  be assigned on the ascending and descending machines 1, 2 respectively. Assume that a job  $y$ ,  $t_y = \theta' \neq \theta$ , is scheduled on 1 after type  $\theta$ . Because  $s$  is a PNE, job  $x_2$  does not decrease its completion time if it moves to machine 1; because  $y$  is scheduled after  $x_1$  on 1:

$$c_{x_1}(s) \geq c_{x_2}(s) - \ell_{x_2}, \text{ and } c_y(s) \geq c_{x_1}(s) + w(\theta') + \ell_y \tag{2}$$

If  $y$  switches to machine  $M_2$  then it will be scheduled before type  $\theta$ , thus its completion time will be at most  $c_{x_2}(s) - w(\theta) - \ell_{x_2} + w(\theta') + \ell_y$  if  $\theta' \notin T_2(s)$  (and at most  $c_{x_2}(s) - w(\theta) - \ell_{x_2} + \ell_y$  otherwise). In the worst case, we obtain:

$$c_y(s) \leq c_{x_2}(s) - w(\theta) - \ell_{x_2} + w(\theta') + \ell_y \tag{3}$$

By (2) and (3),  $c_y(s) \leq c_y(s) - w(\theta) < c_y(s)$ , a contradiction, because  $\theta \in \mathcal{T}'$ .  $\square$

The next result identifies upper bounds on the  $PoA$  of AD. A proposition that follows proves tightness, through lower bounds on the Price of Stability, the ratio of the *least* expensive PNE makespan over the optimum makespan. We take  $k \geq 2$ ; AD is identical to **Makespan** for  $k = 1$ .

**Theorem 7.** *The Price of Anarchy of the AD mechanism for the scheduling game with setup times is at most  $\frac{m+1}{2}$  when  $m \leq k$  and at most  $\frac{k+3}{2} - \epsilon$  ( $\epsilon = \frac{k}{m}$  when  $m$  is even and  $\epsilon = \frac{k-1}{m-1}$  otherwise), when  $m > k$ .*

*Proof.* Let  $s$  be a PNE assignment and  $\mathcal{T}' \subseteq \mathcal{T}$  contain types with non-zero setups. Assume  $C(s) = C_1(s) = \max_i C_i(s)$ . Let  $\theta_0$  be the type scheduled last on machine 1 and  $x$  a job with  $t_x = \theta_0$ . Define  $\mathcal{T}'_C \subseteq \mathcal{T}'$  to be types with jobs assigned to both ascending and descending machines under  $s$ . Let  $\mathcal{T}'_A \subseteq \mathcal{T}' \setminus \mathcal{T}'_C$  and  $\mathcal{T}'_D \subseteq \mathcal{T}' \setminus \mathcal{T}'_C$  contain types exclusively assigned to ascending and descending machines respectively. Notice that at most one type  $\theta_1 \in \mathcal{T}'_C$  may appear on at least  $\frac{m}{2} + 1$  machines, when  $m$  even, and  $\frac{m+1}{2}$  machines, when  $m$  odd; thus any type in  $\mathcal{T}'_C \setminus \{\theta_1\}$  appears on at most  $\frac{m}{2}$  or  $\frac{m-1}{2}$  machines respectively. We study 2 cases depending on whether  $\theta_1$  exists and whether it coincides with  $\theta_0$  or not.

**CASE 1:  $\theta_0 = \theta_1$  or  $\theta_1$  does not exist.** Job  $x$  will not decrease its completion time by moving to machine  $p$  for  $p = 2, \dots, m$ . If  $M_{\theta_0}(s)$  are the indices of machines which contain type  $\theta_0$ , then:

$$\begin{aligned} \forall p \in M_{\theta_0}(s), \quad c_x(s) &\leq C_p(s) + \ell_x && \text{and} \\ \forall p \notin M_{\theta_0}(s), \quad c_x(s) &\leq C_p(s) + w(\theta_0) + \ell_x \end{aligned} \tag{4}$$

To obtain the upper bound we sum up (4) for  $p \in \{2, \dots, m\}$ , add  $C_1(s)$  in the left and right hand part, and take  $\sum_{\theta \notin \mathcal{T}'} w(\theta) = 0$ . We will do this analysis below, collectively for cases 1 and 2.

**CASE 2:  $\theta_0 \neq \theta_1$  and  $\theta_1$  exists.** Assume  $\theta_0 < \theta_1$  and let  $R$  contain the indices of ascending machines which have at least one job of type  $\theta_1$  assigned (if  $\theta_0 > \theta_1$ , we consider the indices of descending machines). Let  $R$  be the indices of these machines and  $R' \subseteq R$  be the indices of machines that are also assigned type  $\theta_0$  jobs (note that  $\theta_0 \notin \mathcal{T}'_C$  if  $R' \neq \emptyset$ ). If job  $x$  moves to a machine with index in  $p \in R'$ , the completion time of  $x$  becomes at most  $C_p(s) - w(\theta_1) + \ell_x$  and  $C_p(s) - w(\theta_1) + w(\theta_0) + \ell_x$  if  $p \in R'' = R \setminus R'$ . Since  $s$  is a PNE:

$$\begin{aligned} \forall p \in R', \quad c_x(s) &\leq C_p(s) - w(\theta_1) + \ell_x, \\ \forall p \in R'', \quad c_x(s) &\leq C_p(s) - w(\theta_1) + w(\theta_0) + \ell_x \end{aligned} \tag{5}$$

We will sum up inequalities (4) or (5) for  $p \in \{2, \dots, m\}$  depending on whether  $p \in R$  or not. As in case 1 we add  $C_1(s)$  to left and right hand parts and consider  $\sum_{\theta \notin \mathcal{T}'} w(\theta) = 0$ . Before summing note that when  $m$  is even, each type in  $\mathcal{T}'_A \cup \mathcal{T}'_D$  has jobs assigned to at most  $\frac{r}{2}$  machines, for  $r = m$ . When  $m$  is odd assume w.l.o.g. that there are  $\frac{m+1}{2}$  descending machines. We ignore one of them - different than  $M_1$  - in the summation (we assume  $m \geq 3$ ; otherwise  $m = 1$  and  $C(s) = C(s^*)$ ). Then, in case 2, type  $\theta_1$  appears at most  $\frac{r}{2}$  times,  $r = m - 1$ , in the remaining  $m - 1$  machines.

$$\begin{aligned} rC_1(s) &\leq \frac{r}{2} \left( \sum_{\theta \in \mathcal{T}'_A \cup \mathcal{T}'_D \setminus \{\theta_0\}} w(\theta) + \sum_{\theta \in \mathcal{T}'_C \setminus \{\theta_0\}} w(\theta) \right) + rw(\theta_0) + \sum_{j \in \mathcal{J}} \ell_j + (r - 1)\ell_x \\ &= \frac{r}{2} \left( \sum_{\theta \in \mathcal{T}} w(\theta) + \sum_{j \in \mathcal{J}} \ell_j \right) + \frac{r}{2}w(\theta_0) + (r - 1)\ell_x - \frac{r - 2}{2} \sum_{j \in \mathcal{J}} \ell_j \end{aligned} \tag{6}$$

$$\leq \frac{r}{2} \left( \sum_{\theta \in \mathcal{T}} w(\theta) + \sum_{j \in \mathcal{J}} \ell_j \right) + \frac{r}{2} (w(\theta_0) + \ell_x) \text{ since } \sum_{j \in \mathcal{J}} \ell_j \geq \ell_x \tag{7}$$

When  $k \geq m$  we use (IIa,b) with (7) to obtain  $C_1(s) \leq \frac{m+1}{2}OPT$ . When  $k < m$  we rewrite (6) as:

$$C(s) \leq \frac{1}{r} \left( \sum_{\theta \in \mathcal{T}} w(\theta) + \sum_{j \in \mathcal{J}} \ell_j \right) + \left( \frac{1}{2} - \frac{1}{r} \right) \left( \sum_{\theta \in \mathcal{T}} w(\theta) + \ell_x \right) + \frac{1}{2} (w(\theta_0) + \ell_x) \tag{8}$$

Using (IIb,c), we get  $kC(s^*) \geq \sum_{\theta \in \mathcal{T}} w(\theta) + \ell_x$  and replacing  $r = m$  and  $r = m - 1$  for even and odd  $m$  respectively, yields the stated bounds with respect to  $k$ .  $\square$

**Proposition 1.** *The Price of Stability of the scheduling game with setup times under the AD mechanism is  $\frac{m+1}{2}$  when  $k > m$  and  $\frac{k+3}{2} - \epsilon$  ( $\epsilon = \frac{k}{m}$  when  $m$  is even and  $\epsilon = \frac{k-1}{m-1}$  otherwise) when  $k \leq m$ .*

*Proof.* For  $k > m$  we use the same example as in the proof of theorem 6, but replace the zero length jobs with very small  $\epsilon > 0$  length. For AD the described assignment for  $a, d \geq 1$  applies, and it is a PNE with makespan  $m + 1 + (m - 1)\epsilon$ ;

the socially optimum makespan has length  $2 + m\epsilon$ . In any PNE, all jobs of types 1 and  $2m - 1$  will play exactly the strategies specified in the described PNE assignment, because a lower completion time is not achievable for them in any assignment. Inductively, jobs of types  $i$  and  $2m - i$ ,  $i = 2 \dots m - 1$ , follow the same practice, given the strategies of jobs of types  $i - 1$ ,  $2m - i + 1$ . For the jobs of type  $m$ , the strategies described in the aforementioned assignment are best possible, given the strategies of all other jobs. Therefore the described PNE is unique, hence  $PoS \rightarrow \frac{m+1}{2}$  for  $\epsilon \rightarrow 0$ . The same uniqueness argument holds when  $k \leq m$ , for the instances given below.

**$k \geq 2$  even.** There are  $m = 2k$  machines,  $k$  ascending  $A_i$  and  $k$  descending  $D_i$  for  $i = 1, \dots, k$ . For each type  $\theta \neq \frac{k}{2} + 1$ ,  $w(\theta) = 1$  and there are  $k$  jobs of this type with length  $\epsilon$ . Finally, there are  $k + 1$  jobs of type  $\frac{k}{2} + 1$  with length 1 and  $w(1 + k/2) = 0$ . Consider the state  $s$  where  $A_1$  has a job of each type  $1, \dots, \frac{k}{2} + 1$ , machine  $A_i$ ,  $i = 2, \dots, k$ , has one job of each type  $1, \dots, \frac{k}{2}$ , and finally descending machine  $D_i$ ,  $i = 1, \dots, k$ , has a job of each type  $\frac{k}{2} + 1, \dots, k$ .  $s$  is a PNE and  $C(s) = \frac{k}{2} + 1 + \frac{k}{2}\epsilon$ . A socially optimum assignment  $s^*$  is defined as follows. For each type  $\theta \neq \frac{k}{2} + 1$ , a dedicated machine schedules all jobs of type  $\theta$ . Thus,  $k - 1$  machines are busy and  $k + 1$  are free. Each job of type  $\frac{k}{2} + 1$  is scheduled on a dedicated machine out of the  $k + 1$  free ones. Then  $C(s^*) = 1 + k\epsilon$ . Since  $m = 2k$ , when  $\epsilon$  tends to 0, we get:  $PoA = \frac{k}{2} + 1 = \frac{k+3}{2} - \frac{k}{m}$ .

**$k \geq 3$  odd.** Take  $m = k$  machines:  $\frac{k+1}{2}$  of ascending index and  $\frac{k-1}{2}$  of descending index. Each type has setup time 1 and the length of each job is  $\epsilon$ . There are  $\frac{k+1}{2}$  jobs for each of the first  $\frac{k-1}{2}$  types, assigned to a distinct ascending index machine each. There are  $\frac{k-1}{2}$  jobs for each of the last  $\frac{k-1}{2}$  types, assigned to a distinct descending index machine each. The middle type (with index  $\frac{k+1}{2}$ ) has  $k$  jobs, each assigned to a distinct machine. This assignment is a PNE and has makespan  $\frac{k+1}{2}(1 + \epsilon)$ . In the socially optimum assignment we place all jobs of every type on a dedicated machine and achieve makespan  $1 + k\epsilon$  (type  $\frac{k+1}{2}$ ). The ratio tends to  $\frac{k+1}{2} = \frac{k+3}{2} - \frac{k-1}{m-1}$  as  $\epsilon \rightarrow 0$ . □

## 7 Open Problems

The universe of types  $\mathcal{U}$  is required to be huge (double exponential) in the proof of Theorem 6. This size is non-realistic for most interesting practical settings. Is there a lower size of  $\mathcal{U}$  that yields  $PoA \geq \frac{m+1}{2}$  for type ordering mechanisms? E.g., the proof requires that  $|\mathcal{U}| \geq 9$  when  $m = 2$ , although  $|\mathcal{U}| \geq 3$  suffices. The performance of type ordering mechanisms is not fully characterized by theorem 6: there may be certain sizes of  $|\mathcal{U}|$  below which these mechanisms may perform better. Another interesting issue is when type ordering mechanisms are a priori aware of the subset of types  $\mathcal{T}$  that corresponds to players  $\mathcal{J}$ . What is the impact of such an a priori knowledge to the achievable  $PoA$  by type ordering mechanisms? We have not considered in this paper simply local mechanisms or more challenging machine environments (uniformly related or unrelated machines). All these are interesting aspects for future developments on the subject.

## References

1. Aumann, R.J.: Acceptable points in games of perfect information. *Pacific Journal of Mathematics* 10, 381–417 (1960)
2. Koutsoupias, E., Papadimitriou, C.H.: Worst-case Equilibria. In: *Proc. STACS. LNCS*, vol. 1543, pp. 404–413. Springer, Heidelberg (1999)
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. Koutsoupias, E., Mavronicolas, M., Spirakis, P.G.: Approximate Equilibria and Ball Fusion. *Theory of Computing Systems* 36(6), 683–693 (2003)
5. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms* 3(1) (2007)
6. Christodoulou, G., Gourvès, L., Pascual, F.: Scheduling Selfish Tasks: About the Performance of Truthful Algorithms. In: Lin, G. (ed.) *COCOON 2007. LNCS*, vol. 4598, pp. 187–197. Springer, Heidelberg (2007)
7. Immorlica, N., Li, L., Mirrokni, V.S., Schulz, A.: Coordination Mechanisms for Selfish Scheduling. In: Deng, X., Ye, Y. (eds.) *WINE 2005. LNCS*, vol. 3828, pp. 55–69. Springer, Heidelberg (2005)
8. Azar, Y., Jain, K., Mirrokni, V.S.: (almost) optimal coordination mechanisms for unrelated machine scheduling. In: *Proc. ACM-SIAM SODA*, pp. 323–332 (2008)
9. Caragiannis, I.: Efficient coordination mechanisms for unrelated machine scheduling. In: *Proc. AMC-SIAM SODA*, pp. 815–824 (2009)
10. Durr, C., Nguyen, T.K.: Non-clairvoyant Scheduling Games. In: *Proc. SAGT. LNCS*. Springer, Heidelberg (to appear, 2009)
11. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: *Proc. ACM-SIAM SODA*, pp. 189–198 (2007)
12. Epstein, L., van Stee, R.: The Price of Anarchy on Uniformly Related Machines Revisited. In: Monien, B., Schroeder, U.-P. (eds.) *SAGT 2008. LNCS*, vol. 4997, pp. 46–57. Springer, Heidelberg (2008)
13. Vöcking, B.: Selfish Load Balancing. In: *Algorithmic Game Theory*, ch. 20, pp. 517–542. Cambridge University Press, Cambridge (2007)
14. Finn, G., Horowitz, E.: A linear time approximation algorithm for multiprocessor scheduling. *BIT* 19, 312–320 (1979)
15. Schuurman, P., Vredeveld, T.: Performance guarantees of local search for multiprocessor scheduling. In: Aardal, K., Gerards, B. (eds.) *IPCO 2001. LNCS*, vol. 2081, pp. 370–382. Springer, Heidelberg (2001)
16. Ausiello, G., Crescenzi, P., Gambosi, G., Kann, V., Marchetti-Spaccamela, A., Proti, M.: *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer, Heidelberg (1999)
17. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-completeness*. W.H.Freeman & Co Ltd., New York (1979)
18. Schuurman, P., Wöginger, G.J.: Preemptive scheduling with job-dependent setup times. In: *Proc. ACM-SIAM SODA*, pp. 759–767 (1999)
19. Chen, B.: A better heuristic for preemptive parallel machine scheduling with batch setup times. *SIAM Journal on Computing* 22, 1303–1318 (1993)
20. Wöginger, G.J., Yu, Z.: A heuristic for preemptive scheduling with set-up times. *Computing* 49(2), 151–158 (1992)
21. Chen, B., Ye, Y., Zhang, J.: Lot-sizing scheduling with batch setup times. *Journal of Scheduling* 9(3), 299–310 (2006)
22. Jukna, S.: *Extremal Combinatorics with Applications in Computer Science*. Springer, Heidelberg (2001)



# Computational Aspects of Multimarket Price Wars

Nithum Thain<sup>1,\*</sup> and Adrian Vetta<sup>2,\*\*</sup>

<sup>1</sup> Department of Mathematics and Statistics, McGill University  
nithum.thain@mail.mcgill.ca

<sup>2</sup> Department of Mathematics and Statistics, and School of Computer Science,  
McGill University  
vetta@math.mcgill.ca

**Abstract.** We consider the complexity of decision making with regards to predatory pricing in multimarket oligopoly models. Specifically, we present multimarket extensions of the classical single-market models of Bertrand, Cournot and Stackelberg, and introduce the War Chest Minimization Problem. This is the natural problem of deciding whether a firm has a sufficiently large war chest to win a price war. On the negative side we show that, even with complete information, it is hard to obtain any multiplicative approximation guarantee for this problem. Moreover, these hardness results hold even in the simple case of linear demand, price, and cost functions. On the other hand, we give algorithms with arbitrarily small *additive* approximation guarantees for the Bertrand and Stackelberg multimarket models with linear demand, price, and cost functions. Furthermore, in the absence of fixed costs, this problem is solvable in polynomial time in all our models.

## 1 Introduction

This paper concerns price wars and predatory pricing in markets. We focus on multiple markets (or a single segmentable market) as it allows us to model a broader and more realistic set of interactions between firms. A firm may initiate a price war in order to increase market share or to deter other firms from competing in particular markets. The firm suffers a short-term loss but may gain large future profits, particularly if the price war forces out the competition and allows it to price as a monopolist.

Price wars (and predatory pricing) have been studied extensively from both an economic and a legal perspective. A detailed examination of all aspects of price wars is far beyond the scope of this paper. Rather, we focus on just one important aspect: the complexity of decision making in oligopolies (e.g. duopolies). Specifically, we consider the budget required by a firm in order to successfully launch a price war. This particular question is fundamental in determining the risk and benefits arising from predatory practices. Moreover, it arises naturally in the following two scenarios:

---

\* Supported in part by NSERC CGS grant.

\*\* Supported in part by NSERC grant 288334.



**ENTRY DETERRENCE:** How much of a war chest must a monopolist or cartel have on hand so that they are able to successfully repel a new entrant?

**COMPETITION REDUCTION:** How much money must a firm or cartel have to force another firm out of business? For example, in a duopoly how much does a firm need to save before it can defeat the other to create a monopoly?

The War Chest Minimization Problem, a generalization of the above scenarios, is the problem of determining how large a budget one firm or cartel needs before it can legally win a price war. This paper studies the computational complexity of and approximation algorithms for this more general problem.

## 1.1 Background

Price wars and predatory pricing are tools that have been long associated with monopolies and cartels. The literature on these topics is vast and we touch upon just a small sample in this short background section.

Given the possible rewards for monopolies and cartels engaging in predatory behaviour, it is not surprising that it has been a recurrent theme over time. The late 19th century saw cartels engaging in predation in a plethora of industries. Prominent examples include the use of “fighting ships” by the British Shipping Conferences ([19], [17]) to control trade routes, the setting up of phoney independents by the American Tobacco Company to undercut smaller competitors [6]. Perhaps the most infamous instance, though, of a cartel concerns Standard Oil under the leadership of John D. Rockefeller [16]. More recent examples of price wars include the cigarette industry [9], the airline industry [4], and the retail industry [5]. In the computer industry, Microsoft regularly faces accusations of predatory practices ([12], [15]).

Antitrust legislation has been introduced in many countries to prevent anti-competitive behaviour like predatory pricing or oligopolistic collusion<sup>1</sup>. In the United States, the most important such legislation is the Sherman Act of 1890. One of the Act’s earliest applications came in 1911 when the Supreme Court ordered the break-up of both Standard Oil and American Tobacco; more recently, it was applied when the Court ordered the break-up of American Telephone and Telegraph (AT&T) in 1982<sup>2</sup>.

Given that such major repercussions may arise, there is a need for a cloak of secrecy around any act of predation. This has meant the extent of predatory pricing is unknown and has been widely debated in the literature. Indeed, early

<sup>1</sup> Whilst it is easy to see the negative aspect of cartels, it is interesting to note that there may even be some positive consequences. For example, it has been argued [11] that the predatory actions of cartels may *increase* consumer surplus.

<sup>2</sup> In 2000, a lower court also ordered the breakup of Microsoft for antitrust violations under the Sherman Act. On appeal, this punishment was removed under an agreed settlement in 2002.

economic work of McGee [16] suggested that predatory pricing was not rational. However, in Stigler's seminal work on oligopolies [21], price wars can be viewed as a break-down of a cartel, *albeit* they do not arise in equilibria because collusion can be enforced via punishment mechanisms. Moreover, recent models have shown how price wars can be recurrent in a "functioning" cartel! For example, this can happen assuming the presence of imperfect monitoring [14] or of business cycles [18]. This is particularly interesting as recurrent price wars were traditionally seen as indicators of a healthy competitive market<sup>3</sup>.

Based primarily on the work of McGee, the US Supreme court now considers predatory pricing to be *generally implausible*<sup>4</sup>. As a result of this, and in an attempt to strike a balance between preventing anti-competitive behaviour and overly restricting normal competition, the Court applied the following strict definition to test for predatory practices.

- (a) The predator is pricing below its short-run costs.
- (b) The predator has a strong chance or recouping the losses incurred during the price-war.

The established way for the Court to test for the first requirement is the Areeda-Turner rule of 1975 [1] which established marginal cost (or, as an approximate surrogate, average variable cost) as the primary criteria for predatory pricing<sup>5</sup>. We will incorporate the Areeda-Turner rule as a legal element in our multimarket oligopoly models in Section 2. The second requirement essentially states that the "short-run loss is an investment in prospective monopoly profits" [8]. This requirement is typically simpler to test for in practice, and will be implicit in our models.

Finally, we remark that we are not aware of any other work concerning the complexity of price wars. One interesting related pricing strategy is that of loss-leaders which Balcan et al. [2] examine with respect to profit optimization. For the scale and type of problem we consider, however, using strategies that correspond to "loss-leaders" is illegal. Alternative models for oligopolistic competition and collusion in a single market setting can be found in the papers of Ericson and Pakes [10] and Weintraub et al. [24].

---

<sup>3</sup> Therefore, should such behaviour also arise in practice it would pose intriguing questions for policy makers. Specifically, when is a price war indicative of competition and when is it indicative of the presence of a cartel or a predatory practice?

<sup>4</sup> See the 1986 case *Matsushita Electric Industrial Company vs Zenith Radio Corporation* and the 1993 case *Brooke Group Limited vs Brown and Williamson Tobacco Corporation*.

<sup>5</sup> We note that the Areeda-Turner rule may be inappropriate in high-tech industries because fixed costs there are typically high. Therefore, measures of variable costs may not be reflective of the presence of a price-war. In fact, hi-tech industries may be particularly susceptible to predatory practices as large marginal profits are required to cover the high fixed costs. Consequently, predatory pricing can be used to inflict great damage on smaller firms.

## 1.2 Our Results

A firm with price-making power belongs to an industry that is a monopoly or oligopoly. In Section 2, we develop three multimarket models of oligopolistic competition. We then introduce the War Chest Minimization Problem to capture the essence of the Entry Deterrence and Competition Reduction scenarios outlined above.

In Section 3, we prove that this problem is NP-Hard in all three multimarket models under the legal constraints imposed by the Areeda-Turner rule. We emphasise that decision making is hard even under complete information.

The hardness results of Section 4 imply that no multiplicative approximation guarantee can be obtained for the Minimum War Chest Problem, even in the simple case of linear cost, price, and demand functions. The situation for potential predators is less bleak than this result appears to imply. To see this we present two positive results in Section 4, assuming linear cost, price, and demand functions. First, the problem can be solved in polynomial time if the predator faces no fixed costs. In addition, for the Bertrand and Stackelberg models there is a natural way to separate the markets into two types, those where player one is making a profit and those in which she is truly fighting a price war. Our second result states that in these models, we can solve the problem on the former set of markets exactly and can find a fully polynomial time approximation scheme for the problem on the latter markets. This leads to a polynomial time algorithm with an arbitrarily small additive guarantee.

## 2 Models

### 2.1 Multimarket Models of Oligopoly

In this section, we formulate multimarket versions of the classical models of Bertrand, Cournot, and Stackelberg for oligopolistic competition. Our models allow for the investigation of the numerous and assorted interactions between firms.

**A Multimarket Bertrand Model.** The Bertrand model is a natural model of price competition between firms (henceforth referred to as “players”) in an oligarchy [3]. In this paper, we will focus on the following generalization of the asymmetric Bertrand model to multiple markets [6]. We will describe the model for the duopoly case, but all of the definitions are easily generalizable. Suppose we have two players and  $n$  markets  $m_1, m_2, \dots, m_n$ . Every player  $i$  has a budget  $B_i$  where a negative budget is thought of as the fixed cost for the firm to exist and a positive budget is thought of as a war chest available to that firm in the round. Every market  $m_k$  has a demand curve  $D_k(p)$  and each player  $i$  also has a marginal cost,  $c_{ik}$ , for producing one unit of good in market  $m_k$ . In addition,

---

<sup>6</sup> We remark that this multimarket Bertrand model is also a generalization of the multiple market model used in the facility location game of Vetta [23].

each player  $i$  has a fixed cost,  $f_{ik}$ , for each market  $m_k$  that she pays if and only if she enters the market, i.e. if she sets some finite price.

We model the price war as a game between the two players. A strategy for player  $i$  is a complete specification of prices in all the markets. Both players choose their strategies simultaneously. If  $p_{ik} < \infty$  then we will say that player  $i$  enters market  $m_k$ . If player  $i$  chooses not to enter market  $m_k$ , this is signified by setting  $p_{ik} = \infty$ . The demand for each market then all goes to the player with the lowest price. If the players set the same price, then the demand is shared equally. Thus if player  $i$  participates, then she gets profit  $\Pi_{ik}$  in market  $m_k$  where  $\Pi_{ik}(p_{ik}, p_{jk}) = (p_{ik} - c_{ik})D_{ik}(p_{ik}, p_{jk}) - f_{ik}$  and where  $D_{ik}$  is the demand for player  $i$ 's good in market  $m_k$  and is defined as

$$D_{ik}(p_{ik}, p_{jk}) = \begin{cases} D_k(p_{ik}) & \text{if } p_{ik} < p_{jk} \\ \frac{1}{2}D_k(p_{ik}) & \text{if } p_{ik} = p_{jk} \\ 0 & \text{if } p_{ik} > p_{jk} \end{cases}$$

If player  $i$  chooses not to participate then her revenue and costs are both zero; thus, she gets 0 profit. The sum of these profits over all markets is added to each player's budget. A player is eliminated if her budget is negative at the end of the round.

Our hardness results will apply even when the demand functions are linear. So, unless stated otherwise, for the remainder of this paper we will assume our demand functions are of the form  $D(p) = a - bp$ .

**Multimarket Cournot and Stackelberg Models.** Economists have considered a number of alternative models for competition [22]. One prominent alternative is the Cournot model, formulated by Augustin Cournot in 1838 [7]. We now formulate a multimarket version of this model. Again we will restrict ourselves to the case of the duopoly as the generalization is obvious. In this Cournot model, there are  $n$  independent Cournot markets  $m_1, \dots, m_n$ . Each market  $m_k$  has a price function  $P_k(q) = a_k - q$ . Each player also has a budget  $B_i$ , which serves the same role as in the Bertrand case. Each player also has a cost function in every market  $C_{ik}(q_{ik})$ . We only consider cost functions of the form  $C_i(q_i) = c_i q_i + f_i$  for  $q_i > 0$  where  $c_i$  is a constant marginal cost and  $f_i$  is a fixed cost.

As before we model the price war as a game. This time, a strategy for each player  $i$  is a choice of quantities  $q_{ik}$  for each market  $m_k$ . Again, both players choose a strategy simultaneously. We say that player  $i$  enters market  $m_k$  if  $q_{ik} > 0$ . Player  $i$  then makes a profit in market  $m_k$  equal to  $\Pi_{ik}(q_{ik}, q_{jk}) = q_{ik}P_k(q_{ik} + q_{jk}) - C_{ik}(q_{ik})$ .

Again, if player  $i$  does not enter a market then her profits and revenues are zero. Each player's total profit is added to their budget at the end of the round. A player is eliminated if her resulting budget is negative.

The Stackelberg model was formulated by Heinrich von Stackelberg in 1934 as an adaptation of the Cournot model [20]. For the multimarket Stackelberg model, we define all of the quantities and functions as in the Cournot case. However, we now consider one player to be the leader and one to be the follower.

The game is no longer simultaneous, as the leader gets to commit to a production level in each market before the follower moves.

## 2.2 The War Chest Minimization Problem

We will examine the questions of entry deterrence and competition reduction in the two-firm setting. Thus, we focus on the computational problems facing (i) a monopolist fighting against a potential market entrant (entry deterrence) and (ii) a firm in a duopoly trying to force out the other firm (competition reduction). We model both these situations using the same duopolistic multimarket models of Section 2.1.

We remark that our focus on a firm rather than a cartel does not effect the fundamental computational aspects of the problem. This restriction, however, will allow us to avoid the distraction arising from the strategic complications that occur in ensuring coordination amongst members of a cartel.

Our game is then as follows. We assume that players one and two begin with budgets  $B_1$  and  $B_2$ , respectively. They then play one of our three multimarket games. The goal of firm one is to stay/become a monopoly; if it succeeds it will subsequently be able to act monopolistically in each market. To achieve this goal the firm needs a non-negative payoff at the end of the game whilst its opponent has a negative payoff (taking into account their initial budgets). This gives us the following natural question:

**War Chest Minimization Problem:** *How large a budget  $B_1$  does player one need to ensure that it can eliminate an opponent with a budget  $B_2 < 0$ .*

The players can play any strategy they wish *provided* it is legal, that is, they must abide by the Areeda-Turner Rule. All our results will be demonstrated under the assumptions of this rule, as it represents the current legal environment. However, similar complexity results can be obtained without assuming this rule.

**Areeda-Turner Rule:** *It is illegal for either player to price below their marginal cost in any market.*

Before presenting our results we make a few comments about the problem and what the legal constraints mean in our setting. First, notice that we specify a negative budget for player two but place no restriction on the budget for player one. This is natural for our models. We can view the budget as the money a firm initially has at its disposable minus the fixed costs required for it to operate; these fixed costs are additional to the separate fixed costs required to operate in any individual market. Consequently, if the second firm has a positive budget it cannot be eliminated from the game as it has sufficient resources to operate (cover its fixed costs) even without competing in any of the individual markets; thus we must constrain the second firm to have a negative budget. On the other hand, for the first firm no constraint is needed. Even if its initial budget is negative, it is plausible that it can still eliminate the second firm and end up

with a positive budget at the end of the game, by making enough profit from the individual markets. Specifically, the legal constraints imposed by the Areeda-Turner rule may ensure that the second firm cannot maliciously bankrupt the first firm even if the first firm has a negative initial war-chest.

Second, since we are assuming that player one wishes to ensure success regardless of the strategy player two chooses, we will analyze the game as an asynchronous game where player two may see player one's choices before making her own. Player two will then first try to survive despite player one's choice of strategy. If she cannot do so, she will undercut player one in every market in an attempt to eliminate her also. To win the price war, player one must find strategies that keep herself safe and eliminate player two irrespective of how player two plays. Therefore, an optimal strategy for player one has maximum profit (i.e. minimum negative profit) amongst the collection of strategies that achieve these goals, assuming that player two plays maliciously.

Finally, the Areeda-Turner Rule has a straightforward interpretation in the Bertrand model of price competition, that is, neither player can set the price in any market below their marginal cost in that market! In models of quantity competition, however, the interpretation is necessarily less direct. For the Cournot model of quantity competition, we interpret the rule as saying that neither player can produce a quantity that will result in a price less than their marginal cost assuming the other player produces nothing, in other words  $q_{ik} < a_k - c_{ik}$ . This is the weakest interpretation possible for this simultaneous game. Finally, for the Stackelberg game, we assume that the restriction imposed by the Areeda-Turner rule is the same for player one as in the Cournot model, as she acts first and player two has not set a quantity when player one decides. Player two on the other hand, must produce a quantity so that her marginal price is greater than her marginal cost, given what player one has produced. In other words, for the Stackelberg game  $q_{1k} < a_k - c_{1k}$  and  $q_{2k} < a_k - q_{1k} - c_{2k}$ .

### 3 Hardness Results

We are now in a position to show that the War Chest Minimization Problem is hard in all three models.

**Theorem 1.** *The War Chest Minimization Problem is NP-hard for the multi-market Bertrand, Cournot, and Stackelberg models (the latter assuming player one is the Stackelberg leader), even in the case of linear demand, price and cost functions.*

*Proof.* We only include the proof for the Stackelberg case, due to space limitations. The other proofs are similar and can be found in the full paper.

We give a reduction from the knapsack problem. There we have  $n$  items, each with value  $v_i$  and weight  $w_i$ , and a bag which can hold weight at most  $W$ . In general, it is NP-hard to decide whether we can pack the items into the bag so that  $\sum w_i \leq W$  and  $\sum v_i > V$  for some constant  $V$  (where the sums are taken over packed items).

We will now create a multimarket Stackelberg game based on the above instance. Set  $a_k = 4\sqrt{v_k}$ , and suppose that there are  $n$  markets and each market has price function  $P_k(q) = a_k - q$ . We now set player one's marginal cost in market  $m_k$  to be  $c_{1k} = 0$  and her fixed cost to be  $f_{1k} = 4v_k + w_k$ . Player two's marginal cost in market  $m_k$  is set to be  $c_{2k} = a_k/2 = 2\sqrt{v_k}$  and her fixed cost is set to be  $f_{2k} = 0$ . Finally, set the budgets to be  $B_1 = W$  and  $B_2 = V - \sum_{k=1}^n v_k$ .

Now consider the decision player one faces when deciding whether or not to enter market  $m_k$ . First notice that her monopoly quantity is  $q_{1k}^* = a_k/2 = 2\sqrt{v_k}$  which we can calculate by maximizing  $\Pi_{1k}(q_{1k}, 0)$  through simple calculus. Notice also that  $a_k - q_{1k}^* - c_{2k} = 0$  and so, by the Areeda-Turner rule, player two cannot produce in any market in which player one is producing.

Thus, if player one enters any market then she will produce her monopoly quantity in that market and player two will not enter that market. In this case, player one makes profit

$$\Pi_{1k}(q_{1k}^*, 0) = q_{1k}^*(P_k(q_{1k}^*) - c_{1k}) - f_{1k} = 2\sqrt{v_k}(2\sqrt{v_k} - 0) - (4v_k + w_k) = -w_k$$

and player two makes profit  $\Pi_{2k}(q_{1k}^*, 0) = 0$ . On the other hand, if player one does not enter the market then player two will produce her monopoly quantity,  $q_{2k}^* = a_k/4 = \sqrt{v_k}$ , and will make profit

$$\Pi_{2k}(0, q_{2k}^*) = q_{2k}^*(P_k(q_{2k}^*) - c_{2k}) - f_{2k} = \sqrt{v_k}(3\sqrt{v_k} - 2\sqrt{v_k}) - 0 = v_k$$

Since player one did not enter, she will make profit 0.

Thus, if player one could solve the War Chest Minimization Problem then she could determine whether or not there exists a set of indices  $K$  of markets that she should enter such that both of the following equations hold simultaneously:

$$\begin{aligned} W - \sum_{k \in K} w_k &\geq 0 \\ V - \sum_{k=1}^n v_k + \sum_{k \notin K} v_k &< 0 \end{aligned}$$

Rearranging these equations, we obtain the conditions of the knapsack equations, namely  $\sum_{k \in K} w_k \leq W$  and  $\sum_{k \in K} v_k > V$ . □

## 4 Algorithms

In this section, we explore algorithms for solving the War Chest Minimization Problem. We highlight a case where the problem can be solved exactly and explore the approximability of the problem in general. For the entirety of this section, we assume linear cost, demand, and price functions.

### 4.1 A Polynomial Time Algorithm in the Absence of Fixed Costs

All of the complexity proofs in Section 3 have a similar flavor. We essentially use the fixed costs in the markets to construct weights in a knapsack problem.

It turns out that in the absence of these fixed costs, it is computationally easy for a player to determine if they can win a multimarket price war even under the restrictions of the Areeda-Turner rule. For space reasons, we omit the proof of this result, though it can be found in the full paper.

**Theorem 2.** *In the absence of fixed costs and assuming linear cost, price, and demand functions, the War Chest Minimization Problem in the Cournot, Bertrand, and Stackelberg models can be solved in polynomial time.  $\square$*

## 4.2 An Inapproximability Result

In this section, we will explore approximation algorithms for the War Chest Minimization Problem. A first inspection is disheartening for would-be predators, as demonstrated by the following theorem.

**Theorem 3.** *It is NP-hard to obtain any approximation algorithm for the War Chest Minimization Problem under the Bertrand, Stackelberg, or Cournot model.*

*Proof.* We prove this for the Stackelberg model - the other cases are similar. Let  $n, W, V, w_i$ , and  $v_i$  be an instance of the knapsack problem. Construct markets  $m_1, \dots, m_n$  exactly as in Theorem 1, with identical price functions, fixed costs, and marginal costs. Let  $W^*$  denote the optimal solution to the War Chest Minimization Problem in this case. Notice that  $W^* > 0$  since player one makes a negative profit in all of her markets. We now construct a new market  $m_{n+1}$  as follows. Let  $P_{n+1}(q) = 2\sqrt{W^*} - q$  be the price function. Let player one's fixed and marginal costs be  $c_{1,n+1} = f_{1,n+1} = 0$ . Let player two's marginal cost be  $c_{2,n+1} = 2\sqrt{W^*}$  and let her fixed cost be an arbitrary nonnegative value. Then player one will clearly enter the market and produce her monopoly quantity,  $q_{1,n+1} = \sqrt{W^*}$ , thereby forcing player two to stay out of the market, by the Areeda-Turner rule. Thus player one will earn her monopoly quantity of  $W^*$  in this market. Consequently, the budget required for this War Chest Minimization Problem is zero. Any approximation algorithm would then have to solve this problem, and thereby the knapsack problem, exactly.  $\square$

## 4.3 Additive Approximation Guarantees

Observe that the difficulty in obtaining a multiplicative approximation guarantee arises due to conflict between markets that generate a loss for player and markets that generate a profit. Essentially the strategic problem for player one is to partition the markets into two groups,  $\alpha$  and  $\beta$ , and then conduct a price war in the markets in group  $\alpha$  and try to gain revenue to fund this price war from markets in group  $\beta$ . This is still not sufficient because, in the presence of fixed costs, the optimal way to conduct a price war is not obvious even when the group  $\alpha$  has been chosen. However, in this section we will show how to partition the markets and generate an arbitrarily small additive guarantee in the Bertrand and Stackelberg cases.

Given an optimal solution with optimal partition  $\{\alpha^*, \beta^*\}$ , let  $w_{\alpha^*}$  be the absolute value of the sum of the profits of the markets with negative profit, and



let  $w_{\beta^*}$  be the sum of the profits in positive profit markets. Then the optimal budget for player one is simply  $OPT = w_{\alpha^*} - w_{\beta^*}$ . For both the Bertrand and Stackelberg models, we present algorithms that produce a budget of most  $(1+\epsilon)w_{\alpha^*} - w_{\beta^*}$ , for any constant  $\epsilon$ . Observe this can be expressed as  $OPT + \epsilon w_{\alpha^*}$ , and since  $w_{\alpha^*}$  represents the actual cost of the price war (which takes place in the markets in  $\alpha^*$ ), our solution is then at most  $OPT$  plus epsilon times the optimal cost of fighting the price war. We begin with the Bertrand model.

**Theorem 4.** *There is an algorithm that solves the War Chest Minimization Problem for the Bertrand model within an additive bound of  $\epsilon w_{\alpha^*}$ , and runs in time polynomial in the input size and  $\frac{1}{\epsilon}$ , assuming linear demand functions.*

*Proof.* We begin by proving that we can find the optimal partition  $\{\alpha^*, \beta^*\}$  of the markets. Towards this goal we show that there is a optimal pricing scheme for any market, should player one choose to enter the market. Using this scheme we will be able to see which markets are revenue generating for player one and which are not. This will turn out to be sufficient to obtain  $\{\alpha^*, \beta^*\}$ . This is because, in the Bertrand model, player two cannot make a profit in a market if player one does and vice versa and because player one needs a strategy that maintains a non-negative budget even if player two acts maliciously (but legally).

The pricing scheme for player one should she choose to enter market  $m_k$  is  $p_{1k}^+ = \max\{c_{1k}, \min\{p_{1k}^*, c_{2k} - \gamma\}\}$ , where  $\gamma$  is the minimum increment of price and  $p_{1k}^*$  is player one's monopoly price. She should not price below  $p_{1k}^+$  as either (i) it is illegal by the Areeda-Turner rule or (ii) she cannot increase her profit by doing so (as the profit function for player one is a concave quadratic in  $p_{1k}$ ). She also should not price above  $p_{1k}^+$ . If she did then either (i) she cannot increase her profit (due to concavity) or (ii) player two could undercut her or increase her own existing profits in the market. Indeed, it is certain that player two will try to undercut her if player one succeeds in keeping player two's budget negative.

Given that we have the optimal pricing scheme for player one, we may calculate the profit she could make on entering a market assuming that player two acts maliciously. Let  $\alpha$  be the set of markets where she makes a negative profit under these conditions, and let  $\beta$  be the set of markets where she makes a non-negative profit. Since all markets in  $\beta$  give player one a non-negative profit even if player two is malicious, she will clearly always enter all of them. Consequently, as we are in the Bertrand model, player two cannot make any profit from markets in  $\beta$ . Thus by entering every market in  $\beta$  player one will earn  $w_\beta$  profit, and this must be the optimal for player one if the goal is to put player two out of business. So  $\{\beta, \alpha\} = \{\beta^*, \alpha^*\}$  is an optimal partition.

It remains only to show that there is a fully polynomial time approximation scheme for the markets in  $\alpha$ . We will prove this result by demonstrating an approximation preserving reduction of the War Chest Minimization Problem with only  $\alpha$ -type Bertrand markets to the *Minimization Knapsack Problem*. Define  $w_k$  to be the negative of the profit earned by player one if she enters the market  $m_k$  and assuming player two undercuts if possible. By the above, she will price at  $p_{1k} = p_{1k}^+$  and thus

$$w_k = \begin{cases} -(p_{1k}^+ - c_{1k})D(p_{1k}^+) + f_{1k} & \text{if } c_{1k} < c_{2k} \\ f_{1k} & \text{otherwise.} \end{cases}$$

Recall that  $w_k$  is non-negative for markets in  $\alpha$ . Let  $p_{2k}^*$  be player two’s monopoly price in market  $m_k$  and let  $\Pi_{2k}^*$  be her monopoly profit in that market. We also let  $v_k = \Pi_{2k}^* - \Pi_{2k}(p_{1k}^+)$ , where  $\Pi_{2k}(p_{1k}^+)$  is the maximum profit that player two can achieve in market  $m_k$  if player one enters and prices at  $p_{1k}^+$ . The War Chest Minimization Problem is that of maximizing player one’s profit (i.e. minimizing the negative of her profit) even if player two acts maliciously, while ensuring that player two’s budget is always negative. So it can be expressed as

$$\begin{aligned} & \min \sum_k w_k y_k \\ \text{s.t. } & B_2 + \sum_k (\Pi_{2k}^*(1 - y_k) + \Pi_{2k}(p_{1k}^+) \cdot y_k) \leq 0 \\ & y_k \in \{0, 1\} \end{aligned}$$

Setting the constant  $C$  to be the sum of player two’s budget and her monopoly profit in all of the markets, that is  $C = B_2 + \sum_k \Pi_{2k}^*$ , the problem can be rewritten as

$$\begin{aligned} & \min \sum_k w_k y_k \\ \text{s.t. } & \sum_k v_k y_k \geq C \\ & y_k \in \{0, 1\} \end{aligned}$$

Finally, since the  $w_k$  are non-negative, this formulation is exactly the minimization knapsack problem. The reduction is approximation preserving and so we are done as there is a fully polynomial time approximation scheme for the minimization knapsack problem [13]. □

A slightly more complex idea is needed for the Stackelberg model. Instead of reducing to the minimization knapsack problem, we reduce to a different problem that can be solved via rounding a dynamic program. The proof is omitted here, for space reasons, but can be found in the full paper.

**Theorem 5.** *There is an algorithm that solves the War Chest Minimization Problem for the Stackelberg model within an additive bound of  $\epsilon w_{\alpha^*}$  and with running time polynomial in the input size and  $\frac{1}{\epsilon}$ , assuming linear cost and price functions.* □

The approach taken for the Stackelberg model does not apply directly to the Cournot model as a more subtle rounding scheme is required there when player one is more competitive than player two. We conjecture, however, that a similar type of additive approximation guarantee is possible in the Cournot model.

## References

1. Areeda, P., Turner, D.: Predatory pricing and related issues under Section 2 of the Sherman Act. *Harvard Law Review* 88, 733–997 (1975)
2. Balcan, M., Blum, A., Chan, H., Hajiaghayi, M.: A theory of loss-leaders: making money by pricing below cost. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 293–299. Springer, Heidelberg (2007)

3. Bertrand, J.: Théorie mathématique de la richesse sociale. *Journal des Savants*, 499–508 (1883)
4. Blair, R., Harrison, J.: Airline price wars: competition or predation. *Antitrust Bulletin* 44(2), 489–518 (1999)
5. Boudreaux, D.: Predatory pricing in the retail trade: the Wal-Mart case. In: Coate, M., Kleit, A. (eds.) *The Economics of the Antitrust Process*, pp. 195–215. Kluwer, Dordrecht (1989)
6. Burns, M.: Predatory pricing and the acquisition cost of competitors. *Journal of Political Economy* 94, 266–296 (1986)
7. Cournot, A.: *Recherches sur les Principes Mathématiques de la Théorie des Richesses*, Paris (1838)
8. Elzinga, K., Mills, D.: Testing for predation: Is recoupment feasible? *Antitrust Bulletin* 34, 869–893 (1989)
9. Elzinga, K., Mills, D.: Price wars triggered by entry. *International Journal of Industrial Organization* 17, 179–198 (1999)
10. Ericson, R., Pakes, A.: Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies* 62(1), 53–82 (1995)
11. Fershtman, C., Pakes, A.: A dynamic oligopoly with collusion and price wars. *RAND Journal of Economics* 31(2), 207–236 (2000)
12. Fisher, F.: The IBM and Microsoft cases: what's the difference? *American Economic Review* 90(2), 180–183 (2000)
13. Gens, G.V., Levner, E.V.: Computational complexity of approximation algorithms for combinatorial problems. In: *Mathematical Foundations of Computer Science*. LNCS, vol. 74, pp. 292–300. Springer, Berlin (1979)
14. Green, E., Porter, R.: Noncooperative collusion under imperfect price competition. *Econometrica* 52, 87–100 (1984)
15. Klein, B.: Did Microsoft engage in anticompetitive exclusionary behavior. *Antitrust Bulletin* 46(1), 71–113 (2001)
16. McGee, J.: Predatory price cutting: the Standard Oil (NJ) case. *Journal of Law and Economics* 1, 137–169 (1958)
17. Podolny, J., Scott Morton, F.: Social status, entry and predation: the case of British shipping cartels 1879–1929. *Journal of Industrial Economics* 47(1), 41–67 (1999)
18. Rotemberg, J., Soloner, G.: A supergame-theoretic model of price wars during booms. *American Economic Review* 76(3), 390–407 (1986)
19. Scott Morton, F.: Entry and predation: British shipping cartels 1879–1929. *Journal of Economics and Management Strategy* 6, 679–724 (1997)
20. von Stackelberg, H.: *Marktform und Gleichgewicht*, Julius Springer, Vienna (1934)
21. Stigler, G.: A theory of oligopoly. *Journal of Political Economy* 72(1), 44–61 (1964)
22. Tirole, J.: *The Theory of Industrial Organization*. MIT Press, Cambridge (1988)
23. Vetta, A.: Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. In: *Proceedings of the 43rd Symposium on the Foundations of Computer Science (FOCS)*, pp. 416–425 (2002)
24. Weintraub, G., Benkard, C.L., Roy, B.V.: Markov perfect industry dynamics with many firms. *Econometrica* 76(6), 1375–1411 (2008)

# The Impact of Social Ignorance on Weighted Congestion Games\*

Dimitris Fotakis<sup>1</sup>, Vasilis Gkatzelis<sup>2</sup>, Alexis C. Kaporis<sup>3,4</sup>, and Paul G. Spirakis<sup>4</sup>

<sup>1</sup> School of Electrical and Computer Engineering,  
National Technical University of Athens, 15780 Athens, Greece

<sup>2</sup> Computer Science Department, Courant Institute,  
New York University, 251 Mercer Street, NY 10012, USA

<sup>3</sup> Department of Information and Communication Systems Engineering,  
University of the Aegean, 83200 Samos, Greece

<sup>4</sup> Research Academic Computer Technology Institute,  
N. Kazantzaki Str., University Campus, 26500 Patras, Greece  
fotakis@cs.ntua.gr, gkatz@cims.nyu.edu,  
kaporis@ceid.upatras.gr, spirakis@cti.gr

**Abstract.** We consider weighted linear congestion games, and investigate how social ignorance, namely lack of information about the presence of some players, affects the inefficiency of pure Nash equilibria (PNE) and the convergence rate of the  $\varepsilon$ -Nash dynamics. To this end, we adopt the model of graphical linear congestion games with weighted players, where the individual cost and the strategy selection of each player only depends on his neighboring players in the social graph. We show that such games admit a potential function, and thus a PNE. Our main result is that the impact of social ignorance on the Price of Anarchy (PoA) and the Price of Stability (PoS) is naturally quantified by the *independence number*  $\alpha(G)$  of the social graph  $G$ . In particular, we show that the PoA grows roughly as  $\alpha(G)(\alpha(G) + 2)$ , which is essentially tight as long as  $\alpha(G)$  does not exceed half the number of players, and that the PoS lies between  $\alpha(G)$  and  $2\alpha(G)$ . Moreover, we show that the  $\varepsilon$ -Nash dynamics reaches an  $\alpha(G)(\alpha(G) + 2)$ -approximate configuration in polynomial time that does not directly depend on the social graph. For unweighted graphical linear games with symmetric strategies, we show that the  $\varepsilon$ -Nash dynamics reaches an  $\varepsilon$ -approximate PNE in polynomial time that exceeds the corresponding time for symmetric linear games by a factor at most as large as the number of players.

## 1 Introduction

Congestion games provide a natural model for non-cooperative resource allocation in large-scale systems and have recently been the subject of intensive research. In a (weighted) *congestion game*, a finite set of non-cooperative players, each controlling an unsplittable (weighted) demand, compete over a finite set of resources. All players using a resource experience a delay (or cost) given by a non-negative and non-decreasing

---

\* This work was partially supported by the ICT Programme of the EU, under contract no. ICT-2008-215270 (FRONTS), by NSF grant CCF0830516, and by Andreas Mentzelopoulos Scholarships for the University of Patras.

function of the resource's total demand (or congestion). Among a given set of resource subsets (or strategies), each player selects one selfishly trying to minimize her *individual cost*, that is the sum of the delays on the resources in her strategy. The most natural solution concept is that of a *pure Nash equilibrium* (PNE), a configuration where no player can decrease her individual cost by unilaterally changing her strategy. Rosenthal [20] proved that the PNE of (unweighted) congestion games correspond to the local optima of a potential function, and thus every congestion game admits a PNE. A similar result was recently shown for weighted congestion games with linear delays [15].

**Motivation and Previous Work.** The prevailing questions in recent work on congestion games have to do with quantifying the inefficiency due to the players' non-cooperative and selfish behaviour (see e.g. [19,3,5,12,11,2,9,13]), and bounding the convergence time to (approximate) PNE if the players select their strategies in a selfish and decentralized fashion (see e.g. [14,11,10,21,6]).

*Inefficiency of Pure Nash Equilibria.* It is well known that a PNE may not optimize the system performance, usually measured by the *total cost* incurred by all players. The main tools for quantifying and understanding the performance degradation due to the players' non-cooperative and selfish behaviour have been the *Price of Anarchy* (PoA), introduced by Koutsoupias and Papadimitriou [19], and the *Price of Stability* (PoS), introduced by Anshelevich *et al.* [3]. The (pure) PoA (resp. PoS) is the *worst-case* (resp. *best-case*) ratio of the total cost of a PNE to the optimal total cost.

Many recent contributions have provided tight bounds on the PoA and the PoS for several interesting classes of congestion games, mostly congestion games with linear and polynomial delays. Awerbuch *et al.* [5] and Christodoulou and Koutsoupias [12] proved that the PoA of congestion games is  $5/2$  for linear delays and  $d^{\Theta(d)}$  for polynomial delays of degree  $d$ . Subsequently, Aland *et al.* [2] obtained exact bounds on the PoA for congestion games with polynomial delays. For weighted congestion games with linear delays, Awerbuch *et al.* [5] proved that the PoA is  $(3+\sqrt{5})/2$ . Christodoulou and Koutsoupias [11] and Caragiannis *et al.* [9] proved that the PoS for congestion games with linear delays is  $1+\sqrt{3}/3$ . Recently, Christodoulou *et al.* [13] obtained tight bounds on the PoA and the PoS of approximate PNE for games with linear delays.

*Convergence Time to Pure Nash Equilibria.* The existence of a potential function implies that a PNE is reached in a natural way when the players iteratively select strategies that improve on their individual cost, given the strategies of the rest. Nevertheless, this may take an exponential number of steps, since computing a PNE is PLS-complete for symmetric congestion games and for asymmetric network games with linear delays [14,1]. In fact, the proofs of [14,1] establish the existence of instances where any sequence of players' improvement moves is exponentially long.

A natural approach to circumvent the strong negative results of [14,1] is to resort to *approximate* PNE, where no player can *significantly* improve her individual cost by changing her strategy. Chien and Sinclair [10] considered symmetric congestion games with a weak restriction on the delay functions, and proved that several natural families of sequences of significant improvement moves converge to an approximate PNE in polynomial time. On the other hand, Skopalik and Vöcking [21] proved that computing an approximate PNE for asymmetric congestion games is PLS-complete, and that even with the restriction of [10] on the delay functions, there are instances where any

sequence of significant improvement moves leading to an approximate PNE is exponentially long. However, Awerbuch *et al.* [6] showed that for unweighted games with polynomial delays and for weighted games with linear delays, many families of sequences of significant improvement moves reach an approximately optimal configuration in polynomial time, where the approximation ratio is arbitrarily close to the PoA.

*Social Ignorance in Congestion Games.* Most of the recent work on congestion games focuses on the full information setting, where each player knows the precise weights and the actual strategies of all players, and her strategy selection takes all this information into account. In many typical applications of congestion games however, the players have incomplete information not only about the weights and the strategies, but also about the mere existence of (some of) the players with whom they compete for resources<sup>1</sup> (see also e.g. [16][18][17][8]). In fact, in many applications, it is both natural and convenient to assume that there is a *social context* associated with the game, which essentially determines the information available to the players. In particular, one may assume that each player has complete information about the set of players in her *social neighborhood*, and limited (if any) information about the remaining players.

Our motivation is to investigate how such social-context-related information considerations affect the inefficiency of PNE and the convergence rate to approximate PNE. To come up with a manageable setting that allows for some concrete answers, we make the simplifying assumption that each player has complete information about the players in her social neighborhood, and no information whatsoever about the remaining players. Therefore, since each player is not aware of the players outside her social neighborhood, her individual cost and her strategy selection are not affected by them. In fact, this is the model of *graphical congestion games*, introduced by Bilò, Fanelli, Flammini, and Moscardelli [7]. The new ingredient in the definition of graphical congestion games is the *social graph*, which represents the players' social context. The social graph is defined on the set of players and contains an edge between each pair of players that know each other. The basic idea (and assumption) behind graphical congestion games is that the individual *presumed cost* of each player only depends on the players in her social neighborhood, and thus her strategy selection is only affected by them.

Bilò *et al.* [7] considered unweighted graphical congestion games, and proved that such games with linear delays and undirected social graphs admit a potential function, and thus a PNE. For unweighted linear graphical games, Bilò *et al.* proved that the PoS is at most  $n$ , and the PoA is at most  $n(\deg_{\max} + 1)$ , where  $n$  is the number of players and  $\deg_{\max}$  is the maximum degree of the social graph, and presented certain families of instances for which these bounds are tight. To the best of our understanding, the fact that these bounds are tight for some instances illustrates that expressing the PoA and the PoS as functions of  $n$  and  $\deg_{\max}$  only does not provide an accurate picture of the impact of social ignorance (see also [7] Section 1.2). In particular, the bound on the PoA conveys the message that the more the players know (or learn) about other players, the worse the PoA becomes, and fails to capture that as the social graph tends to the complete graph, the PoA should become a small constant that tends to  $5/2$ .

<sup>1</sup> In many applications, information considerations have to do not only with what the players actually know or are able to learn about the game, but also with how much information they are able or willing to handle in their strategy selection process.



**Contribution.** Adopting graphical linear congestion games as a model, we investigate whether there is a natural parameter of the social graph that completely characterizes the impact of social ignorance on the inefficiency of PNE and on the convergence rate of the  $\varepsilon$ -Nash dynamics. We restrict our attention to graphical linear games with *undirected* social graphs. We consider *weighted players*, so as to investigate the impact of different weights (i.e. could the PoA and the PoS become worse, and if yes, by how much, when many “small” players ignore a few “large” ones compared against the same social situation with all players having the same weight?). With a single exception, the PoA (resp. PoS) is defined with respect to the *actual total cost* of the worst (resp. best) PNE, while equilibria are defined with respect to the players’ *presumed cost*, which is an underestimation of their actual individual cost due to limited social knowledge.

We prove that the impact of social ignorance on the PoA and the PoS is naturally quantified by the *independence number*  $\alpha(G)$  of the social graph  $G$ , i.e. by the cardinality of the largest set of players that do not know each other. In particular, we show that the PoA is at most  $\alpha(G)(\alpha(G) + 2)$ , which is essentially tight as long as  $\alpha(G) \leq n/2$ , and that the PoS lies between  $\alpha(G)$  and  $2\alpha(G)$ .

More specifically, we first show that graphical linear games with weighted players admit a potential function, and thus a PNE (Theorem 1). Our potential function generalizes the potential function of [15, Theorem 3.2], where the social graph is complete, and the potential function of [7, Theorem 1], where the players are unweighted.

To bound the inefficiency of PNE, we show that the total actual cost in any configuration is an  $\alpha(G)$ -approximation of the total presumed cost in the same configuration (Lemma 1). Then, we prove that the PoA of any graphical linear congestion game with weighted players is at most  $\alpha(G)(\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})/2$ , which varies from  $(3 + \sqrt{5})/2$ , when the social graph is complete, to  $\alpha(G)(\alpha(G) + 2)$ , when  $\alpha(G)$  is large (Theorem 2). This bound is essentially tight, even for unweighted players, as long as  $\alpha(G) \leq n/2$  (Theorem 3). For games with unweighted players, we show that the PoA is also bounded from above by  $2n(n - \alpha(G) + 1)$  (Theorem 4), which is tight (up to a small constant factor) when  $\alpha(G) \geq n/2$  (Theorem 5). Furthermore, we prove that the upper bound of  $\alpha(G)(\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})/2$  remains valid if the PoA is calculated with respect to the total presumed cost (Theorem 6), and that this bound is essentially tight as long as  $\alpha(G) \leq \sqrt{n/2}$  (Theorem 7). As for the PoS, we prove that it is at most  $\frac{2n\alpha(G)}{n + \alpha(G)}$  (Theorem 8) and at least  $\alpha(G) - \varepsilon$ , for any  $\varepsilon > 0$  (Theorem 9).

It is rather surprising that the upper bounds on the PoA and the PoS only depend on the *cardinality* of the largest set of players that do not know each other, not on their weights. In addition, the fact that all our lower bounds are established for the case of unweighted players implies that as long as the worst-case PoA and PoS are concerned, considering players with different weights does not make things worse.

As for the convergence time to approximately optimal configurations, we show that it does not directly depend on the structure of the social graph, only the approximation ratio does. In particular, using the techniques of Awerbuch *et al.* [6], we show that the *largest improvement*  $\varepsilon$ -Nash dynamics reaches an approximately optimal configuration

---

<sup>2</sup> Moreover, we can show that the PoA for unweighted linear games is at most  $\frac{3\alpha(G)+7}{3\alpha(G)+1}\alpha^2(G)$ . We omit the details, since this bound is better than the one stated only if  $\alpha(G) \in \{1, 2\}$ .

in a polynomial number of steps (Theorem 10). The approximation ratio is arbitrarily close to the PoA, so it is roughly  $\alpha(G)(\alpha(G) + 2)$ , while the convergence time is linear in  $n$  and in the logarithm of the initial potential value, with the only dependence on the social graph hidden in the latter term. For graphical linear games with unweighted players and symmetric strategies, we use the techniques of Chien and Sinclair [10], and show that the largest improvement  $\varepsilon$ -Nash dynamics converges to an  $\varepsilon$ -PNE in a polynomial number of steps (Theorem 11). Compared to the bound of [10, Theorem 3.1] for symmetric linear congestion games, the convergence time increases by a factor up to  $n$  due to social ignorance. Both results can be extended to the unrestricted  $\varepsilon$ -Nash dynamics, which proceeds in rounds of bounded length, and the only requirement is that each player gets a chance to move in every round.

A subtle point about our results is that they refer to a *static* social information context, an assumption questionable in many settings. This is especially true for the results on the convergence rate of the  $\varepsilon$ -Nash dynamics, since during the convergence process, players can become aware of some initially unknown players. However, our results convey the message that the more the available social information, the better the situation becomes due to the fact that  $\alpha(G)$  tends to decrease. So as the players increase their social neighborhood, the  $\varepsilon$ -Nash dynamics keeps reaching better and better configurations. The entire process takes a polynomial number of steps, since there are  $O(n^2)$  edges to be added to the social graph, and for each fixed social graph, the  $\varepsilon$ -Nash dynamics reaches a configuration with the desired properties in polynomially many steps.

**Other Related Work.** To the best of our knowledge, Gairing, Monien, and Tiemann [16] were the first to investigate the impact of incomplete social knowledge on the basic properties of weighted congestion games. They adopted a Bayesian approach, and mostly focused on parallel-link games.

Our information model can be regarded as a simplified version of the information model considered by Koutsoupias, Panagopoulou, and Spirakis [18]. Their model is based on a directed social graph, where each player knows the precise weights of the players in her social neighborhood, and only a probability distribution for the weights of the rest. Koutsoupias *et al.* obtained upper and lower bounds on the PoA for a very simple game with just two identical parallel links.

An alternative information model was suggested by Karakostas *et al.* [17]. In [17], a fraction of the players are totally ignorant of the presence of other players, and thus oblivious to the resource congestion when selecting their strategies, while the remaining players have full knowledge. Karakostas *et al.* considered non-atomic congestion games, and investigated how the PoA depends on the fraction of ignorant players.

After introducing graphical congestion games in [7], Bilò *et al.* [8] considered the PoA and the PoS of graphical multicast cost sharing games, and proved that one can dramatically decrease the PoA by enforcing a carefully selected social graph.

In an orthogonal approach, Ashlagi, Krysta, and Tennenholtz [4] associated the social graph not with the information available to the players, but with their individual cost. They suggested that the individual cost of each player is given by an *aggregation function* of the delays of the players in her social neighborhood (including herself). The aggregation function is also part of the social context, since it represents the players' attitude towards their neighbors.



## 2 Model and Preliminaries

For an integer  $k \geq 1$ , we let  $[k] \equiv \{1, \dots, k\}$ . For a vector  $x = (x_1, \dots, x_n)$ , we denote  $x_{-i} \equiv (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$  and  $(x_{-i}, x'_i) \equiv (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$ .

**Weighted Congestion Games.** A *congestion game* with weighted players is defined by a set  $V = [n]$  of players, a positive integer weight  $w_i$  associated with each player  $i$ , a set  $R$  of resources, a strategy space  $\Sigma_i \subseteq 2^R \setminus \{\emptyset\}$  for each player  $i$ , and a non-decreasing delay function  $d_e : \mathbb{N} \mapsto \mathbb{R}_{\geq 0}$  associated with each resource  $e$ . The game is (or the players are) unweighted if  $w_i = 1$  for all  $i \in [n]$ . A congestion game has *symmetric strategies* if all players share a common strategy space. We consider *linear* congestion games, where the delay of each resource  $e$  is given by  $d_e(x) = a_e x + b_e$ ,  $a_e, b_e \geq 0$ .

A *configuration* is a tuple  $s = (s_1, \dots, s_n)$  consisting of a strategy  $s_i \in \Sigma_i$  for each player  $i$ . We let  $s_e = \sum_{i:e \in s_i} w_i$  denote the congestion induced on each resource  $e$  by  $s$ , and let  $c_i(s) = \sum_{e \in s_i} w_i(a_e s_e + b_e)$  denote the (actual) cost of player  $i$  in  $s$ .

**Graphical Congestion Games.** The new ingredient in the definition of graphical congestion games is the *social graph*  $G(V, E)$ , which is defined on the set of players  $V$  and contains an edge  $\{i, j\} \in E$  between each pair of players  $i, j$  that know each other. We consider graphical games with weighted players and simple undirected social graphs.

Given a graphical congestion game with a social graph  $G(V, E)$ , let  $\alpha(G)$  be the independence number of  $G$ , i.e. the cardinality of the largest set of players that do not know each other. For a configuration  $s$  and a resource  $e$ , let  $V_e(s) = \{i \in V : e \in s_i\}$  be the set of players using  $e$  in  $s$ , let  $G_e(s)(V_e(s), E_e(s))$  be the social subgraph of  $G$  induced by  $V_e(s)$ , and let  $n_e(s) = |V_e(s)|$  and  $m_e(s) = |E_e(s)|$ . For each player  $i$  (not necessarily belonging to  $V_e(s)$ ), let  $\Gamma_e^i(s) = \{j \in V_e(s) : \{i, j\} \in E\}$  be  $i$ 's social neighborhood among the players using  $e$  in  $s$ . In any configuration  $s$ , a player  $i$  is aware of a *presumed congestion*  $s_e^i = w_i + \sum_{j \in \Gamma_e^i(s)} w_j$  on each resource  $e$ , and of her *presumed cost*  $p_i(s) = \sum_{e \in s_i} w_i(a_e s_e^i + b_e)$ . We note that the presumed cost coincides with the actual cost if the social graph is complete.

For graphical congestion games, a configuration  $s$  is a *pure Nash equilibrium* (PNE) if no player can improve her *presumed cost* by unilaterally changing her strategy. Formally,  $s$  is a PNE if for every player  $i$  and all strategies  $\sigma_i \in \Sigma_i$ ,  $p_i(s) \leq p_i(s_{-i}, \sigma_i)$ .

**Social Cost, Price of Anarchy, and Price of Stability.** We evaluate configurations using the objective of total (actual) cost. The *total cost*  $C(s)$  of a configuration  $s$  is the sum of players' actual costs in  $s$ , i.e.  $C(s) = \sum_{i=1}^n c_i(s) = \sum_{e \in E} (a_e s_e^2 + b_e s_e)$ . The *optimal configuration*, denoted  $o$ , minimizes the total cost among all configurations.

The (pure) *Price of Anarchy* (PoA) of a graphical congestion game  $\mathcal{C}$  is the maximum ratio  $C(s)/C(o)$  over all PNE  $s$  of  $\mathcal{C}$ . The (pure) *Price of Stability* (PoS) of  $\mathcal{C}$  is the minimum ratio  $C(s)/C(o)$  over all PNE  $s$  of  $\mathcal{C}$ .

**Other Notions of Cost.** We also consider the *total presumed cost*  $P(s)$  of a configuration  $s$ , defined as  $P(s) = \sum_{i=1}^n p_i(s)$ . We note that  $P(s) \leq C(s)$ , which holds with equality if the social graph is complete. We let  $o'$  denote the configuration of minimum total presumed cost. The Price of Anarchy with respect to the total presumed cost is the maximum ratio  $P(s)/P(o')$  over all PNE  $s$ .

Moreover, it is helpful to define  $U(s) = \sum_{i=1}^n \sum_{e \in s_i} w_i (a_e w_i + b_e)$ . We note that  $U(s) \leq P(s)$ , which holds with equality if the social graph is an independent set.

**Potential Functions.** A function  $\Phi : \Sigma_1 \times \cdots \times \Sigma_n \mapsto \mathbb{R}_{\geq 0}$  is a *potential function* for a (graphical congestion) game if when a player  $i$  moves from her current strategy in a configuration  $s$  to a new strategy  $s'_i$ , the difference in the potential value equals the difference in the (presumed) cost of player  $i$ , i.e.  $\Phi(s_{-i}, s'_i) - \Phi(s) = p_i(s_{-i}, s'_i) - p_i(s)$ .

**Improvement Moves and Approximate Equilibria.** A strategy  $s'_i$  is a *best response* of a player  $i$  to a configuration  $s$  if for all  $\sigma_i \in \Sigma_i$ ,  $p_i(s_{-i}, s'_i) \leq p_i(s_{-i}, \sigma_i)$ . We let  $\Delta(s) = \sum_{i=1}^n (p_i(s) - p_i(s_{-i}, s'_i))$  denote the sum of the improvements on the presumed cost if each player  $i$  moves from her current strategy in a configuration  $s$  to her best response  $s'_i$ . A strategy  $\sigma_i$  is an *improvement move* of player  $i$  in a configuration  $s$  if  $p_i(s_{-i}, \sigma_i) < p_i(s)$ . Given an  $\varepsilon \in (0, 1)$ , a strategy  $\sigma_i$  is an (improvement)  $\varepsilon$ -*move* of player  $i$  in  $s$  if  $p_i(s_{-i}, \sigma_i) < (1 - \varepsilon)p_i(s)$ , i.e. if player  $i$  moving from her current strategy in  $s$  to  $\sigma_i$  improves her presumed cost by a factor more than  $\varepsilon$ .

For a (graphical congestion) game that admits a potential function, every improvement move decreases the potential. Hence, the Nash dynamics, i.e. any sequence of improvement moves, converges to a PNE in a finite number of steps, and the  $\varepsilon$ -Nash dynamics, i.e. any sequence of  $\varepsilon$ -moves, converges to a *pure Nash  $\varepsilon$ -equilibrium* ( $\varepsilon$ -PNE), i.e. a configuration where no player has an  $\varepsilon$ -move available. Formally, a configuration  $s$  is a  $\varepsilon$ -PNE if for every player  $i$  and all strategies  $\sigma_i \in \Sigma_i$ ,  $(1 - \varepsilon)p_i(s) \leq p_i(s_{-i}, \sigma_i)$ .

### 3 Potential Function and Cost Approximation

**Potential Function.** We first show that graphical linear congestion games with weighted players admit a potential function.

**Theorem 1.** *Every graphical linear congestion game with weighted players admits a potential function, and thus a pure Nash equilibrium.*

*Proof.* Let  $s$  be a configuration of a graphical linear game with weighted players. Then,

$$\Phi(s) = \sum_{e \in R} \left[ a_e \left( \sum_{i \in V_e(s)} w_i^2 + \sum_{\{i,j\} \in E_e(s)} w_i w_j \right) + b_e \sum_{i \in V_e(s)} w_i \right] = \frac{P(s) + U(s)}{2}$$

is a potential function for such a game, since when a player  $i$  switches from her current strategy in  $s$  to a strategy  $s'_i \in \Sigma_i$ ,  $\Phi(s_{-i}, s') - \Phi(s) = p_i(s_{-i}, s') - p_i(s)$ .  $\square$

**Actual Cost vs Presumed Cost.** Next we show that for any configuration  $s$ , the total actual cost  $C(s)$  is an  $\alpha(G)$ -approximation of the total presumed cost  $P(s)$ .

**Lemma 1.** *Let  $G(V, E)$  be the social graph associated with a graphical linear congestion game with weighted players, and let  $\alpha(G)$  be the independence number of  $G$ . Then for any configuration  $s$ ,  $C(s) \leq \alpha(G)P(s)$ .*

*Proof.* For any configuration  $s$  and any resource  $e$ , let  $G_e(s)(V_e(s), E_e(s))$  be the subgraph induced by the players in  $V_e(s)$ , and let  $P_e(s) = \sum_{i \in V_e(s)} w_i (a_e s_e^i + b_e)$  and  $C_e(s) = a_e s_e^2 + b_e s_e$ . It suffices to show that  $C_e(s) \leq \alpha(G_e(s))P_e(s)$ . For simplicity and since we consider a fixed configuration  $s$  throughout this proof, we omit the dependence of the social subgraph and its parameters on  $s$ .

We first establish the lemma for unweighted players, and then reduce the weighted case to the unweighted one. For the unweighted case, it suffices to show a lower bound on the number of edges  $m_e$ , since  $P_e(s) = a_e(2m_e + n_e) + b_en_e$ , i.e.  $P_e(s)$  only depends on the number of edges in  $G_e$ , not on which players are connected by them.

We let  $k = n_e/\alpha(G_e)$  and  $r = k - \lfloor k \rfloor$  (resp.  $k - r$ ) be  $k$ 's fractional (resp. integral) part. We partition the vertices of  $G_e$  into a sequence of at least  $\lceil k \rceil$  independent sets of non-increasing cardinality as follows: We begin with  $\ell = 1$  and  $G_e^{(1)} = G_e$ . As long as  $G_e^{(\ell)}$  is non-empty, we find a maximum independent set  $I_\ell$  of  $G_e^{(\ell)}$ , obtain the next graph  $G_e^{(\ell+1)}$  by removing the vertices of  $I_\ell$  from  $G_e^{(\ell)}$ , increase  $\ell$  by one, and iterate.

Let  $q \geq \lceil k \rceil$  be the number of independent sets obtained by the decomposition above. Since  $I_\ell$  is a maximum independent set of  $G_e^{(\ell)}$ , for every  $j = \ell + 1, \dots, q$ , each vertex  $u \in I_j$  is connected by an edge to at least one vertex in  $I_\ell$ . Otherwise,  $I_\ell \cup \{u\}$  would be an independent set larger than  $I_\ell$ . Hence, for every  $j = 2, \dots, q$ , each vertex  $u \in I_j$  is incident to at least  $j - 1$  edges connecting it to vertices in the independent sets  $I_1, \dots, I_{j-1}$ , and  $m_e \geq \sum_{j=2}^q (j - 1)|I_j|$ . Since  $\sum_{j=2}^q |I_j| = n_e - \alpha(G_e)$  and  $|I_j| \leq \alpha(G_e)$  for all  $j \in [q]$ , the lower bound on the total number of edges is minimized when  $q = \lceil k \rceil$ ,  $|I_j| = \alpha(G_e)$  for all  $j = \{1, \dots, \lceil k \rceil\}$ , and  $|I_{\lceil k \rceil}| = r\alpha(G_e)$ . Therefore,

$$m_e \geq \frac{(k - r)(k - r - 1)}{2}\alpha(G_e) + (k - r)r\alpha(G_e) = \frac{(k - r)(k + r - 1)}{2}\alpha(G_e)$$

Using that  $r \geq r^2$ , since  $r \in [0, 1)$ , we conclude the proof for the unweighted case:

$$\begin{aligned} \alpha(G_e)P_e(s) &\geq a_e[(k - r)(k + r - 1) + k]\alpha^2(G_e) + b_ek\alpha^2(G_e) \\ &\geq a_ek^2\alpha^2(G_e) + b_ek\alpha(G_e) = C_e(s) \end{aligned}$$

For weighted players, we create a new graph  $G'_e$  by replacing each player  $i \in V_e$  with a clique  $Q_i$  of  $w_i$  unweighted players that use the same strategy  $s_i$ . For each edge  $\{i, j\} \in E_e$ , we add  $w_iw_j$  edges to  $G'_e$  that connect every vertex in the clique  $Q_i$  to every vertex in the clique  $Q_j$ . The lemma follows since  $G'_e$  consists of unweighted players,  $\alpha(G_e) = \alpha(G'_e)$ , and the transformation does not affect  $P_e(s)$  and  $C_e(s)$ .  $\square$

### 4 The Price of Anarchy and the Price of Stability

The following lemma is useful both in bounding the PoA and in establishing the fast convergence of the  $\varepsilon$ -Nash dynamics to approximately optimal configurations.

**Lemma 2.** *For any configuration  $s$  of a graphical linear congestion game with weighted players arranged in a social graph  $G$ ,*

$$P(s) \leq \frac{\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)}}{2}C(o) + 2\Delta(s) \tag{1}$$

*Proof.* We follow the general approach of [6, Lemma 4.1] and [5, Theorem 3.1]. The presumed cost of each player  $i$  if she switches to her best response strategy  $s'_i$  is at most her presumed cost if she switches to her optimal strategy  $o_i$ . Thus,

$$p_i(s_{-i}, s'_i) \leq p_i(s_{-i}, o_i) \leq \sum_{e \in o_i} w_i(a_e(s_e + w_i) + b_e)$$

Summing up over all players, and using the Cauchy-Schwarz inequality and Lemma 1 we obtain that:

$$\sum_{i=1}^n p_i(s_{-i}, s'_i) \leq \sqrt{C(o)C(s)} + C(o) \leq \sqrt{C(o)\alpha(G)P(s)} + C(o) \quad (2)$$

Adding  $\Delta(s)$  to both sides of (2) and dividing by  $C(o)$ , we obtain that :

$$P(s)/C(o) \leq \sqrt{\alpha(G)}\sqrt{P(s)/C(o)} + 1 + \Delta(s)/C(o) \quad (3)$$

Setting  $\beta = \sqrt{P(s)/C(o)}$  and  $\gamma = \Delta(s)/C(o)$ , (3) becomes  $\beta^2 \leq \sqrt{\alpha(G)}\beta + 1 + \gamma$ . By simple algebra, we obtain that  $2\beta \leq \sqrt{\alpha(G)} + \sqrt{\alpha(G) + 4(1 + \gamma)}$ . The lemma follows by taking the squares of both sides of the previous inequality and observing that  $\sqrt{\alpha^2(G) + 4\alpha(G)(1 + \gamma)} \leq \sqrt{\alpha^2(G) + 4\alpha(G) + 2\gamma}$ .  $\square$

**The Price of Anarchy.** In Lemma 2 when  $s$  is a PNE, we get  $\Delta(s) = 0$ . Therefore, Lemma 1 and Lemma 2 immediately imply the following upper bound on the PoA.

**Theorem 2.** *For graphical linear congestion games with weighted players, the Price of Anarchy is at most  $\alpha(G)(\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})/2$ .*

Therefore, the PoA of graphical linear games with weighted players is less than  $\alpha(G)(\alpha(G) + 2)$ . Bilò *et al.* [7, Theorem 13] present a family of unweighted graphical games with  $n$  players arranged in a bipartite social graph  $G$  with  $\alpha(G) = n/2$ , for which the PoA can be as large as  $\alpha^2(G)$ . Next we present a different family of unweighted graphical games for which the ratio  $\alpha(G)/n$  can take any value in  $(0, \frac{1}{2}]$  and the PoA is at least  $\alpha(G)(\alpha(G) + 1)$ . This implies that as long as  $\alpha(G) \leq n/2$ , the upper bound of Theorem 2 is essentially tight.

**Theorem 3.** *For any integers  $\ell \geq 1$  and  $n \geq 2\ell$ , there is a graphical linear congestion game with  $n$  unweighted players arranged in a social graph  $G$  with  $\alpha(G) = \ell$ , for which the PoA is  $\ell(\ell + 1)$ .*

*Proof sketch.* For simplicity, we focus on the case where  $k = n/\ell$  is an integer. The social graph  $G$  is the complete  $k$ -partite graph with  $\ell$  vertices in each part. There are  $k(\ell + 1)$  resources  $e_i^j, j \in [k], i \in \{0\} \cup [\ell]$ , with delay function  $d(x) = x$ .

Each player has two strategies, the “short” and the “long” one. In particular, for the  $i$ -th player in the  $j$ -th part, the “short” strategy is  $\{e_i^j\}$ , and the “long” strategy is  $\{e_0^{(j \bmod k)+1}, e_1^{(j \bmod k)+1}, \dots, e_\ell^{(j \bmod k)+1}\}$ . The optimal configuration  $o$  assigns all players to their “short” strategies and has  $C(o) = k\ell$ . On the other hand, there is a PNE  $s$  with  $C(s) = k\ell^2(\ell + 1)$  where all players use their “long” strategies.  $\square$

Next we establish a stronger upper bound on the PoA of graphical linear games with unweighted players and social graphs with a very large independence number, and show that this bound is tight (up to a small constant factor) when  $\alpha(G) \geq n/2$ . The proofs of the following two theorems are omitted due to lack of space.

**Theorem 4.** *For graphical linear congestion games with  $n$  unweighted players, the Price of Anarchy is at most  $2n(n - \alpha(G) + 1)$ .*

**Theorem 5.** *For any integers  $\ell \geq 1$  and  $\ell \leq n \leq 2\ell$ , there is a graphical linear game with  $n$  unweighted players arranged in a social graph  $G$  with  $\alpha(G) = \ell$ , for which the Price of Anarchy is at least  $\ell(n - \ell + 1) - (n - \ell)$ .*

In addition, we observe that Lemma 1 and Lemma 2 imply that the bound of Theorem 2 remains valid if the PoA is calculated with respect to the total presumed cost.

**Theorem 6.** *For a graphical congestion game with weighted players, let  $o'$  be the configuration that minimizes the total presumed cost, and let  $s$  be any PNE. Then,*

$$P(s)/P(o') \leq \alpha(G)(\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})/2$$

Moreover, we show that the bound of Theorem 6 is essentially tight for social graphs  $G$  with  $\alpha(G) \leq \sqrt{n/2}$ , even for unweighted players. We emphasize that such a lower bound is best possible, since the PoA with respect to the total presumed cost is at most  $n$  (see e.g. [7] Theorem 9), which can be easily generalized to the weighted case).

**Theorem 7.** *For any integers  $\ell \geq 1$  and  $n \geq 2\ell^2$ , there is a graphical linear congestion game with  $n$  unweighted players arranged in a social graph  $G$  with  $\alpha(G) = \ell$ , for which the Price of Anarchy with respect to the total presumed cost is  $\ell^2$ .*

*Proof sketch.* For simplicity, we focus on the case where  $k = n/\ell^2$  is an integer. The social graph  $G$  consists of  $k$  groups  $G_i$ ,  $i \in [k]$ , where each  $G_i$  consists of  $\ell$  disjoint independent sets  $I_i^j$ ,  $j \in [\ell]$ , with  $\ell$  vertices each. The vertices within each part  $G_i$  are also partitioned into  $\ell$  cliques of cardinality  $\ell$ , with each clique including one vertex from each independent set  $I_i^j$ ,  $j \in [\ell]$ . The edges between the vertices in the same clique are the only edges between vertices in the same group. All pairs of vertices from different groups are connected to each other by an edge. There are  $k\ell$  resources  $e_i^j$ ,  $i \in [k]$ ,  $j \in [\ell]$ , one for each independent set  $I_i^j$ , with delay function  $d(x) = x$ .

Each player has two strategies, the “short” and the “long” one. More specifically, for a player in the independent set  $I_i^j$ , the “short” strategy is  $\{e_i^j\}$ , and the “long” strategy is  $\{e_{(i \bmod k)+1}^1, \dots, e_{(i \bmod k)+1}^\ell\}$ . The configuration  $o'$  of minimum total presumed cost assigns each player to her “short” strategy and has  $P(o') = k\ell^2$ . On the other hand, there is a PNE  $s$  with  $P(s) = k\ell^4$  where all players use their “long” strategies.  $\square$

**The Price of Stability.** An upper bound on the PoS follows easily from the potential function of Theorem 1 and Lemma 1

**Theorem 8.** *For graphical linear congestion games with  $n$  weighted players, the Price of Stability is at most  $\frac{2n\alpha(G)}{n+\alpha(G)}$ .*

*Proof sketch.* Let  $s$  be a minimizer of the potential function  $\Phi$ . Clearly,  $s$  is a PNE. We observe that  $C(o) \geq \Phi(o) \geq \Phi(s) = P(s)/2 + U(s)/2$ . The lemma follows by observing that (i)  $P(s) \geq C(s)/\alpha(G)$ , by Lemma 1 and (ii)  $U(s) \geq C(s)/n$ , by the Cauchy-Schwarz inequality.  $\square$

The following theorem, whose proof is omitted due to lack of space, shows that the upper bound of Theorem 8 is essentially tight.

**Theorem 9.** *For any positive integers  $\ell$  and  $n \geq \ell$ , and any  $\varepsilon > 0$ , there is a graphical linear congestion game with  $n$  unweighted players arranged in a social graph  $G$  with  $\alpha(G) = \ell$ , for which the Price of Stability is  $\ell - \varepsilon$ .*

## 5 Convergence Rate of the $\varepsilon$ -Nash Dynamics

**Convergence to Near Optimal Configurations.** Employing the techniques of Awerbuch *et al.* [6], we show that the *largest improvement*  $\varepsilon$ -Nash dynamics reaches an approximately optimal configuration in a polynomial number of steps, where the approximation ratio is arbitrarily close to the PoA of the graphical congestion game. In each step of the largest improvement  $\varepsilon$ -Nash dynamics, among all players with an  $\varepsilon$ -move available, the player with the largest improvement on his presumed cost moves. The proof of the following theorem is omitted due to lack of space.

**Theorem 10.** *For a graphical linear game with  $n$  weighted players arranged in a social graph  $G$ , let  $s^*$  be a minimizer of the potential function  $\Phi$ , and let  $\frac{1}{8} \geq \delta \geq \varepsilon > 0$ . Starting from a configuration  $s_0$ , the largest improvement  $\varepsilon$ -Nash dynamics reaches a configuration  $s$  with  $C(s) \leq \frac{\alpha(G)}{2}(\alpha(G) + 2 + \sqrt{\alpha^2(G) + 4\alpha(G)})(1 + 8\delta)C(o)$  in  $O(\frac{n}{\delta} \log \frac{\Phi(s_0)}{\Phi(s^*)})$  steps.*

Following the approach of [6, Theorem 3.3], we can establish a similar convergence time for the unrestricted  $\varepsilon$ -Nash dynamics, which proceeds in rounds of bounded length, and the only requirement is that each player gets a chance to move in every round.

**Convergence to Approximate Equilibria.** For graphical linear games with unweighted players and symmetric strategies, we employ the techniques of Chien and Sinclair [10] and show that the largest improvement  $\varepsilon$ -Nash dynamics converges to an  $\varepsilon$ -PNE in polynomial time. The proof of the following theorem is omitted due to lack of space.

**Theorem 11.** *For a graphical linear congestion game with symmetric strategies and  $n$  unweighted players, let  $s^*$  be a minimizer of the potential function  $\Phi$ , and let  $\varepsilon \in (1, 0)$ . Starting from a configuration  $s_0$ , the largest improvement  $\varepsilon$ -Nash dynamics converges in  $O(\frac{n^2}{\varepsilon} \log \frac{\Phi(s_0)}{\Phi(s^*)})$  steps.*

Following the approach of [10, Theorem 4.1], we can establish a similar convergence time to an  $\varepsilon$ -PNE for the unrestricted  $\varepsilon$ -Nash dynamics, where the only requirement is that each player gets a chance to move in every round.

**Acknowledgements.** The first author thanks Piotr Krysta for helpful discussions that greatly contributed to shaping his view on the topic of this work.

## References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the Impact of Combinatorial Structure on Congestion Games. *Journal of the ACM* 55(6) (2008)
2. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact Price of Anarchy for Polynomial Congestion Games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 218–229. Springer, Heidelberg (2006)
3. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The Price of Stability for Network Design with Fair Cost Allocation. In: Proc. of the 45th IEEE Symp. on Foundations of Computer Science (FOCS 2004), pp. 295–304 (2004)



4. Ashlagi, I., Krysta, P., Tennenholtz, M.: Social Context Games. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 675–683. Springer, Heidelberg (2008)
5. Awerbuch, B., Azar, Y., Epstein, A.: The Price of Routing Unsplittable Flow. In: Proc. of the 37th ACM Symp. on Theory of Computing (STOC 2005), pp. 57–66 (2005)
6. Awerbuch, B., Azar, Y., Epstein, A., Mirrokni, V., Skopalik, A.: Fast Convergence to Nearly Optimal Solutions in Potential Games. In: Proc. of the 9th ACM Conf. on Electronic Commerce (EC 2008), pp. 264–273 (2008)
7. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: Graphical Congestion Games. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 70–81. Springer, Heidelberg (2008)
8. Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: When Ignorance Helps: Graphical Multicast Cost Sharing Games. In: Ochmański, E., Tyszkiewicz, J. (eds.) MFCS 2008. LNCS, vol. 5162, pp. 108–119. Springer, Heidelberg (2008)
9. 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)
10. Chien, S., Sinclair, A.: Convergece to Approximate Nash Equilibria in Congestion Games. In: Proc. of the 18th Symp. on Discrete Algorithms (SODA 2007), pp. 169–178 (2007)
11. 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)
12. Christodoulou, G., Koutsoupias, E.: The Price of Anarchy of Finite Congestion Games. In: Proc. of the 37th ACM Symp. on Theory of Computing (STOC 2005), pp. 67–73 (2005)
13. Christodoulou, G., Koutsoupias, E., Spirakis, P.: On the Performance of Approximate Equilibria in Congestion Games. In: Proc. of the 17th European Symposium on Algorithms (ESA 2009). LNCS, vol. 5757. Springer, Heidelberg (2009)
14. Fabrikant, A., Papadimitriou, C., Talwar, K.: The Complexity of Pure Nash Equilibria. In: Proc. of the 36th ACM Symp. on Theory of Computing (STOC 2004), pp. 604–612 (2004)
15. Fotakis, D., Kontogiannis, S., Spirakis, P.: Selfish Unsplittable Flows. *Theoretical Computer Science* 348, 226–239 (2005)
16. Gairing, M., Monien, B., Tiemann, K.: Selfish Routing with Incomplete Information. *Theory of Computing Systems* 42, 91–130 (2008)
17. Karakostas, G., Kim, T., Viglas, A., Xia, H.: Selfish Routing with Oblivious Users. In: Prencipe, G., Zaks, S. (eds.) SIROCCO 2007. LNCS, vol. 4474, pp. 318–327. Springer, Heidelberg (2007)
18. Koutsoupias, E., Panagopoulou, P., Spirakis, P.: Selfish Load Balancing Under Partial Knowledge. In: Kučera, L., Kučera, A. (eds.) MFCS 2007. LNCS, vol. 4708, pp. 609–620. Springer, Heidelberg (2007)
19. Koutsoupias, E., Papadimitriou, C.: Worst-case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
20. Rosenthal, R.W.: A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
21. Skopalik, A., Vöcking, B.: Inapproximability of Pure Nash Equilibria. In: Proc. of the 40th ACM Symp. on Theory of Computing (STOC 2008), pp. 355–364 (2008)

# The Complexity of Models of International Trade

Christopher A. Wilkens\*

Computer Science Division,  
University of California at Berkeley, CA, 94720 USA  
cwilkens@cs.berkeley.edu

**Abstract.** We show a range of complexity results for the Ricardo and Heckscher-Ohlin models of international trade (as Arrow-Debreu production markets). For both models, we show three types of results:

1. When utility functions are Leontief and production functions are linear, it is NP-hard to decide if a market has an equilibrium.
2. When utility functions and production functions are linear, equilibria are efficiently computable (which was already known for Ricardo).
3. When utility functions are Leontief, equilibria are still efficiently computable when the diversity of producers and inputs is limited.

Our proofs are based on a general reduction between production and exchange equilibria. One interesting byproduct of our work is a generalization of Ricardo's Law of Comparative Advantage to more than two countries, a fact that does not seem to have been observed in the Economics literature.

## 1 Introduction

How does production in an economy affect the computability of equilibria? A wave of research has shown a broad spectrum of results for pure exchange economies (e.g. [5,6,2]); however, only a handful of papers approach equilibria in the presence of production (e.g. [11,8,9]). The papers that do consider production typically construct sophisticated algorithms to compute equilibria, and they do not present negative results.

We take a different approach: in the spirit of Jain and Mahdian's reduction for the Fisher market [7], we reduce production economies to exchange economies. The reduction yields a variety of complexity results for two classical models of trade: the Ricardo model and the Heckscher-Ohlin model. Mathematically, both are special cases of the Arrow-Debreu production market [1]. Economists use them because they represent different motivations for international trade: differentiation in production technology and differentiation in raw materials. For our purposes, their formulations are conveniently simple: the Ricardo model uses linear production functions with a single raw material, and the Heckscher-Ohlin model specifies that agents have identical production functions.

---

\* Supported by a fellowship from the University of California at Berkeley and NSF ITR grants CCR-0121555 and CCF-0515259.



Our reductions will leverage the plentiful literature on computing equilibria in pure exchange economies. The simplest results reduce the Ricardo and Heckscher-Ohlin models to exchange economies with linear utilities. A wide variety of algorithms already exist for this case — for example, Devanur et al. use a primal dual approach [5] and Garg and Kapoor use an auction algorithm [6].

Our hardness results are based on the NP-hardness result of Codenotti et al. for pure exchange economies [2]. Revisiting their proof yields a convenient tool for showing NP-hardness for our production economies. It is noteworthy that the production model need not be complicated — we show that linear production functions with a single production input suffice to preserve hardness.

Our most interesting computational result is that equilibria may become easier to compute when there are only a few types of producers or raw materials. Devanur and Kannan [4] show that for exchange markets, equilibria in a Leontief exchange economy become easier to compute when there are few goods or agents. We use their result to show that equilibria in the Ricardo model are easy to compute when there are few types of producers. Mathematically, this translates to a type of low-rank constraint on the production coefficients in the Ricardo economy. Similarly, for the Heckscher-Ohlin model, we show that equilibria are efficiently computable under Leontief utilities and production functions independent of the number of goods, provided the number of raw materials is small.

The previous two results are interesting in a broader context because real economies have less variation in technologies and raw materials than they do in consumers, goods, and preferences. For example, typical uses of the Heckscher-Ohlin model [12] employ very few raw materials: labor, land, capital, etc. Thus, our results concern economies which may be closer to reality or to patterns studied by economists.

Our complexity results are summarized in Table 1.

**Table 1.** A summary of the results in this paper

Model	Production	Utilities	Complexity	Note
Ricardo	Linear	Leontief	NP-hard	Already known, e.g. [11] Similar producers
		Linear	P	
		Leontief	P	
Heckscher-Ohlin	Linear	Leontief	NP-hard	$O(1)$ raw materials
	Linear	Linear	P	
	Leontief	Leontief	P	

As a bonus, we encounter a novel generalization of a classical theorem of economics: Ricardo’s law of comparative advantage. This law states that each of two trading countries will specialize in the production of goods for which its relative labor efficiency is larger, with the ratio of the equilibrium price of labor (wage) as the cut-off point. We establish a multi-dimensional generalization (from the interval  $[0, 1]$  to the simplex, see Figure 1).

## 2 Markets, Equilibria, and Production

We will use four special cases of Arrow and Debreu's market model [1]: the exchange economy, the pairing exchange economy, the Ricardo production economy, and the Heckscher-Ohlin production economy.

### 2.1 Exchange Economies

An exchange economy consists of  $n$  agents and  $m$  divisible goods. Each agent  $i$  is initially given  $e_{ij}$  units of good  $j$  and has a utility function  $u_i(x)$  mapping a bundle of goods  $x = \{x_1, \dots, x_m\}$  to a nonnegative utility. Agents trade goods to improve their utilities.

We will use both linear and Leontief utilities in this paper. Linear utility functions take the form

$$u_i(x_1, \dots, x_m) = \sum_j \phi_{ij} \cdot x_j .$$

A player with Leontief utilities desires goods in fixed proportions. The utility functions take the form

$$u_i(x_1, \dots, x_m) = \min_j \frac{x_j}{\phi_{ij}} .$$

Let  $\Phi = [\phi_{ij}]$  be the matrix of coefficients  $\phi_{ij}$ .

An equilibrium in an exchange economy is an allocation  $x$  and a set of prices  $\pi$  such that  $x$  maximizes the utility of each agent subject to the budget constraint

$$\sum_j \pi_j \cdot x_{ij} \leq \sum_j \pi_j \cdot e_{ij} .$$

*The Pairing Leontief Economy.* In the pairing model (Ye [13]), agent  $i$  is endowed with exactly 1 unit of good  $i$  and nothing else. When the agents have Leontief utilities, we call it a pairing Leontief economy. Since endowments are fixed, the pairing Leontief economy is completely specified by  $\Phi$ .

Codenotti et al. used pairing Leontief economies to show that it is NP-hard to decide whether a general Leontief exchange economy has an equilibrium [2]. In fact, the pairing constraint is not violated by their proof, yielding:

**Theorem 1.** (Derived from Codenotti et al. [2]) *It is NP-hard to decide whether a pairing Leontief exchange economy has an equilibrium.* (Proof omitted.)

### 2.2 Production Economies

We will restrict Arrow and Debreu's production model. We say that each agent  $i$  has one production function  $f_{ij}$  for each tradable good  $j$  (of  $m$  total). Each function  $f_{ij}(l)$  maps a bundle  $l$  of  $K$  non-tradable raw materials (indexed by  $k$ ) to  $f_{ij}(l)$  units of the  $j$ -th good. An agent is endowed with a bundle of raw materials

$l_i$  (for which he has no utility). For our purposes, the production functions will be either linear or Leontief, parameterized by coefficients  $a_{ij}$  with matrix form  $A = [a_{ij}]$ . As before, each agent has a utility function  $u_i$ .

Such a production economy may be understood to operate in two stages. First, agent  $i$  chooses a production plan to turn his endowment  $l_i$  into a bundle of tradable goods  $x_i$  using the functions  $f_{ij}$ . Second, the agents exchange goods as in an exchange economy.

We will use  $w_{ik}$  to denote the effective price of raw material  $k$  of agent  $i$ .

*The Ricardo Model.* The special case with a single raw material ( $K = 1$ ) and linear production technologies was used by economist David Ricardo and is commonly known as the Ricardo model. In this restricted setting, the production functions take the form

$$f_{ij}(l) = a_{ij} \cdot l$$

where  $l$  is a scalar. We use  $l_i$  to denote the amount of raw material possessed by agent  $i$  and  $w_i$  the price for agent  $i$ 's raw material. Historically, the raw material  $l$  represents labor and the price  $w_i$  represents wages.

*The Heckscher-Ohlin Model.* The case where there are many inputs but production technologies are identical is known as the Heckscher-Ohlin model. In this model, the form of the production functions is not specified.

### 3 The Upside-Down Reduction

Many of our theorems reduce a production economy to an *upside-down* exchange economy. In an upside-down economy, trade precedes production — agents trade raw materials, then produce their optimal bundles given the raw materials they acquire. To preserve the possibilities of the original economy, raw materials retain the production technology of their original agent. As a result, the production functions are absorbed into the utilities, as each player's utility function for a bundle of raw materials  $l$  will be

$$u_i(l) = \max_{x \in X} u_i(x)$$

where  $X$  is the set of all bundles agent  $i$  can produce given  $l$ . This type of reduction was used by Jain and Mahdian in the context of the Fisher model [7], but we use it more broadly.

When the production functions exhibit constant returns to scale, the production possibilities in the upside-down economy are identical to those in the original production economy. Thus, the equilibria are also identical. We use the fact that both linear and Leontief functions exhibit constant returns to scale.

We denote functions and variables in the upside-down economy with a ( $'$ ). In general, an upside-down economy will have  $n' = n$  agents and  $m' = (n \times K)$  goods. (Since raw materials carry technology, the raw materials of two agents

are different goods.) We index goods as  $(ik)$  and use  $x'_{(ik)}$  to refer to an amount of raw material  $k$  that has the production technology of agent  $i$ .

The following lemmas give three cases where the reduction behaves nicely — the Leontief/Leontief, linear/Leontief, and linear/linear cases respectively. The technique is similar, so we only prove the Leontief/Leontief case.

**Lemma 2.** *When all agents have identical production functions, and both production functions and utilities are Leontief, then*

1. *the utility functions in the upside-down economy are also Leontief with easily computable parameters, and*
2. *we can recover equilibrium prices as*

$$\pi_j = \sum_k \frac{\pi(k)}{a_{jk}}$$

*Proof.* Since all agents have identical production functions, there will be  $K$  distinct goods in the upside-down economy.

Consider the behavior of a single agent, Alice, and drop her subscripts for clarity. Let  $x'_{(k)j}$  be the amount of raw material  $k$  that Alice uses to produce good  $j$ . We can write the amount of good  $j$  that Alice produces as

$$\min_k \frac{x'_{(k)j}}{a_{jk}}$$

and Alice’s subsequent utility as

$$u(x') = \min_j \frac{\min_k \frac{x'_{(k)j}}{a_{jk}}}{\phi_j} = \min_k \min_j \frac{x'_{(k)j}}{a_{jk} \cdot \phi_j} .$$

In order to maximize her utility, Alice will distribute each input  $x_{(k)}$  over goods so as to maximize  $\min_j \frac{x'_{(k)j}}{a_{jk} \cdot \phi_j}$ . This will occur when all terms are equal, so we know that

$$\min_j \frac{x'_{(k)j}}{a_{jk} \cdot \phi_j} = \frac{1}{m} \sum_j \frac{x'_{(k)j}}{a_{jk} \cdot \phi_j} = \frac{x'_{(k)}}{m} \sum_j \frac{1}{a_{jk} \cdot \phi_j} .$$

Substituting gives Alice’s utility function:

$$u(x') = \min_k \left( \frac{x'_{(k)}}{m} \sum_j \frac{1}{a_{jk} \cdot \phi_j} \right) .$$

As claimed, this is Leontief. Moreover, the coefficients  $\phi'$  may be easily computed from  $\phi$ ,  $a$ , and  $m$ .

Since there is only one production technology for each good, we can compute the price of good  $j$  as the total cost of the inputs required to make one unit:

$$\pi_j = \sum_k \frac{\pi(k)}{a_{jk}} .$$

□

**Lemma 3.** *When there is a single type of raw material, production functions are linear, utilities are Leontief, and for all goods  $j$  we are told that agents use the raw material of agent  $i_j$  to produce good  $j$ , then*

1. *the utility functions in the upside-down economy are Leontief and easily computable, and*
2. *we can recover equilibrium prices as*

$$\pi_j = \frac{\pi(i_j)}{a_{i_j j}} .$$

**Lemma 4.** (Like Jain and Mahdian with multiple raw materials [7].) *When the production functions and utilities in the production economy are linear, then*

1. *the utility functions in the upside-down economy are linear and easily computable, and*
2. *equilibrium prices  $\pi$  in the original economy may be recovered from equilibrium prices  $\pi'$  in the upside-down economy as*

$$\pi_j = \min_{i,k} \frac{\pi'_{(ik)}}{a_{i j k}} .$$

## 4 Computability in the Ricardo Model

We show a broad range of computational results for the Ricardo model. Computability with linear and Leontief utilities parallels the exchange economy. Interestingly, we find that with Leontief utilities, equilibria are efficiently computable when producers are sufficiently similar.

### Linear Utilities

As a warm-up, we use an upside-down reduction to show that Ricardo equilibria are efficiently computable when the utility functions are linear. (The computability was already known, e.g. the auction algorithm of Kapoor et al. [11].) Note that Jain and Mahdian use the same proof for the Fisher model in [7].

**Theorem 5.** *Equilibria in the Ricardo model are efficiently computable when agents' utility functions are linear.*

*Proof.* The Ricardo model, with linear production functions and one raw material, is a special case of the linear production economy reduced in Lemma 4. Following this lemma, the upside-down counterpart to the linear Ricardo economy has linear utility functions that are efficiently computable from the original utilities. Furthermore, we can recover equilibrium prices from the upside-down equilibrium and use them to compute demands (given prices, it is easy to compute demands under linear utilities.)

To complete the proof, we note that many algorithms exist to compute equilibria in linear exchange economies, e.g. [5,6]. Thus, linear Ricardo equilibria are efficiently computable. □

**Leontief Utilities**

While it seems nontrivial to reduce a general exchange economy to a Ricardo economy, it is easy to reduce a pairing one — this yields the following hardness result:

**Theorem 6.** *It is NP-hard to decide whether a Ricardo model economy with Leontief utility functions has an equilibrium.*

*Proof.* Let  $\Phi$  represent the preferences in a pairing Leontief exchange economy. Observe that choosing  $A = I$  and  $l_i = 1$ ,  $i$ -th country can only produce the  $i$ -th good and can produce at most 1 unit of it. Since it has no value for the raw material, it may be assumed to produce 1 unit of the  $i$ -th good.

Since each agent  $i$  has the same goods for trade and the same utilities in both economies, the equilibria must also be the same. NP-hardness follows from Theorem □

**Similar Producers**

We show that equilibria are efficiently computable when the utility functions are Leontief provided that the producers are similar. Specifically, we will require a low-rank-like requirement on the matrix of production parameters  $A$ .

First, we make the following common observation about the Ricardo model:

**Observation 7.** *In equilibrium, agent  $i$  may produce good  $j$  only if  $\pi_j = \frac{w_i}{a_{ij}} = \min_i \frac{w_i}{a_{ij}}$ . Alternatively, country  $i$  may produce a good if and only if  $\frac{w_i}{w_{i'}} \leq \frac{a_{ij}}{a_{i'j}}$  for all other countries  $i'$ .*

Intuitively, this holds because when  $\pi_i < \frac{w_i}{a_{ij}}$ , then country  $i$  loses by producing good  $j$ . On the other hand, if  $\pi_i > \frac{w_i}{a_{ij}}$ , then a buyer would resist buying and force the price down.

A key insight is that given prices (which may be completely specified by either the  $\pi_i$ 's or the  $w_i$ 's,) the pattern of production is fixed. This will allow us to prove the following lemma decomposing the price space into production patterns:

**Lemma 8.** *In a Ricardo economy with  $n$  producers and  $m$  goods, there are at most  $O(m^{O(n^2)})$  distinct production patterns. Moreover, each production pattern occurs in a convex polytope in the price space.*

*Proof.* Observation □ implies that if

$$\frac{a_{i1}}{a_{i'1}} \geq \dots \geq \frac{a_{ik}}{a_{i'k}} > \frac{w_i}{w_{i'}} > \frac{a_{i(k+1)}}{a_{i'(k+1)}} \geq \dots \geq \frac{a_{im}}{a_{i'm}} ,$$

then agent  $i$  cannot produce any good for which  $\frac{w_i}{w_{i'}} > \frac{a_{ij}}{a_{i'j}}$  while agent  $i'$  cannot produce any good for which  $\frac{a_{ij}}{a_{i'j}} > \frac{w_i}{w_{i'}}$ . Thus, we may give a combinatorial specification of the pattern of production between two countries by specifying where  $\frac{w_i}{w_{i'}}$  appears in the ordering of goods. Note that there are  $(2m + 1)$  possibilities.

Extending this idea to  $n$  agents, we want to show that specifying the pairwise combinatorial production patterns will specify the overall pattern of production. For a given good  $j$ , either

1. There is some agent  $i$  such that  $\frac{a_{ij}}{a_{i'j}} > \frac{w_i}{w_{i'}}$  for all other agents  $i'$ , or
2. there is a cycle of agents  $i, i_2, \dots, i_r$  such that  $\frac{a_{ij}}{a_{i_2j}} > \frac{w_i}{w_{i_2}}, \dots, \frac{a_{i_{r-1}j}}{a_{ij}} > \frac{w_{i_{r-1}}}{w_i}$ .

However, option (2) is impossible: multiplying the first  $r - 1$  inequalities gives  $\frac{a_{ij}}{a_{i_rj}} > \frac{w_i}{w_{i_r}}$ , which contradicts the final inequality. Thus, for each good, there is some producer who can produce it. It follows that specifying all the combinatorial pairwise patterns must specify the overall pattern of production.

Since there are  $O(n^2)$  pairs of agents and  $(2m + 1)$  patterns for each pair, there are at most  $O(m^{O(n^2)})$  different combinatorial characterizations and therefore production patterns in the economy.

Finally, note that the production pattern will be specified by  $O(n^2)$  inequalities of the form

$$\frac{a_{ik}}{a_{i'k}} > \frac{w_i}{w_{i'}} > \frac{a_{i(k+1)}}{a_{i'(k+1)}}$$

or an equality of the form

$$\frac{a_{ik}}{a_{i'k}} = \frac{w_i}{w_{i'}}$$

Each equality/inequality bounds the equilibrium prices between a pair of hyperplanes. The union of the hyperplanes defines the convex polytope in the price space in which this production pattern occurs.  $\square$

**Lemma 9.** *If the rows of the production matrix  $A$  are scalar multiples of  $K = O(1)$  different vectors, then computing equilibria in the Leontief Ricardo economy is as easy as finding equilibria in a Leontief exchange economy with  $K$  goods restricted to a convex polytope in the price space.*

*Proof.* Briefly, the  $K = O(1)$  bound dictates that there are  $K = O(1)$  interesting raw materials. Combined with Lemma 8, we will conclude that there are a polynomial number of distinct production patterns. This permits an upside-down reduction for each production pattern using Lemma 3 to reduce to a Leontief exchange economy.

First, Observation 7 and our restriction on  $A$  will imply that agents see  $K$  distinct producers in the economy. Let  $A_i$  denote the  $i$ -th row of  $A$ . Let  $i$  and  $i'$  be agents whose production vectors are scalar multiples, i.e.  $A_i = c \cdot A_{i'}$  for some constant  $c$ . We claim that in equilibrium,  $w_i = c \cdot w_{i'}$ , and therefore agents are ambivalent between having one unit of  $i$ 's raw material and  $c$  units of  $i'$ 's raw material.

Assume the contrary, i.e.  $w_i \neq c \cdot w_{i'}$ . If  $w_i < c \cdot w_{i'}$ , then  $\frac{w_i}{a_{ij}} < \frac{w_{i'}}{a_{i'j}}$  for all goods  $j$ . By Observation 7, this implies that agent  $i'$  does not produce anything. Similarly,  $w_i > c \cdot w_{i'}$  would imply that agent  $i$  does not produce anything. This can only happen in equilibrium if neither agent  $i$  nor agent  $i'$  produce anything, in which case it must be that  $w_i = w_{i'} = 0$ .

Thus, the raw materials of  $i$  and  $i'$  are indistinguishable. It follows that from a computational perspective, we need only consider an economy with  $K$  distinct producers (we can normalize so that  $A_i = A_{i'}$ .)

According to Lemma 8, this implies that there are at most  $O(m^{O(K^2)}) = O(m^{O(1)})$  different production patterns. Since we have one raw material, linear production functions, Leontief utilities, and knowledge of the production pattern, we apply Lemma 3 to reduce the problem to a Leontief exchange economy with  $K$  goods in a polytope in the price space. (The relationship between prices in the Ricardo and upside-down economies tells us how to transform the Ricardo price polytope to the price space of the upside-down economy.)  $\square$

**Theorem 10.** *If the rows of the production matrix  $A$  are scalar multiples of  $K = O(1)$  different vectors in a Leontief Ricardo economy, then equilibria are efficiently computable.*

*Proof.* We use the method of Devanur and Kannan 4 to compute equilibria in a polytope for a Leontief exchange economy. Combining this with Lemma 9 gives a polynomial time algorithm.  $\square$

### 4.1 Ricardian Comparative Advantage

The price-space decomposition implied by Observation 7 suggests a new generalization of Ricardo’s law of comparative advantage. A well-known theorem in economics, Ricardo’s law of comparative advantage for two agents (originally countries) is as follows:

**Theorem 11.** *(David Ricardo) In equilibrium for a two agent Ricardo economy, if the goods are ordered by relative production factors  $a_{ij}$  and equilibrium raw material prices  $w_i$  as*

$$\frac{a_{11}}{a_{21}} \geq \dots \geq \frac{a_{1k}}{a_{2k}} > \frac{w_1}{w_2} > \frac{a_{1(k+1)}}{a_{2(k+1)}} \geq \dots \geq \frac{a_{1m}}{a_{2m}}$$

*then agent 1 produces goods 1 through  $k$  and agent 2 produces goods  $(k + 1)$  through  $m$ . If, for some good  $j$  we have  $\frac{w_1}{w_2} = \frac{a_{1j}}{a_{2j}}$ , then either country may produce good  $j$ .*

Interestingly, previous attempts to generalize comparative advantage failed to produce a nice theory 10,3. However, hyperplane-partitioning leads to the following intuitive generalization:

**Theorem 12.** *Comparative advantage in an  $n$ -agent Ricardo economy may be understood as a partition of an  $(n - 1)$ -dimensional simplex by the price vector  $w$  into  $n$  convex polytopes. A good  $j$  is produced by country  $i$  if and only if its relative production technologies map to a point in  $i$ ’s polytope.*

*Proof.* Observation 7 tells us that in equilibrium, country  $i$  produces all goods for which  $\frac{w_i}{w_{i'}} < \frac{a_{ij}}{a_{i'j}}$  for all other countries  $i'$ , and that it may produce goods for which  $\frac{w_i}{w_{i'}} \leq \frac{a_{ij}}{a_{i'j}}$  for all  $i'$ .

Consider the material prices and production coefficients as vectors

$$w = (w_1, \dots, w_n), \quad a_j = (a_{1j}, \dots, a_{nj})$$




and normalize them according to their  $L_1$  norm,

$$w' = \frac{w}{|w|_1}, \quad a'_j = \frac{a_j}{|a_j|_1} .$$

Since all  $a_{ij}$  are positive, this maps the wage and production vectors to points on the  $(n - 1)$ -dimensional simplex.

The points  $\pi$  in the price space where  $\frac{\pi_i}{\pi_{i'}} = \frac{w_i}{w_{i'}}$  form a hyperplane. By Observation 7, this hyperplane partitions the space (and therefore the simplex) between goods possibly produced by  $i$  and goods possibly produced by  $i'$ . Together, the hyperplanes partition the simplex into  $n$  convex polytopes  $P_i$  where country  $i$  produces those goods whose normalized production technology  $a'_j$  falls inside  $P_i$ .

Figure 1 illustrates the generalization. □



**Fig. 1.** Comparative advantage for 3 agents in the Ricardo model. The space of relative production ratios is visualized as a 2-dimensional simplex (i.e. triangle) following Theorem 12. If  $w$  represents the equilibrium price vector for the raw material, then good 1 will be produced by country  $B$ , 3 will be produced by  $C$ , 4 will be produced by  $A$ , and 2 may be produced by either  $A$  or  $C$ .

## 5 Computability in the Heckscher-Ohlin Model

The Heckscher-Ohlin model stipulates that agents’ production functions are identical. Again, we show a variety of results and, most interestingly, see that when the number of raw materials is small ( $K = O(1)$ ), the number of goods may not matter (see Corollary 13).

**Theorem 13.** *It is NP-hard to determine if a Heckscher-Ohlin economy with linear production functions and Leontief utilities has an equilibrium.*

*Proof.* Like Theorem 6, it is easy to simulate a pairing Leontief exchange economy. Let  $\Phi$  parameterize a pairing Leontief exchange economy. Construct a

Heckscher-Ohlin economy with  $n$  raw materials,  $n$  outputs, and production functions parameterized by

$$a_{jk} = \begin{cases} 1 & i = j \\ 0 & \text{otherwise} \end{cases}$$

Endow agent  $i$  with one unit of raw material  $i$  and nothing else, i.e.

$$l_{ik} = \begin{cases} 1, & i = k \\ 0, & \text{otherwise} \end{cases}$$

Agent  $i$  can produce exactly one unit of good  $i$  and nothing else, so the goods for trade are identical to the pairing Leontief economy. Thus, the equilibria must be the same.  $\square$

Our next two results will be corollaries of the following theorem:

**Theorem 14.** *When the utility and production functions are both linear (or both Leontief), computing equilibria in the Heckscher-Ohlin model reduces to computing equilibria in an exchange economy with linear (Leontief) utilities and  $K$  goods.*

*Proof.* The linear and Leontief cases are straightforward applications of Lemmas [4](#) and [2](#) respectively. In both cases, the reductions are efficiently computable.  $\square$

**Corollary 15.** *When the utility and production functions are both Leontief and there are  $K = O(1)$  raw materials, equilibria in the Heckscher-Ohlin model are efficiently computable.*

*Proof.* It is sufficient to compute equilibria in an exchange economy with Leontief utilities and  $m = O(1)$  goods. Devanur and Kannan show that such equilibria are efficiently computable [4](#).  $\square$

**Corollary 16.** *When the utility and production functions are both linear, equilibria in the Heckscher-Ohlin model are efficiently computable.*

*Proof.* It is sufficient to compute equilibria in an exchange economy with linear utilities, a problem for which many efficient algorithms exist, e.g. [5,6](#).  $\square$

## Acknowledgments

The author would like to thank Christos Papadimitriou for his helpful suggestions and comments.

## References

1. Arrow, K.J., Debreu, G.: Existence of an equilibrium for a competitive economy. *Econometrica* 22(3), 265–290 (1954)
2. Codenotti, B., Saberi, A., Varadarajan, K., Ye, Y.: Leontief economies encode nonzero sum two-player games. In: *SODA 2006: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp. 659–667. ACM, New York (2006)

3. Deardorff, A.V.: The general validity of the law of comparative advantage. *The Journal of Political Economy* 88(5), 941–957 (1980)
4. Devanur, N.R., Kannan, R.: Market equilibria in polynomial time for fixed number of goods or agents. In: *FOCS 2008: Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, Washington, DC, USA, pp. 45–53. IEEE Computer Society, Los Alamitos (2008)
5. Devanur, N.R., Papadimitriou, C.H., Saberi, A., Vazirani, V.V.: Market equilibrium via a primal-dual-type algorithm. In: *FOCS 2002: Proceedings of the 43rd Symposium on Foundations of Computer Science*, Washington, DC, USA, pp. 389–395. IEEE Computer Society, Los Alamitos (2002)
6. Garg, R., Kapoor, S.: Auction algorithms for market equilibrium. *Math. Oper. Res.* 31(4), 714–729 (2006)
7. Jain, K., Mahdian, M.: Computing equilibria in a fisher market with linear single-constraint production units. In: Deng, X., Ye, Y. (eds.) *WINE 2005. LNCS*, vol. 3828, pp. 788–792. Springer, Heidelberg (2005)
8. Jain, K., Varadarajan, K.: Equilibria for economies with production: constant-returns technologies and production planning constraints. In: *SODA 2006: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp. 688–697. ACM, New York (2006)
9. Jain, K., Vazirani, V.V., Ye, Y.: Market equilibria for homothetic, quasi-concave utilities and economies of scale in production. In: *SODA 2005: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, Philadelphia, PA, USA, pp. 63–71. Society for Industrial and Applied Mathematics (2005)
10. Jones, R.W.: Comparative advantage and the theory of tariffs: A multi-country, multi-commodity model. *The Review of Economic Studies* 28(3), 161–175 (1961)
11. Kapoor, S., Mehta, A., Vazirani, V.: An auction-based market equilibrium algorithm for a production model. *Theor. Comput. Sci.* 378(2), 153–164 (2007)
12. Leamer, E.E.: *The heckscher-ohlin model in theory and practice*. Princeton Studies in International Economics 77, International Economics Section, Department of Economics Princeton University (1995)
13. Yinyu, Y.: Exchange market equilibria with leontief’s utility: Freedom of pricing leads to rationality. *Theor. Comput. Sci.* 378(2), 134–142 (2007)

# The Geometry of Truthfulness<sup>\*</sup>

Angelina Vidali

Max-Planck Institut für Informatik, Saarbrücken, Germany  
angelina@mpi-inf.mpg.de

**Abstract.** We study the geometrical shape of the partitions of the input space created by the allocation rule of a truthful mechanism for multi-unit auctions with multidimensional types and additive quasilinear utilities. We introduce a new method for describing the allocation graph and the geometry of truthful mechanisms for an arbitrary number of items(/tasks). Applying this method we characterize all possible mechanisms for the case of three items.

Previous work shows that Monotonicity is a necessary and sufficient condition for truthfulness in convex domains. If there is only one item, monotonicity is the most practical description of truthfulness we could hope for, however for the case of more than two items and additive valuations (like in the scheduling domain) we would need a global and more intuitive description, hopefully also practical for proving lower bounds. We replace Monotonicity by a geometrical and global characterization of truthfulness.

Our results apply directly to the scheduling unrelated machines problem. Until now such a characterization was only known for the case of two tasks. It was one of the tools used for proving a lower bound of  $1 + \sqrt{2}$  for the case of 3 players. This makes our work potentially useful for obtaining improved lower bounds for this very important problem.

Finally we show lower bounds of  $1 + \sqrt{n}$  and  $n$  respectively for two special classes of scheduling mechanisms, defined in terms of their geometry, demonstrating how geometrical considerations can lead to lower bound proofs.

## 1 Introduction

Mechanism design is the branch of game theory that tries to implement social goals taking into account the selfish nature of the individuals involved. Mechanism design constructs allocation algorithms that together with appropriate payments elicit from the players their secret values or preferences. In this paper we give a characterization result that reveals the exact geometry of truthful mechanisms. The goal of this paper is to understand and visualize truthful mechanisms better. We realized the need for such a result while trying to improve the lower bound for the scheduling selfish unrelated machines problem [14,7,10],

---

<sup>\*</sup> Work done while the author was at the University of Athens and was supported in part by IST-15964 (AEOLUS), IST-2008-215270 (FRONTS), and the Greek GSRT.

however the result is more broadly applicable and interesting from itself continuing a line of research attempting to grasp truthfulness better [7,9,12,14]. What differentiates our work from this line of research is that we fully exploit the linearity in the geometry of additive valuations.

There exists a simple necessary and sufficient condition for truthfulness in convex domains and a finite number of outcomes, the Monotonicity Property. In single parameter domains, like for example in an auction where there is only one item, monotonicity is exactly the monotonicity we know from calculus and the most practical description of truthfulness we could hope for. The allocation should be a monotone (for the case of auctions an increasing, while for the case of scheduling a decreasing) function of the player's valuation for the item. However for the case of two or more items Monotonicity is a local condition that should be satisfied by any pair of instances of the problem and does not give us any clue about the global picture of the mechanism, when considering the whole space of inputs together. We would instead need a global and more intuitive description, hopefully also practical for proving lower bounds. We replace Monotonicity by a geometrical and global characterization of truthfulness, for the case when the valuations are additive.

Until now such a characterization was known in the context of the scheduling unrelated machines problem only for the easy case of two tasks [7] and it turned out to be a quintessential element of the characterization proof in [6] and the lower bound in [7]. We believe that our result here can be used for obtaining new lower bounds. The only discouraging fact is that even for the case of 3 tasks the different mechanisms are too many and geometrically complicated.

No matter how many are the players participating in a mechanism, determining whether a mechanism is truthful boils down to a single-player case. Truthfulness requires that for fixed values of the other players, a player should not be able to increase his utility by lying. Studying the mechanism for fixed values of the other players is like studying a single-player case. Consequently in our setting there is a single player and  $m$  different indivisible items (or tasks). The player's type is denoted by the vector  $t = (t_1, \dots, t_m)$ , where  $t_i$  is the valuation of the bidder for the  $i$ -th item/task and the allocation is denoted by  $a = (a_1, \dots, a_m)$  where  $a_i \in \{0, 1\}$ .

We assume that the bidder has additive valuations and hence the bidder's valuation function when his type is  $t$  and his allocation  $a$  is  $v_t(a) = a \cdot t$ . In fact it is easy to see that our results also apply if the valuations are of the form  $v_t(a) = \lambda(a \cdot t) + \gamma_a$  for some constants  $\lambda, \gamma_a$  (we can have one different  $\gamma_a$  for each different allocation  $a$ ). The reason for this is simple namely these valuations also satisfy the Monotonicity Property and moreover the possible truthful mechanisms for such valuations are like in Figure 1 (this would not be the case for valuations with  $v_t(11) = t_1 \cdot t_2$  or  $v_t(11) = 2t_1 + t_2$  as the sloped hyperplane would not be  $45^\circ$ ). A mechanism consists of an allocation algorithm  $a$  and a payment algorithm  $p$ . We make the standard assumption that the utilities are quasilinear, that is the utility of the player is  $u(a, t) = v_t(a) - p(a)$ .

The allocation part of the mechanism gives a partition of the space  $\mathbb{R}^m$  of possible values of a player to  $2^m$  different regions, one for each possible different allocation  $a$  of the player. But which are exactly the possible partitions of the space the mechanism creates? This is exactly the question we address in this paper.

We know [17] that a mechanism is truthful if and only if its allocation part satisfies the monotonicity property.

**Definition 1 (Monotonicity Property).** *An allocation algorithm is called monotone if it satisfies the following property: for every two sets of tasks  $t$  and  $t'$  the associated allocations  $a$  and  $a'$  satisfy  $(a - a') \cdot (t - t') \leq 0$ , where  $\cdot$  denotes the dot product of the vectors, that is,  $\sum_{j=1}^m (a_j - a'_j)(t_j - t'_j) \leq 0$ .*

Notice that the Monotonicity Property is a necessary and sufficient condition for truthfulness and that it only involves the allocation part of the mechanism. Consequently by determining the possible partitions of the input space created by the allocation part of the mechanism we will eventually give a characterization of truthfulness.

As it has already been noticed in [9] in the case of additive valuations the boundaries of the mechanism are hyperplanes of a very specific form, every region created by this partition is a convex polyhedron. In this paper we show exactly which (rather few) polytopes are involved in such a partition. For proving our results we reduce the problem to that of determining the allocation graph of the mechanism, i.e. which of the regions share a common boundary. We can then determine the exact geometrical shape of the mechanism because the hyperplane that separates two regions can be easily derived from the monotonicity property.

Our results apply directly to the scheduling unrelated machines problem:

**Definition 2 (The scheduling unrelated machines problem).** *The input to the scheduling problem is a nonnegative matrix  $t$  of  $n$  rows, one for each machine-player, and  $m$  columns, one for each task. The entry  $t_{ij}$  (of the  $i$ -th row and  $j$ -th column) is the time it takes for machine  $i$  to execute task  $j$ . Let  $t_i$  denote the times for machine  $i$ , which is the vector of the  $i$ -th row. The output is an allocation  $a = a(t)$ , which partitions the tasks into the  $n$  machines. We describe the partition using indicator values  $a_{ij} \in \{0, 1\}$ :  $a_{ij} = 1$  iff task  $j$  is allocated to machine  $i$ . We should allocate each task to exactly one machine, or more formally  $\sum_{j=1}^n a_{ij} = 1$ . The goal is to minimize the makespan, i.e. to minimize the total processing time of the player that finishes last.*

## 1.1 Our Tools

Besides the potential applications of our characterization, we believe that also the method we introduce for studying the allocation graph is of particular interest as it provides a very simple way to handle a very complicated partition of the space. We propose a method for determining all possible allocation graphs and the geometrical shapes of the mechanism: For each region  $R_a$  of the mechanism instead of considering its complicated geometrical shape we define a box that

contains the region. The signs of distances between parallel to each other boundaries of the mechanism determine whether two of these boxes intersect. If two boxes intersect then the corresponding regions share a common boundary. Alternatively if two boxes intersect then there is an edge between the corresponding edges in the allocation graph. These distances however are not independent from each other. Applying cycle-monotonicity for appropriately chosen zero-length cycles allows us to determine how these constants relate. As boundaries between regions that differ only in one allocation always exist we will concentrate on the subgraph of the allocation graph that consists of the edges corresponding to Hamming distance-1 boundaries.

## 1.2 Related Work

Myerson [13] gave a characterization of truthful algorithms for one-parameter problems, in terms of a monotonicity condition, which was rediscovered by Archer and Tardos [2]. For the case of multidimensional types Bikchandani et al. [5] prove that a simple necessary monotonicity property of the allocations of different inputs (and without any reference to payments) is also sufficient for truthful mechanisms, while Gui, Müller, and Vohra [9] extend this to a greater variety of domains (this work is rather close to ours as it also follows a geometrical approach). Saks and Yu [17] generalize this result to cover all convex domains of finitely many outcomes. Monderer [12] showed that this result cannot be essentially extended to a larger class of domains. Both these results concern domains of finitely many outcomes. There are however cases, like the fractional version of the scheduling problem, when the set of all possible allocations is infinite. For these, Archer and Kleinberg [1] provided a necessary and sufficient condition for truthfulness. Very recently Berger et al. [4] generalize all these results for the case of convex valuations.

Nisan and Ronen introduced the mechanism-design version of the scheduling problem on unrelated machines in the paper that founded the algorithmic theory of Mechanism Design [14,15]. They showed that the well-known VCG mechanism, which is a polynomial-time algorithm and truthful, has approximation ratio  $n$ . They conjectured that there is no deterministic mechanism with approximation ratio less than  $n$ . They also showed that no mechanism (polynomial-time or not) can achieve approximation ratio better than 2. This was improved to  $1 + \sqrt{2}$  [7], and further to  $1 + \varphi$  in [10]. For the case of two machines [8] Dobzinski and Sundararajan characterized all mechanisms with finite approximation ratio, while [6] gave a characterization of all (regardless of approximation ratio) decisive truthful mechanisms in terms of affine minimizers and threshold mechanisms. In a very recent paper [3] Ashlagi, Dobzinski and Lavi prove a lower bound of  $n$  for a special class of mechanisms, which they call *anonymous*. Lavi and Swamy [11] considered another special case of the same problem when the processing times have only two possible values low or high; the use of cycle monotonicity played a central role in this work as well.

## 2 Preliminaries

We denote by  $R_a$  the closure of the subset of  $\mathbb{R}^m$  where the mechanism gives assignment  $a$  and we will call it a *region of the mechanism*. For any two different assignments  $a, b$  for player  $i$  we define  $f_{a:b} := \sup\{(a - b) \cdot t \mid t \in R_a\}$ .

We define the Hamming Distance  $\text{Hd}(a, b)$  between two vectors  $a, b$ , as the number of positions in which the two vectors are different. The *Minkowski sum* of two sets  $A, B \subseteq \mathbb{R}^m$  is  $A \oplus B = \{a + b \mid a \in A, b \in B\}$ . Let also  $B_a := \{t \mid (-1)^{a_j} t_j \geq 0, j = 1, \dots, m\}$ . For  $m = 2$  each  $B_a$  is a quadrant of  $\mathbb{R}^2$ .

### Lemma 1

a) If a point  $b$  belongs to region  $R_a$  of a truthful mechanism, then also  $b \oplus B_a \subseteq R_a$ .

b) Regions  $R_a$  and  $R_{a'}$  are separated by the hyperplane  $(a - a') \cdot t = f_{a:a'}$  and each region is bounded by a convex polytope.

c) For  $F_a := (f_{a:a_{-1}, 1-a_1}, \dots, f_{a:a_{-m}, 1-a_m})$  every region  $R_a$  satisfies  $R_a \subseteq F_a \oplus B_a$ . In other words region  $R_a$  is included in the box we get by shifting the box  $B_a$  so that it has its vertex at the point  $F_a$ .

This means that every region  $R_a$  is included in a box defined by the boundaries of  $R_a$  with all regions  $R_b$  such that  $\text{Hd}(a, b) = 1$ . The proof is immediate by the monotonicity property and the definition of  $f_{a:b}$ .

### 2.1 The Allocation Graph of Each Player

We define an edge-weighted directed graph  $G$ , the *allocation graph*, whose vertex set are all possible allocations of the player. For each two allocations  $a, b$  the weight of the edge from  $a$  to  $b$  is  $f_{a:b}$ .

The following property is necessary and sufficient for truthfulness [16].

**Definition 3 (Cycle monotonicity).** An allocation algorithm satisfies cycle monotonicity if for every integer  $K$  and cycle  $a_1, \dots, a_K, a_{K+1} = a_1$  on the allocation graph  $\sum_{k=1}^K f_{a_k:a_{k+1}} \leq 0$ .

The following Lemma is an essential tool for our proofs.

### Lemma 2


a) Two regions  $R_a, R_{a'}$  that share at least one common boundary point satisfy  $p(a) - p(a') = f_{a:a'} = -f_{a':a}$ .

b) Any cycle on the allocation graph in which each pair of consecutive nodes corresponds to a pair of regions sharing a common boundary point has length zero.

## 3 New Tools for the Case of $m$ Items

The mechanism consists of sloped hyperplanes, as well as hyperplanes vertical to some axis, which we will call Hd-1 boundaries (because they separate regions that





**Fig. 1.** The two possible ways to partition the positive orthant for the case of 2 tasks and the threshold mechanism as a degenerate case of both

have Hamming distance equal to 1, i.e. differ in only one task). The trouble with the sloped hyperplanes is that they do appear as boundaries in all possible shapes of the mechanism, so we have to take cases. Luckily the hyperplanes vertical to some axis appear in all possible shapes. We will use the distance between these hyperplanes in order to describe the allocation graph of the mechanism. The sign of these distances determines exactly which of the sloped lines appear in the geometrical picture of the mechanism. Knowing the allocation graph we can then easily draw the picture of the mechanism.

The idea of our approach is best depicted if we apply it for the easy case of two tasks (for which we already know that the two possible mechanisms are depicted in Figure [II](#)). A purely algebraic way to obtain this description is to just apply once cycle monotonicity. Taking the cycle  $00 \rightarrow 01 \rightarrow 11 \rightarrow 10 \rightarrow 00$  we get  $f_{11:01} + f_{00:10} = f_{11:10} + f_{00:01}$ . If we define  $c_{12} := f_{11:01} + f_{00:10}$  (This is the distance between the two lines vertical to the axis  $t_1$ .) then by the previous cycle it turns out that the distance between the two lines vertical to the axis  $t_2$ , which can be expressed as  $f_{11:01} + f_{00:10}$  is also equal to  $c_{12}$ . Region  $R_{11}$  is contained in the box defined by  $f_{11:01}, f_{11:10}$  and  $R_{11}$  and  $R_{00}$  share a common boundary line if and only if the boxes that contain them intersect i.e. if and only if  $c_{12} > 0$ . That is the sign of  $c_{12}$  determines which of the two possible shapes has the mechanism.

We proceed to define some constants that express the distances between regions, generalizing this idea we demonstrated for the case of two tasks. The constant  $c_{ij|a_{-\{i,j\}}}$  measures the distance between the two parallel hyperplanes  $t_i = f_{11a_{-\{i,j\}}:01a_{-\{i,j\}}}$  and  $t_i = f_{10a_{-\{i,j\}}:00a_{-\{i,j\}}}$ , which correspond to separating hyperplanes of the mechanism between regions in Hamming distance 1. This constant fully describes the geometry of the mechanism if the allocation of all tasks, except for tasks  $i, j$ , is fixed to  $a_{-\{i,j\}}$ . To provide some intuition why we choose these consider that in a decisive mechanism this would give an asymptotic picture of the mechanism: If the values of only two tasks  $i, j$  are allowed to be variables, while the remaining tasks with allocation 1 are fixed to the biggest possible value ( $+\infty$ ) and the tasks with allocation 0 are fixed to the smallest possible value, this constant describes the geometry of the mechanism that allocates tasks  $i, j$ .

**Definition 4.** For all  $i, j$  and all possible  $m - 2$ -tuples (/allocations)  $a_{-\{i,j\}}$  we define

$$\begin{aligned}
 c_{ij|a_{-\{i,j\}}} &:= f_{11a_{-\{i,j\}}:01a_{-\{i,j\}}} + f_{00a_{-\{i,j\}}:10a_{-\{i,j\}}} \\
 &= f_{11a_{-\{i,j\}}:10a_{-\{i,j\}}} + f_{00a_{-\{i,j\}}:01a_{-\{i,j\}}} \quad (1)
 \end{aligned}$$

But are these constants independent from each other? As the following Lemma shows, the answer is no and the relation between these constants is derived from Cycle Monotonicity.

**Lemma 3.** If a mechanism is truthful then the constants  $c_{ij|a_{-\{i,j\}}}$  satisfy the following equation  $c_{ij|1a_{-\{i,j,k\}}} - c_{ij|0a_{-\{i,j,k\}}} = c_{ik|1a_{-\{i,j,k\}}} - c_{ik|0a_{-\{i,j,k\}}}$ .

By Lemma □ each region  $R_a$  of the mechanism is contained in a box formed by the separating hyperplanes between  $R_a$  and all regions with assignment in Hamming distance 1 from  $a$ . If we concentrate on a pair of intersecting regions, then the boxes that contain them have a non-empty intersection. But it is also the other way round:

**Lemma 4.** If the boxes corresponding to two regions intersect then the regions share a common boundary hyperplane.

We proceed to define  $d_{a;b}^i$  as the difference of the Hd-1 boundaries on axis  $i$  corresponding to two distinct regions  $R_a, R_b$ . We have  $d_{a;1-a}^i > 0$  for all  $i = 1, \dots, m$  if and only if regions  $R_a$  and  $R_{1-a}$  intersect.

Even though the geometry of the mechanism is complicated it turns out that we can derive a general formula for the  $d_{a;b}^i$ s using now a more complicated zero-length cycle on the allocation graph.

**Definition 5.** We define the distance  $d_{a;b}^i := f_{a:1-a_i,a_{-i}} + f_{b:1-b_i,b_{-i}}$ .

**Lemma 5.** We have  $d_{a;b}^i = d_{b;a}^i$  (symmetry) and  $d_{a;b}^i = -d_{1-a_i,a_{-i};1-b_i,b_{-i}}^i$ .

**Lemma 6.** The distance  $d_{a;1-a}^i$  can be expressed as the following sum of constants:  $d_{a;1-a}^i := \sum_{j \neq i, j \in \{1, \dots, m\}} (-1)^{a_i+a_j} c_{ij|b_{-\{i,j\}}}$ , where the  $k$ -th coordinate

of the allocation  $b_k$  is  $b_k = \begin{cases} 1 - a_k & \text{if } k < j \\ a_k & \text{if } k > j. \end{cases}$

## 4 Characterization of 3-Dimensional Mechanisms

### 4.1 Calculating the Distances

We believe that the tools we have developed in the preceding section are useful for the study of the allocation graph for an arbitrary number of tasks  $m$ . We demonstrate this by using them in order to determine the allocation graphs and the corresponding geometrical shapes a truthful mechanism can take for the case  $m = 3$ .

For the case of 3 tasks we will apply Lemmas 6 and 5 in order to compute the distances  $d_{a:1-a}^i$  with respect to the constants  $c_{i,j|a-\{i,j\}}$ . For simplicity of notation we will write  $d^j$  instead of  $d_{111:000}^j$ , for  $j = 1, 2, 3$  and it turns out that all other distances  $d_{a:b}^j$ , between regions  $R_a$  and  $R_b$ , can be expressed using the three distances  $d^1, d^2, d^3$  between regions  $R_{111}$  and  $R_{000}$ . We define the constant  $e$  as  $e := c_{12|0} - c_{12|1} = c_{13|0} - c_{13|1} = c_{23|0} - c_{23|1}$ . then  $c_{12|0} = c_{12|1} + e$  and we can rewrite the equalities in the following way:

$$\begin{aligned} d_{111:000}^1 &= c_{13|1} + c_{12|0} = c_{13|1} + c_{12|1} + e = c_{13|0} + c_{12|0} - e = d^1 \\ d_{111:000}^2 &= c_{12|1} + c_{23|0} = c_{12|1} + c_{23|1} + e = c_{12|0} + c_{23|0} - e = d^2 \\ d_{111:000}^3 &= c_{13|1} + c_{23|0} = c_{13|1} + c_{23|1} + e = c_{13|0} + c_{23|0} - e = d^3 \end{aligned}$$

Similarly for the rest of the distances we have  $d_{011:100}^1 = -d^1, d_{011:100}^2 = -c_{12|1} + c_{23|1} = d^3 - d^1, d_{011:100}^3 = -c_{13|1} + c_{23|1} = d^2 - d^1$  and then  $d_{101:010}^1 = -c_{12|1} + c_{13|1} = d^3 - d^2, d_{101:010}^2 = -d^2, d_{101:010}^3 = -(d^2 - d^1)$  and  $d_{110:001}^1 = -(d^3 - d^2), d_{110:001}^2 = -(d^3 - d^1), d_{110:001}^3 = -d^3$ .

### 4.2 Properties Satisfied by the Allocation Graph

**Lemma 7.** *There always exist two regions  $R_a, R_b$  in  $\text{Hd} = 3$  such that  $d_{a:b}^i \geq 0$  for  $i = 1, 2, 3$ .*

In what follows we will make the assumption that this pair of regions  $R_a, R_b$  in  $\text{Hd} = 3$  such that  $d_{a:b}^i \geq 0$  for  $i = 1, 2, 3$ , guaranteed to exist by Lemma 7 are  $R_{111}$  and  $R_{000}$ .

For any mechanism we present here you can get another truthful mechanism by applying the following rotations: Think of the mechanism as a partition of the cube, if you rotate one of the possible partitions so that the faces of the cube go to faces of the cube after the rotation (and the center of axes goes to another vertex of the cube), you also get a truthful mechanism. The reason is that the slope of the separating hyperplane between two regions only depends on their Hamming Distance, i.e. on the number of tasks on which they differ. The characteristic of the rotation we described is that it respects the Hamming distances.


**Lemma 8.** *If  $R_{111}$  and  $R_{000}$  intersect then a) if  $e < 0$  then at least two of the constants  $c_{12|1}, c_{13|1}, c_{23|1}$  are strictly positive,*

*b) if  $e > 0$  then at least two of the constants  $c_{12|0}, c_{13|0}, c_{23|0}$  are strictly positive.*

**Lemma 9.** *If a pair of regions  $R_a, R_{1-a}$  share a common Hd-3 boundary then no other pair  $R_b, R_{1-b}$  of regions share a common Hd-3 boundary.*

### 4.3 All Possible Mechanisms

**Definition 6.** *A degenerate version of a mechanism  $M$  is a mechanism for which some of the constants  $c_{ij|0}, c_{ij|1}, d_{a:b}^k$ , for some  $i, j, k \in \{1, 2, 3\}$  and some*



**Fig. 2.** 3D models of the possible partitions (up to rotation). Looking just at the blue projections you can determine the constants  $c_{ij|0}$  and from the red projections the constants  $c_{ij|1}$ .

*allocations  $a, b$ , become 0, while all other such constants retain the same sign as in the non-degenerate mechanism.*

We describe the possible shapes of the mechanism when a Hd-3 boundary exists and thanks to Lemma 7 any other mechanism is a degenerate version of a mechanism with a Hd-3 boundary. Summarizing all restrictions to the shape of the mechanism we obtained in the previous section we get the following characterization:

**Theorem 1.** *The possible truthful mechanisms are the following five possible partitions of the space and all their rotations. (In Figure 2 you can see their geometrical shapes.)*

*As for any mechanism we give here we also include in our characterization all its rotations, we suppose without loss of generality that  $R_{111}, R_{000}$  share a common boundary, that  $e < 0$  and that the two constants guaranteed to be positive by Lemma 8 are  $c_{12|1} > 0, c_{23|1} > 0$ .*

1.  $c_{12|1} > 0, c_{13|1} > 0, c_{23|1} > 0, c_{12|0} > 0, c_{13|0} > 0, c_{23|0} > 0$
2.  $c_{12|1} > 0, c_{13|1} > 0, c_{23|1} > 0, c_{12|0} < 0, c_{13|0} < 0, c_{23|0} < 0$
3.  $c_{12|1} > 0, c_{13|1} > 0, c_{23|1} > 0, c_{12|0} < 0, c_{13|0} < 0, c_{23|0} > 0$
4.  $c_{12|1} > 0, c_{13|1} > 0, c_{23|1} > 0, c_{12|0} > 0, c_{13|0} < 0, c_{23|0} > 0$
5.  $c_{12|1} > 0, c_{13|1} < 0, c_{23|1} > 0, c_{12|0} > 0, c_{13|0} < 0, c_{23|0} > 0$ .

## 5 Lower Bounds for Some Scheduling Mechanisms

Observing the figures we got from our characterization we see that many of the regions have the shape of a box, for some of these cases the region that has the shape of the box is  $R_{1\dots 1}$ . Threshold(/additive) mechanisms [6,14] are the special case of these mechanisms, when all regions are boxes. Even though these mechanisms are much more general, we can still show the same lower bound of  $1 + \sqrt{n}$  using an argument very similar to the one used in [6]. For these cases we can prove a lower bound of  $1 + \sqrt{n}$ .

**Theorem 2.** *Every mechanism for which  $R_{1\dots 1}$  is a box has approximation ratio at least  $1 + \sqrt{n}$ .*

Finally there is a non-trivial geometrically defined class of mechanisms for which we can provide an  $n$  lower bound. We say that a mechanism is *non-penalizing* if in the allocation graph no pair of regions of the form  $R_{a10}, R_{b01}$ , where  $a, b$  are  $(m - 2)$ -dimensional allocation vectors, share a common boundary. The first mechanism in Figure 2 is an example of such a mechanism. The intuition behind these mechanisms is that, if for fixed values of the other players, a player lowers one of his values he only gets more tasks (regardless of his initial allocation for the tasks he lowers), in other words a machine never loses a job, just because it becomes faster for another job.

**Theorem 3.** *Every non-penalizing mechanism has approximation ratio at least  $n$ .*

## 6 Concluding Remarks and Open Problems

Our characterization is only for the case of 3 tasks, the tools we have developed to obtain this characterization are however for the general case of  $m$  tasks. Can we find a succinct way to describe all possible allocation graphs for the general case?

We would like to stress the connection of our results with the scheduling unrelated machines problem. The lower bounds in the last section show that many mechanisms have bad approximation ratio just because of the geometrical shape of their projections. Finally we believe that the characterization for the case of three tasks can be used to improve the existing [10] lower bound of 2.465 for the case of 4 machines to a better constant.

## Acknowledgements

I would like to thank Christos Athanasiadis, Ioannis Emiris and Elias Koutsoupias for helpful discussions and my brother Aris Vidalis for making the nice shaded 3D models I include in this paper (and also for bothering to imagine these complicated partitions of the space!).

## References

1. Archer, A., Kleinberg, R.: Truthful germs are contagious: A local to global characterization of truthfulness. In: EC (2008)
2. Archer, A., Tardos, É.: Truthful mechanisms for one-parameter agents. In: FOCS, pp. 482–491 (2001)
3. Ashlagi, I., Dobzinski, S., Lavi, R.: An optimal lower bound for anonymous scheduling mechanisms. In: EC (2009)
4. Berger, A., Müller, R., Naeemi, S.H.: Characterizing incentive compatibility for convex valuations. In: SAGT (2009)

5. Bikhchandani, S., Chatterji, S., Lavi, R., Mu'alem, A., Nisan, N., Sen, A.: Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica* 74(4), 1109–1132 (2006)
6. Christodoulou, G., Koutsoupias, E., Vidali, A.: A characterization of 2-player mechanisms for scheduling. In: Halperin, D., Mehlhorn, K. (eds.) *Esa 2008*. LNCS, vol. 5193, pp. 297–307. Springer, Heidelberg (2008)
7. Christodoulou, G., Koutsoupias, E., Vidali, A.: A lower bound for scheduling mechanisms. *Algorithmica* (2008)
8. Dobzinski, S., Sundararajan, M.: On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In: *EC (2008)*
9. Gui, H., Müller, R., Vohra, R.V.: Dominant strategy mechanisms with multidimensional types. In: *Computing and Markets (2005)*
10. Koutsoupias, E., Vidali, A.: A lower bound of  $1+\phi$  for truthful scheduling mechanisms. In: Kučera, L., Kučera, A. (eds.) *MFCS 2007*. LNCS, vol. 4708, pp. 454–464. Springer, Heidelberg (2007)
11. Lavi, R., Swamy, C.: Truthful mechanism design for multi-dimensional scheduling via cycle monotonicity. In: *EC*, pp. 252–261 (2007)
12. Monderer, D.: Monotonicity and implementability. In: *EC (2008)*
13. Myerson, R.B.: Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73 (1981)
14. Nisan, N., Ronen, A.: Algorithmic mechanism design (extended abstract). In: *STOC*, pp. 129–140 (1999)
15. Nisan, N., Ronen, A.: Algorithmic mechanism design. *Games and Economic Behavior* 35, 166–196 (2001)
16. Rochet, J.-C.: A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics* 16, 191–200 (1987)
17. Saks, M.E., Yu, L.: Weak monotonicity suffices for truthfulness on convex domains. In: *EC*, pp. 286–293 (2005)

# Optimal Incentives for Participation with Type-Dependent Externalities<sup>\*</sup>

Michal Feldman<sup>1</sup>, Ran Tessler<sup>2</sup>, and Yoav Wilf<sup>3</sup>

<sup>1</sup> School of Business Administration and Center for the Study of Rationality,  
Hebrew University of Jerusalem

`mfeldman@huji.ac.il`

<sup>2</sup> Mathematics Division, Hebrew University of Jerusalem

`ran.tessler@mail.huji.ac.il`

<sup>3</sup> Computer Science Division, Hebrew University of Jerusalem

`yoavwilf@gmail.com`

**Abstract.** We study a “principal-agent” setting in which a principal motivates a team of agents to participate in her project (e.g., friends in a social event or store owners in a shopping mall). A key element in our model is the *externalities* among the agents; i.e., the benefits that the agents gain from each others’ participation. Bernstein and Winter [6] devised a basic model for this setting and characterized the optimal incentive mechanism inducing full participation as a unique Nash equilibrium. Here we suggest and embark on several generalizations and extensions to the basic model, which are grounded in real-life scenarios. First, we study the effect of side payments among the agents on the structure of the optimal mechanism and the principal’s utility. Second, we study the optimal partition problem in settings where the principal operates multiple parallel projects.

## 1 Introduction

Should you pay a celebrity to attend your party? Should you have special admissions for the “rich and famous”? Would it make sense for you to reduce the rental fees for anchor stores? These are some of the questions that motivate us in this study.

Suppose you wish to organize a social event and you desire as many people to attend. Obviously, the happiness of the participants would be directly affected by the identity of the other participants, which will in turn determine your cost to attract them to join your party. Similarly, if you wish to open up a shopping mall, the willingness of chain stores you may want to attract depends heavily on the identity of the other stores located at your shopping mall and how likely they are to attract potential customers. This is the case in many other social events, economic ventures, academic conferences, and commercial projects,

---

<sup>\*</sup> This work was partially supported by the Israel Science Foundation (grant number 1219/09) and by the Leon Recanati Fund of the Jerusalem school of business administration.

whose success depends to a large extent on the identity of their participants. In many cases attracting your participants is a costly task, but you can use the benefit your potential participants obtain from each other to construct a well-designed incentive scheme that will attract your wishful participant list at the lowest possible cost.

Bernstein and Winter [6] devised a principal-agent model (henceforth the *basic model*) of multi-agent projects in which a “principal” runs a project, whose success depends on the identities of the “agents” participating in the project. In this setting, the principal motivates the agents to participate by offering them a set of payments, which together put the agents in an equilibrium of the induced game. At the heart of the model lies the externalities structure, representing the benefits and losses the agents gain or incur from each others’ participation. The principal’s problem, our problem in this paper, is that of designing the optimal set of payments; i.e., the payments that maximize the principal’s net benefit.

The design of optimal incentives consist of two stages, namely the *selection* problem, in which the principal selects the set of agents for her project, and the *participation* problem, in which the principal introduces a set of payments in order to induce the participation of the selected group. The optimal solution can be obtained by first characterizing the optimal incentive scheme that induces the participation of a given set of agents (i.e., the participation problem), then working backward to determine the optimal set of agents (i.e., the selection problem). In this paper, as in Bernstein and Winter [6], the focus is on the participation problem. An optimal solution to the participation problem is a vector of payments offered by the principal that sustain full participation of the given set at minimal total cost.

The externalities among the  $n$  agents are represented by an  $n$  by  $n$  matrix  $w$ , where entry  $w_{i,j}$  represents the extent to which agent  $i$  is attracted to the project when agent  $j$  participates. We denote this entry by  $w_i(j)$ . Under additive externalities, the benefit to a participant  $i$  when the set  $S$  participates is  $\sum_{j \in S} w_i(j)$ . Bernstein and Winter characterized the optimal incentive mechanism sustaining full participation in a unique Nash equilibrium (NE) under this setting.

In this paper we study two natural extensions to the basic model, namely side payments and parallel projects. Some of the proofs are deferred to the full version. The basic model [6] assumes that payments transfer only between the principal and the agents while no communication exists among the agents themselves. Here we consider the scenario in which the agents can encourage or discourage each other to participate through payments (in addition to the offers made by the principal). The motivation for this extension is very natural; if agent  $i$  earns or loses from  $j$ ’s participation, he might be willing to pay in order to affect it. Thus, agents may be willing to transfer side payments among themselves to induce a particular outcome. For example, an agent running a store at a shopping center might be willing to pay an anchor store to co-locate at the same center in order to attract potential consumers. Alternatively, the store owner might be willing to pay another store for not participating if he has a reason to believe it might cause him some economic loss due to competition.



Additionally, when agents coordinate their actions, Nash equilibrium is not a well-suited solution concept. A more natural solution concept to consider in this case is the strong equilibrium (SE) solution concept, where no coalition can deviate such that each coalition member is strictly better-off.

We find that if the principal wishes to induce full participation as a unique strong equilibrium, the principal's utility in scenarios with side payments never exceeds her utility in the basic model, and it is equal if and only if all the externalities are positive. Thus, if some agents incur losses from others' participation, the principal strictly loses if side payments can take place. This is interesting in light of the fact that side payments can never be realized in an equilibrium. It is their virtual existence which affects the agents' behavior in equilibrium.

We also consider side payment in dynamic settings. We find that the principal's utility when inducing a subgame perfect equilibrium (SPE) of full participation is the same in the basic model and under side payments, but when side payments take place, a wide range of optimal mechanisms exist, compared to a unique optimal mechanism in the basic model. Within dynamic settings, we also consider a perturbed game, which demonstrates scenarios in which some agents might get "cold feet" and decide irrationally not to participate with a small probability. We compute the optimal mechanism in the perturbed game under side payments, and show that it can be attained only if side payments are allowed.

Finally, we consider the case in which the principal runs multiple parallel projects and wishes to partition the agents among the different projects in an optimal manner. We show that if all the externalities are positive, it is optimal for the principal to concentrate all the agents in a single project. In contrast, under negative externalities, it is always optimal for the principal to spread out the agents among all the projects, but it is NP-hard to compute the optimal partition. On the positive side, based on well-known approximation results, the optimal partition can be approximated up to a constant factor.

The general approach of contracting with externalities is closely related to [9,10]. Our approach is also related to the analysis of optimal incentive schemes under hidden action of the agents considered in [12] and later in [3,4]. Like in the papers above, the principal offers payments to the agents in order to motivate them to follow some desired outcome in equilibrium. Type-dependent externalities were also considered in [7] in the context of single-item auctions.

## 2 Model and Preliminaries

The basic model of multi-agent projects is given by a tuple  $\langle N, w, c \rangle$ .  $N$  is a set of  $n$  agents,  $w$  is the externalities matrix. Each agent has a binary decision whether to participate in the principal's project or not, where  $c = (c_1, \dots, c_n)$  is the outside-option vector of the agents if they choose not to participate. For simplicity of presentation we assume that the outside option is uniform across agents and equals zero, but all of our results can be easily generalized.

The externalities structure,  $w$ , is given by an  $n$  by  $n$  matrix specifying the bilateral externalities among the agents, such that the entry  $w_{i,j}$  represents the

benefit agent  $i$  obtains from agent  $j$ 's participation; i.e., the extent to which agent  $i$  is attracted to the project when agent  $j$  is participating. We denote this entry  $w_i(j)$ , and also denote the subset of agents affected negatively by agent  $i$  by  $N(i) = \{j | w_j(i) < 0\}$ . Under additive preferences, agent  $i$ 's utility from participating jointly with a set  $S$  of agents is  $\sum_{j \in S} w_i(j)$  for every  $S \subseteq N$ . The externalities structure  $w$  is fixed and exogenous. In addition, we assume that agents gain no additional benefit from their own participation (i.e.,  $w_i(i) = 0$  for every  $i$ ), and that if an agent gains the same benefit if he participates or not he prefers to participate. Similarly, if an agent  $i$  gains the same benefit if agent  $j$  participates or not, he prefers that agent  $j$  participates.

In the basic model [6], the set of contracts can be described by a payment vector  $v = (v_1, v_2, \dots, v_n)$ , where agent  $i$  receives a payment of  $v_i$  if he decides to participate and zero otherwise (where  $v_i$  is not constrained in sign; i.e., the principal can either pay or charge each agent). Given a mechanism  $v$  agents face a normal form game  $G(v, w, c)$ , and each agent chooses whether to participate or not. The decision is made by all the agents simultaneously. For a given set  $S$  of participating agents, the utility of an agent  $i \in S$  is  $u_i(v, S) = \sum_{j \in S} w_i(j) + v_i$ , and the utility of an agent  $i \notin S$  is his outside option.

We consider both simultaneous and sequential models with side payments. In the simultaneous case, the principal offers the agents a payment vector  $v = (v_1, v_2, \dots, v_n)$  as above, but the agents can also offer payments to each other for participating or not participating. The agents simultaneously decide whether or not to participate in the project and how much to offer every other agent for participating or not. In the sequential case the offers are made sequentially, and consequently agents make their decisions (both their participation and their side payments decisions) sequentially. We assume a model with complete information, where each agent fully observes the history of the decisions made thus far.

### 3 Side Payments in Simultaneous Games

A mechanism in the simultaneous game is denoted by  $v = (v_1, \dots, v_n)$ , representing the payments the principal offers to the agents. A strategy for agent  $j$ , denoted  $a_j$ , consists of the agent's decision of participating or not participating and the payment for each player in order for him to participate or not. The participation decision of agent  $j$  is denoted by  $e_j \in \{0, 1\}$ , where  $e_j = 1$  iff agent  $j$  participates. The payment agent  $j$  offers agent  $i$  is based on agent  $i$ 's participation decision  $e_i$  and is denoted  $t_{j,i}(e_i)$ . Under a mechanism  $v$ , and a profile of strategies  $a = (a_1, \dots, a_n)$  the set of agents who decide to participate is denoted by  $S(v, a)$ . When clear in the context, we denote it simply by  $S$ .

The utility of agent  $i$  under mechanism  $v$ , profile  $a$  and the participation of the set  $S$  is given by

$$u_i(S, v, a) = e_i \cdot \left( v_i + \sum_{j \in S} w_i(j) \right) - \sum_{j \neq i} (t_{i,j}(e_j) - t_{j,i}(e_i)) \quad (1)$$

An optimal mechanism for the principal is a mechanism that achieves full participation in equilibrium at lowest possible cost.

A trivial, yet important observation, is that side payments can never be realized in a Nash equilibrium of a simultaneous game. Clearly, a unilateral deviation does not change the participation decision of the other agents, therefore, not paying is always a beneficial deviation.

**Observation 1.** *Side payments can never occur as part of a Nash equilibrium in a simultaneous game.*

It is easy to see that by the above observation, side payments have no effect on the optimal mechanism that induces full participation as a unique Nash equilibrium. However, when coordination among the agent is possible, a more natural solution concept to consider is a *strong equilibrium*. A strong equilibrium, introduced by Aumann [2], is a strategy profile from which no *coalition* can deviate and improve the utility of each member of the coalition.

One can easily verify that in the basic model the optimal mechanism that induces full participation as a strong equilibrium is  $v_i = -\sum_{j \neq i} w_i(j)$  for every  $i$ . In contrast, the following theorem shows that under side payments, the structure of the optimal mechanism is different.

**Theorem 2.** *The mechanism given by  $v_i = -\sum_{j \neq i} w_i(j) - \sum_{j \in N(i)} w_j(i)$  for every  $i$  is the optimal mechanism which induces full participation as a strong equilibrium. Moreover, under this mechanism, full participation is a unique strong equilibrium.*

*Proof.* We first show that in any mechanism that induces full participation as a strong equilibrium, it holds that for every  $i$ :

$$v_i \geq -\sum_{j \neq i} w_i(j) - \sum_{j \in N(i)} w_j(i). \tag{2}$$

Suppose towards contradiction that there is a mechanism  $v$  that induces full participation as a SE, in which for some agent  $i$ ,  $v_i < -\sum_{j \neq i} w_i(j) - \sum_{j \in N(i)} w_j(i)$ . The obtained profile is in particular a NE, thus by Observation 1 has no side payments. Thus, it follows from Equation 1 that  $u_i < -\sum_{j \neq i} w_i(j) - \sum_{j \in N(i)} w_j(i) + \sum_{j \neq i} w_i(j) = -\sum_{j \in N(i)} w_j(i)$ . Let  $R = -u_i - \sum_{j \in N(i)} w_j(i)$ , and  $|N(i)| = k$ . Clearly  $R > 0$ , and the sum of influences of  $i$  on the set  $N(i)$  is exactly  $\sum_{j \in N(i)} w_j(i)$ . We claim that the coalition  $T = N(i) \cup \{i\}$  has a beneficial deviation to a new profile in which agent  $i$  does not participate and each agent  $j \in N(i)$  pays  $i$  a value of  $-w_j(i) - R/(k + 1)$ . The utility of each member of  $N(i)$  after the deviation is increased by  $R/(k + 1) > 0$ . The same is true for  $i$ 's utility, which becomes  $-\sum_{j \in N(i)} w_j(i) - kR/(k + 1) = u_i + R/(k + 1)$  by the definition of  $R$ ; thus reaching a contradiction.

It is now left to show that under the above payment structure, full participation is indeed a strong equilibrium. Assume towards contradiction that this is not a strong equilibrium, and let  $C \subseteq N$  be a coalition with a beneficial deviation. Clearly, it is not beneficial for any of its members to pay any amount to an

agent that is not part of the coalition. Let us divide the set  $C$  into two subsets,  $C = C^+ \cup C^-$ , such that  $C^-$  denotes the set of agents that leave the project, and  $C^+$  denotes the set of agents staying in the project (and perhaps compensating the members of  $C^-$ ). We next show that the total gain of the coalition's members decreases, hence at least one of its members must be worse-off by joining the coalitional deviation. The total benefit before the deviation is:

$$\sum_{i \in C} u_i = - \sum_{i \in C^+} \sum_{j \in N(i)} w_j(i) - \sum_{i \in C^-} \sum_{j \in N(i)} w_j(i).$$

After the deviation, the total benefit of the members of  $C^-$  is 0, and the benefit of every player  $i$  in  $C^+$  is:

$$\begin{aligned} & - \sum_j w_i(j) - \sum_{j \in N(i)} w_j(i) + \sum_{j \notin C^-} w_i(j) \\ &= - \sum_{j \in N(i)} w_j(i) - \sum_{j \in C^-} w_i(j) \\ &\leq - \sum_{j \in N(i)} w_j(i) - \sum_{j \in C^- \wedge i \in N(j)} w_i(j). \end{aligned}$$

Summing over all the members of  $C^+$ , their total benefit is given by:

$$\begin{aligned} & - \sum_{i \in C^+} \sum_{j \in N(i)} w_j(i) - \sum_{i \in C^+} \sum_{j \in C^- \wedge i \in N(j)} w_i(j) \\ &= - \sum_{i \in C^+} \sum_{j \in N(i)} w_j(i) - \sum_{j \in C^-} \sum_{i \in C^+ \cap N(j)} w_i(j) \\ &\leq - \sum_{i \in C^+} \sum_{j \in N(i)} w_j(i) - \sum_{j \in C^-} \sum_{i \in N(j)} w_i(j), \end{aligned}$$

as required.

We next show that under the above scheme full participation is the unique strong equilibrium. Assuming towards contradiction that there is an additional SE in the game, without full participation. It must also be a NE. Let  $S \subset N$  be the set of players that participate. As we are in a NE, according to Observation □ no side payments are made, hence the total net benefit of the players is exactly the sum of benefits of  $S$ 's elements. This yields:

$$\begin{aligned} & \sum_{i, j \in S} w_i(j) - \sum_{i \in S} \left( \sum_j w_i(j) + \sum_{j \in N(i)} w_j(i) \right) \\ &= - \sum_{i \in S} \sum_{j \in N(i)} w_j(i) - \sum_{i \in S, j \notin S} w_i(j) \\ &\leq - \sum_{i \in S} \sum_{j \in N(i)} w_j(i) - \sum_{j \notin S} \sum_{i \in S \cap N(j)} w_i(j) \end{aligned}$$

$$\begin{aligned} &\leq - \sum_{i \in S} \sum_{j \in N(i)} w_j(i) - \sum_{j \notin S} \sum_{i \in N(j)} w_i(j) \\ &= - \sum_i \sum_{j \in N(i)} w_j(i) \end{aligned}$$

This shows that if everyone participates, the total utility of all the agents does not decrease. Consequently, they can redistribute the money such that all the agents are better off. Note that if their total utility remains the same, the agents are better-off in a full participation profile by the assumption on the agents' preferences.  $\square$

### 4 Side Payments in Sequential Games

Static notions like Nash or strong equilibrium are not rich enough to convey the full power of side payments. In what follows we study the participation problem in a dynamic setting, and concentrate on the subgame perfect equilibrium and a variant of the trembling hand equilibrium solution concepts. Under the sequential model, the principal induces some order on the agents, and each agent, in his turn, decides whether to participate or not and how much to pay each one of the other agents for participating. We assume that each agent, at the time of decision, knows the history of the decisions made thus far. A strategy is now a function from the decisions of the previous players and the principal's compensation scheme to the decision on participation and side payments. We restrict attention to positive externalities. In this case, paying an agent who has already decided to participate is a dominated strategy since agents cannot change their decision after their turn. The strategy of an agent is denoted as  $a_j(v) = (e_j, t_{j,j+1}, \dots, t_{j,n})$ , where  $e_j \in \{0, 1\}$  (i.e.,  $e_j = 1$  iff agent  $j$  participates), and  $t_{j,i}$  is  $j$ 's non-negative offer to agent  $i$  for participating.

#### 4.1 Subgame Perfect Equilibrium

**Definition 1.** *Given an order on the agents, a mechanism  $v$  is called the stable vector with respect to this order if  $v_1 = 0$ , and  $v_i = - \sum_{j < i} (w_i(j) + w_j(i))$  for every  $i > 1$ .*

The following theorem shows that the stable vector with respect to any order induces full participation as a subgame perfect equilibrium (SPE). Moreover, this is the optimal mechanism that induces full participation as an SPE.

**Theorem 3.** *For every order on the agents, the stable vector with respect to that order is an optimal mechanism that induces full participation as a subgame perfect equilibrium.*

*Proof (Sketch).* Let  $M = \sum_{i,j} w_i(j)$ . First observe that in any SPE of full participation the principal must pay at least  $-M$ ; otherwise the sum of the agents'

utilities is negative, thus there must be an agent who prefers to deviate. Therefore, it is sufficient to show that the stable vector induces full participation as an SPE. Consider an arbitrary order of the players and the stable vector with respect to that order. We claim that the strategy profile defined below is an SPE:

The first player's strategy is to participate and to offer each player  $i > 1$  a payoff  $w_1(i)$  to participate. The strategy of the  $i^{\text{th}}$  player for  $i > 1$  is the following: If every agent  $k < i$  entered and offered each player  $j$  such that  $j > k$  a payoff  $w_k(j)$ , the  $i^{\text{th}}$  agent does the same (i.e., participates and offers every player  $j$  such that  $j > i$  a payoff  $w_i(j)$ ). It will be shown to be an optimal action. Otherwise, player  $i$  acts according to an optimal strategy; i.e., a strategy that maximizes his net benefit (with ties broken arbitrarily). It can be shown that while the utility function is not continuous and hence compactness arguments cannot be directly applied, an optimal strategy does exist, and thus the definition is well defined.

Either if all the agents decided to participate or only a subset, it can be shown that their total utility is not positive. If there exists an agent with negative utility, clearly it cannot be an SPE. Otherwise, all the agents gain zero. In this case, however, based on the assumption, each agent prefers full participation. Using backward induction it follows that each agent participates and makes offers as described above, which ensures a full participation with zero utility.  $\square$

The following proposition considers the optimal mechanism in the basic model which induces full participation as an SPE.

**Proposition 1.** *Under positive externalities, there is a unique optimal mechanism which induces an SPE of full participation in the basic model, and the principal's utility under this mechanism is identical to the principal's utility in the optimal mechanism in the simultaneous side payments model.*

*Proof (Sketch).* An argument similar to the one given in the proof of Theorem 3 (stating that the stable vector allows an SPE of full participation), can be carried out for the vector described in Theorem 2 as well. This vector guarantees a total gain of 0 to all the group. If it does not hold that all the agents participate, at least one of the participating agents will have a negative utility. Thus this mechanism induces an SPE of full participation in the side payments model. As the variety of possible strategies in the basic model is narrower (there are less possible action options), it follows easily from a simple backward induction argument that this mechanism induces an SPE in that model too. Moreover, in the basic model this is the only optimal mechanism that induces an SPE, as a cheaper mechanism cannot even induce a NE of full participation, since otherwise the sum of the agents' utilities will be negative, and thus at least one of these utilities will be negative. It can be easily seen that the total payments under the above compensation scheme and the stable vectors is the same. They both equal  $-\sum_{i,j} w_i(j)$ .  $\square$

We conclude that under positive externalities, the principal's utility under the optimal mechanisms achieving full participation as an SPE with and without

side payments is the same. Yet, while in the basic model, there is a unique optimal mechanism, side payments allow for a wide range of optimal mechanisms, e.g. the stable vectors with respect to any order of the agents. The following observation, which is used in the sequel, considers the utility of the agents under these mechanisms:

**Observation 4.** *In every optimal mechanism that induces an SPE of full participation the utility of each agent is zero.*

### 4.2 Perturbation-Proof Equilibrium

The above analysis assumes full rationality of the agents. Yet, if one of the agents gets cold feet and decides irrationally not to participate, it might be in the best interest of some of the other agents to quit as well. We next come up with an equilibrium notion that is stable against such situations. We first define a perturbed strategy. Given a strategy  $a_j$  of player  $j$  in the sequential game, we denote by  $a_j^\varepsilon$  an  $\varepsilon$ -perturbed version of  $a_j$ , where agent  $j$  plays according to  $a_j$  with probability  $1 - \varepsilon$ , and with probability  $\varepsilon$  he does not participate and offers zero payments to the other agents, independent of their actions and the original strategy  $a$ . We also denote  $a^\varepsilon = (a_1^\varepsilon, \dots, a_n^\varepsilon)$ . With this we are ready to define the notion of perturbation-proof equilibrium.

**Definition 2.** *Given a mechanism  $v$ , a strategy profile  $a = (a_1, \dots, a_n)$  is a perturbation proof equilibrium (PPE) with respect to  $v$  if  $a$  is an SPE with respect to  $v$ , and there exists a positive value  $c$  such that for every  $0 < \varepsilon < c$ , and for every agent  $i$ , if every agent  $k < i$  played according to  $a_k$ , then if every agent  $k > i$  plays according to  $a_k^\varepsilon$ , playing  $a_i$  is an optimal strategy for agent  $i$ .*

We wish to find a mechanism  $v$  under which there is an equilibrium of full participation even if every agent does not enter the project (irrationally) with some small probability. That is, a mechanism  $v$  under which there exists a PPE with full participation.

Given the stable vector with respect to some order on the agents, consider the following strategy profile: For every agent  $j$ , if every player  $k < j$  acted according to the strategy defined in the proof of Theorem 3, agent  $j$  also follow this strategy. If there exists an agent  $k < j$  who did not follow that strategy, agent  $j$  plays an optimal strategy (as computed using backward induction by computing expected benefits given  $\varepsilon$ ). Note that agents cannot change their strategies after their turn.

**Theorem 5.** *Under the above strategies full participation is a perturbation proof equilibrium. Moreover, every optimal mechanism which allows a perturbation proof equilibrium of full participation must be a stable vector with respect to some order.*

Note that the optimal mechanism (i.e., the one achieving full participation at minimum cost) cannot be attained in the absence of side payments (yet can get arbitrarily close to the minimum payment).

## 5 Operating Multiple Projects

In this section we study scenarios in which the principal operates  $k$  parallel projects, for an exogenously given  $k$ , and wishes to optimally divide the agents among them; i.e., find the mechanism with the cheapest total payments that induces full participation as a Nash equilibrium. We assume that the principal can charge (or pay) differently for different projects (even from the same agent). In addition, the principal does not necessarily have to operate all  $k$  projects, but he cannot operate more than  $k$  projects. This is the case, for example, in situations where an economic entity has a concession for opening  $k$  shopping malls.

As before, a mechanism is denoted by  $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_n)$ , but here  $\mathbf{v}_i$  is itself a vector of payments for agent  $i$  for participating in the different projects; i.e.,  $\mathbf{v}_i = (v_i^1, \dots, v_i^k)$ , where  $v_i^j$  for  $1 \leq j \leq k$  is the payment the principal offers agent  $i$  for participating in project  $j$ . Given a mechanism  $v$ , each agent chooses whether or not to participate and in which project. The participation vector is denoted by  $\mathbf{S} = (S_1, \dots, S_k)$ , where  $S_j$  is the set of agents participating in project  $j$ . We also denote by  $n_j$  the number of agents participating in project  $j$ ; i.e.,  $n_j = |S_j|$  for  $1 \leq j \leq k$ .

We assume that the principal wishes to induce full participation of the agents at the lowest possible cost. That is, the principal wishes to minimize the following expression:

$$P = \sum_{j=1}^k \sum_{i \in S_j} v_i^j.$$

The utility of agent  $i$  under mechanism  $v$  and participation vector  $S$  if he chooses to participate in project  $j$  is given by  $v_i^j + \sum_{l \in S_j} w_i(l)$ .

In what follows we present some observations regarding the cases of positive externalities and negative externalities.

*Positive Externalities.* We first observe that in the case of positive externalities, it is optimal for the principal to induce an equilibrium in which all the agents participate in the same project.

**Observation 6.** *Under positive externalities, the optimal mechanism that induces full participation as a NE motivates all the agents to participate in the same project.*

By the last observation, the problem of finding the optimal mechanism can be solved by solving  $k$  copies of the optimal mechanism problem in the basic model [6]. The last problem can be shown to be NP-hard by a reduction from the feedback arc set problem.

*Negative Externalities.* The exact opposite case is where the externalities are all negative. In this case the principal wishes to distribute the agents as much as possible, as stated in the following observation.



**Observation 7.** *Under negative externalities, in the optimal mechanism that induces full participation as a NE, if  $n \geq k$  then  $n_i > 0$  for every  $j \in \{1, \dots, k\}$ .*

Given some partition of the agents into  $k$  subsets, the principal should pay each agent the sum of the externalities he incurs from all the agents in his subset. Given the last observation, the problem can be described as an undirected graph in which the agents are the vertices and the edge between any two vertices has weight equals to the absolute value of the sum of their corresponding mutual externalities. The problem of finding the optimal mechanism that sustains full participation as a NE is actually that of finding a partition of the set of vertices into  $k$  subsets of minimal total weight. This problem is equivalent to the MAX-K-CUT problem: finding a partition of the vertices into  $k$  subsets with maximal total weight of crossing edges.

The MAX-K-CUT problem is known to be NP-hard. It can be approximated within a constant factor [8] but does not admit a polynomial-time approximation scheme (PTAS) unless  $P=NP$  [11]. In some special cases (e.g., dense graphs) the problem admits a PTAS [5,11].

## References

1. Arora, S., Karger, D., Karpinski, M.: Polynomial Time Approximation Schemes for Dense Instances of NP-Hard Problems. *Journal of Computer and System Sciences* 58(1), 193–210 (1999)
2. Aumann, R.: Acceptable Points in General Cooperative n-Person Games. In: *Contributions to the Theory of Games*, vol. 4 (1959)
3. Babaioff, M., Feldman, M., Nisan, N.: Combinatorial Agency. In: *ACM Electronic Commerce* (2006)
4. Babaioff, M., Feldman, M., Nisan, N.: Mixed Strategies in Combinatorial Agency. In: *2nd international Workshop on Internet and Network Economics* (2006)
5. Fernandez de la Vega, W.: Max-Cut Has a Randomized Approximation Scheme in Dense Graphs. *Random Struct. Algorithms* 8(3), 187–198 (1996)
6. Winter, E., Bernstein, S.: How to Throw a Party: Multi-Agent Initiatives with Type Dependent Externalities. *Center for the study of rationality: discussion paper* (2008)
7. Jehiel, P., Moldovanu, B.: Strategic Nonparticipation. *RAND Journal of Economics* 27, 84–98 (1996)
8. Frieze Alan, M., Mark, J.: Improved Approximation Algorithms for Max k-Cut and Max Bisection. In: *Proceedings of the 4th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pp. 1–13. Springer, London (1995)
9. Segal, I.: Contracting with Externalities. *Quarterly Journal of Economics* 114, 337–388 (1999)
10. Segal, I.: Coordination and Discrimination in Contracting with Externalities: Divide and Conquer? *Journal of Economic Theory* 113, 337–388 (2004)
11. Viggo, K., Sanjeev, K., Jens, L., Alessandro, P.: On the Hardness of Approximating Max k-Cut and its Dual. *Technical report* (1997)
12. Winter, E.: Incentives and Discrimination. *American Economic Review* 94, 764–773 (2004)

# Nash Dynamics in Congestion Games with Similar Resources

Anand Bhargat, Tanmoy Chakraborty, and Sanjeev Khanna

Department of Computer and Information Science  
University of Pennsylvania

**Abstract.** We study convergence of  $\epsilon$ -Nash dynamics in congestion games when delay functions of all resources are similar. Delay functions are said to be similar if their values are within a polynomial factor at every congestion level. We show that for any  $\epsilon > 0$ , an  $\epsilon$ -Nash dynamics in symmetric congestion games with similar resources converge in steps polynomial in  $1/\epsilon$  and the number of players are resources, yielding an FPTAS. Our result can be contrasted with that of Chien and Sinclair [3], which showed polynomial convergence result for symmetric congestion games where the delay functions have polynomially bounded jumps. Our assumption of similar delay functions is orthogonal to that of bounded jumps in that neither assumption implies the other. Our convergence result also hold for several natural variants of  $\epsilon$ -Nash dynamics, including the most general *polynomial liveness* dynamics, where each player is given a chance to move frequently enough. We also extend our positive results to give an FPTAS for computing equilibrium in asymmetric games with similar resources, in which players share  $k$  distinct strategy spaces for any constant  $k$ .

We complement our positive results by showing that computing an exact pure Nash equilibrium in symmetric congestion game with similar resources is PLS-complete. Furthermore, we show that for any  $\epsilon > 0$ , all sequences of  $\epsilon$ -Nash dynamics takes exponential steps to reach an approximate equilibrium in general congestion games with similar resources, as well as in symmetric congestion games with two groups of similar resources.

## 1 Introduction

A fundamental problem in algorithmic game theory is the computational complexity of computing a *Nash equilibrium* for various classes of *non-cooperative games*. Nash [8] showed that a *mixed* Nash equilibrium always exists in any *normal-form* game with a finite number of players and strategies. In a mixed Nash equilibrium, players are allowed to play randomized strategies, and they wish to maximize their expected payoff. In contrast, a pure Nash equilibrium, where players play deterministic strategies only, may not exist in all games. However, there are natural classes of games which always have a pure Nash equilibrium. A prominent class with this property is *congestion games*, defined by Rosenthal [9].

**Nash Dynamics:** Nash equilibrium is especially interesting in systems where selfish agents make their own decisions, in a decentralized fashion. The joint strategy of all the players is said to define a *state* of the game. Given a state of the game, a *better response* of a player is any of its strategies that has a higher payoff (or less cost) if the strategies of the other players remain unchanged. A pure Nash equilibrium is a state where no player has a better response. *Nash dynamics* is an evolution of the game over time that starts from some *initial state*, and in each step, a player switches to a better response given the current state. We refer to each such step as an *improving move*. Clearly, Nash dynamics stops (converges) when no player has a better response, that is, the dynamics has reached a pure Nash equilibrium. Nash dynamics is a well-studied model of selfish decentralized decision-making in a game, where players always play deterministic strategies (eg. [12]).

**Congestion Games:** A *congestion game* is an  $n$ -player game with  $m$  resources, and each player is assigned several strategies that it may choose from, where a strategy of a player is a subset of the resources. A player playing a strategy  $s$  is said to be *using* the resources in  $s$ . Each resource  $e$  is also associated with a non-negative increasing integral *delay function*  $d_e(f_e)$ , where  $f_e$  is the number of players using  $e$ , also called the *congestion* of  $e$ . Each player that uses  $e$  suffers a delay of  $d_e(f_e)$  on  $e$ , and the total delay of a player playing strategy  $s$  is  $\sum_{e \in s} d_e(f_e)$ . Every player must play one strategy, and selfishly seeks to *reduce* its own delay. Rosenthal [9] showed that every congestion game has a pure Nash equilibrium.

**Computational Complexity:** Nash dynamics can be viewed as a natural algorithm to compute pure Nash equilibrium in congestion games. However, the number of steps required to converge to an equilibrium from some initial state may not be polynomially bounded, and so this algorithm need not run in polynomial time. Johnson et. al. [7] introduced a complexity class PLS (polynomial-time local search) related to local search problems, and Nash dynamics implies that computing a pure Nash equilibrium for congestion games belong to PLS. In fact, Fabrikant et. al. [5] proved that the computing a pure Nash equilibrium in congestion games is PLS-complete. Moreover, they showed that computing an equilibrium in *symmetric* congestion games, where every player has the same strategy space, is PLS-complete. The PLS-completeness proofs for these problems also show that there exists initial states such that any sequence of improving moves from the initial state to an equilibrium state is exponential in length.

**Approximate Nash Equilibrium and Greedy Dynamics:** Since it is considered to be unlikely that  $PLS = P$ , much research has focused on computing approximate equilibria of congestion games. For any  $\epsilon > 0$ , an  $\epsilon$ -*Nash equilibrium* (see, eg. [10, 4, 1, 6, 3, 11, 2]) is defined to be a state where no player can unilaterally change its own strategy so that her delay decreases by at least an  $\epsilon$  fraction of its current delay. In other words, a player switches strategy only if there is a significant improve in delay. We note here that the definition of approximate equilibrium has varied in literature, though most of them are equivalent, with different parameters. In particular,  $\epsilon$ -Nash equilibrium in the language of Chien

and Sinclair [3] is a  $\frac{1}{1-\epsilon}$ -approximate equilibrium in the language of Skopalik and Vöcking [11]. In this paper, we follow the terminology of [3]; precise definitions are given in Section 2. When searching for  $\epsilon$ -Nash equilibrium using Nash dynamics, it is natural to assume that all players shall make  $\epsilon$ -moves only, that is, a player may switch to a better response during Nash dynamics if the move reduces its delay by at least an  $\epsilon$  fraction of its current delay. Such greedy dynamics, called  $\epsilon$ -Nash dynamics, were first studied in [4].

A resource  $e$  is said to satisfy the *bounded jump condition* if its delay increases by a factor of at most  $\beta$  with the addition of a new player using  $e$ , for some parameter  $\beta$ . Chien and Sinclair [3] showed that while computing exact equilibrium in these games remains PLS-complete, however, starting from any state  $s$ , an  $\epsilon$ -Nash dynamics converges in time that is polynomial in the size of the input,  $\beta$ , and  $1/\epsilon$ , thus yielding an FPTAS (fully-polynomial time approximation scheme) for polynomially bounded  $\beta$ . Note that Nash dynamics does not specify the method to select a move if several players can make an  $\epsilon$ -move. Chien and Sinclair [3] showed that their fast convergence result holds for several natural methods, such as largest gain, largest relative gain, heaviest first, and the most general *polynomial liveness*, where it is only required that every player is given an opportunity to make an  $\epsilon$ -move, if available, in every polynomially long sequence of steps. Finally, Chien and Sinclair [3] also showed fast convergence of  $\epsilon$ -Nash dynamics for symmetric congestion games where only a constant number of edges have unbounded jump, again yielding an FPTAS. The last result easily implies fast convergence in games where the players share a constant number of distinct strategy spaces, a generalization of symmetric games.

Progress has also been made on the inapproximability of equilibria in congestion games. Skopalik and Vöcking [11] showed that the fast convergence results in [3] does not extend to arbitrary asymmetric congestion games with bounded jumps, by constructing such a game where there exists an initial state such that any sequence of  $\epsilon$ -moves from this state takes exponential number of steps to converge. They also showed that for asymmetric congestion games with unbounded jumps (that is, general congestion games), computing an  $\epsilon$ -Nash equilibrium is PLS-complete for any polynomially computable  $\epsilon < 1$ . The last result is true for symmetric congestion games as well.

**Our Results and Techniques:** In this paper, we study Nash dynamics in congestion games with similar resources. While the bounded jump assumption of [3] is well-motivated, it disallows a natural class of resources, namely, resources with a *threshold behavior*, that is, delay of a resource may suddenly jump once a certain congestion is reached. Since the results of [11] show that computing an approximate equilibrium in symmetric congestion games is hard, elimination of the bounded jump condition makes it necessary to impose some other restriction. We make a novel assumption that the delay functions of the resources are *similar*. A set of delay functions are said to be *similar*, with a *similarity factor* of  $Q$ , if for any positive integer  $k$ , the maximum and minimum delay values at congestion  $k$  for all resources are within a factor  $Q$ . We assume that  $Q$  is polynomial in the size of the input. In particular,  $Q = 1$  implies that the delay functions are

identical. Note that similarity of resources does not imply that the jumps are bounded, and vice versa. The similar resources assumption is of interest in many applications, for example, when the roads in a traffic routing game have similar dimensions, or the wires in a routing network are made of the same material. Moreover, constructions of hard instances such as those in [11] use resources whose delay functions have very different properties, and the use of such distinct delay functions seem intrinsic to their proof. So it is interesting from a theoretical perspective that very different delay functions are in fact necessary for proving the hardness results. Intuitively, we require that any threshold behavior in the delay values, should happen at a similar congestion value for all resources.

We show fast convergence of  $\epsilon$ -Nash dynamics in symmetric congestion games with similar resources. We show that several natural dynamics, including the most general condition of polynomial liveness, converge in time that is polynomial in the size of the input,  $Q$  and  $1/\epsilon$ , thus giving an FPTAS. We extend our FPTAS result to the case when there are only  $k$  distinct *types* of players for some constant  $k$ . Our technique is based on splitting the delay functions into multiple levels. Within the same level, the delay function has polynomially bounded jump for each increase in congestion. When the delay function jumps by a large factor, the level changes. Intuitively, the level change corresponds to the threshold behavior of the delay function.

We match our positive results with appropriate negative results showing that computing exact equilibrium in these games in PLS-complete. Moreover, if we relax any of the assumptions, then we show that one cannot find even approximate equilibrium by Nash dynamics. That is,  $\forall 0 < \epsilon < 1$ , there exists a game and a state in the game from which all sequences of improving moves take exponential number of steps to reach an  $\epsilon$ -Nash equilibrium. In particular, we show this last result for general congestion games (asymmetric, with arbitrarily many types of players) with similar resources, as well as symmetric congestion games with only *two types of resources*, that is, the resources can be divided into two groups of similar resources. This result is derived by a modification of the construction in [11] to show such a result for asymmetric games with bounded jumps. Thus, with respect to these assumptions, our results are tight. A summary of our results is given in Table 1.

**Table 1.** Summary of our results

Class of game	Computing exact equilibrium	FPTAS	Arbitrary approximation factor
Symmetric game, similar resources	PLS-complete	Yes	Yes
Constant types of players, similar resources	PLS-complete	Yes	Yes
Arbitrary types of players, similar resources	PLS-complete	All sequences exponential	All sequences exponential
Symmetric game, two types of resources	PLS-complete	All sequences exponential	All sequences exponential

**Organization.** The rest of the paper is organized as follows: we formally define congestion games and its various restrictions in Section 2. In Section 3, we give an efficiently computable sequence of  $\epsilon$ -moves that converges fast in symmetric games with similar resources. In Section 4, we show that several natural dynamics converge fast in these games. In Section 5, we generalize our positive result to games with constant types of players and similar resources. Finally, we present our negative results in Section 6.

## 2 Preliminaries

A *game* consists of a finite set of players  $p_1, p_2 \dots p_n$ . Each player  $p_i$  has a finite set of strategies  $S_i$  and a *cost or delay* function  $c_i : S_1 \times \dots \times S_i \times \dots \times S_n \rightarrow \mathbb{N}$  that it wishes to minimize. We refer to  $S_i$  as the *strategy space* of  $p_i$ . A game is called *symmetric* if all strategy spaces  $S_i$  are identical. Two players are said to be of the same *type* if they have the same strategy space. An  $n$ -tuple of strategies  $s = (s_1, s_2 \dots s_n) \in S_1 \times \dots \times S_i \times \dots \times S_n$  is called a *state* of the game, where player  $p_i$  plays strategy  $s_i \in S_i$ . A state  $s$  is a *pure Nash equilibrium* if for each player  $p_i$ ,  $c_i(s_1, \dots, s_i, \dots, s_n) \leq c_i(s_1, \dots, s'_i, \dots, s_n)$  for every  $s'_i \in S_i$ . Thus in an equilibrium state, no player can improve its cost by unilaterally changing its strategy.

*Congestion games* is a class of games where players' costs are based on the shared usage of a common set of resources  $R = \{r_1, r_2 \dots r_m\}$ . The strategy set of a player  $p_i$  is  $S_i \subseteq 2^R$ , an arbitrary collection of subsets of  $R$ . Each resource  $r \in R$  has a non-decreasing delay function  $d_r : \mathbb{N} \rightarrow \mathbb{N}$  associated with it. If  $j$  players are using a resource  $r$ , each of these players incurs a delay of  $d_r(j)$  on resource  $r$ . The *delay* incurred by a player  $p_i$  in a state  $s = (s_1, \dots s_n)$  is the sum of the delays it incurs on each resource in its strategy, that is,  $c_i(s) = \sum_{r \in s_i} d_r(f_s(r))$ , where  $f_s(r)$  is the number of players using resource  $r$  in state  $s$ , that is,  $f_s(r) = |\{j : r \in s_j\}|$ .

Given any state  $s = (s_1 \dots s_i \dots s_n)$  of a congestion game, the *potential function*  $\phi$  of the game ([9]) is defined to be  $\phi(s) = \sum_r \sum_{i=1}^{f_s(r)} d_r(i)$ . Note that if  $p_i$  unilaterally changes its strategy from  $s_i$  to  $s'_i$ , then the potential of the resulting state  $s' = (s_1 \dots s'_i \dots s_n)$  is  $\phi(s') = \phi(s) + (c_i(s') - c_i(s))$ , that is, the change in potential is equal to the change in the delay of  $p_i$ . Thus the potential decreases whenever any player makes a move which decreases its delay.

Given a state  $s = (s_1 \dots s_n)$ , a *better response* strategy of a player  $p_i$  is any strategy  $s'_i \in S_i$  such that if the player switches its strategy from  $s_i$  to  $s'_i$  its delay decreases. A *best response* strategy is a better response strategy that maximizes this decrease. Changing the state by making a player switch to a better response strategy is called an *improving move*. A *Nash dynamics* starting from some initial state  $s$  refers to a sequence of states such that only one player changes strategy in each step, and each such change is an improving move for the player with respect to the preceding state.

**Definition 1.** ([3]) *An  $\epsilon$ -Nash equilibrium is a state  $s$  such that for any strategy  $s'_i \in S_i$ , if  $s' = (s_1, s_2, \dots, s'_i, \dots s_n)$ , then  $c_i(s') > (1 - \epsilon)c_i(s)$ , for all  $1 \leq i \leq n$  i.e. no player can decrease its delay at least by a factor of  $\epsilon$  by unilaterally*

changing its strategy. An  $\epsilon$ -move is an improving move where the change of strategy causes the player's delay to decrease by at least  $\epsilon$  fraction, that is, if state  $s$  changed to state  $s'$ ,  $c_i(s') \leq (1 - \epsilon)c_i(s)$ . An  $\epsilon$ -Nash dynamics is Nash dynamics where each improving move is an  $\epsilon$ -move.

### 2.1 Similar Resources

We now formalize the notion of *similar* resources.

**Definition 2.** A congestion game is said to have similar resources if there exists a function  $Q$  that is polynomial in  $n, m$ , and a function  $d$  such that given an instance of the game with  $n$  players and  $m$  resources for any  $m, n \geq 1$ , for every resource  $r$ ,  $d_r(k) \in [d(k), d(k)Q]$  for  $1 \leq k \leq n$ . The function  $d$  is said to be a common delay function for the resources. The function  $Q$  is the similarity factor for the delay functions of the resources.

The concept of similar resources is a generalization of identical resources. In the case of identical resources i.e. when all resources have same delay function,  $Q = 1$ . The function  $Q$  is the measure of how closely related delay functions for the resources are. For the sake of smooth presentation of our proofs, which involve a recursive argument, we shall prove our algorithmic results for a slightly more general concept of *synchronized* resources.

**Definition 3.** A congestion game is said to have synchronized resources if there exists  $Q$  which is polynomial in  $(n, m)$  and a function  $d$  such that such that  $\forall m, n \geq 1$ , in a game with  $n$  players and  $m$  resources, for every resource  $r$ , there exists an integer  $\text{shift}(r)$  such that  $\forall k \geq 1$ ,  $d_r(k) \in [d(k + \text{shift}(r)), d(k + \text{shift}(r))Q]$ . The function  $d$  is said to be a common function for the resources.

Thus a congestion game with similar resources is a special case of congestion game with synchronized resources with  $\text{shift}(r) = 0$  for every resource  $r$ . The concept of synchronized resources allow resources to synchronize with the common delay function at a different congestion value.

**Definition 4.** Let  $s^1, s^2 \dots s^k$  be a sequence of states derived by an  $\epsilon$ -Nash dynamics. We say that a good event occurs in the sequence if  $\phi(s^1) - \phi(s^k) \geq \frac{\epsilon \phi(s_1)}{32m^2Q^2n}$ , where  $Q$  is polynomial in  $(m, n)$ . If a good event occurs in a sequence comprising only of 2 states, we call the associated  $\epsilon$ -move a good  $\epsilon$ -move.

As in [3], the following simple lemma will be useful in our analysis.

**Lemma 1.** Consider a sequence of states derived by  $\epsilon$ -Nash dynamics. Consider any partition of the sequence into segments such that a good event occurs in each segment. If the first state of the sequence is  $s^{\text{begin}}$ , then there can be at most  $O(m^2Q^2n\epsilon^{-1} \log \phi(s^{\text{begin}}))$  segments.

We now list some simple lemmas about *congestion games with synchronized resources*, whose proofs are deferred to the full version.



**Definition 5.** Let  $G$  be a congestion game with synchronized resources with common function  $d$ . We define levels of congestion based on the jumps in value of  $d$ . Let  $\text{level}(1) = 1$ , and if  $d(i+1) \leq 2mQd(i)$ , then  $\text{level}(i+1) = \text{level}(i)$ , otherwise  $\text{level}(i+1) = \text{level}(i) + 1$ . Note that  $\text{level}()$  is a non-decreasing function. In a state  $s$ , if  $j$  players are using a resource  $r$ , then we define the level of  $r$  to be  $\text{level}(j + \text{shift}(r))$ . We define the level of a player  $p_i$  in state  $s$  to be  $\max_{r \in s_i} \text{level}(r)$ .

**Lemma 2.** For any  $\ell > 1$ , the minimum delay incurred by any player in level  $\ell$ , is at least twice the maximum delay that can be incurred by any player in level  $(\ell - 1)$  or below.

**Corollary 1.** Whenever a player  $p_i$  lowers its level by making a move, then the move decreases the delay of player  $p_i$  by at least a factor of  $1/2$ .

**Corollary 2.** Suppose a player  $p_i$  at level  $\ell$  in current state  $s$  makes an improving move, and let  $s'$  be the state of the game after the move. Then for every player  $p_j$  at a level above  $\ell$  in  $s$ , its level remains the same in  $s'$ . For every player  $p_j$  in level  $\ell$  in  $s$ , its level does not increase in going to  $s'$ .

**Lemma 3.** Let  $p_i$  and  $p_j$  be two players with the same strategy space such that  $p_i$  has level  $\ell$  and  $p_j$  has level at most  $(\ell - 2)$  in some state  $s$ , for some integer  $\ell$ . Then  $p_i$  has an  $\epsilon$ -move from  $s_i$  to  $s_j$ , for  $0 < \epsilon \leq 1/2$ .

**Lemma 4.** Let  $0 < \epsilon \leq 1/2$ . Let  $p_j$  be a player in level  $\ell$  in some state  $s$ , and let  $p_i$  be another player in level  $\ell$  that has the same strategy space as  $p_j$ , and has an  $\epsilon$ -move. Then there exists an  $\epsilon$ -move by some player in level  $\ell$  with current delay at least  $(c_j(s))/(8m^2Q^2)$ . Consequently, such a move decreases the potential by at least  $\epsilon(c_j(s))/(8m^2Q^2)$ . If  $p_j$  were a player with maximum delay in state  $s$ , then this move would be a good  $\epsilon$ -move.

*Proof.* Note that if  $p_j$  is a player with maximum delay in state  $s$ , then  $c_j(s) \geq \phi(s)/n$ . If  $c_i(s) \geq \frac{c_j(s)}{8m^2Q^2}$ , then the said  $\epsilon$ -move of  $p_i$  is itself the required move. Otherwise, we show that  $p_j$  can make an  $\epsilon$ -move to  $s_i$ .

There is a resource in  $s_j$  with delay at least  $8mQ^2c_j(s)$ . Also, delay of each resource in  $s_i$  is at most  $c_i(s)$ . Here we will show that resources in level  $\ell$  in  $s_i$  will not increase their level by increase in congestion by one. Since level of both  $p_i$  and  $p_j$  is  $\ell$ , there is a delay value  $\geq 8mQ^2c_i(s)$  which is in level  $\ell$  as well as there is a delay value  $\leq c_i(s)$  which is in level  $\ell$ . It implies increasing congestion by one for a resource in level  $\ell$  in  $s_i$  will not increase the level of that resource beyond  $\ell$ . Consider the delay on each resource in  $s_i$ , if the congestion on them were increased by one more player. It follows from the definition of levels that if such a resource were in level  $\ell$  in  $s$ , then its delay increases by a factor of at most  $2mQ^2$ , and so is at most  $2mQ^2c_i(s)$ . If a resource were in a lower level in  $s$ , then it either remains in a lower level or just reaches level  $\ell$ . In either case, the fact that the resources are synchronized implies that the delay on the resource is at most  $Qc_i(s)$ , since there is a resource in  $s_i$  with level  $\ell$ , and delay at most  $c_i(s)$ .



Since  $s_i$  has at most  $m$  resources, it follows that the delay faced by  $p_j$  when it switches to  $s_i$  is at most  $2m^2Q^2c_i(s) + mQc_i(s) < c_j/2 \leq (1 - \epsilon)c_j$ , completing the proof.

Our algorithms are recursive, for which we shall define a *reduced game*.

**Definition 6.** *Let  $P$  be an arbitrary subset of players in a congestion game  $G$ , and let  $s$  be a state of the game. Then a reduced game  $G(s, P)$  is defined as a game derived from  $G$  by removing all players in  $P$  and modifying the delay function of each resource  $r$  from  $d_r$  to  $d'_r$  such that if  $x$  is the number of players in  $P$  using  $r$  in  $s$ , then  $d'_r(k) = d_r(k + x)$ .*

**Lemma 5.** *If  $G$  is a congestion game with synchronized resources, then any reduced game  $G(s, P)$  is also a congestion game with synchronized resources, with the same common function. Moreover, if the players  $p_i \notin P$  play  $s_i$  in  $G(s, P)$ , then the level of each player and resource is the same as that in state  $s$  of the game  $G$ .*

### 3 FPTAS for Symmetric Games with Similar Resources

In this section, we give an FPTAS for computing an  $\epsilon$ -Nash equilibrium for symmetric congestion games with similar resources.

**Theorem 1.** *For any  $\epsilon \in (0, 1/2]$ , given a symmetric congestion game with similar resources in some initial state  $s^{begin}$ , there exists a poly-time computable sequence of  $O((nm^2Q^2 \log(\phi(s^{begin}))/\epsilon^2)) \epsilon$ -moves leading to an  $\epsilon$ -Nash equilibrium.*

We shall now prove Theorem 1. As mentioned earlier, we shall prove the result for congestion games with synchronized resources. Let  $d$  be the common function for the resources. By Lemma 1, it suffices to show that the behavior of the algorithm can be partitioned into segments each of which contains  $O(nm^2Q^2\epsilon^{-1} \log(\phi(s^{begin}))) \epsilon$ -moves and terminates with a good  $\epsilon$ -move.

Consider a segment with starting state  $s$ . If  $\ell$  is the maximum level of any player in  $s$  and there is a player at level  $(\ell - 2)$  or below, then Lemma 3 implies that the heaviest player has a good  $\epsilon$ -move, and we complete the segment with this move. If all players are either in level  $\ell$  or  $(\ell - 1)$ , and some player in level  $\ell$  has an  $\epsilon$ -move, then some player again has a good  $\epsilon$ -move, by Lemma 4.

Suppose that no player in  $\ell$  has an  $\epsilon$ -move, but some player in level  $(\ell - 1)$  does. In this case, we recursively consider the reduced game  $G(s, P)$ , where  $P$  is the set of players in level  $\ell$ . We terminate the recursion if at any time some player in  $P$  has an  $\epsilon$ -move, since we can then complete the segment by making a good  $\epsilon$ -move. We refer to the above mentioned condition as the *termination condition*. From Corollary 2, it follows that no player in level  $(\ell - 1)$  can get their level increased and no player in  $P$  can have its level changed as a result of  $\epsilon$ -Nash dynamics in the reduced game. Note that, from Lemma 3, termination condition for reduced game also implies that no player in reduced game drops its level to  $(\ell - 2)$  or below. Thus dynamics in the reduced game

keeps the level of all players unchanged, and since the reduced game is also a game with synchronized resources (by Lemma 5), we apply Lemma 4 to observe that whenever there exists an  $\epsilon$ -move, there also exists a good  $\epsilon$ -move in the reduced game. We shall always choose such a good  $\epsilon$ -move, and so, by Lemma 1, the reduced game either reaches an  $\epsilon$ -Nash equilibrium or is terminated within  $O(m^2Q^2n\epsilon^{-1}\log(\phi(s))) = O(m^2Q^2n\epsilon^{-1}\log(\phi(s^{begin})))$  steps. If the reduced game reaches an  $\epsilon$ -Nash equilibrium without getting terminated, then the original game itself has reached an  $\epsilon$ -Nash equilibrium. This completes the proof of Theorem 1.

## 4 Fast Convergence of Natural Dynamics

In this section, we show that several natural dynamics also converge fast.

**Largest gain dynamics:** The largest gain  $\epsilon$ -Nash dynamics [3] is defined as follows: among all the  $\epsilon$ -moves that can be made in the current state, the dynamics chooses a move with the largest absolute decrease in delay for the player who moves (thus causing the largest drop in potential). The proof of the theorem below is deferred to the full version.

**Theorem 2.** *For any  $0 < \epsilon \leq 1/2$ , given a symmetric congestion game with similar resources in some initial state  $s^{begin}$ , a largest gain  $\epsilon$ -Nash dynamics converges to an  $\epsilon$ -Nash equilibrium in  $O((nm^2Q^2\log(\phi(s^{begin}))/\epsilon)^2)$   $\epsilon$ -moves.*

**Heaviest First Dynamics:** The heaviest first dynamics [3] is defined as follows: among all the players which have an  $\epsilon$ -move, a player with maximum delay in the current state moves. The proof of the theorem below is deferred to the full version.

**Theorem 3.** *For any  $0 < \epsilon \leq 1/2$ , given a symmetric congestion game with similar resources in some initial state  $s^{begin}$ , a heaviest first  $\epsilon$ -Nash dynamics converges to an  $\epsilon$ -Nash equilibrium in  $O((nm^2Q^2\log(\phi(s^{begin}))/\epsilon)^2)$   $\epsilon$ -moves.*

**Unrestricted Dynamics with Liveness Condition:** Now we study convergence properties of  $\epsilon$ -Nash dynamics in a much more general setting, namely, unrestricted  $\epsilon$ -Nash dynamics with liveness condition [3]. Under these dynamics, every player is given at least one opportunity to make a move in each interval of  $T$  steps for some polynomially bounded parameter  $T$ . When an opportunity to move is given to a player  $p$ , it does not imply that  $p$  can make an  $\epsilon$ -move: it makes an  $\epsilon$ -move if it can, otherwise it does not move. Independent asynchronous updates by players according to most random processes, such as Poisson clocks, belong to this class of dynamics almost surely. The analysis of this dynamics requires new ideas, since unlike Theorem 1, we cannot always reduce the analysis to the case when all players are in only two levels.

The following lemmas are variations of Lemma 4.2 of Chien and Sinclair [3], and their proofs are deferred to the full version.

**Lemma 6.** *In any congestion game  $G$  with  $\epsilon$ -Nash dynamics, where  $G$  progresses from some state  $s^1$  to another state  $s^2$ , then for any player  $p_i$ , we have  $\phi(s^1) - \phi(s^2) \geq |\epsilon(c_i(s^1) - c_i(s^2))|$ . Moreover, if a player  $p_i$  has made at least one  $\epsilon$ -move as  $G$  progressed from  $s^1$  to  $s^2$ , then we can make a stronger inference, namely  $\phi(s^1) - \phi(s^2) \geq \epsilon(\max\{c_i(s^1), c_i(s^2)\})$ .*

**Corollary 3.** *In any congestion game  $G$  with  $\epsilon$ -Nash dynamics, where  $G$  progresses from some state  $s^1$  to another state  $s^2$ , if a player  $p_i$  has made an  $\epsilon$ -move between  $s^1$  and  $s^2$ , and there exists a state  $s^3$  between  $s^1$  and  $s^2$  with  $c_j(s^3) \geq \phi(s^1)/C$  for some  $C \geq 1$ , then  $\phi(s^1) - \phi(s^2) \geq \epsilon\phi(s^1)/2C$ .*

**Theorem 4.** *For any  $0 < \epsilon \leq 1/2$ , given a symmetric congestion game with similar resources in some initial state  $s^{begin}$ , an unrestricted  $\epsilon$ -Nash dynamics where every player gets an opportunity to move in every  $T$  steps, converges to an  $\epsilon$ -Nash equilibrium in  $O((nm^2Q^2 \log(\phi(s^{begin}))/\epsilon)^2T)$  steps.*

*Proof.* Consider an interval of  $T$  steps, where all players have got at least one chance to move. Let  $s^{start}$  and  $s^{end}$  be the first and last states of the game in this interval. Let  $p_i$  be a player with highest delay in the state  $s^{start}$ , then  $c_i(s^{start}) \geq \phi(s^{start})/n$ . Let  $\ell$  be its level in  $s^{start}$ . Then by Corollary 2, the maximum level of any player is at most  $\ell$  in all states that the dynamics can reach from  $s^{start}$ . We consider two cases.

*Case 1:* Suppose that at least one player, which is in level  $\ell$  in state  $s^{start}$ , makes an  $\epsilon$ -move in the interval (the player need not belong to level  $\ell$  when it makes a move). We will show that in this case, a good event occurred in this interval.

If  $p_i$  made an  $\epsilon$ -move in the interval, then by Lemma 6,  $\phi(s^{start}) - \phi(s^{end}) \geq \epsilon c_i(s^{start}) \geq \epsilon\phi(s^{start})/n$ . Hence a good event happened in this interval. We next consider the case that  $p_i$  did not make any  $\epsilon$ -move in the interval.

If  $c_i(s) < c_i(s^{start})/2$  for some state  $s$  in the interval, then by Lemma 6, the drop in potential in the interval is at least  $\epsilon(c_i(s^{start}) - c_i(s)) \geq \epsilon c_i(s^{start})/2 \geq \epsilon\phi(s^{start})/2n$ , so a good event has occurred. Now suppose that for all states  $s$  between  $s^{start}$  and  $s^{end}$ ,  $c_i(s) \geq c_i(s^{start})/2$ . Then by Lemma 2, it follows that the level of  $p_i$  is  $\ell$  in all states in this interval.

Let  $p_j$  be the first player, which is in level  $\ell$  in state  $s^{start}$ , that makes an  $\epsilon$ -move in the interval. If  $c_j(s) \geq \phi(s^{start})/16nm^2Q^2$  for any state  $s$  in the interval, then Corollary 3 implies that a good event occurs in the interval.

We now consider the only remaining case, where for every state  $s$  in the interval, we have  $c_i(s) \geq c_i(s^{start})/2$  and  $c_j(s) < \phi(s^{start})/16nm^2Q^2$ . Thus for every state  $s$ , we have  $c_i(s) \geq 8m^2Q^2c_j(s)$ , that is, delays of  $p_i$  and  $p_j$  are widely separated. Let  $s'$  be the first state when  $p_i$  was given a chance to move. Then we show that  $p_i$  can make an  $\epsilon$ -move from  $s'_i$  to  $s'_j$  in state  $s'$ , a contradiction to the assumption that  $p_i$  did not make an  $\epsilon$ -move in the interval.

If  $p_j$  is in level  $\ell$  in  $s'$ , then since  $p_i$  is also in level  $\ell$  and has delay at least  $8m^2Q^2$  times that of  $p_j$ , it follows (since the resources are synchronized) that increasing congestion on the resources in  $s'_j$  by one player does not change their level, and in fact their delays change by a factor of at most  $2mQ$ , so  $p_i$  can make

an  $\epsilon$ -move to  $s'_j$ . If  $p_j$  is in level less than  $\ell$  in  $s'$ , then using the fact that  $p_j$  was in level  $\ell$  in  $s^{start}$  and the fact that the resources are synchronized, we know that the maximum delay any resource  $r$  can have when it first enters level  $\ell$  is at most  $Qc_j(s^{start})$ . Thus, if  $p_i$  switches to  $s'_j$ , either the resources in  $s'_j$  still remain below level  $\ell$ , or just reached level  $\ell$ , and have delays of at most  $Qc_j(s^{start})$ . Hence  $p_i$  faces a delay of at most  $mQc_j(s^{start}) < c_i(s^{start})/4 \leq c_i(s')/2$  by moving to  $s'_j$ , so the move is an  $\epsilon$ -move.

*Case 2:* Now suppose that no player in level  $\ell$  in state  $s^{start}$  made an  $\epsilon$ -move in the entire interval. In this case, if no player which is in level  $(\ell - 1)$  in state  $s^{start}$  has also not been able to make an  $\epsilon$ -move in the interval, then by Corollary [2](#), the sets of players in levels  $(\ell - 1)$  and  $\ell$  remain unchanged in the entire interval. If there is no player in a lower level, the dynamics has already reached an  $\epsilon$ -Nash equilibrium. If there is a player, say  $p_k$ , below level  $(\ell - 1)$  in state  $s^{start}$ , then when an opportunity to move is given to player  $p_i$ , player  $p_k$  is still below level  $(\ell - 1)$ . So if  $p_i$  switches to the strategy of  $p_k$ , it would make the level of  $p_i$  drop to  $(\ell - 1)$  or less, thus giving an  $\epsilon$ -move, which is a contradiction.

So the only remaining case is that some player, say  $p_j$ , belonging to level  $(\ell - 1)$  in state  $s^{start}$  has been able to make an  $\epsilon$ -move in the interval. Let  $P$  be the set of players which are in level  $\ell$  in  $s^{start}$ , and consider the reduced game  $G(s^{start}, P)$ . It can be argued, as we did in Case 1, that a good event happens for the reduced game in this interval.

In conclusion, in every interval of  $T$  steps where approximate equilibrium has not been reached, either a good event happens in the original game or a good event happens in the reduced game. By Lemma [1](#), at most  $O(nm^2Q^2 \log(s^{begin})/\epsilon)$  good events can happen in the reduced game before it reaches an equilibrium, so a good event must happen in the original game in every  $O(nm^2Q^2 \log(s^{begin})/\epsilon)$  intervals of  $T$  steps each. Applying Lemma [1](#) on the original game, we get that the number of steps required for convergence is  $O((nm^2Q^2 \log(\phi(s^{begin}))/\epsilon)^2T)$ .

## 5 FPTAS for Constant Types of Players

We now generalize the FPTAS result of Theorem [1](#) to congestion games which allow *constant types of players*: two players are said to be of the same *type* if they have the same strategy space. Thus we shall assume that the  $n$  players share  $k$  distinct strategy spaces for some constant  $k$ . Note that  $k = 1$  implies that the game is symmetric. We defer the proof of the theorem below to full version.

**Theorem 5.** *For any  $0 < \epsilon \leq 1/2$ , given a symmetric congestion game with similar resources in some initial state  $s^{begin}$ , there exists a poly-time computable sequence of  $O((m^2Q^2n \log(\phi(s^{begin}))/\epsilon)^{2k})$   $\epsilon$ -moves leading to an  $\epsilon$ -Nash equilibrium, thus yielding a FPTAS when  $k$  is a constant.*

## 6 Hardness Results

In this section, we present our negative results.

**Theorem 6.** *Computing an exact pure Nash equilibrium in symmetric congestion games with similar resources is PLS-complete.*

Moreover, the assumptions of symmetry among players and similarity among resources are both necessary to get any non-trivial approximation using Nash dynamics. The following theorem states that similarity of resources is not a sufficient condition, and its proof is essentially a modification of the construction used to prove Theorem 3 in [11]. We also show that the game being symmetric does not suffice either, even if there are only two types of resources.

**Theorem 7.** *For every  $\epsilon \in (0, 1)$ ,  $\exists n_0$  such that for all  $n \geq n_0$ , there is a congestion game  $G(n)$  with similar resources and a state  $s$  with following properties: the description of  $G(n)$  is polynomial in  $n$ , and every sequence of improving moves from  $s$  to an  $\epsilon$ -Nash equilibrium is exponential in  $n$ .*

**Theorem 8.** *For every  $\epsilon \in (0, 1)$ ,  $\exists n_0$  such that for every  $n \geq n_0$ , there is a symmetric congestion game  $G(n)$  and a state  $s$  with following properties: the description of  $G(n)$  is polynomial in  $n$ , its resources can be divided into two groups such that all resources in a group are similar, and every sequence of improving moves from  $s$  to  $\epsilon$ -Nash equilibrium is exponential in  $n$ .*

## References

- [1] Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. In: FOCS, pp. 613–622 (2006)
- [2] Awerbuch, B., Azar, Y., Epstein, A., Mirrokni, V.S., Skopalik, A.: Fast convergence to nearly optimal solutions in potential games. In: ACM Conference on Electronic Commerce, pp. 264–273 (2008)
- [3] Chien, S., Sinclair, A.: Convergence to approximate nash equilibria in congestion games. In: SODA, pp. 169–178 (2007)
- [4] Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and approximation in potential games. In: STACS, pp. 349–360 (2006)
- [5] Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The complexity of pure nash equilibria. In: STOC, pp. 604–612 (2004)
- [6] Goemans, M.X., Li, L., Mirrokni, V.S., Thottan, M.: Market sharing games applied to content distribution in ad hoc networks. IEEE Journal on Selected Areas in Communications 24(5), 1020–1033 (2006)
- [7] Johnson, D.S., Papadimitriou, C.H., Yannakakis, M.: How easy is local search? J. Comput. Syst. Sci. 37(1), 79–100 (1988)
- [8] Nash, J.F.: Noncooperative games. Annals of Mathematics 54, 289–295 (1951)
- [9] Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. International Journal of Game Theory 2, 65–67 (1973)
- [10] Roughgarden, T., Tardos, É.: How bad is selfish routing? J. ACM 49(2), 236–259 (2002)
- [11] Skopalik, A., Vöcking, B.: Inapproximability of pure nash equilibria. In: STOC, pp. 355–364 (2008)
- [12] Tardos, E., Wexler, T.: Network formation games and the potential function method. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) Algorithmic Game Theory, ch. 19. Cambridge University, Cambridge (2007)

# Online Ad Assignment with Free Disposal

Jon Feldman<sup>1</sup>, Nitish Korula<sup>2</sup>, Vahab Mirrokni<sup>1</sup>, S. Muthukrishnan<sup>1</sup>,  
and Martin Pál<sup>1</sup>

<sup>1</sup> Google Inc., 76 9th Avenue, New York, NY 10011

{jonfeld,mirrokn,muthu,mpal}@google.com

<sup>2</sup> Dept. of Computer Science, University of Illinois, Urbana, IL 61801

Work done while at Google Inc., NY

nkorula2@illinois.edu

**Abstract.** We study an online weighted assignment problem with a set of fixed nodes corresponding to advertisers and online arrival of nodes corresponding to ad impressions. Advertiser  $a$  has a contract for  $n(a)$  impressions, and each impression has a set of weighted edges to advertisers. The problem is to assign the impressions online so that while each advertiser  $a$  gets  $n(a)$  impressions, the total weight of edges assigned is maximized.

Our insight is that ad impressions allow for *free disposal*, that is, advertisers are indifferent to, or prefer being assigned more than  $n(a)$  impressions without changing the contract terms. This means that the value of an assignment *only* includes the  $n(a)$  highest-weighted items assigned to each node  $a$ . With free disposal, we provide an algorithm for this problem that achieves a competitive ratio of  $1 - 1/e$  against the offline optimum, and show that this is the best possible ratio. We use a primal/dual framework to derive our results, applying a novel exponentially-weighted dual update rule. Furthermore, our algorithm can be applied to a general set of assignment problems including the *ad words* problem as a special case, matching the previously known  $1 - 1/e$  competitive ratio.

## 1 Introduction

*Motivation: Display Ads Allocation.* Many web publishers (e.g., news sites) have multiple pages (sports, arts, real estate, etc) where they show image, video or text ads. When a visitor to such a web site is exposed to an ad, this is called an “impression.” Advertisers typically buy blocks of impressions ahead of time via contracts, choosing blocks carefully to target a particular market segment, typically as part of a more general advertising campaign across various web sites and other media outlets. Once the contract is agreed upon, the advertiser expects a particular number of impressions to be delivered by the publisher over an agreed-upon time period.

The publisher enters all such impression contracts into an ad delivery system. Such systems are typically provided as a service by third party companies, but sophisticated publishers may develop their own software. When a user views one of the pages with ad slots, this system determines the set of eligible ads for that

slot, and selects an ad to be shown, all in real time. Because traffic to the site is not known beforehand, it must solve an online matching problem to satisfy the impression contracts. However, before committing to a set of contracts, it would have been already determined using traffic forecasts that the contracts are likely to be fulfillable. Thus, if this were purely a cardinality matching problem, it would typically be easy to solve; what makes the problem challenging is the fact that not all impressions are of equal value to an advertiser (e.g., top vs. side slots, sports vs. arts pages). The publisher is interested not only in filling the impression contracts, but also delivering well-targeted impressions to its advertisers (as measured, e.g., by click-throughs). Thus the ADS, when deciding which ad to serve, has the *additional* goal of maximizing the overall quality of impressions used to fill the contracts. We formulate and study this online optimization problem.

*Online Ad Allocation Problem.* We have a set of advertisers  $A$  known in advance, together with an integer impression contract  $n(a)$  for each advertiser  $a \in A$ . Each  $a \in A$  corresponds to a node in one partition of the bipartite graph we define. The set of impressions  $I$  forms the nodes of the other partition and they arrive online. When an impression  $i \in I$  arrives, its value  $w_{ia} \geq 0$  to each advertiser  $a$  becomes known (some of the  $w_{ia}$ 's are possibly zero). The value  $w_{ia}$  might be a prediction of click-through probability, an estimate of targeting quality, or even the output of a function given by the advertiser; we treat this abstractly for the purposes of this work. The impression  $i$  must be assigned immediately to some advertiser  $a \in A$ .

Let  $I^a \subseteq I$  be the set of impressions assigned to  $a$  during the run of the algorithm. The goal of the algorithm is to maximize overall advertiser satisfaction, i.e.,  $\sum_{a \in A} S(a, I^a)$  for some satisfaction function  $S$ . To encode the impression contracts  $n(a)$  as part of  $S$ , one possible choice is to say  $S(a, I^a) = \sum_{i \in I^a} w_{ia}$  if  $|I^a| \leq n(a)$  (and  $S(a, I^a) = -\infty$  otherwise). In other words, maximize overall quality without exceeding any of the contracts  $n(a)$ . As stated, no bounded competitive ratio can be obtained for this problem: just consider the simple case of a single advertiser,  $n(a) = 1$ , and two items arriving. The first item that arrives has value 100. If it is assigned, then the next item has value 10000; if it is not assigned, the next item has value 1. (In both cases the algorithm achieves less than 1/100th the value of the optimal solution.)

The main insight that inspires our model is that the strict enforcement of the impression contract as an upper bound is inappropriate, since impressions exhibit what is known as the property of *free disposal* in Economics. That is, in the presence of a contract for  $n(a)$  impressions, the advertiser is only pleased — or is at least indifferent to — getting *more* than  $n(a)$  impressions. Therefore, a more appropriate formulation of the problem is the following. We let  $I_k^a$  be the  $k$  impressions  $i \in I^a$  with the largest  $w_{ia}$ . Then, define

$$S(a, I^a) = \sum_{i \in I_{n(a)}^a} w_{ia}.$$

In other words, each advertiser draws its value from its top  $n(a)$  impressions, and draws zero value from its remaining impressions (yielding free disposal).



We call this the *display ads* (DA) problem. Free disposal makes the problem tractable; e.g., for the counterexample above with a single advertiser  $a$ , the trivial algorithm that assigns all the impressions to that advertiser is optimal. (The general problem with multiple advertisers is, of course, nontrivial.) This choice of  $S$  also allows us to tradeoff between quality and contract fulfillment by adding a constant  $W$  to each  $w_{ia}$ ; for large  $W$  the problem becomes closer to a pure maximum-cardinality matching.

*Our Results and Techniques.* Our main technical contribution is an online algorithm for the DA problem with competitive ratio of  $1 - 1/e$ , as long as  $n(a) \rightarrow \infty$ . Further, this is the best possible for any (even randomized) online algorithm.

We generalize our algorithm to the case of non-uniform item sizes, the so-called *Generalized Assignment Problem* (GAP). More specifically, we can add “sizes”  $s_{ia}$  to the model, where the contract then refers to the total *size* of impressions assigned to an advertiser (and the function  $S$  is defined appropriately; see Section 3 for more details). This generalization captures both the DA problem as well as the well-studied *ad words* (AW) problem [19], where the advertisers express budgets  $B_a$  (simply set  $s_{ia} = w_{ia}$ ). Our bound of  $1 - 1/e$  when sizes are “small” matches the best known ratio for the AW problem. Furthermore, GAP is a unifying generalization that can handle hybrid instances where some advertisers are budget-constrained, and some are inventory constrained.

Our algorithm for the DA problem is inspired by the techniques developed for the online Ad Words (AW) allocation problem in [19], as well as the general primal-dual framework for online allocation problems [4]. The key element of this technique is to develop a *dual update rule* that will maintain dual feasibility as well as a good bound on the gap between the primal and dual solutions. Previous algorithms for related online packing problems such as AW [5] typically update dual (covering) variables by multiplying them by a small factor (such as  $1 + 1/n$ ) at each step, and adding a term proportional to the increase in primal value. By contrast, our update rule sets the dual variable for each advertiser  $a$  to be a carefully weighted average of the weights of the top  $n(a)$  impressions currently assigned to  $a$ . In fact, the value of the dual variable for an advertiser with a set of impressions  $I^a$  is the same as it would be if we *re-ordered the impressions* in increasing order of weight and used the update rules of previous algorithms (as in [4]) on  $I^a$ . By choosing our exponentially-weighted update rule, we balance the primal and dual objectives effectively and obtain an optimal algorithm for the DA problem.

*Related Work.* The related AW problem discussed above is NP-Hard in the offline setting, and several approximations have been designed [6,22,2]. For the online setting, it is typically assumed that every weight is very small compared to the corresponding budget, in which case there exist  $(1 - 1/e)$ -factor online algorithms [19,4,15,11], and this factor is tight. In order to go beyond the competitive ratio of  $1 - \frac{1}{e}$  in the adversarial model, stochastic online variants of the problem have been studied, such as the random order and i.i.d models [15]. In particular, for any  $\varepsilon$ , a primal-dual  $1 - \varepsilon$ -approximation has been developed for



this problem in the random order model with the assumption that  $opt$  is larger than  $O(\frac{n^2}{\epsilon^3})$  times each bid [9]. Moreover, a 0.67-competitive algorithm has been recently developed for the (unweighted) max-cardinality version of this problem in the i.i.d. model (without any extra assumption) [12]. Previously, a randomized  $(1 - \frac{1}{e})$ -competitive algorithm for the max-cardinality problem was known in the adversarial model [16]. The online maximum weighted  $b$ -matching problem *without free disposal* in the random permutation model has also been studied, and a  $\frac{1}{8}$ -approximation algorithm has been developed for this problem [17].

Prior to the development of the  $(1 - \frac{1}{e})$ -approximation algorithm for the offline GAP, various  $\frac{1}{2}$ -approximation algorithms had been obtained for this problem [8, 21, 13]. It has been observed that beating the approximation ratio  $1 - \frac{1}{e}$  for more general packing constraints is not possible unless  $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$ . However, for GAP with simple knapsack constraints, an improved  $1 - \frac{1}{e} + \delta$ -approximation (with  $\delta \approx 10^{-180}$ ) was developed by Feige and Vondrak [11]. In the online model with small sizes, our approximation factor of  $1 - \frac{1}{e}$  is tight.

The offline variants of DA, AW, and GAP are special cases of the problem of maximizing a monotone submodular function subject to a matroid constraint [13]. Recently, the approximation factor for this problem has been improved from  $\frac{1}{2}$  to  $1 - \frac{1}{e}$  [23], but these algorithms do not work in the online model. The algorithm in [18], although studied for the offline setting, works for the online DA problem and gives a  $\frac{1}{2}$ -competitive algorithm (discussion below).

## 2 The Display Ads Problem

In this section, we provide online algorithms for the DA problem with small competitive ratios. Recall that the competitive ratio of an online algorithm for a maximization problem is defined as the minimum, over all possible input sequences, of the ratio between the value obtained by the algorithm and the optimum value on that sequence. We first give a simple upper bound:

**Lemma 1.** *No deterministic algorithm for the Display Ads problem achieves a competitive ratio better than  $1/2$ .*

*Proof.* Consider an instance in which there are two advertisers  $a_1, a_2$  each with capacity 1, and two impressions  $i_1, i_2$ . Impression  $i_1$  has value  $w$  for both advertisers, and arrives first. Once it has been assigned,  $i_2$  arrives, and has value  $w$  for the same advertiser to which  $i_1$  was assigned. Thus we obtain a value of  $w$ , while the optimal solution has value  $2w$ .  $\square$

In this section, we show that a greedy algorithm is always  $1/2$ -competitive, matching the bound of Lemma 1. On real instances of the Display Ads problem, though, advertisers request far more than a single impression, and so a natural question is whether one can obtain better deterministic algorithms if  $n(a)$  is large for each advertiser  $a$ . Also in this section, we answer this question affirmatively, giving an algorithm that achieves a competitive ratio tending to  $1 - 1/e$  as  $n(a)$  tends to infinity.

*The Greedy Algorithm.* Consider an algorithm for the DA problem, assigning impressions online. When impression  $i$  arrives, what is the benefit of assigning it to advertiser  $a$ ? This impression can contribute  $w_{ia}$  to the value obtained by the algorithm, but if advertiser  $a$  already has  $n(a)$  impressions assigned to it, one of these impressions cannot be counted towards the value. Let  $v(a)$  denote the value of the least valuable impression currently assigned to  $a$  (if there are fewer than  $n(a)$  such impressions,  $v(a) = 0$ ). Clearly, if  $w_{ia} \leq v(a)$ , there is no benefit to assigning impression  $i$  to advertiser  $a$ . Let  $A_i = \{a: w_{ia} > v(a)\}$ ; any algorithm should only assign  $i$  to an impression in  $A_i$ .

Perhaps the simplest algorithm is to assign an impression  $i$  to the advertiser  $a \in A_i$  that maximizes  $w_{ia}$ . The competitive ratio of this naive algorithm is arbitrarily bad: Consider a set of advertisers  $\{a^*, a_1, a_2, \dots, a_n\}$  each with capacity 1, and impressions  $\{i_1, i_2, \dots, i_n\}$  that appear in that order. Impression  $i_j$  has value  $1 + j\varepsilon$  for  $a^*$ , and value 1 for  $a_j$ . The algorithm above obtains value  $1 + n\varepsilon$ , while the optimal solution has value  $n + n\varepsilon$ .

One can do better by noticing that the increase in value by assigning impression  $i$  to  $a$  is  $w_{ia} - v(a)$ , and therefore greedily assigning  $i$  to the advertiser  $a$  maximizing this quantity, which we call the *marginal gain* from assigning  $i$  to  $a$ .

The following theorem shows that the greedy algorithm (maximizing the marginal gain at each step) is 1/2-competitive:

**Theorem 1.** *The greedy algorithm is  $\frac{1}{2}$ -competitive for display ad allocation.*

This theorem is a special case of Theorem 8 in [18] which studies combinatorial allocation problems with submodular valuation functions. This follows from the fact that the valuation function of each advertiser in the online DA problem is submodular in terms of the set of impressions assigned to it, i.e.,  $\sum_{i \in I_{n(a)}^a} w_{ia}$  is submodular in  $I^a$ . Though [18] studied this problem in the offline setting, their greedy algorithm can be implemented as an online algorithm. Other offline  $\frac{1}{2}$  and  $1 - \frac{1}{e}$ -approximation algorithms for a more general problem of submodular maximization under matroid constraints are known [13][23], but these offline algorithms do not provide an online solution.

When  $n(a)$  is large for each advertiser  $a$ , the upper bound of Lemma 1 does not hold; it is possible to achieve competitive ratios better than 1/2. However, even in this setting, the performance of the greedy algorithm does *not* improve.

**Lemma 2.** *The competitive ratio of the greedy algorithm is 1/2 even when  $n(a)$  is large for each advertiser  $a \in A$ .*

*Proof.* Let each of advertisers  $a_1, a_2$  have capacity  $n$ ; suppose there are  $n$  copies of impression  $i_1$  with value  $w$  to  $a_1$  and  $w - 1/n$  to  $a_2$ . The greedy algorithm assigns all of these impressions to  $a_1$ , obtaining value  $wn$ . Subsequently,  $n$  copies of impression  $i_2$  arrive, with value  $w$  to  $a_1$  and 0 to  $a_2$ . Thus, the optimal solution has value  $2nw - 1$ , while the greedy algorithm only obtains a value of  $nw$ .  $\square$

The greedy algorithm does badly on the instance in Lemma 2 because it does not take the capacity constraints into account when assigning impressions.

*Primal-Dual algorithms for the DA problem.* We write a linear program where for each we have variables  $x_{ia}$  to denote whether impression  $i$  is one of the  $n(a)$  most valuable impressions assigned to advertiser  $a$ .

<p><b>Primal:</b> <math>\max \sum_{i,a} w_{ia}x_{ia}</math></p> <p><math>\sum_a x_{ia} \leq 1 \quad (\forall i)</math></p> <p><math>\sum_i x_{ia} \leq n(a) \quad (\forall a)</math></p>	<p><b>Dual:</b> <math>\min \sum_a n(a)\beta_a + \sum_i z_i</math></p> <p><math>\beta_a + z_i \geq w_{ia} (\forall i, a)</math></p> <p><math>[x_{ia}, \beta_a, z_i \geq 0]</math></p>
--	--

The algorithms we consider simultaneously construct feasible solutions to the primal and dual LPs, using the following outline:

- Initialize the dual variables  $\beta_a$  to 0 for each advertiser.
- Subsequently, when an impression  $i$  arrives online, assign  $i$  to the advertiser  $a' \in A$  that maximizes  $w_{ia} - \beta_a$ . (If this value is negative for each  $a$ , leave impression  $i$  unassigned.)
- Set  $x_{ia'} = 1$ . If  $a'$  previously had  $n(a')$  impressions assigned, let  $i'$  be the least valuable of these; set  $x_{i'a'} = 0$ .
- In the dual solution, set  $z_i = w_{ia'} - \beta_{a'}$  and increase  $\beta_{a'}$  using an appropriate *update rule* (see below); different update rules give rise to different algorithms/assignments.

The outline above results in a valid integral assignment (primal solution) and a feasible dual solution; to completely describe such an algorithm, we only need to specify the update rule used. We consider the following update rules:

1. **Greedy:** For each advertiser  $a$ ,  $\beta_a$  is the weight of the lightest impression among the  $n(a)$  heaviest impressions currently assigned to  $a$ . That is,  $\beta_a$  is the weight of the impression which will be discarded if  $a$  receives a new high-value impression.
2. **Uniform Weighting:** For each advertiser  $a$ ,  $\beta_a$  is the average weight of the  $n(a)$  most valuable impressions currently assigned to  $a$ . If  $a$  has fewer than  $n(a)$  assigned impressions,  $\beta_a$  is the ratio between the total weight of assigned impressions and  $n(a)$ .
3. **Exponential Weighting:** For each advertiser  $a$ ,  $\beta_a$  is an “exponentially weighted average” (see Def. □) of the  $n(a)$  most valuable impressions.

It is easy to see that the Greedy rule simply gives rise to the greedy algorithm that assigns each impression to the advertiser that maximizes marginal gain. Using Uniform Weighting, one can obtain an improved ratio  $\approx 3/4$  on the instance of Lemma □, as the first  $n$  copies of impression  $i_1$  are split evenly between advertisers  $a_1$  and  $a_2$ , and thus half the copies of impression  $i_2$  can be assigned to  $a_1$ . We state and analyze the Exponential Weighting rule in more detail below, but as a warm-up, we use the primal-dual technique to show that the Uniform Weighting rule gives a 1/2-competitive algorithm.

**Lemma 3.** *The primal-dual algorithm with Uniform Weighting is  $\frac{1}{2}$ -competitive.*

*Proof.* We show that the value of the feasible dual solution constructed by the algorithm is at most twice the value of the assignment; by weak duality, this implies that the algorithm is  $1/2$ -competitive. It suffices to show that in any step, the increase in value of the assignment is at least  $1/2$  of the increase in value of the dual solution. If impression  $i$  is assigned to advertiser  $a$ , let  $v$  be the value of the least valuable impression among the best  $n(a)$  impressions previously assigned to  $a$ . Thus, the increase in value of the assignment is  $w_{ia} - v$ . We set  $z_i = w_{ia} - \beta_a \leq w_{ia} - v$ , as the least valuable impression is worth no more than the average. The increase in  $\beta_a$  is precisely  $\frac{1}{n}(w_{ia} - v)$ , and hence the total increase in the dual objective function is at most  $2(w_{ia} - v)$ .  $\square$

Using the Greedy Rule,  $\beta_a$  is simply the weight of the edge/impression that will be discarded, while with Uniform Weighting,  $\beta_a$  is the average of all the best  $n(a)$  weights currently assigned to  $a$ . The disadvantage of the first approach is that it only takes into account the *least* valuable impression, ignoring how much capacity is unused. For Uniform Weighting, Lemma 3 showed that the increase in dual value is  $(w_{ia} - v) + (w_{ia} - \beta_a)$ , but as one can only use the fact that  $v \leq \beta_a$ , we get a ratio of 2. To obtain a  $(1 - 1/e)$ -competitive algorithm, we use an intermediate exponentially-weighted average in which the less valuable impressions are weighted more than the more valuable ones, as follows:

**Definition 1 (Exponential Weighting).** *Let  $w_1, w_2, \dots, w_{n(a)}$  be the weights of impressions currently assigned to advertiser  $a$ , sorted in non-increasing order.*

*Let  $\beta_a = \frac{1}{n(a) \cdot ((1+1/n(a))^{n(a)} - 1)} \sum_{j=1}^{n(a)} w_j \left(1 + \frac{1}{n(a)}\right)^{j-1}$ .*

**Theorem 2.** *The primal-dual algorithm with the Exponential Weighting update rule has a competitive ratio of  $(1 - 1/e)$  as  $n(a) \rightarrow \infty$  for each advertiser  $a$ .*

*Proof.* Let  $e_n = (1 + 1/n)^n$ ; we have  $\lim_{n \rightarrow \infty} e_n = e$ . Analogous to the proof of Lemma 3, it suffices to show that at each impression/step of the algorithm, the increase in the value of the assignment is at least  $(1 - 1/e_{n(a)})$  times the increase in value of the feasible dual solution, where  $a$  is the advertiser to which this impression is assigned.

As before, let impression  $i$  be assigned to advertiser  $a$ , and let  $v$  be the value of the least valuable impression among the best  $n(a)$  impressions previously assigned to  $a$ . Thus, the increase in value of the assignment is  $w_{ia} - v$ , and we set  $z_i = w_{ia} - \beta_a$ . It remains to bound the increase in  $\beta_a$ , which we do as follows.

Let  $\beta_o, \beta_n$  denote the old and new values of  $\beta_a$  respectively. Suppose that after  $i$  is assigned to  $a$ , it becomes the most valuable impression assigned to  $a$ . Then, we have  $\beta_n = (1+1/n)\beta_o - \frac{ve_n}{n(e_n-1)} + \frac{w_{ia}}{n(e_n-1)}$ . Thus,  $n(\beta_n - \beta_o) = \beta_o - \frac{ve_n}{e_n-1} + \frac{w_{ia}}{e_n-1}$ . Therefore, the total dual increase, which is the sum of  $z_i$  and  $n$  times the increase in  $\beta_a$  is  $(w_{ia} - \beta_o) + \beta_o - \frac{ve_n}{e_n-1} + \frac{w_{ia}}{e_n-1} = \frac{(w_{ia}-v)e_n}{e_n-1}$ . Therefore, the ratio between the increase in assignment value and dual objective function is  $1 - 1/e_n$ .

We assumed above that  $i$  became the most valuable impression assigned to  $a$ ; what if this is not true? It is not difficult to verify that in this case, the

increase in  $\beta_a$  is *less* than otherwise; to see this, note that if it is the  $j$ th most valuable impression, the contribution of  $w_{ia}$  to  $\beta_a$  must be multiplied by a factor of  $(1 + 1/n)^{j-1}$  compared to the previous case, but the contributions of  $j - 1$  more valuable impressions will be decreased by a factor of  $(1 + 1/n)$ .  $\square$

**Theorem 3** ([19]). *No algorithm achieves a competitive ratio of greater than  $1 - 1/e$  for the display ad allocation problem. This is true even with weights in  $\{0, 1\}$ , and for randomized algorithms against oblivious adversaries.*

The lower bound of Theorem 3 was proved by [19] for the *Ad words* problem; the example they give is a valid instance of the Display Ads problem, and hence the same lower bound applies. Thus, our primal-dual algorithm with the Exponential Weighting update rule is optimal for the DA problem.

### 3 The Generalized Assignment Problem

In the Generalized Assignment Problem (GAP), a set  $A$  of bins/machines and a set  $I$  of items/jobs is given. Each bin  $a \in A$  has a capacity  $C_a$ ; for each item  $i$  and bin  $a$ , we have a size  $s_{ia}$  that item  $i$  occupies in bin  $a$  and a weight/profit  $w_{ia}$  obtained from placing  $i$  in  $a$ . (Alternately, one can think of GAP as a scheduling problem with  $s_{ia}$  as the processing time job  $i$  takes on machine  $a$ , and with  $w_{ia}$  being the value gained from scheduling job  $i$  on machine  $a$ .) Note that the special case of GAP with a single bin/machine is simply the Knapsack problem.

We first note that GAP captures both the Display Ads problem and the Ad Words problem as special cases, where bins correspond to advertisers and items to impressions. The DA problem is simply the special case in which  $s_{ia} = 1$  for all  $i, a$ , and the AW problem is the special case in which  $w_{ia} = s_{ia}$  for all  $i, a$ .

For the offline GAP, the best approximation ratio known is  $1 - 1/e + \delta$ , where  $\delta \approx 10^{-180}$  [11]; this improves on the previous  $(1 - 1/e)$ -approximation of [14]. In an online instance of GAP, the set of bins  $A$  is known in advance, together with the capacity of each bin. Items arrive online, and when item  $i$  arrives,  $w_{ia}$  and  $s_{ia}$  are revealed for each  $a \in A$ . The only previous work on online GAP appears to have been for the special case corresponding to the Knapsack problem [3].

Recall that without free disposal, the online Display Ads problem was intractable. We make a similar assumption to solve GAP online; here, we assume that we can assign items of total size more than  $C_a$  to bin  $a$ , but that the total value derived by bin  $a$  is given by the most profitable set of assigned items that actually fits within capacity  $C_a$ . (Note that such an assumption is not necessary for the easier Ad Words problem, in which the value/weight of an item in a bin is equal to its size; thus, there is never a need for over-assignment.) Thus, an online algorithm for GAP immediately gives algorithms with the same competitive ratio for the DA and AW problems. In fact, an algorithm for GAP allows one to simultaneously handle ad allocation problems in which some bidders have budget constraints and others have inventory constraints. Unfortunately, we have:

**Lemma 4.** *No deterministic online algorithm for GAP with free disposal can achieve a competitive ratio better than  $n^{-1/2}$ .*

Given this lower bound, for the rest of this section, we consider the case of *small items*; that is, we assume that for each item  $i$  and bin  $a$  such that  $w_{ia} > 0$ ,  $s_{ia} \leq \varepsilon C_a$ .<sup>1</sup> This is a reasonable assumption for both the DA and AW problems, where contracts are for large numbers of impressions or individual bids are small compared to budgets. We refer to GAP restricted to such instances – where no individual item can occupy more than an  $\varepsilon$  fraction of any bin – as  $\varepsilon$ -GAP. Let  $e_{1/\varepsilon} = (1 + \varepsilon)^{1/\varepsilon}$ ; we prove the following theorem:

**Theorem 4.** *There is a  $(1 - 1/e)$ -competitive algorithm for  $\varepsilon$ -GAP as  $\varepsilon \rightarrow 0$ .<sup>2</sup>*

*Proof Sketch.* We construct a feasible dual solution (primal and dual linear programs for GAP are given below) as in the proof of Theorem 2, but a problem arises in dealing with non-uniform sizes. It may sometimes be necessary for the algorithm to place an item in a bin even when doing so would *decrease* the value of the solution; this holds even when item sizes are all less than  $\varepsilon$  times the bin capacities. The intuition is as follows: Suppose an item  $i$  arrives with value/size ratio significantly better than the average for a given bin  $a$ ; it is clear that we should take it, and discard the existing items. (The inability to do this provides the lower bound of Lemma 4.) But if the items already in the bin are larger than the new item, one may lose value by discarding the existing items. This difficulty appears because in integral solutions an item cannot continuously move from being in the bin to outside. We deal with this issue by having the algorithm act as though it *could* derive value from such fractional solutions, in which the item of lowest value/size ratio is partly in the bin, and the value obtained from this item depends on how much of it is in the bin. Under this metric, we show the algorithm’s (fractional) value is at least  $(1 - 1/e_{1/\varepsilon})$  times that of a feasible dual solution. Since the algorithm does not truly obtain any integral value from such partially assigned items, it loses at most the value of these items, which is an  $\varepsilon$  fraction of its overall value. Thus, we obtain an integral solution which achieves an approximation ratio of  $(1 - 1/e_{1/\varepsilon})(1 - \varepsilon)$ .

<p><b>Primal:</b>    <math>\max \sum_{i,a} w_{ia} x_{ia}</math></p> <p style="margin-left: 40px;"><math>\sum_a x_{ia} \leq 1 \quad (\forall i)</math></p> <p style="margin-left: 40px;"><math>\sum_i s_{ia} x_{ia} \leq C_a \quad (\forall a)</math></p>		<p><b>Dual:</b>    <math>\min \sum_a C_a \beta_a + \sum_i z_i</math></p> <p style="margin-left: 40px;"><math>s_{ia} \beta_a + z_i \geq w_{ia} (\forall i, a)</math></p> <p style="margin-left: 100px;"><math>[x_{ia}, \beta_a, z_i \geq 0]</math></p>
--	--	---

<sup>1</sup> This lower bound does not apply to *randomized* algorithms; see Section 4.

<sup>2</sup> More formally, we obtain a ratio of  $(1 - 1/e_{1/\varepsilon})(1 - \varepsilon)$  for  $\varepsilon$ -GAP. This is greater than 1/2 for  $\varepsilon \leq 0.17$ .

## 4 Extensions and Future Work

*Randomized Algorithms and Lower Bounds.* For the basic Display Ads problem, we showed an upper bound of  $(1 - 1/e)$  on the competitive ratio of all algorithms, and a deterministic algorithm that matches this bound when  $n(a)$  is large. Further, Lemma 1 shows that no deterministic algorithm has competitive ratio larger than  $1/2$  when  $n(a)$  is small; does this bound also apply to randomized algorithms? The randomized algorithm of [16] gets a competitive ratio of  $1 - 1/e$  for the unweighted case. Extending this result to the weighted case seems difficult; a new approach may be necessary.

Similarly, Lemma 4 shows that no deterministic online algorithm for GAP has a competitive ratio better than  $n^{-1/2}$ . One can avoid this bound using randomization: Toss a coin to determine whether bins should accept only *large* items (that occupy more than  $1/3$  the bin), or only small items (that occupy at most  $1/3$  the bin.) In the latter case, use the algorithm of Theorem 4; in the former case, have each bin accept a single item. Since each bin can accept only two big items, we obtain a constant-competitive algorithm in both cases. (A similar observation was also made in [3] for the easier Knapsack problem.) Optimizing constants, we obtain the following theorem:

**Theorem 5.** *There is a 0.15-competitive randomized online algorithm for GAP.*

Extending these results for the online GAP to more general packing problems is an interesting subject of study. In particular, this idea may be applicable to packing problems with *sparse* constraint matrices; see [20,7] for recent work on the offline versions of these problems.

*General non-linear valuation functions.* The display ad business is performed through a set of pre-determined contracts. Hence, in many settings, the *number* of impressions assigned to an advertiser is an important quality measure in addition to the total valuation (or total weight) of the impressions. In other words, the valuation (or utility) of an advertiser  $a$  for receiving a set  $I^a$  of impressions is  $v_a(I^a) = \sum_{i \in I^a} w_{ia} + f_a(|I^a|)$  where  $f_a : N \rightarrow N$  is a non-decreasing function of the number of impressions assigned to  $a$ . We may also assume that  $f_a(x) = f_a(n(a))$  for any  $x \geq n(a)$ . The corresponding online ad allocation problem here is to assign impressions to advertisers and maximize  $\sum_{a \in A} v_a(I^a)$ .

Depending on various quality measures, this function  $f_a$  could be concave or convex. A convex function  $f_a$  models the guaranteed delivery property of advertisers in that receiving a number of impression close to  $n(a)$  is very important. A concave function  $f_a$ , on the other hand, captures the diminishing return property of extra impressions for advertisers. We observe that for convex functions  $f$ , the ad allocation problem becomes inapproximable, even in the offline case; this hardness result uses a reduction from a banner ad allocation problem with penalties studied in [10]. On the other hand, if all functions  $f_a$  are concave, the problem becomes a special case of submodular valuation and the greedy



algorithm gives a  $\frac{1}{2}$ -competitive algorithm. An interesting question is whether the competitive ratio of  $\frac{1}{2}$  can be improved to  $1 - \frac{1}{e}$ .

*“Underbidding” and Incentives.* One disadvantage of using the free disposal property is that it may incentivize advertisers to declare smaller  $n(a)$ , in the hope of getting more impressions in the final allocation. We can partially address this concern by modifying the algorithm slightly so that the sum of weights of *all* impressions assigned to  $a$  is at most twice the sum of weights of the top  $n(a)$  impressions:

**Theorem 6.** *There is a  $\frac{1-1/e}{2}$ -competitive algorithm for the DA problem such that for each advertiser,  $\sum_{i \in I^a} w_{ia} \leq 2 \sum_{i \in I_{n(a)}^a} w_{ia}$ .*

To prove this theorem, one simply needs to use the Exponential Weighting update rule but double  $\beta_a$  for each  $a$ ; we omit details from this extended abstract.

**Concluding Remarks:** We have used free disposal to solve the online DA problem with a competitive ratio of  $1 - 1/e$ . An outstanding issue is to understand how free disposal affects the incentives of advertisers, who may be led to speculate. (Note that even the sub-optimal algorithm of Theorem 6 only bounds the total weight of impressions assigned to an advertiser, not the number of impressions received.) A model for incentives must simultaneously handle contract selection/pricing and the online ad allocation problem; this is an interesting subject of future research.

## References

1. Alaei, S., Malekian, A.: Maximizing sequence-submodular functions (2009) (manuscript)
2. Azar, Y., Birnbaum, B., Karlin, A.R., Mathieu, C., Nguyen, C.T.: Improved Approximation Algorithms for Budgeted Allocations. In: Proc. Automata, Languages and Programming (2008)
3. Babaioff, M., Hartline, J., Kleinberg, R.: Selling ad campaigns: Online algorithms with cancellations. In: ACM EC (2009)
4. 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)
5. Buchbinder, N., Naor, J.: The Design of Competitive Online Algorithms via a Primal-Dual Approach. Foundations and Trends in Theoretical Computer Science 3(2-3), 93–263 (2007)
6. Chakrabarty, D., Goel, G.: On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. In: Proc. FOCS, pp. 687–696 (2008)
7. Chekuri, C., Ene, A., Korula, N.: Unsplittable flow in paths and trees and column-restricted packing integer programs. In: Dinur, I., et al. (eds.) APPROX and RANDOM 2009. LNCS, vol. 5687, pp. 42–55. Springer, Heidelberg (2009)
8. Chekuri, C., Khanna, S.: A PTAS for the multiple knapsack problem. In: 11th ACM-SIAM Symp. on Discrete Algorithms (SODA), pp. 213–222 (2000)



9. Devanur, N., Hayes, T.: The adwords problem: Online keyword matching with budgeted bidders under random permutations. In: Proceedings of the 10th ACM Conference on Electronic Commerce, pp. 71–78 (2009)
10. Feige, U., Immorlica, N., Mirrokni, V., Nazerzadeh, H.: A combinatorial allocation mechanism for banner advertisement with penalties. In: WWW (2008)
11. Feige, U., Vondrak, J.: Approximation algorithms for allocation problems: Improving the factor of  $1-1/e$ . In: FOCS (2006)
12. Feldman, J., Mehta, A., Mirrokni, V., Muthukrishnan, S.: Online stochastic matching: Beating  $1 - 1/e$ . In: FOCS (to appear, 2009)
13. Fisher, M., Nemhauser, G., Wolsey, L.: An analysis of the approximations for maximizing submodular set functions II. *Math. Prog. St.* 8, 73–87 (1978)
14. Fleischer, L., Goemans, M., Mirrokni, V.S., Sviridenko, M.: Tight approximation algorithms for maximum general assignment problems. In: Proc. SODA (2006)
15. Goel, G., Mehta, A.: Online budgeted matching in random input models with applications to adwords. In: SODA, pp. 982–991 (2008)
16. Karp, R.M., Vazirani, U.V., Vazirani, V.V.: An optimal algorithm for online bipartite matching. In: Proc. STOC (1990)
17. Korula, N., Pal, M.: Algorithms for secretary problems on graphs and hypergraphs. In: Albers, S., et al. (eds.) ICALP 2009, Part II. LNCS, vol. 5556, pp. 508–520. Springer, Heidelberg (2009)
18. Lehman, Lehman, Nisan, N.: Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behaviour*, 270–296 (2006)
19. Mehta, A., Saberi, A., Vazirani, U., Vazirani, V.: Adwords and generalized online matching. In: FOCS (2005)
20. Pritchard, D.: Approximability of Sparse Integer Programs. In: Proceedings of the 17th Annual European Symposium on Algorithms, pp. 83–94 (2009)
21. Shmoys, D., Tardos, E.: An approximation algorithm for the generalized assignment problem. *Mathematical Programming* 62(3), 461–474 (1993)
22. Srinivasan, A.: Budgeted Allocations in the Full-Information Setting. In: Goel, A., Jansen, K., Rolim, J.D.P., Rubinfeld, R. (eds.) APPROX and RANDOM 2008. LNCS, vol. 5171, pp. 247–253. Springer, Heidelberg (2008)
23. Vondrak, J.: Optimal approximation for the submodular welfare problem in the value oracle model. In: STOC (2008)

# Truthful and Quality Conscious Query Incentive Networks

Devansh Dikshit and Narahari Yadati

Department of Computer Science and Automation  
Indian Institute of Science, Bangalore, India  
{devansh,hari}@csa.iisc.ernet.in

**Abstract.** Query incentive networks capture the role of incentives in extracting information from decentralized information networks such as a social network. Several game theoretic models of query incentive networks have been proposed in the literature to study and characterize the dependence, of the monetary reward required to extract the answer for a query, on various factors such as the structure of the network, the level of difficulty of the query, and the required success probability. None of the existing models, however, captures the practical and important factor of *quality* of answers. In this paper, we develop a complete mechanism design based framework to incorporate the quality of answers, in the monetization of query incentive networks. First, we extend the model of Kleinberg and Raghavan [2] to allow the nodes to modulate the incentive on the basis of the quality of the answer they receive. For this *quality conscious model*, we show the existence of a unique Nash equilibrium and study the impact of quality of answers on the growth rate of the initial reward, with respect to the branching factor of the network. Next, we present two mechanisms, the direct comparison mechanism and the peer prediction mechanism, for truthful elicitation of quality from the agents. These mechanisms are based on *scoring rules* and cover different scenarios which may arise in query incentive networks. We show that the proposed quality elicitation mechanisms are incentive compatible and ex-ante budget balanced. We also derive conditions under which ex-post budget balance can be achieved by these mechanisms.

## 1 Introduction

We consider the scenario where a person is seeking some information from a social network. She formulates a query and asks her friends in the social network. If they know the answer, they reply her back, otherwise they forward the query to their friends and like that the query propagates through the social network. Similarly, when someone answers the query, the answer propagates back to the original person. In the real world, however, the picture is not so simplistic. Since every person is an intelligent and rational agent and since forwarding the query (and then reporting back the answer) requires a certain amount of effort on her part, she may not be willing to do so.

At present, the concept of incentive based queries is used in various QA networks such as Yahoo! Answers, Orkut's Ask Friends, LinkedIn, etc. However, in these cases, only the person who answers the query is rewarded, with no reward for the intermediaries. The net result is lack of enough exposure and propagation of the query over the social network with, generally, only the immediate friends viewing/answering the query. While offering an appropriate incentive to the intermediate nodes will increase the total reward which must be offered by the person posing the query, it will also increase the exposure of the query. This concept was captured by the model of Kleinberg and Raghavan [2]. Their model however, does not take into account the relevance or quality of the answer.

As an example, consider a query that seeks a complete proof of existence of mixed strategy Nash equilibria in a finite strategic form game and someone answers with only a sketch of the proof but not with full mathematical details. Then the root agent (the agent who originally posed the query) may not be willing to give the full incentive promised, as the answer is not complete. However, an appropriate fractional incentive is required to be paid to the answering agent and the intermediate agents. This motivates the need to develop a systematic scheme where the quantum of reward depends on the quality of the answer offered. An important research gap in this setting is the absence of a proper mechanism for eliciting and determining the quality of an answer, in a highly decentralized setting of social networks.

## 1.1 Relevant Work

The branching process model for query incentive networks, proposed by Kleinberg and Raghavan [2], is as follows. First, without loss of generality, assume that the query is originated at the root node  $v_{root}$  of a tree. Let  $n$  denote the rarity of the answer implying that one out of  $n$  nodes, on an average, holds the answer to the question being posed. This means that every node has the answer to the query, independently, with probability  $\frac{1}{n}$ . Now take the case of any general node  $v$ . Let  $r$  be the reward offered to  $v$  by its parent and let  $f_v(r)$  be the reward which  $v$  offers to its children. Here,  $f_v(\cdot)$  denotes the reward function for node  $v$ . Then the payoff for node  $v$  is  $(r - f_v(r) - 1)$ . The factor of  $-1$  is introduced to account for the effort put in by the agent in forwarding the query. Kleinberg and Raghavan [2] proved that in the Nash equilibrium profile, each agent will offer a reward  $x$  to her children where  $x$  maximizes  $(r - x - 1)\alpha_v(f, x)$ . Here,  $\alpha_v(f, x)$  is the probability that node  $v$  will receive the answer from its children if it offers the reward  $x$  and  $f$  is the common reward function. They proved that this Nash equilibrium is in fact unique under fairly weak technical conditions. Further, they investigated the relation between the growth rate of required reward and the branching factor of the tree network.

Arcaute, Kirsch, Kumar, Liben-Nowell, and Vassilvitskii [3] generalized the results of [2] to an arbitrary branching process model and proved some additional bounds. The combined results of [2] and [3] show that for any constant failure probability  $\sigma$ , the growth rate of reward at  $v_{root}$  is linear in the expected depth of the search tree if  $b > 2$  where  $b$  is the branching factor. For  $1 \leq b \leq 2$ ,

the growth rate of reward is exponential. However when  $\sigma$  is a polynomial in  $1/n$ , then this threshold effect disappears and the growth rate is exponential for all values of the branching factor. It should be mentioned here that the above analysis is not applicable for  $b < 1$  since the tree becomes finite in nature.

## 1.2 Contributions and Outline

The related work in the area does not capture quality or relevance of answers in any way. In this paper, we design a complete framework for incorporating the factor of quality of answers in determining the payments in a query incentive network. Our specific contributions are as follows.

We first non-trivially extend the strategic form game model of query incentive networks proposed by Kleinberg and Raghavan [2] to capture a key factor, namely, the quality of answers into the model. We prove the existence of a unique Nash equilibrium for the reward functions employed by the nodes in our model. We next investigate the impact of quality of answers on the growth rate of reward in the proposed model.

Our next contribution in this paper involves development of two mechanisms, namely the Direct Comparison (DC) and the Peer Prediction (PP) mechanisms, for honest elicitation of quality from the agents. These two mechanisms are based on the concept of scoring rules [4]. We show that the DC and PP mechanisms are incentive compatible. We then investigate the issue of budget balance in these two mechanisms.


## 2 Quality Conscious Model: Formulation and Analysis

Throughout this paper, we will consider the underlying network to be a branching process. As will be evident in the following analysis, the branching process assumption does not affect the analysis much and similar analysis can be repeated for other network structures also. Also, during the analysis, we will use the terms agents and nodes interchangeably.

### 2.1 Design of the Model

Consider an infinite  $d$ -ary tree where every node is active independently with a probability  $q$ . A node is said to be active when it is willing, as well as, able to participate in the query propagation process. For such a network, we will have the branching factor as  $b = qd$ . Also let  $v_{root}$  be the root node of the tree where the query is originated. Let  $v_i$  denote any general node at level  $i$ ; we assume that the root is at level 0. Now the process of incentive based query propagation will progress in three steps as follows (see Figure 1).

(1) The query originates at the root and an initial reward of  $r_1$  is offered to all the nodes at level 1. This query propagates down the tree. At each level  $i$ , node  $v_i$  is offered a reward  $r_i$  for answering the query. If  $v_i$  does not have the answer, it keeps a fraction of reward  $r_i$  for itself and offers the rest of it as reward



**Fig. 1.** The model for quality conscious query incentive networks

$r_{i+1} (= f_{v_i}(r_i))$  to its children for answering the query. Here  $f_{v_i}(\cdot)$  is the reward function employed by  $v_i$ . The propagation continues until either (a) some node comes up with the answer, or, (b) the reward becomes zero, or, (c) a node is not active at the time of propagation.

(2) In step 2, The answer propagates back to the root node through appropriate intermediate nodes. With respect to the formulation above, let us assume that node  $v_m$  answered the query. Then the intermediate nodes will be  $v_{m-1}, v_{m-2}, \dots, v_2, v_1$ . We say that these nodes form a *Link Chain* for the given query answering process. Further, in this step, all the intermediate nodes along with  $v_{root}$  and  $v_m$  declare the quality of the answer as they perceive it to be. For node  $v_i$ , the quality of the answer which it perceives will be denoted by  $\theta_i^* (\in \Delta(T))$ , with  $T = \{1, 2, \dots, t\}$  representing the set of possible levels of quality. If  $\theta_i^* = (0, 0, \dots, 0, 1)$ , then the implication is that the answer is perceived to be perfect by  $v_i$ . Declaration of quality by the nodes in the above manner implies that, due to the tree structure, only nodes  $v_{i-1}$  and  $v_{i+1}$  are aware of the quality reported by the agent  $v_i$ . For now, we will assume that the agents declare their quality assessment truthfully. Also, we will make the reasonable assumption that, while reporting the quality, the agents are not aware of the quality announcements of the other agents. This restriction will be helpful in Section 3, where we will design mechanisms which will make honest reporting, a best response strategy for the agents.

(3) In the final step, the rewards will propagate down from the root node along the sequence of agents which acted as intermediaries, all the way to the final agent that answered the query. The actual reward  $\hat{r}_i$  for agent  $v_i$  will depend only upon the initial reward offered  $r_i$  and the reported types  $\theta_{i-1}$  and  $\theta_i$  of agents  $v_{i-1}$  and  $v_i$ , respectively. This is because agent  $v_i$  is not aware of existence any of the nodes  $v_j (j < i - 1)$ . For agent  $v_i$ ,  $v_{i-1}$  posed the question. Similarly,  $v_{i-1}$  is not aware of any of the children of  $v_i$ , for him  $v_i$  answered the query. Thus the only common information between  $v_i$  and  $v_{i-1}$  is the promised reward  $r_i$  and their respective reported qualities.

It might happen that at a particular node, two or more of its children might reply with an answer. Selection criteria of answers, for example on the basis of the quality, will not have any impact on the working of our model.

### 2.2 Quality Aggregation Function

We consider a general function of the form  $\varphi : \Delta(T) \times \Delta(T) \rightarrow [0, 1]$  which takes as input reported quality values from two nodes and outputs a representative quality of the answer. This function can be viewed as an agreement by the agents over the quality of answer when each one of them independently perceives a different quality. We will not use any particular function for  $\varphi$  till Section 3 where we show how  $\varphi$  can be modeled for honest quality elicitation. We will call  $\varphi$  as the *quality aggregation function*. Using this quality aggregation function  $\varphi$ , we can model the incentives as follows: Reward promised to  $v_i$  by its parent =  $r_i$ . Reward promised by  $v_i$  to its children =  $r_{i+1}$ . Actual reward ( $\hat{r}_i$ ) received by  $v_i = r_i\varphi(\theta_{i-1}, \theta_i)$ . Actual reward ( $\hat{r}_{i+1}$ ) given by  $v_i = r_{i+1}\varphi(\theta_i, \theta_{i+1})$ . Thus the actual incentive for agent  $v_i$  is  $r_i\varphi(\theta_{i-1}, \theta_i) - r_{i+1}\varphi(\theta_i, \theta_{i+1}) - 1$ . The factor of  $-1$  is the cost incurred by agent  $v_i$  in forwarding the query to its children and then reporting back the answer to its parent.

### 2.3 Nash Equilibrium Analysis

Let us take  $\alpha_v(f, x)$  to be the probability that node  $v$  will receive the answer from its children, if it offers the reward  $x$  and  $f$  is the common reward function. Let  $\beta_v(f, x)$  be the corresponding complementary probability. Then it can be recursively defined as:

$$\beta_v(f, x) = \prod_{w \in \text{child}(v)} (1 - q(1 - p\beta_w(f, f_w(x))))$$

where  $p$  is the constant probability that a node does not hold the answer,  $q$  is the probability that the node is active at the time the query is asked and  $\text{child}(v)$  is the set of all children of  $v$ . Note that, the probability with which an agent  $v$  receives the answer, increases with the reward  $f_v$  offered by  $v$ . This happens because a higher level of reward leads to deeper propagation of the query within the network and hence higher exposure. With this setup, we state the following lemmas which can be proved as an extension of results in [2]. Due to lack of space, we have omitted several proofs from the paper. These can be found in [8].

**Lemma 1.** *A set of payoff functions  $f$  is a Nash equilibrium if it maximizes  $E_\theta[(r\varphi(\theta_{i-1}, \theta_i) - f_v(r)\varphi(\theta_i, \theta_{i+1}) - 1)\alpha_v(f, f_v(r))]$ .*

Let  $e_{i,j} = E[\varphi(\theta_i, \theta_j)]$ . Then, using elementary probability, we can rewrite the term for the Nash equilibrium as  $f$  which maximizes  $(re_{i-1,i} - f_v(r)e_{i,i+1} - 1)\alpha_v(f, f_v(r))$ . For proving the uniqueness, we will have to show that there cannot be two reward functions  $f$  and  $g$ , both of which attain the maximum explained above. We will take  $g$  to be the set of payoff functions corresponding to a Nash equilibrium as computed in Lemma 1. Also, we will assume  $g_v(2) = 1$  because  $v$  will anyway get zero reward whether  $g_v(2) = 1$  (1 unit for her effort) or  $g_v(2) = 0$  (query does not propagate forward).

**Lemma 2.** *If  $p$  is generic with respect to  $q$  and  $f$  is a Nash Equilibrium in which  $f_v(2) = 1$  for all nodes  $v$ , then  $f_v(r) = g_v(r) \quad \forall v, \forall r$ , where  $g_v(r)$  is any Nash equilibrium as calculated above.*

### 2.4 Breakpoint and Growth-Rate Analysis

In this section, we investigate the impact of the quality factor on the growth rate of the initial reward which must be offered by the root agent, in order to get an answer with a certain probability. We begin by introducing some notation. Let  $R_\sigma(n, b)$  denote the minimum reward, which must be offered by the root node, in order to receive the answer with probability  $\sigma$ , given the branching factor  $b$  and rarity of the answer  $n$ . The rarity of the answer can be expressed in terms of  $p$ , the probability that a particular node holds the answer as  $n = 1/(1 - p)$ . Also, let  $\hat{\phi}_j$  denote the probability that none of the nodes in the first  $j$  levels from the root, holds the answer to the query. To establish a relation between  $R_\sigma(n, b)$  and  $\hat{\phi}_j$ , we need to look at the equilibrium behavior of each node.

**Breakpoint Analysis.** Consider a typical node  $v$ , which has been offered an incentive  $r$  for finding an answer to a particular query. In case  $v$  does not possess the answer, it will offer a reward  $g(r)$  to its children to come up with the answer to the query, where  $g(\cdot)$  is the equilibrium reward function. Let  $\delta(r)$  denote the number of times we have to iterate  $g$  on  $r$  in order to reduce it to zero. Then  $\delta(r)$  will represent the depth up to which the query will further propagate from  $v$ . We further define  $u_j$  as the minimum reward  $r$  for which  $\delta(r)$  is at least  $j$  implying that  $u_j$  is the minimum reward for which the query will propagate for at least  $j$  more levels. Note that the growth rate of reward is actually a step function of  $\sigma$ . To see this, suppose by offering a reward of  $u_j$ , the query will be pushed up to  $j$  levels. Then to push the query up to  $j + 1$  levels,  $v$  needs to offer, by definition, at least a reward of  $u_{j+1}$ . Thus for any reward value between any two such levels, the reward can be reduced to the next highest break point without affecting the probability of success and in the process increasing the payoff.

With respect to the above analysis, the expected payoff for a particular node  $v_i$  which has been offered reward  $u_j$  will be  $(u_j e_{i-1,i} - u_{j-1} e_{i,i+1} - 1)(1 - \hat{\phi}_{j-1})$ . Let  $\Delta_j$  be the gap between the breakpoints, that is,  $\Delta_j = u_j - u_{j-1}$ . We will assume that the agents are identical, which implies that they have a common prior about each other's beliefs. Using the properties of the quality aggregation function  $\varphi$ , one can easily show that  $e_{i-1,i} = e_{i,i+1}$ . Let us use the notation  $\mu = e_{i-1,i} = e_{i,i+1}$ . Here,  $\mu$  signifies the expected value of the quality agreement between any two agents in the network. We are now in a position to prove the following result relating the breakpoints of the reward to the success probability.

**Lemma 3.** *For a quality conscious query incentive network,  $\frac{\Delta_{j+1}}{\Delta_j} \geq \frac{\mu}{\frac{1-\hat{\phi}_{j+1}}{1-\hat{\phi}_j} - 1}$*

**Growth-Rate Analysis.** In this section, we will show as to how varying the incentives on the basis of the quality helps in reduction of the required reward at

the root level, when the branching factor  $b$  is less than 2. Throughout this section we will assume that  $\sigma \gg n^{-1}$ . This is not a serious restriction as practically the model will be extensively used only for queries with rare answers. Thus the value of  $n$  will be usually very large.

We begin by formulating a relation between  $\hat{\phi}_j$  and  $\hat{\phi}_{j-1}$ . Let  $\hat{\phi}_j = t(\hat{\phi}_{j-1})$ . Then by definition of the probability of finding an answer within the subtree rooted at a particular node explained in Section 2.3, we have  $t(x) = (1 - q(1 - px))^d$ . Recall that the setting corresponds to a  $d$ -ary tree. The next step in the analysis involves bounding of the number of iterations of  $t$  required to reduce the initial failure probability ( $= 1 - \Theta(n^{-1})$ ) to the final required failure probability ( $= 1 - \Theta(1)$ ). This can be summed up in the form of the following results.

**Theorem 1.** *If  $b < 2$  and  $\hat{\phi}_j < 1 - \sigma$  for any  $j$ , then  $R_\sigma(n, b) \geq \mu^{\log n} n^c$ , where  $c > 1$ .*

Since we have  $0 \leq \mu \leq 1$ , this is a significant improvement over the earlier bound of  $\Theta(n^c)$ . The effect of this factor will be even more significant in the case of rare/tough queries where practically the low expected quality of the answer will negate the effect of large  $n$ .

**Theorem 2.** *If  $b > 2$  and  $\hat{\phi}_j < 1 - \sigma$  for any  $j$ , then  $R_\sigma(n, b)$  is  $O(\log n)$ .*

Without the incorporation of the quality factor, the growth rate of reward is  $O(\log n)$ . Thus although there is no change in the asymptotic rate or growth, the factor  $\mu$ , which is independent of  $n$ , reduces the constant term involved. Again we can have the intuitive argument of low quality answers, in the case of tough queries, having significant effect on the behavior of the growth rate in this case. Also, the significant point to note here is that since the agents are assumed to be rational, therefore each agent involved as an intermediary must be offered some reward. Since one out of every  $n$  persons, on an average, has the answer, in the worst case there will be  $O(\log n)$  intermediaries in a tree structure. So asymptotic reduction in the growth rate of  $O(\log n)$  really seems to be out of question.

### 3 Honest Quality Elicitation

In this section, we address the key issue that is necessary to ensure practical viability of quality conscious query incentive networks. As mentioned, the final payment to any agent  $v_i$  is dependent on the quality of the answer reported by her. This induces strategic behavior among the agents. Therefore it is important to design an incentive compatible quality aggregation function. It is easy to see that naive functions, such as mean of quality levels, lead to a scenario where the payment is directly related to the reported quality type.

Before going into details of designing a quality elicitation mechanism, we would like to emphasize two technical assumptions about the process. First, an agent does not know the reported quality of other agents, in particular that of her child and parent, until after she declares her own perceived quality. Secondly,



all the monetary transfers are via a social planner, in this case the social networking site. One can see that these two assumptions are quite logical, not very restrictive, and are also easy to implement.

### 3.1 Quality Elicitation Mechanism

To summarize the scenario at the most basic level, we have an agent  $v_i$  along with her parent  $v_{i-1}$  and child  $v_{i+1}$ . The observed quality valuation, henceforth referred to as the observed type, of agent  $v_i$  is  $\theta_i^*$  with the type variable  $\theta_i \in \Delta(T)$  where  $T$  is the set of possible quality levels. The reported type of agent  $v_i$  will be represented by  $\hat{\theta}_i$  which can be different from her true observed type. The target is to define the quality aggregation function  $\varphi()$  in a way that the agent maximizes her profit by reporting her true type. For this we will use **scoring rules** which are explained next.

**Scoring Rules.** A scoring rule  $S$  is a sequence of scoring functions,  $S_1, S_2, \dots, S_t$ , such that  $S_i$  assigns a score  $S_i(z)$  to every  $z \in \Delta(T)$  [4]. Note that  $z = (z_1, z_2, \dots, z_t)$ . We will consider only real valued scoring rules. Scoring rules are primarily used for comparing the predicted distribution with the true observed one. Suppose  $w \in \Delta(T)$  is the true observed distribution. Then the *expected score* of any general distribution  $z \in \Delta(T)$  against  $w$  is defined as  $V(z|w) = \sum_{i=1}^t w_i S_i(z)$ . The *expected score loss* is defined as  $L(z|w) = V(w|w) - V(z|w)$ . A scoring rule  $S$  is called *strictly proper* or *incentive compatible* if  $\forall z, w \in \Delta(T)$  with  $z \neq w$ ,  $L(z|w) > 0$ . There are three popular proper scoring rules in the literature [4]:

- Quadratic scoring rule:  $S_i(z_1, \dots, z_i, \dots, z_t) = 2z_i - \sum_{j=1}^t z_j^2$
- Logarithmic scoring rule:  $S_i(z_1, \dots, z_i, \dots, z_t) = \ln z_i$
- Spherical scoring rule:  $S_i(z_1, \dots, z_i, \dots, z_t) = \frac{z_i}{\sqrt{\sum_{j=1}^t z_j^2}}$

We will consider only the quadratic scoring rule in this paper. A strong motivation for this is that the quadratic scoring rule and its derivatives or variants alone follow the property of *neutrality*, along with other desirable properties. A scoring rule is *neutral* if, for any two given distributions  $z$  and  $w$ ,  $L(z|w) = L(w|z)$ . Thus in scenarios like social networks, where there is no actual/observed distribution and agents are valued against each other’s reports, *neutrality* is a desirable property. The continuous version of the quadratic scoring rule, which will be required in the later part of this paper, is given by:

$$S(z(x)) = 2z(x) - \int_{-\infty}^{+\infty} z(x)^2 dx \quad ; x \in \mathfrak{R}$$

where  $z(x)$  is the probability density at  $x$ .

We now formulate two mechanisms, based on scoring rules: (1) Direct Comparison (DC) mechanism (2) Peer Prediction (PP) mechanism.

### 3.2 Direct Comparison Mechanism

This mechanism is applicable in the scenario where the agents have absolutely no prior information about each other’s possible behavior. The only way an agent can predict another agent’s perceived quality of an answer is through her own assessment of it. Technically speaking, given any pair of nodes  $v_i$  and  $v_j$ , the conditional probability distribution functions  $P(\theta_j^*|\theta_i^*)$  over the space of unit- $t$  simplex corresponding to  $\theta_j$  has all the probability mass concentrated at  $\theta_j = \theta_i^*$ .

In the direct comparison mechanism, the monetary transfer between two agents  $i$  and  $j$  is dependent on the difference between the reported quality types  $\hat{\theta}_i$  and  $\hat{\theta}_j$ , a measure of which is calculated by a proper scoring rule. For a particular agent  $v_i$ , the final amount received would be

$$= r_i \sum_{j=1}^t \hat{\theta}_{i-1,j} S_j(\hat{\theta}_i) - r_{i+1} \sum_{j=1}^t \hat{\theta}_{i+1,j} S_j(\hat{\theta}_i) - 1$$

where  $\hat{\theta}_{i,j}$  is the  $j^{th}$  component of the probability vector  $\hat{\theta}_i$ .

**Theorem 3.** *All agents reporting true quality is a Nash equilibrium in the DC mechanism when the agents are identical and the root agent is truthful.*

*Proof.* Let us assume a dummy amount  $v_{m+1}(= 0)$  to be paid by the last node  $v_m$  of the *Link chain*. Now for any node  $v_i$ , the actual amount she receives is  $r_i V(\hat{\theta}_{i-1}|\hat{\theta}_i) - r_{i+1} V(\hat{\theta}_{i+1}|\hat{\theta}_i) - 1$ . However, at the time of announcing her quality, the reported quality types of other agents are unknown to her. Thus the expected amount before reporting her quality is given by:

$$\left( \int_{\theta_{i-1}} r_i P(\theta_{i-1}|\theta_i^*) V(\hat{\theta}_{i-1}|\hat{\theta}_i) d\theta_{i-1} \right) - \left( \int_{\theta_{i+1}} r_{i+1} P(\theta_{i+1}|\theta_i^*) V(\hat{\theta}_{i+1}|\hat{\theta}_i) d\theta_{i+1} \right) - 1$$

where both the integrals represent shorthand for  $t$  integrations, one over each dimension of the unit- $t$  simplex corresponding to the two qualities  $\theta_{i+1}$  and  $\theta_{i-1}$ . Using the identical agents assumption, this can be simplified to:

$$\left( \int_{\theta_j} (r_i - r_{i+1}) P(\theta_j|\theta_i^*) V(\hat{\theta}_j|\hat{\theta}_i) d\theta_j \right) - 1$$

where  $\theta_j$  is the quality distribution corresponding to any general node  $j$ . Now when other agents, except  $v_i$ , are truthful, it implies that  $\hat{\theta}_j = \theta_j^*$ . However, in the above expression, since it represents the expected payoff, truthfulness amounts to  $\theta_j^* = \theta_j$  in the score function  $V()$ . Also since  $V()$  corresponds to a strictly proper scoring rule,  $V(\hat{\theta}_j|\hat{\theta}_i)$  maximizes when  $\hat{\theta}_i = \hat{\theta}_j$ . Combining this with the probability mass assumption, the expected payment for  $v_i$  is maximized when  $\hat{\theta}_i = \theta_i^*$ . Thus reporting true quality value is a Nash equilibrium for the agents.

Note that the above arguments are not valid for the root agents, as the root node has no monetary incentive. Hence we require the assumption of her being truthful. This is a reasonable assumption as the root node will be deriving some value from the answer which will contribute towards her overall payoff.

### 3.3 Peer Prediction Mechanism

The DC mechanism is applicable in a scenario where agents cannot form beliefs about each other’s assessment. However, in practice, that may not be the case always. In fact, it would be more beneficial if we can formulate a mechanism for a more general setting where, given her own assessment of the quality, an agent can form a belief about other agents’ assessments. The proposed peer prediction mechanism aims to achieve exactly this.

In the proposed PP mechanism, we ask agents to predict the assessment of other agents, specifically her parent and child. The actual payment made is based on the score computed by a scoring rule rather than the expected score as in DC mechanism. The score is calculated using the conditional probability of quality announcement rather than based on the actual reported qualities. That is, the payment to agent  $v_i$  is  $r_i S(P(\hat{\theta}_{i-1}|\hat{\theta}_i)) - r_{i+1} S(P(\hat{\theta}_{i+1}|\hat{\theta}_i)) - 1$ . Note that here an agent is asked for a common prediction for both her child and parent, which is consistent with our assumption of identical agents. Also, since the score is now calculated upon the probability function, we will have to use a continuous scoring rule for PP mechanisms.

**Theorem 4.** *All agents reporting true quality is a Nash equilibrium in the PP mechanism when the agents are identical and the root agent is truthful.*

*Proof.* The expected payoff for agent  $v_i$  is:

$$= r_i \left( \int_{\theta_{i-1}} P(\theta_{i-1}|\theta_i^*) S(P(\theta_{i-1}|\hat{\theta}_i)) d\theta_{i-1} \right) - r_{i+1} \left( \int_{\theta_{i+1}} P(\theta_{i+1}|\theta_i^*) S(P(\theta_{i+1}|\hat{\theta}_i)) d\theta_{i+1} \right) - 1$$

Note that inside the scoring function, the probability used is conditioned on the announced type of  $v_i$ , while the expectation is calculated on the probability conditioned on the actual type of  $v_i$ . Also there is an underlying assumption in the above expression that other agents are truthful. Now, using the symmetric agents assumption, the above expression can be written as:

$$= (r_i - r_{i+1}) \left( \int_{\theta_j} P(\theta_j|\theta_i^*) S(P(\theta_j|\hat{\theta}_i)) d\theta_j \right) - 1$$

The first part is precisely the expression for the *expected score* of the scoring rule while the second part is a constant. Therefore the expected payoff is uniquely maximized when  $\hat{\theta}_i = \theta_i^*$ . Thus reporting true quality value is a Nash equilibrium strategy. The arguments for  $v_{root}$  and  $v_m$  are similar as in the case of DC mechanism.

### 3.4 Weighted Scoring Rules

Note that the expected score calculated by a proper scoring rule is high when the two probability distributions being compared are identical to each other. While this partially serves our purpose, it does not capture our original intent of modulating payments on the basis of quality of answer. Specifically, we would

like to have a high score when the quality of answer is high and in addition, the agents agree on that. For this, we introduce the notion of weighted scoring rules with higher quality levels having more weight in calculation of scores. To the best of our knowledge, there does not seem to be any relevant work in this direction. Also, in this section we address the issue of normalizing the scoring rule such that its value lies in the interval  $(0, 1)$ . We propose the weighted discrete quadratic scoring rule (WDQSR) as follows:

$$S_i(z_1, \dots, z_i, \dots, z_t) = \frac{2z_i i - \sum_{j=1}^t z_j^2 j}{t}$$

**Lemma 4.** *The WDQSR is a proper scoring rule.*

For the continuous case, we need to define weights which are applicable at each point in the unit- $t$  simplex and capture the notion of importance of dimensions as well. For this, we will use the generalized Euclidean distance of a point,  $x = (x_1, x_2, \dots, x_t)$ , from the origin, as the weight for the corresponding point. We define the weighted continuous quadratic scoring rule (WCQSR) as follows:

$$S(p(x)) = \frac{2z(x)\rho(x) - \int_{\Delta(T)} z^2(x)\rho(x)dx}{t} \quad ; \quad \rho(x) = \sqrt{\sum_{j=1}^t x_j^2 j}$$

**Lemma 5.** *The WCQSR is a proper scoring rule.*

### 3.5 Budget Balance

In this final part, we consider the issue of budget balance. Note that the quality aggregation function introduced in Section 2.2 was inherently budget balanced because the same function was used for agents  $v_i$  and  $v_{i+1}$  for all  $i$ . However while designing the quality elicitation mechanism, the quality aggregation function was modified in a way that the payment by a particular agent might not always be equal to the receipt by her child. However, for the viability of the model, it is necessary that the payment is always greater than or at least equal to the receipts. This issue is important from the social planner’s point of view which in this case happens to be the social networking site.

**Budget Balance of DC Mechanism.** Consider an agent  $v_i$  and her child  $v_{i+1}$ . In the DC mechanism, the payment by agent  $v_i$  would be  $r_{i+1}V(\hat{\theta}_{i+1}|\hat{\theta}_i)$ . The receipt by agent  $v_{i+1}$  for the corresponding transaction would be  $r_{i+1}V(\hat{\theta}_i|\hat{\theta}_{i+1})$ . We have the following results for the budget balance in DC mechanism.

**Lemma 6.** *The DC mechanism is ex-ante strictly budget balanced.*

**Lemma 7.** *A DC mechanism based on WDQSR is ex-post budget balanced if  $\sum_{j=1}^t \hat{\theta}_{i,j}^2 j^2 \leq \sum_{j=1}^t \hat{\theta}_{i+1,j}^2 j^2 \forall i \in \{2, \dots, m\}$ .*

**Lemma 8.** *A DC mechanism which is based on a scoring rule and is ex-post budget balanced, need not be incentive compatible.*

**Budget Balance of PP Mechanism.** In the PP mechanism, the payment by an agent  $v_i$  would be  $r_{i+1}S(P(\hat{\theta}_{i-1}|\hat{\theta}_i))$ . The receipt by agent  $v_{i+1}$  for the corresponding transaction would be  $r_{i+1}S(P(\hat{\theta}_i|\hat{\theta}_{i-1}))$ . To this end, we can prove following results for budget balance in a PP mechanism.

**Lemma 9.** *The PP mechanism is ex-ante strictly budget balanced.*

**Lemma 10.** *The PP mechanism is ex-post budget balanced if and only if, for each agent, the observed quality type has a uniform density over a unit- $t$  simplex.*

## 4 Conclusion

We have shown in this paper how the classical query incentive network model can be extended to incorporate quality of answers. Our work opens up several possibilities for future research. One can explore the formulation of similar models in other network settings such as arbitrary branching process, power law networks, etc. As already mentioned in the paper, the effect of the selection criteria, when more than one child replies with the answer, on the answer quality and the reward structure, can be explored. In the scoring rule technique that we have employed, there is a monetary incentive for agents to report their true valuations. We believe that similar objectives can be attained by using reputation based mechanisms also.

## References

1. Guha, R., Kumar, R., Raghavan, P., Tomkins, A.: Propagation of trust and distrust. In: Proc. International Conference on the World Wide Web (WWW), pp. 403–412 (2004)
2. Kleinberg, J., Raghavan, P.: Query incentive networks. In: Proc. Symposium on Foundations of Computer Science (FOCS), pp. 132–141 (2005)
3. Arcaute, E., Kirsch, A., Kumar, R., Liben-Nowell, D., Vassilvitskii, S.: On Threshold Behavior in Query Incentive Networks. In: ACM Conference on Electronic Commerce (EC), pp. 66–74 (2007)
4. Selten, R.: Axiomatic Characterization of the Quadratic Scoring Rule. Discussion Paper Serie B 390, University of Bonn, Germany (1996)
5. Jurca, R., Faltings, B.: Confess: Eliciting honest feedback without independent verification authorities. In: Proceedings of Sixth International Workshop on Agent Mediated Electronic Commerce (AMEC VI), pp. 59–72 (2004)
6. Xia, L., Conitzer, V.: Generalized scoring rules and the frequency of coalitional manipulability. In: Proceedings of ACM Conference on Electronic Commerce, pp. 109–118 (2008)
7. Lambert, N.S., Pennock, D.M., Shoham, Y.: Eliciting properties of probability distributions. In: Proceedings of ACM Conference on Electronic Commerce, pp. 129–138 (2008)
8. Dikshit, D.: Quality Conscious and Truthful Query Incentive Networks. M.E. project report. Department of Computer Science and Automation, Indian Institute of Science, Bangalore (July 2009)

# An Analysis of Troubled Assets Reverse Auction

Saeed Alaei and Azarakhsh Malekian

Dept. of Computer Science, University of Maryland, College Park, MD 20742  
saeed@cs.umd.edu, malekian@cs.umd.edu

**Abstract.** In this paper we study the Nash-equilibrium and equilibrium bidding strategies of the Pooled Reverse Auction for troubled assets. The auction was described in (Ausubel & Cramton 2008[1]). We further extend our analysis to a more general class of games which we call *Summation Games*. We prove the existence and uniqueness of a Nash-equilibrium in these games when the utility functions satisfy a certain condition. We also give an efficient way to compute the Nash-equilibrium of these games. We show that then Nash-equilibrium of these games can be computed using an ascending auction. The aforementioned reverse auction can be expressed as a special instance of such a game. We also, show that even a more general version of the well-known oligopoly game of Cournot can be expressed in our model and all of the previously mentioned results apply to that as well.

## 1 Introduction

In this paper, we primarily study the equilibrium strategies of the pooled reverse auction for troubled assets which was described in [1]. The US Treasury is purchasing the troubled assets to infuse liquidity into the market to recover from the current financial crisis. Reverse auctions in general have been a powerful tool for injecting liquidity into the market in places where it will be most useful. As explained in [1] a simple and naive approach for the government could be to run a single reverse auction for all the holders of toxic assets as follows. The auctioneer(government) then sets a total budget to be spent. The auctioneer starts at a price like 100¢ on a dollar. All the holders, bid the quantity of their shares that they are willing to sell at the current prices. There can be excess supply. The auctioneer then lowers the price in steps e.g. 95¢, 90¢, etc. and bidders indicate the quantities that they are willing to sell at each price. At some point (for example at 30¢ on a dollar) the total supply offered by all the holders for sale equals or falls below the specified budget of the treasury. At that point the auction concludes and the auctioneer buys the securities offered at the clearing price. As explained in [1], this simple approach is flawed as it leads to a severe *adverse selection problem*. Note that at the clearing price the securities that are offered are only the ones that are actually worth less than 30¢ on each dollar of face value. They could as well worth far below 30¢. In other words, the government would pay most of its budget to buy the worst of the securities.

In [1], the authors propose the following two type of auctions.

- A Security by Security Reverse Auction
- A Pooled Reverse Auction

They are both part of a two phase plan. The first one can be used to extract private information of holders about the true value of the securities to give an estimate on how much each security and similar securities are actual worth of. Later, that information can be used to establish reference prices in the Pooled Reverse Auction.

In this paper we focus our attention on the second class of auction. In a Pooled Reverse Auction, different securities are pooled together. The government puts a reference price on each security and then runs a reverse auction on all of them together. We explain this auction in more detail in section 2.

In section 3, we study the Nash-equilibrium and the bidding strategies of the Pooled Reverse Auctions in detail. We then create a more abstract model of it at the end of section 2. In section 4 we describe a general class of games that can be used to model the Pooled Reverse Auction as well as other problems. In section 4, we give some exciting result on these games. We give a condition which is sufficient for the existence of a Nash-equilibrium. We further explain how the Nash-equilibrium can be computed efficiently using a an ascending auction-like mechanism. Later in section 5, we show how we can apply our result of section 4 to Pooled Reverse Auctions. section 6 explains how a more general version of the Cournot's oligopoly game can be expressed in our model.

## 1.1 Related Work

We partition the related works to two main groups. The first group that is closely related to our model, are computing equilibrium in Cournot and public good provision games. The second one with similar model but different objective are the works related to bandwidth sharing problems and the efficiency of computed equilibria.

One well known problem that can be considered as an example of our model is the Cournot's oligopoly game. It can be described as an oligopoly of firms producing a homogeneous good. The strategy of firm  $i$  is to choose  $q_i$  which is the quantity it produces. Assuming that the production cost is  $c_i$  per item, the utility of firm  $i$  is  $(p(Q) - c_i)q_i$  for which  $Q = \sum_i q_i$  is the total production and  $p(Q)$  is the global price of the good based on the total production. There is a vast amount of literature on Cournot games (e.g. [7]). Different aspect of Cournot equilibrium has been studied (For example, in [3] Bergstrom and Varian, studied the effect of taxation on Cournot equilibrium and also showed some characteristics of the Cournot equilibrium.)

Another set of results, with similar model, but with different criteria are the works related to bandwidth sharing problem. At a high level, the problem is to allocate a fixed amount of an infinitely divisible good among rational competing users. [8] studies this problem from pricing perspective. Kelly [6], considered a generalized variant of this problem in the context of routing and charging

(However the equilibrium point of his mechanism was not fully efficient) His model, for a single resource with fixed supply, is to give each person proportional to his bid from the resource and charge him his bid. Later, Johari et al in [4], showed that Kelly’s mechanism is at least 75% efficient at the equilibrium point. In another work, Johari et al show that, Kelly’s model minimizes efficiency loss (at the equilibrium point) when price discrimination is not allowed and then they present a class of mechanisms that has an efficient outcome at the equilibrium point assuming that price discrimination is allowed ([5]).

## 2 Model for Pooled Reverse Auction

In this section, we explain the basic model for the reverse auction of pooled securities. We will use this model throughout the rest of this paper. We start by explaining our notations:

- There are  $n$  bidders  $N = \{1, \dots, n\}$ , and  $m$  securities.
- Government has evaluated a reference price of  $r_j$  for each security  $j$ . Also let  $\mathbf{r} = (r_1, \dots, r_m)$  denote the vector of reference prices for all the  $m$  securities. The reference prices are public information. These prices are in the form of the ratio of the evaluated price to the face price and are expressed in cents per dollar. For example  $r_j = 0.75$  means every dollar of the face value of the security is actually worth 75¢.
- Each bidder  $i$  holds  $\bar{q}_{i,j}$  shares of security  $j$ . Also let  $\bar{\mathbf{q}}_i = (\bar{q}_{i,1}, \dots, \bar{q}_{i,m})$  denote the vector of the quantities of shares that bidder  $i$  holds from each security. The shares are expressed in quantity of the face value.
- Each bidder has a private valuation function  $v_i(l)$  for receiving a liquidity amount of  $l$ . In a quasi-linear setting, we would assume that  $v_i(l) = l$ . In our model, we assume the  $v_i$  could be an arbitrary function.  $v_i$  can capture the bonus for acquiring a needed amount of liquidity or can be negative to account for the cost incurred by the shortage thereof. For example consider the following:

$$v_i(l) = \begin{cases} l + (l - L_i) & l \leq L_i \\ l & l \geq L_i \end{cases} \tag{2.1}$$

We could interpret the above  $v_i$  as the following. Bidder  $i$  has a liquidity need of  $L_i$  dollars. She incurs a cost of  $L_i - l$  dollars if she raises only  $l$  dollars where  $l < L_i$ . Her value for any liquidity that she receives beyond  $L_i$  is just the same as the amount that she receives. The experimental study of reverse auction for troubled assets in [2] considers two cases for  $v_i$ . In the first case, each bidder  $i$  has a liquidity need  $L_i$  and  $v_i(l) = 2l$  for  $l \leq L_i$  and  $v_i(l) = l + L_i$  for  $l > L_i$ . In the second case, bidders don’t have liquidity needs, so  $v_i(l) = l$ . In this paper, we consider arbitrary  $v_i$  under some constraints as we will see later.

- Each bidder  $i$  has a private value of  $w_{i,j}$  for each dollar of security  $j$ . Also let  $\mathbf{w}_i = (w_{i,1}, \dots, w_{i,m})$  denote the vector of the valuations of bidder  $i$



for different securities. In reality, we should have assumed a single common value for each security which is unknown and can only be computed by aggregating all the private information of all bidders. However, that model is prohibitively hard to analyze in case of non trivial valuation functions for liquidity (i.e. when  $v_i(l)$  is not the identity function). Therefore, we assume that  $\mathbf{w}_i$  is the private values of bidder  $i$  for the securities.

Next, we briefly explain the reverse auction mechanism for pooled securities as described in [1].

**Auction 1 (Pooled Reverse Auction).** *Initially, the auctioneer (government) establishes the reference prices for all the securities. These reference prices are supposed to be the best estimate of the government about the true value of the securities. The reference prices are announced publicly.*

*The auction uses a single descending clock  $\alpha$  which specifies the current prices as a percentage of the reference prices. For example,  $\alpha = 110\%$  means the current price of each security is 110% of its reference price. As the clock goes down, participants update their bids. Bidder  $i$  submits a bid  $\mathbf{b}_i = (b_{i,1}, \dots, b_{i,m})$ , where  $b_{i,j}$  is the quantity of shares from security  $j$  that bidder  $i$  would like to sell at the current prices. These quantities are specified in terms of dollars of face value. The auctioneer collects all the bids and computes the activity points for each bidder  $i$  as  $a_i = \mathbf{r} \cdot \mathbf{b}_i$  (remember  $\mathbf{r}$  is the vector of reference prices). In other words, the activity points of each bidder is her bid quantity for each security times the reference price of that security summed over all the securities. The auctioneer also computes the total activity point  $A = \sum_i a_i$ . Assuming that  $M$  is the total budget of the government, the clock keeps going down for as long as  $A\alpha > M$ . In practice, the clock goes down in discrete steps. At each step the auctioneer collects all the bids and computes the aggregate activity point. At the first step that  $A\alpha$  becomes less than or equal to  $M$ , the clock stops and the auction concludes. The auctioneer then buys from each bidder the quantity of shares specified in her bid. Bidders are paid at the current prices (i.e. the reference price scaled by the current value of the clock). Assuming that  $\alpha^*$  was the final value of the clock and for each bidder  $i$ ,  $\mathbf{b}_i^*$  was the final bid of bidder  $i$ , the amount of liquidity that bidder  $i$  receives is  $\alpha^* \mathbf{r}_i \cdot \mathbf{b}_i^*$ .*

In the next section we study the equilibrium of the above auction.

### 3 The Equilibrium of Pooled Reverse Auction

In this section, we study the Nash-equilibrium of [Auction. 1](#) and propose a method that can be used to efficiently compute that. We also develop a bidding strategy that leads to the Nash-equilibrium.

First, we show how to compute the utility of each bidder  $i$ . Assume that  $\mathbf{b}_i$  is the bid of bidder  $i$  and  $\alpha$  is the current value of the clock. Also, as we defined in [section 2](#),  $v_i(l)$  is the valuation of bidder  $i$  for receiving amount  $l$  of liquidity and  $\mathbf{w}_i = (w_{i,1}, \dots, w_i)$  is the vector of her valuations for different securities. We

denote by  $u_i$ , the tentative utility of bidder  $i$  which is her utility if the auction stops at the current value of the clock.  $u_i$  can be computed as the following:

$$u_i = v_i(\alpha \mathbf{r} \cdot \mathbf{b}_i) - \mathbf{w}_i \cdot \mathbf{b}_i \tag{3.1}$$

Before we start with the bidding strategies, we restate some of the definitions from [Auction. 1](#).

- For a bidder  $i$  with current bid  $\mathbf{b}_i$ , we use  $a_i$  to define her activity point which is defined as:

$$a_i = \mathbf{r} \cdot \mathbf{b}_i \tag{3.2}$$

- The total activity point of all bidders is defined as:

$$A = \sum_{i=1}^n a_i \tag{3.3}$$

- The auction clock,  $\alpha$ , keeps going down for as long as  $\alpha A > M$  where  $M$  is the total budget of the auctioneer. If we denote the value of the clock when the auction stops by  $\alpha^*$ , then  $\alpha^* A \leq M$ . Note that, to simplify the analysis, we assume quantities do not need to be integers. We also assume that the clock changes continuously and bidders update their bids continuously as well. Respectively, we may assume that when the auction concludes, the auctioneers budget constraint is met with equality so:

$$\alpha^* A = M \tag{3.4}$$

Next, we show that the best strategy for each player  $i$  can be described by just specifying the activity points that she needs to generate. In other words, the only thing that bidder  $i$  has to decide is how much activity point to generate and her best bid vector can be specified as a function of that.

**Lemma 1.** *In order to play her best strategy, bidder  $i$  only needs to choose her activity points  $a_i$  and then among all the bid vectors  $\mathbf{b}_i \in [0, \bar{\mathbf{q}}_i]$ <sup>1</sup> such that  $\mathbf{r} \cdot \mathbf{b}_i = a_i$  her best strategy is to submit a bid  $\mathbf{b}_i$  that minimizes  $\mathbf{w}_i \cdot \mathbf{b}_i$ . We will refer to one such bid vector as  $\mathbf{b}_i(a_i)$ . Formally:*

$$\mathbf{b}_i(a) = \operatorname{argmin}_{\mathbf{b}} \mathbf{w}_i \cdot \mathbf{b} : \mathbf{b} \in [0, \bar{\mathbf{q}}_i] \wedge \mathbf{r} \cdot \mathbf{b} = a \tag{3.5}$$

*Proof.* See the full version.

Based on [Lemma 1](#) to describe a best strategy for a bidder  $i$  we only need to specify the activity points  $a_i$  that she should bid and then [Lemma 1](#) tells us what condition the corresponding bid vector should satisfy. The next lemma describes how we can efficiently compute  $\mathbf{b}_i(a_i)$  for any given  $a_i$ .

---

<sup>1</sup> We use the notation  $[\mathbf{a}, \mathbf{b}]$  to denote all the vectors that are componentwise greater than or equal to  $\mathbf{a}$  and less than or equal to  $\mathbf{b}$ .

**Lemma 2.** For any given  $a_i \in [0, \mathbf{r} \cdot \bar{\mathbf{q}}_i]$  we can compute  $\mathbf{b}_i(a_i)$  by using the following procedure.

Without loss of generality, assume securities are sorted in decreasing order of  $\frac{r_j}{w_{i,j}}$  so that  $\frac{r_j}{w_{i,j}} \geq \frac{r_{j+1}}{w_{i,j+1}}$ . To find the bid vector, we start from an initial zero bid vector and increase each  $q_{i,j}$  up to  $\bar{q}_{i,j}$  starting at  $j = 1$  until the generated activity point reaches  $a_i$ . The following is a more formal definition of  $\mathbf{b}_i(a)$ :

$$\mathbf{b}_i(a) = (\bar{q}_{i,1}, \dots, \bar{q}_{i,y-1}, b_{i,y}, 0, \dots, 0) \tag{3.6}$$

such that:

$$r_y b_{i,y} + \sum_{j=1}^{y-1} r_j \bar{q}_{i,j} = a \tag{3.7}$$

*Proof.* See the full version.

Intuitively, Lemma 2 is saying that a strategic bidder should never sell any shares of a security  $j$  unless for any other security  $k$  for which  $\frac{r_k}{w_{i,k}} > \frac{r_j}{w_{i,j}}$  she has already sold all of her shares of security  $k$ .

**Definition 1.** We can define a cost function  $c_i(a) : [0, \mathbf{w}_i \cdot \bar{\mathbf{q}}_i] \rightarrow \mathbb{R}$  for each bidder  $i$  which only depends on her activity points:

$$c_i(a) = \mathbf{w}_i \cdot \mathbf{b}_i(a) \qquad 0 \leq a \leq \mathbf{r} \cdot \bar{\mathbf{q}}_i \tag{3.8}$$

Intuitively, for bidder  $i$ ,  $c_i(a)$  is the minimum cost of generating 'a' activity points.

At this point, we can define the bid vectors and all the equations only in terms of  $a_i$ . Bidders only need to specify their activity point  $a_i$ . We denote the final activity points of bidder  $i$  when the auction concluded by  $a_i^*$  and the final total activity point by  $A^*$ . The utility of each bidder  $i$  can now be written as the following:

$$u_i = v_i(\alpha^* a_i^*) - c_i(a_i^*) \tag{3.9}$$

Also, the auction concludes at the highest clock  $\alpha^*$  such that:

$$\sum_{i=1}^n \alpha^* A^* = M \tag{3.10}$$

Next, we define the Nash-equilibrium. Before that, notice we can write the utility of each bidder  $i$  as  $u_i(a, A)$  which is a function of her own bid and the total aggregate bid. Formally:

$$u_i(a, A) = v_i\left(\frac{a}{A} M\right) - c_i(a) \tag{3.11}$$

Now we are ready to describe the Nash-equilibrium. Suppose  $a_1^*, \dots, a_n^*$  are the activity points at which the auction has concluded. We say the outcome of

the auction is *stable* or is a *Nash-equilibrium* if for every bidder  $i$ ,  $a_i^*$  is a best response to  $a_{-i}^*$ . For a Nash equilibrium, the first order and boundary conditions are sufficient. Assume that  $\bar{a}_i$  denotes the maximum possible activity points that bidder  $i$  can generate (i.e.,  $\bar{a}_i = \mathbf{r} \cdot \bar{q}_i$ ). The first order and boundary conditions of the Nash-equilibrium are the following:

$$\forall i \in N : \begin{cases} \frac{d}{da_i^*} u_i(a_i^*, A^*) = 0 & \text{and } 0 < a_i^* < \bar{a}_i \\ \text{or} \\ \frac{d}{da_i^*} u_i(a_i^*, A^*) \leq 0 & \text{and } a_i^* = 0 \\ \text{or} \\ \frac{d}{da_i^*} u_i(a_i^*, A^*) \geq 0 & \text{and } a_i^* = \bar{a}_i \end{cases} \quad (3.12)$$

$$A^* = \sum_{i=1}^n a_i^* \quad (3.13)$$

Note that, to use the first order conditions, we need  $u_i(a, A)$  to be a continuous and differentiable function in its domain. We can however relax the differentiability requirement and allow  $u_i(a, A)$  to have different left and right derivatives at a finite number of points. In that case, if assume that  $\rho_i^-$  is the left derivative of  $\frac{d}{da_i^*} u_i(a_i^*, A^*)$  and  $\rho_i^+$  is its right derivative, then in the first condition, we can replace  $\frac{d}{da_i^*} u_i(a_i^*, A^*) = 0$  with  $\rho_i^- \leq 0 \leq \rho_i^+$ . To keep the proofs simple, we do not use this general form but we will refer to it later when we explain how to compute the equilibrium.

We further expand the first order and boundary conditions. Notice that  $\frac{d}{da_i^*} u_i(a_i^*, A^*) = \frac{\partial}{\partial a} u_i(a_i^*, A^*) + \frac{\partial}{\partial A} u_i(a_i^*, A^*) \frac{d}{da_i^*} A^*$ . Because we always have  $\frac{d}{da_i^*} A^* = 1$ , we can write  $\frac{d}{da_i^*} u_i(a_i^*, A^*) = \frac{\partial}{\partial a} u_i(a_i^*, A^*) + \frac{\partial}{\partial A} u_i(a_i^*, A^*)$ . So the first order and boundary conditions can be rephrased as:

$$\forall i \in N : \begin{cases} \frac{\partial}{\partial a} u_i(a_i^*, A^*) + \frac{\partial}{\partial A} u_i(a_i^*, A^*) = 0 & \text{and } 0 \leq a_i^* \leq \bar{a}_i \\ \text{or} \\ \frac{\partial}{\partial a} u_i(a_i^*, A^*) + \frac{\partial}{\partial A} u_i(a_i^*, A^*) \leq 0 & \text{and } a_i^* = 0 \\ \text{or} \\ \frac{\partial}{\partial a} u_i(a_i^*, A^*) + \frac{\partial}{\partial A} u_i(a_i^*, A^*) \geq 0 & \text{and } a_i^* = \bar{a}_i \end{cases} \quad (3.14)$$

$$A^* = \sum_{i=1}^n a_i^* \quad (3.15)$$

Next, we state the main theorem of this section which gives a sufficient condition for the existence of a Nash-equilibrium and provides a method for computing it as well as a bidding strategy.

**Theorem 1.** Consider the Auction. 1, in which as explained before, each bidder's utility is given by  $u_i = v_i(\alpha \mathbf{r} \cdot \mathbf{b}_i) - \mathbf{c}_i \cdot \mathbf{b}_i$  if the auction stops at the current

clock  $\alpha$ . Assuming that the valuation functions  $v_i$  are continuous, differentiable<sup>2</sup> and concave, there exists a unique Nash-equilibrium that satisfies the first order and boundary conditions of (3.14). Furthermore, there are bid functions  $g_i(\alpha)$ , such that for every  $i$  if bidder  $i$  bids  $\mathbf{b}_i^*(a_i)$  where  $a_i = g_i(\alpha)$ , then the outcome of the auction coincides with the unique Nash-equilibrium.  $g_i(\alpha)$  is given below ( $v'_i$  and  $c'_i$  are the derivatives of  $v_i$  and  $c_i$ ):

$$g_i(\alpha) = \operatorname{argmin}_{a \in [0, \bar{a}_i]} \left| v'_i(\alpha a) \frac{M - \alpha a}{M} \alpha - c'_i(a) \right| \tag{3.16}$$

Furthermore, the  $g_i(\alpha)$  can be computed efficiently using binary search on 'a' (the parameter of the argmin) because the expression inside the absolute value is a decreasing function of  $a$ .

Note that the requirement of  $v_i$  functions being concave is quite natural. It simply means that the derivative of  $v_i$  should be decreasing which can be interpreted as the marginal value of the first dollar received being more than the marginal value of the last dollar.

It is worth mentioning that the bid function  $g_i(\alpha)$ , as described in (3.16) is not necessarily an increasing function of  $\alpha$ . In other words, as the clock goes down,  $g_i(\alpha)$  may increase at some points which means bidder  $i$  is actually offering more for sale although the prices are going down. This phenomenon is in fact quite common when bidder  $i$  has liquidity needs as we will explain in section 5.

We defer the proof of Theorem 1 to section 5. Instead of proving Theorem 1 directly, we prove a more general theorem in the next section. Later, in section 5, we show that Theorem 1 is a special case of that.

## 4 Summation Games

In this section, we describe a general class of games which we will refer to as summation games. Later, we show that the reverse auction explained in the previous section and some well known problems like the Courant-Nash equilibrium of an oligopoly game [7] can be expressed in this model. Next, we define a *Summation Game*:

**Definition 2 (Summation Game).** *There are  $n$  players  $N = \{1, \dots, n\}$ . Each player can choose a number  $a_i$  from the interval  $[0, \bar{a}_i]$  where  $\bar{a}_i$  is a constant. The utility of each bidder depends only on her own number as well as the sum of all the numbers. In other words, assuming that  $A = \sum_{i=1}^n a_i$ , the utility of each bidder  $i$  is given by  $u_i(a_i, A)$ .*

We next show that if the utility functions  $u_i(a, A)$  meet a certain requirement, the summation game has a unique Nash-equilibrium that can be computed efficiently. Before that, we define the following notation

---

<sup>2</sup> We may relax this to allow  $v_i$  to have different left and right derivatives at a finite number of points.

**Definition 3.** For each player  $i$ , assuming that  $u_i(a, A)$  is her utility function, define her characteristic function  $h_i(x, T)$  as the following:

$$h_i(x, T) = \frac{\partial}{\partial a} u_i(xT, T) + \frac{\partial}{\partial A} u_i(xT, T) \tag{4.1}$$

**Theorem 2.** If all the characteristic functions  $h_i(x, T)$  are strictly decreasing functions in both  $x$  and  $T$ , then the game has a unique Nash-equilibrium and in that equilibrium, the bid of each player  $i$  is  $a_i = x_i(A)A$  where  $x_i$  is defined as the following:

$$x_i(T) = \operatorname{argmin}_{x \in [0, \min(1, \frac{a_i}{T})]} |h_i(x, T)| \tag{4.2}$$

Furthermore, because  $h_i(x, T)$  is decreasing in both  $x$  and  $T$ ,  $x_i(T)$  is also decreasing in  $T$  and the equilibrium can be computed efficiently using two nested binary searches or using an auction-like mechanism with an ascending clock  $T$  in which the each bidder  $i$  submits  $a_i = x_i(T)T$  and the clock  $T$  keeps going up as long as  $\sum_i a_i > T$ .

*Proof.* See the full version.

To find the values of  $a_i$ 's at the Nash-equilibrium we can use the following algorithm:

**Algorithm 1**

- Start with  $T = 0$  (or a sufficiently small positive  $T$ ).
- Keep increasing  $T$  for as long as  $\sum_{i=1}^n x_i(T) > 1$ .
- Stop as soon as  $\sum_{i=1}^n x_i(T) \leq 1$  and then set each bid  $a_i = x_i(T)T$

*Proof.* See the full version.

Note that **Alg. 1** can be implemented either using binary search on  $T$  or as an ascending auction-like mechanism in which each player submits the bid  $a_i = x_i(T)T$  where  $T$  is the ascending clock and in which the clock stops once  $\sum_i a_i \leq T$ .

In the next section we finish our analysis of the pooled reverse auction of **1**. Later, in **section 6**, we give example of a well-known problem that can be expressed in our model and its Nash-equilibrium can be computed using **Alg. 1**.

<sup>3</sup> Note that we allow  $h_i(x, T)$  to be discontinuous at a finite number of points (e.g. a step function).

<sup>4</sup> If we relax the requirement of  $h_i$ 's being strictly decreasing to just being *non-increasing* then there is a continuum of Nash-equilibria in which there is one Nash-equilibrium that is strictly preferred by some players and is just as good as other Nash equilibria for other players.

<sup>5</sup> If  $\sum_{i=1}^n x_i(T^*) < 1$  then arbitrarily choose each  $a_i^*$  from the interval  $[\lim_{\epsilon \rightarrow 0+} x_i(T^* - \epsilon)T^*, x_i(T^*)T^*]$  such that  $\sum_{i=1}^n a_i^* = T^*$  (It is easy to show that each player  $i$  is indifferent to all  $a_i^* \in [\lim_{\epsilon \rightarrow 0+} x_i(T^* - \epsilon)T^*, x_i(T^*)T^*]$ ).

## 5 Back to Pooled Reverse Auction

In the previous section we described a more general class of games and in [Theorem 2](#) we gave sufficient conditions for the existence of a Nash-equilibrium. We explained when it is unique and how to compute it. In this section we continue our analysis of [Auction. 1](#). We first give a proof for [Theorem 1](#) by reducing it to a special case of [Theorem 2](#).

Next, we give a proof for [Theorem 1](#) which is based on a reduction to [Theorem 2](#).

*Proof (Proof of [Theorem 1](#)).* To be able to apply [Theorem 2](#), we first need to show that the utility function of each bidder in [Auction. 1](#) meets the requirement of [Theorem 2](#). More specifically, we have to show that  $h_i(x, T) = \frac{\partial}{\partial a} u_i(xT, T) + \frac{\partial}{\partial A} u_i(xT, T)$  is a decreasing function in both  $x$  and  $T$ . Remember that in our model for [Auction. 1](#), we can write the utility of bidder  $i$  as  $u_i(a, A) = v_i(\frac{a}{A}M) - c_i(a)$ . First, we show that  $c_i(a)$  is a convex function.

**Lemma 1.** *The cost function  $c_i(x)$  as defined in [\(3.8\)](#) is always a convex function and has an non-decreasing first order derivative in  $[0, \mathbf{r}.\bar{q}]$  although its derivative might be discontinuous in at most  $m$  points.*

*Proof.* See the full version.

**Lemma 2.** *The  $h_i(x, T)$  functions for bidders in [Auction. 1](#) are decreasing in both  $x$  and  $T$ :*

$$h_i(x, T) = \frac{\partial}{\partial a} u_i(xT, T) + \frac{\partial}{\partial A} u_i(xT, T) \tag{5.1}$$

$$= v'_i(xM) \frac{1-x}{T} + c'_i(Tx) \tag{5.2}$$

*Proof.* See the full version.

Since in [Lemma 2](#) we proved that  $h_i(x, T)$  is decreasing in both  $x$  and  $T$ , we can now apply [Theorem 2](#) and all of the claims of [Theorem 1](#) follow from [Theorem 2](#). Also note that  $g_i(\alpha)$  which was defined in [\(3.16\)](#) is actually the same as  $x_i(T)T$ , where  $T = \frac{M}{\alpha}$ .

It is interesting to notice that the auction-like mechanism of [Theorem 2](#) and [Auction. 1](#) are actually equivalent. In fact, the  $x_i(T)$  where  $T = M/\alpha$ , has a very natural interpretation in [Auction. 1](#). It specifies the fraction of the budget of the auctioneer that the bidder  $i$  is demanding at the clock  $\alpha$ . In fact we may modify the auction to ask the bidders to submit the amount of liquidity that they are demanding directly at each step of the clock and then the auction stops when the demand becomes less than or equal to the budget of the auctioneer. Then, each bidder will be required to sell enough quantity of her shares at the current prices to pay for the liquidity that she had demanded.

It is easy to see that the liquidity that each bidder demands may only decrease as the  $\alpha$  increases. However, the value of the bid,  $x_i(T)T$ , may actually increase because bidder  $I$  may want to maintain her demand for the liquidity.

## 6 Application to Cournot's Oligopoly

In this section, we show how the well-known problem of *Cournot's Oligopoly* can be expressed in our model of a summation game and all the results of [Theorem 2](#) can therefore be applied:

### Definition 4 (Cournot's Oligopoly)

- There are  $n$  firms. The firms are oligopolist suppliers of a homogenous good.
- At each period, each firm chooses a quantity  $q_i$  to supply.
- The total supply  $Q$  on the market is the sum of all firms' supplies:

$$Q = \sum_i q_i \quad (6.1)$$

- All firms receive the same price  $p$  per unit of the good. The price  $p$  on the market depends on the total supply  $Q$  as:

$$p(Q) = p_0(Q_{\max} - Q) \quad (6.2)$$

- Each firm  $i$  incurs a cost  $c_i$  per unit of good. These costs can be different for different firms and are private information
- Each firm  $i$ 's profit is given by:

$$u_i(q_i, Q) = (p(Q) - c_i)q_i \quad (6.3)$$

- After each market period, firms are informed of the total quantity  $Q$  and the market price  $p(Q)$  of the previous period.

If we write down the  $h_i(x, T)$  for each firm  $i$  we get:

$$h_i(x, T) = \frac{\partial}{\partial a} u_i(xT, T) + \frac{\partial}{\partial A} u_i(xT, T) \quad (6.4)$$

$$= p(T) - c_i + p'(T)Tx \quad (6.5)$$

$$= p_0(Q_{\max} - T) - c_i - p_0Tx \quad (6.6)$$

Notice that clearly the above  $h_i(x, T)$  is a decreasing function of both  $x$  and  $T$  and therefore all of the nice results of [Theorem 2](#) can be applied. Notice in fact that as long as  $p(Q)$  is concave and a decreasing function of  $Q$ ,  $h_i(x, T)$  is still a decreasing function of both  $x$  and  $T$  and all of the results of [Theorem 2](#) still holds.

## 7 Conclusion

In this paper we studied the Nash-equilibrium and equilibrium bidding strategies of the troubled assets reverse auction. We further generalized our analysis to a more general class of games with non quasi-linear utilities. We proved the existence and uniqueness of a Nash-equilibrium in those games and we also gave



an efficient way to compute the equilibrium of those games. We also showed that finding the Nash equilibrium can be implemented using an ascending mechanism so that the participants don't need to reveal their utility functions. We also, showed that even a more general version of the well-known problem of Cournot's Oligopoly can be expressed in our model and all of the previously mentioned results apply to that as well.

## References

1. Ausubel, L., Cramton, P.: A troubled asset reverse auction. Working paper (2008)
2. Ausubel, L.M., Cramton, P., Filiz-Ozbay, E., Higgins, N., Ozbay, E., Stocking, A.: Common-value auctions with liquidity needs: an experimental test of a troubled assets reverse auction. University of maryland, economics working paper series, Department of Economics, University of Maryland (2008)
3. Bergstrom, T., Varian, H.: Two remarks on cournot equilibria. University of california at santa barbara, economics working paper series, Department of Economics, UC Santa Barbara (February 1985)
4. Johari, R., Mannor, S., Tsitsiklis, J.N.: Efficiency loss in a network resource allocation game. *Mathematics of Operations Research* 29, 407–435 (2004)
5. Johari, R., Tsitsiklis, J.: Efficiency of scalar-parameterized mechanisms. *Operation Research* (2008)
6. Kelly, F.: Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8, 33–37 (1997)
7. Mas-colell, A., Whinston, M.D., Green, J.R.: *Microeconomics Theory*. Oxford University Press, Oxford (1995)
8. Shenker, S., Clark, D., Estrin, D., Herzog, S.: Pricing in computer networks: Reshaping the research agenda. *ACM Computer Communication Review* 26, 183–201 (1996)

# Direction Preserving Zero Point Computing and Applications

## (Extended Abstract)

Xiaotie Deng<sup>1</sup>, Qi Qi<sup>2</sup>, and Jie Zhang<sup>1</sup>

<sup>1</sup> Department of Computer Science,  
City University of Hong Kong, Hong Kong SAR, P.R. China  
deng@cs.cityu.edu.hk, zhangjie3@student.cityu.edu.hk

<sup>2</sup> Department of Management Science and Engineering,  
Stanford University, Stanford, California, USA  
kaylaqi@stanford.edu

**Abstract.** We study the connection between the direction preserving zero point and the discrete Brouwer fixed point in terms of their computational complexity. As a result, we derive a PPAD-completeness proof for finding a direction preserving zero point, and a matching oracle complexity bound for computing a discrete Brouwer’s fixed point.

Building upon the connection between the two types of combinatorial structures for Brouwer’s continuous fixed point theorem, we derive an immediate proof that TUCKER is PPAD-complete for all constant dimensions, extending the results of Pálvölgyi for 2D case [20] and Papadimitriou for 3D case [21]. In addition, we obtain a matching algorithmic bound for TUCKER in the oracle model.

## 1 Introduction

Fixed point theorems have been fundamental in the literatures of Economics, such as equilibrium analysis of markets [1]. Fixed point computation has also been important in computational complexity study of equilibria. There is however a choice to make when we discretize the original concept of fixed point from the continuous domain.

Clearly, Sperner’s simplex is a natural candidate for a discrete version of the fixed point concept. Scarf, on the other hand, proposed a primitive set structure in his study of fixed point computation [23]. Iimura [13] defined the direction preserving function as a discrete version for the continuous function, and developed theorems for the existence of a zero point for any bounded direction preserving function.

The direction preserving function is defined on hypergirards, consisting of numerous hypercubes aligned along hyperplanes perpendicular to the coordinates. Informally, values of a direction preserving function will not have opposite signs on any pair of neighboring nodes. Based on such a concept in modeling continuous functions, Chen and Deng proved a matching algorithmic bound for approximate fixed point computation under the oracle complexity model [2].

Daskalakis, et al., on the other hand, developed a different discrete version of fixed point, defining it to be a unit hypercube on which all  $d + 1$  colors appear, for a hypergrid colored with  $d + 1$  colors. Its computation lays down the foundation for the settlement of the computational complexity of Nash equilibrium computation [8,5].

In this paper, we establish the equivalence of the two concepts in terms of their computation, both under the oracle function model, and the polynomial time computable function model. Therefore, it derives two immediate new results: PPAD-completeness for finding a zero point for any direction preserving function, and a matching algorithmic bound for finding a unit hypercube that contains all colors.

The computational equivalence of those two important models for discrete fixed point solutions has important implications. Based on the PPAD-complete result for zero point computation for direction preserving functions, we establish a PPAD-completeness proof of Tucker's problem for all constant dimensions. In the seminal paper of Papadimitriou introducing the **PPAD** class [21], he proved the PPAD-completeness of Tucker for 3D. It was also noticed that there is a similarity in difficulties to attempt PPAD-complete proofs for 2D SPERNER and 2D TUCKER. For 2D SPERNER, the difficulty was solved later by Chen and Deng [4]. The recent proof of PPAD-completeness for 2D Tucker by Pálvölgyi [20] exploits the techniques developed by Chen and Deng for 2D Sperner's PPAD proof. Our proof unifies the proof for all constant dimensions, through a modular approach. The connection also allows us to derive a matching oracle complexity bound for the computation of Tucker, for all constant dimensions.

The connection between DPZP and BROUWER requires a combinatorial structure that relates an integer function of values from 0 through  $d$  to a direction preserving function of values  $\{0, \pm e_i : i = 1, 2, \dots, d\}$ . For the oracle complexity of TUCKER, we develop a combinatorial parity lemma that relates the existence of complementary edges on the boundary with certain parity property of the fully colored simplexes on the boundary of the polyhedron.

Integer fixed point computation has important applications in Economics [13]. The PPAD-complete proof for direction preserving functions further reinforces our understanding of the correct complexity class it belongs to. The connection of the two different discrete fixed points may help in the studies of other related applications.

Tucker's lemma characterizes the fundamental combinatorial property underlying important mathematical problems such as the Borsuk-Ulam theorem, and Lovász's theorem on Kneser Graph (by Matoušek [18]). It has also applications to a type of the fair division problem called consensus-halving problem (Simmons and Su [24]) It has also been applied to the necklace problem and consensus- $\frac{1}{k}$ -division problem (Longueville and Živaljević [16]).

In section 2, we review related concepts such as BROUWER, Direction Preserving Zero Point (DPZP for short) and TUCKER. In section 3, we prove DPZP is PPAD-complete and derive a matching bound for BROUWER. In section 4, we prove TUCKER is PPAD-complete for all constant dimension  $d$ , which

extends the result for 2-D Tucker by [20] and 3-D Tucker by [21]. In section 5, we derive an oracle computational matching bound for finding a complementary edge in TUCKER.

## 2 Preliminaries

In this section, we formally introduce three problems: the BROUWER, DPZP and the TUCKER.

Those three problems have solutions as guaranteed by mathematical theorems, often with an embedded parity argument. Complexities for finding such solutions of search problems are characterized by the classes of **FNP** and **TFNP** [19]. **PPAD** is a subclass of **TFNP** that is defined by a complete problem, the **End of Line** problem (or called **LEFFE** as in [3]). Several important fundamental problems are shown to be complete in this class, including SPERNER [214], NASH [8,9,3,5], BROUWER, KAKUTANI, BORSUK-ULAM [21], approximate-NASH [6], Exchange Economy [21].

### 2.1 BROUWER

We should start with one class of those **PPAD**-complete classes, BROUWER, defined as follows.

Consider a  $d$  dimension hypercube ( $d$ -hypercube for short) of size  $N$ . A  $d$  dimension hypergrid ( $d$ -hypergrid for short) of scale  $N$  places  $N - 1$  equally spaced hyperplanes parallel to each of the boundary faces of the hypercube and divides the hypercube into  $N^d$  base hypercubes.

Let

$$V_N^d = \{\mathbf{p} = (p_1, p_2, \dots, p_d) \in \mathcal{Z}^d \mid \forall i : 0 \leq p_i \leq N\}$$

Its boundary  $B_N^d$  consists of point  $\mathbf{p} \in V_N^d$  with  $p_i \in \{0, N\}$  for some  $i$ :

$$B_N^d = \{\mathbf{p} \in V_N^d : p_i \in \{0, N\} \text{ for some } i\}.$$

For each  $\mathbf{p} \in \mathcal{Z}^d$ , let

$$\mathcal{K}_{\mathbf{p}} = \{\mathbf{q} : q_i \in \{p_i, p_i + 1\}\}.$$

Let  $m(\mathbf{p}) = 0$  if  $\forall i, p_i > 0$  and  $m(\mathbf{p}) = \min\{i : p_i = 0\}$ , otherwise. A color function  $g : V_N^d \Rightarrow \{0, 1, \dots, d\}$  is valid if  $g(\mathbf{p}) = m(\mathbf{p}), \forall \mathbf{p} \in B_N^d$ .

**Definition 1.** *The input of BROUWER is a pair  $(G, V_N^d)$  where  $G$  generates a valid  $d$ -coloring  $\mathbf{g}$  on  $V_N^d$ . That is, given any  $\mathbf{p} \in V_N^d$ ,  $G$  takes time polynomial in  $N$  to compute the color  $g(\mathbf{p})$  of a point  $\mathbf{p}$ .*

*The output of BROUWER is a point  $\mathbf{p} \in V_N^d$  such that  $\mathcal{K}_{\mathbf{p}}$  is fully colored, that is,  $\mathcal{K}_{\mathbf{p}}$  has all  $d + 1$  colors.*

It is known that BROUWER is **PPAD**-complete [214] when  $G$  is a polynomial time algorithm. As it is closely related to the Sperner's problem SPERNER [4], it is known to have a matching bound for the oracle model [2] as a result of the same lower bound for SPERNER [10].

## 2.2 Direction Preserving Zero Point

We should introduced another problem closely related to the fixed-point problem but for a restricted class of functions, the direction preserving functions.

**Definition 2.** A function  $f : V_N^d \Rightarrow \{0, \pm e_i, i = 1, 2, \dots, d\}$  is direction preserving if for any  $\mathbf{p} \in V_N^d$ ,  $s, t \in \mathcal{K}_{\mathbf{p}}$ ,  $f(s) + f(t) \neq 0$ .

We call a direction preserving function  $f$  on  $V_N^d$  is bounded if  $\forall \mathbf{p} \in V_N^d, \mathbf{p} + f(\mathbf{p}) \in V_N^d$ .

We define the DPZP (Direction Preserving Zero Point) problem as follows:

**Definition 3.** The input of DPZP is a pair  $(F, V_N^d)$  where  $F$  computes the value  $f(\mathbf{p})$  of a bounded direction preserving function for each  $\mathbf{p} \in V_N^d$ . That is, given any  $\mathbf{p} \in V_N^d$ ,  $F$  takes time polynomial in  $N$  to compute the function value  $f(\mathbf{p})$  of a point  $\mathbf{p}$ . The output of DPZP is a point  $\mathbf{p} \in V_N^d$  such that  $f(\mathbf{p}) = 0$ , which is called a zero point.

It was known that a bounded direction preserving function  $f$  on  $V_N^d$  always has a zero point [13,14], and there is a matching algorithmic bound [2] for the oracle model.

## 2.3 TUCKER

The complexity of TUCKER was first considered by Papadimitriou in [21], it was proved that TUCKER is in **PPAD**, and 3D TUCKER is **PPAD**-complete. The problem can be defined as follows:

**Definition 4.** (TUCKER) Consider the  $d$  dimension hypercube which is triangulated to be base simplexes. All vertices  $V$  of the triangulated hypercube are colored by a function  $f : V \Rightarrow \{\pm 1, \dots, \pm d\}$  and the boundary vertices are colored antipodal preserving, that is,  $f(-\mathbf{v}) = -f(\mathbf{v})$  for any boundary vertex  $\mathbf{v}$ . The problem is to find out a complementary 1-simplex in the hypercube, i.e., find out an edge in the triangulated hypercube so that the coloring of its two endpoints  $\mathbf{p}, \mathbf{q}$  satisfying  $f(\mathbf{p}) = -f(\mathbf{q})$ .

The existence of the solution for this problem is guaranteed by a combinatorial topology theorem, Tucker’s lemma [26]. The original problem is built on a  $d$  dimension ball, here we cast it into a  $d$ -hypercube to simplify of discussion.

## 3 PPAD-Completeness of Direction Preserving Zero Point

In this section, we assume there is a polynomial time algorithm to compute  $F$  in the DPZP problem. We prove DPZP is PPAD-complete and give a matching bound for BROUWER.

We will first introduce a reduction from DPZP to BROUWER. We then introduce a reduction from BROUWER to DPZP. As both reduction are polynomial

time with respect to the algorithm  $F$  and  $G$ , it shows that DPZP is **PPAD**-complete for the case  $F$  is a polynomial time algorithm. Since DPZP has a matching algorithmic bound of  $\theta(N^{d-1})$  [2], it shows that BROUWER has a same matching algorithmic bound of  $\theta(N^{d-1})$ .

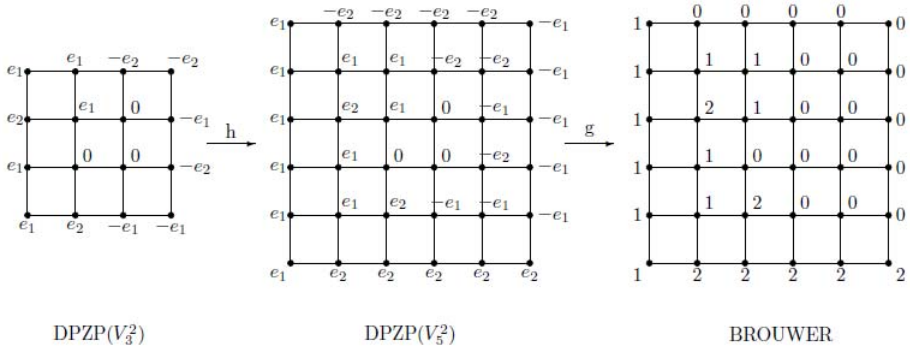
**Lemma 1.** *DPZP is in PPAD class.*

*Proof.* We prove it in two steps. First we add one more layer to any given input of DPZP to achieve a specific boundary condition. Then we reduce the expanded DPZP to BROUWER.

Given an input of DPZP  $(F, V_N^d)$ , first we add one more layer on each face of its boundary which expands  $V_N^d$  to  $V_{N+2}^d$ . Next we define a function  $H$ : for each vertex  $\mathbf{p} \in V_N^d$ ,  $h(\mathbf{p}) = F(\mathbf{p})$  and for any boundary point  $\mathbf{p} \in V_{N+2}^d \setminus V_N^d$ , either  $h(\mathbf{p}) = \mathbf{e}_i$  if  $p_i = 0$  and  $\forall j < i : p_j > 0$ , or  $h(\mathbf{p}) = -\mathbf{e}_i$  if  $\mathbf{p} \gg 0$  (all coordinates of  $\mathbf{p}$  are positive),  $p_i = N + 2$  and  $\forall j < i : p_j < N + 2$ . It is easy to verify that  $(H, V_{N+2}^d)$  is still a DPZP.

Next, we reduce DPZP  $(H, V_{N+2}^d)$  to BROUWER  $(G, V_{N+2}^d)$  as follows: define  $g(\mathbf{p}) = i$  if  $H(\mathbf{p}) = \mathbf{e}_i$ ;  $g(\mathbf{p}) = 0$  otherwise.

We illustrate the processes in Fig.1.



**Fig. 1.** The Proof of Lemma 1

Clearly, a direction preserving with a bounded specific boundary condition function  $h$  for DPZP translates into a valid coloring for BROUWER. For a solution of BROUWER,  $\mathcal{K}_{\mathbf{p}}$ , we should have  $\mathbf{q} \in \mathcal{K}_{\mathbf{p}}$  such that  $g(\mathbf{q}) = 0$ . Then  $h(\mathbf{q}) \in \{0, -\mathbf{e}_1, -\mathbf{e}_2, \dots, -\mathbf{e}_d\}$  by the process of our reduction. For  $\mathbf{t} \in \mathcal{K}_{\mathbf{p}}$  and  $g(\mathbf{t}) = i \neq 0$ , we have  $h(\mathbf{t}) = \mathbf{e}_i$ . Since all colors appear in  $\mathcal{K}_{\mathbf{p}}$ , by direction preserving property of  $h$ , we must have  $h(\mathbf{q}) = 0$ , and hence  $\mathbf{q}$  is a solution to DPZP.

**Lemma 2.** *DPZP is PPAD-hard.*

*Proof.* Given an input of BROUWER  $(G, V_N^d)$ , we define an input for DPZP on an expanded hypergrid. The expanded hypergrid obtained by placing  $2N$

more hyperplanes parallel to each of the boundary faces which refine each base hypercube in  $V_N^d$  into  $3^d$  smaller hypercubes.

The assignment of values to a grid point depends on the minimum dimension original unit hypercube that contains it. The main idea is maintain the following properties (see the full version for details):

- (I) A valid coloring  $g$  in BROUWER is translated into a bounded direction preserving function  $f$  for DPZP;
- (II) Once finding a zero point in DPZP, we can get a corresponding fixed point set in BROUWER.

Combining the above two lemmas, we have

**Theorem 1.** *Direction Preserving Zero Point(DPZP) is **PPAD**-complete.*

Since DPZP has a matching algorithmic bound of  $\theta(N^{d-1})$  [2], by the above reduction processes from DPZP to BROUWER and from BROUWER to DPZP, we obtain the following corollary.

**Corollary 1.** *BROUWER has a matching algorithmic bound of  $\theta(N^{d-1})$ .*

## 4 The Complexity of TUCKER

In this section, we prove that  $d$ -D TUCKER is **PPAD**-complete for all constant dimension  $d$ , and hence extending the results for 2-D Tucker by [20] and 3-D Tucker by [21].

First of all, it is known that TUCKER is in PPAD.

**Lemma 3.** [21]  *$d$ -D TUCKER is in **PPAD**.*

We prove TUCKER is **PPAD**-hard for any dimension  $d$ . The reduction is based on DPZP and Kuhn’s triangulation [15]. We should introduce Kuhn’s triangulation first before giving the formal proof of **PPAD**-hardness.

**Definition 5.** (*Kuhn’s Triangulation*) *Let a  $d$ -hypercube of side length  $N$  be located in the first quadrant and one of its corner point at the origin, denote this vertex as  $\mathbf{v}_0 = (0, 0, \dots, 0)_{1 \times d}$  and its diagonal vertex as  $(N, N, \dots, N)_{1 \times d}$ . We apply hyperplanes of side length 1 which parallel to coordinate axes to cut the big hypercube into  $N^d$  unit  $d$ -hypercubes. Let  $\mathbf{e}_i$  be a  $d$ -dimension vector that its  $i$ -th coordinate be 1 and others be 0. Let  $\pi = (\pi(1), \pi(2), \dots, \pi(d))$  be any permutation of integers  $0, 1, \dots, d - 1$ . Next we triangulate each unit hypercube in the following way: Let the vertex which is closest to  $\mathbf{v}_0$  in each unit hypercube be the base point. Each permutation  $\pi$  corresponds to one base simplex whose vertices are given by  $\mathbf{v}_\pi^i = \mathbf{v}_\pi^{i-1} + \mathbf{e}_{\pi(i)}$ , where  $\mathbf{v}_\pi^0$  is the base point of that unit hypercube. Then each of the unit hypercube is triangulated to  $d!$  base simplexes.*

We illustrate Kuhn’s triangulation for a 3-dimension hypercube with side length 1 in Fig. 2. It is easy to see that all vertices of the base simplexes are the vertices of the hypercubes. We claim that these simplexes all have disjoint interiors and




Fig. 2. Kuhn’s triangulation in 3 dimension

the union of them is the  $d$ -hypercube. In fact, Kuhn’s triangulation can be acquired by an equivalent cutting approach, that is, first using the  $d!$  permutations  $\pi$  to triangulate the big hypercube to  $d!$  simplexes, and then triangulate each big simplex into  $N^d$  base simplexes.

**Lemma 4.**  $d$ -D TUCKER is **PPAD-hard**.

*Proof.* To prove TUCKER is **PPAD-hard**, we reduce DPZP to it. For any input of DPZP  $(F, V_N^d)$ , we first add one more layer to make DPZP  $(H, V_{N+2}^d)$  satisfy the antipodal constraint for the function values. This can be done by defining  $h$  as: for each vertex  $\mathbf{p} \in V_N^d$ ,  $h(\mathbf{p}) = F(\mathbf{p})$  and for any boundary point  $\mathbf{p} \in V_{N+2}^d \setminus V_N^d$ , either  $h(\mathbf{p}) = \mathbf{e}_i$  if  $p_i = 0$  and  $\forall j < i : N + 2 > p_j > 0$ , or  $h(\mathbf{p}) = -\mathbf{e}_i$  if  $p_i = N + 2$  and  $\forall j < i : 0 < p_j < N + 2$ . It is easy to verify that  $(H, V_{N+2}^d)$  is still a DPZP.

Next we reduce DPZP  $(H, V_{N+2}^d)$  to TUCKER  $(G, V_{N+2}^d)$  by defining function  $g$  as:  $g(\mathbf{p}) = i$  if  $h(\mathbf{p}) = \mathbf{e}_i$ ;  $g(\mathbf{p}) = -i$  if  $h(\mathbf{p}) = -\mathbf{e}_i$ ;  $g(\mathbf{p}) = -1$  if  $h(\mathbf{p}) = 0$ . Obviously, the function values of  $g$  satisfy TUCKER’s antipodal boundary condition.

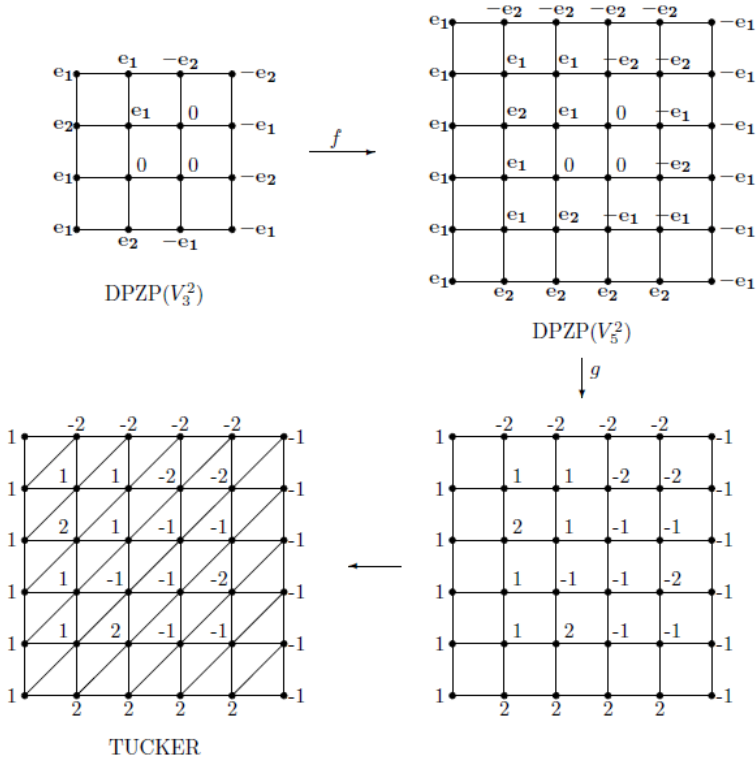
The last step is to triangulate each small hypercube by Kuhn’s triangulation. Since Kuhn’s triangulation will not add any extra vertices, the output of Kuhn’s triangulation combining with the function  $g$  is an instance of TUCKER  $(G, V_{N+2}^d)$ . By TUCKER’s lemma, it has a complementary edge. Because of the direction preserving property of the original function  $h$ , the complementary edge has to be  $(1, -1)$ . Again because of the direction preserving property, that node  $\mathbf{p}$  with value  $-1$  in the complementary  $(1, -1)$  must corresponding to  $h(\mathbf{p}) = 0$  in the original problem. The result follows. We illustrate the reduction process in Fig. 3.

**Theorem 2.**  $d$ -D TUCKER is **PPAD-complete**.

## 5 Matching Bound of TUCKER in Oracle Model

In this section, we derive the oracle matching bound for TUCKER.





**Fig. 3.** The Proof of Lemma 4

**Theorem 3.** (Lower Bound) For any instance of  $d$ -D TUCKER, finding a complementary 1-simplex takes time  $\Omega(N^{d-1})$ .

*Proof.* Using the same reduction as in PPAD-hardness proof of Tucker, we notice that if there is an algorithm that solves a problem of size  $N^d$  for TUCKER, it solves a problem of size  $N^d$  for DPZP. Since there is a lower bound of  $\Omega(N^{d-1})$  of DPZP [2], it results in a same lower bound  $\Omega(N^{d-1})$  for TUCKER.

The upper bound can be derived in two steps. First, we check if there is a complementary edge on the boundary. This takes time complexity  $O(N^{d-1})$ . If the answer is no, we prove that the number of  $(d - 1)$ -simplexes which are fully colored on the boundary must be odd. Then by using binary search based on the parity argument, one can find a complementary edge in time  $O(N^{d-1})$ . Hence, the total time complexity is  $O(N^{d-1})$ .

For  $d = 2$ , the following boundary lemma is easy to check.

**Lemma 5.** For any instance of 2D TUCKER, if there is no complementary edge on the boundary, then the number of  $\{1, 2\}$  edges on the boundary must be odd.

For higher dimensions, we employ the parity argument in Cohen’s proof of Tucker’s lemma [7] to prove the boundary lemma. A proof for the 3D case is presented here.

**Lemma 6.** *For any instance of 3-D TUCKER, if there is no complementary edge on the boundary, then there must exist one kind of  $\{a, b, c\}$ -simplex, where  $a, b, c$  are of distinct different absolute values, whose total number on the boundary is odd.*

*Proof.* For contradiction, we assume that the number of all  $\{a, b, c\}$ -simplexes on the boundary are even. Let  $A$  and  $A'$  be the center of two opposite faces. Let  $P$  be a path on the boundary from  $A$  to  $A'$ , and  $P'$  is the antipodal path of  $P$ . Then  $P$  and  $P'$  separate the boundary of the cube to two parts. Let us denote one part as  $S$  and the other part denoted as  $S'$ .

Let  $S(a, b, c)$  be the number of  $\{a, b, c\}$ -simplex on  $S$ , and  $S'(a, b, c)$  be the number of  $\{a, b, c\}$ -simplex on  $S'$ . According to the assumption, we have

$$S(a, b, c) + S'(a, b, c) = 0 \pmod 2$$

since the number of  $\{a, b, c\}$ -simplex on  $S'$  is equal to the number of  $\{-a, -b, -c\}$ -simplex on  $S$ , i.e.,

$$S'(a, b, c) = S(-a, -b, -c)$$

we get

$$S(a, b, c) + S(-a, -b, -c) = 0 \pmod 2$$

On the other hand, let  $P(a, b)$  denotes the number of  $\{a, b\}$  edges on  $P$ . As  $A$  and  $A'$  are of different signs,  $P$  must has odd number of  $\{a, b\}$  edges where  $a$  and  $b$  possess different signs. Since there is no complementary edge on the boundary, there are six possible this kind of edges, i.e.,  $\{1, -2\}, \{1, -3\}, \{2, -1\}, \{2, -3\}, \{3, -1\}$  and  $\{3, -2\}$ . So, we have

$$P(1, -2) + P(1, -3) + P(2, -1) + P(2, -3) + P(3, -1) + P(3, -2) = 1 \pmod 2$$

Now we going to derive a contradiction by parity argument. See Fig 4 for illustration.

Take  $S$  into consideration, and consider the 2-simplexes on the boundary starting from edges  $\{1, -2\}$  lying on  $P \cup P'$ , since there is no complementary edge on  $S$ , those  $\{1, -2\}$  edges will either reach another  $\{1, -2\}$  edge on  $P \cup P'$ , or will terminate with a  $\{1, -2, \pm 3\}$ -simplex on  $S$ . Starting from  $\{1, -2\}$  edges lying on  $S \setminus (P \cup P')$ , they will be pairised in two  $\{1, -2, \pm 3\}$  simplexes, or one in a  $\{1, -2, \pm 3\}$  simplex while another on  $P \cup P'$ . This gives a parity argument that

$$S(1, -2, 3) + S(1, -2, -3) = P(1, -2) + P'(1, -2) \pmod 2$$

which implies

$$S(1, -2, 3) + S(1, -2, -3) = P(1, -2) + P(-1, 2) \pmod 2$$




Fig. 4. The proof of Lemma 6

Similarly, we have

$$S(1, 2, -3) + S(-1, 2, -3) = P(2, -3) + P(-2, 3) \pmod 2$$

and

$$S(-1, 2, 3) + S(-1, -2, 3) = P(-1, 3) + P(1, -3) \pmod 2$$

Summing up these three equations we get a contradiction with the assumption, that gives a proof to the lemma.

We can transfer the kind of  $\{a, b, c\}$ -simplex whose number is odd to  $\{1, 2, 3\}$ -simplex as follows: for example, if it is  $\{1, -2, 3\}$ -simplex whose number is odd, then we exchange labels of 2 and -2, then the numbers of  $\{1, -2, 3\}$ -simplexes and  $\{1, 2, 3\}$ -simplexes are exchanged and this process does not affect the complementary edges in the instance.

**Proposition 1.** *If there is no complementary edge on the boundary, then the number of  $\{1, 2, \dots, d\}$ -simplexes on the boundary of  $d$ -dimensional hypergrid triangulated by Kuhn's triangulation has the same parity as the number of  $\{-i, 1, 2, \dots, d\}$ -simplexes,  $i = 1, 2, \dots, d$ , in the hypergrid.*

*Proof.* We change the labels of all nodes with negative labels to zero. Then the claim follows from standard theorems on indexes(see, e.g., [2]).

**Theorem 4.** (Upper Bound) *The upper bound of  $d$ -TUCKER is  $O(N^{d-1})$  under Kuhn's triangulation when the utility functions are given by oracle.*

*Proof.* First, by checking boundary we can see whether there is a kind of  $\{v_1, v_2, \dots, v_d\}$ -simplex whose number is odd on the boundary. If not, there must exist complementary edges on the boundary, finding it can be completed in time  $O(N^{d-1})$ . If we did not find any complementary edge on the boundary,

we relabel all nodes of negative labels by changing them to 0. Then the divide-&-conquer approach applies. We will end with a simplex of labels  $\{0, 1, \dots, d\}$ . As 0 was transferred from some negative labelled node, say one labelled with  $-i$ . Then  $(i, -i)$  in this simplex will be the complementary edge we set off to find.

Since the size of the hypercube decrease geometrically, the complexity of checking which case occurred will not exceed  $O(N^{d-1})$ .

**Theorem 5.** *For any  $d$ -hypercube which is an instance of TUCKER and it is triangulated by Kuhn's triangulation, the complementary 1-simplex can be found in time  $\theta(N^{d-1})$  under oracle model.*

## 6 Conclusion

Our study builds a computational complexity connection between two types of discrete versions of fixed point concepts, which implies two new results: a PPAD-completeness proof for computing direction preserving zero point, and a matching oracle complexity bound for discrete BROUWER's fixed point. Furthermore, the connection allows for a clear proof that TUCKER is PPAD-complete for all constant dimensions, that extend the results of Pálvölgyi for 2D case [20] and Papadimitriou for 3D case [21]. At the same time, a matching algorithmic bound for TUCKER is obtained in the oracle model.

As fixed point analysis is fundamental for many economic equilibrium problems, our study, though purely out curious minds for the fundamentals, could be of application values for other algorithmic game theory problems, in one way or another. We strongly believe that simplicity carries a value in its own right in theory building, and could play an important role in the development of a field. We demonstrate this value by presenting a simple and clear proof of complexities of TUCKER, both in the oracle function model and the polynomial time function model.

**Acknowledgement.** Work reported in this work is sponsored by a matching research grant from Comet Electronics (HK) Ltd and Research Grants Council of Hong Kong SAR (Proj. No. 9220046).

## References

1. Arrow, K., Debreu, G.: Existence of an equilibrium for a competitive economy. *Econometrica* 22, 265–290 (1954)
2. Chen, X., Deng, X.: On algorithms for discrete and approximate brouwer fixed points. In: STOC, pp. 323–330 (2005); Journal version appeared at: Chen, X., Deng, X.: Matching algorithmic bounds for finding a Brouwer fixed point. *J. ACM* 55(3) (2008)
3. Chen, X., Deng, X.: 3-NASH is PPAD-Complete. *Electronic Colloquium on Computational Complexity (ECCC)* (134) (2005)
4. Chen, X., Deng, X.: On the Complexity of 2D Discrete Fixed Point Problem. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 489–500. Springer, Heidelberg (2006)

5. Chen, X., Deng, X.: Settling the Complexity of Two-Player Nash Equilibrium. In: FOCS, pp. 261–272 (2006)
6. Chen, X., Deng, X., Teng, S.-H.: Computing Nash Equilibria: Approximation and Smoothed Complexity. In: FOCS, pp. 603–612 (2006)
7. Cohen, D.I.A.: On the Combinatorial Antipodal-Point Lemmas. *Journal of Combinatorial Theory, Series B* (27), 87–91 (1979)
8. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. In: STOC (2006)
9. Daskalakis, C., Papadimitriou, C.H.: Three-Player Games Are Hard. *Electronic Colloquium on Computational Complexity (ECCC)* (139) (2005)
10. Deng, X., Qi, Q., Saberi, A.: On the Complexity of Envy-Free Cake Cutting, <http://arxiv.org/abs/0907.1334>
11. Freund, R.M.: Variable dimension complexes. Part I: Basic theory. *Math. Oper. Res.* 9(4), 479–497 (1984)
12. Freund, R.M., Todd, M.J.: A constructive proof of Tucker’s combinatorial lemma. *Journal of Combinatorial Theory, Series A* 30(3), 321–325 (1981)
13. Iimura, T.: A discrete fixed point theorem and its applications. *Journal of Mathematical Economics* 39, 725–742 (2003)
14. Iimura, T., Murota, K., Tamura, A.: Discrete fixed point theorem reconsidered. *METR* 9 (2004)
15. Kuhn, K.W.: Some combinatorial lemmas in topology. *IBM J. Res. Dev* 4(5), 518–524 (1960)
16. Longueville, M., Živaljević, R.T.: The Borsuk-Ulam-property, Tucker-property and constructive proofs in combinatorics. *Journal of Combinatorial Theory, Series A* (113), 839–850 (2006)
17. Matoušek, J.: *Using the Borsuk-Ulam Theorem*. Springer, Heidelberg (2003)
18. Matoušek, J.: A combinatorial proof of Kneser’s conjecture. *Combinatorica* 24, 163–170 (2004)
19. Megiddo, N., Papadimitriou, C.H.: On Total Functions, Existence Theorems and Computational Complexity. *Theor. Comput. Sci.* 81(2), 317–324 (1991)
20. Pálvolgyi, D.: 2D-TUCKER is PPAD-complete (Preprint in ECCC)
21. Papadimitriou, C.H.: On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 498–532 (1994)
22. Prescott, T., Su, F.E.: A constructive proof of Ky Fan’s generalization of Tucker’s lemma. *Journal of Combinatorial Theory, Series A* (111), 257–265 (2005)
23. Scarf, H.E.: The approximation of fixed points of a continuous mapping. *SIAM J. Applied Mathematics* 15, 997–1007 (1967)
24. Simmons, F.W., Su, F.E.: Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Mathematical Social Science* (45), 15–25 (2003)
25. Su, F.E.: Rental Harmony: Sperner’s Lemma in Fair Division. *Amer. Math. Monthly* (106), 930–942 (1999)
26. Tucker, A.W.: Some topological properties of disk and sphere. In: *Proceedings of the First Canadian Mathematical Congress*, pp. 285–309 (1945)
27. Ziegler, G.M.: *Lectures on Polytopes*. Graduate Texts in Mathematics, vol. 152. Springer, New York (1995)

# Continuity Properties of Equilibria in Some Fisher and Arrow-Debreu Market Models

Vijay V. Vazirani\* and Lei Wang

College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280

**Abstract.** Following up on the work of Megiddo and Vazirani [10], who determined continuity properties of equilibrium prices and allocations for perhaps the simplest market model, Fisher’s linear case, we do the same for:

- Fisher’s model with piecewise-linear, concave utilities
- Fisher’s model with spending constraint utilities
- Arrow-Debreu’s model with linear utilities
- Eisenberg-Gale markets.

## 1 Introduction

Three basic properties that a desirable model of an economy should possess are existence, uniqueness, and continuity of equilibria (see [3], Chapter 15, “Smooth preferences”). These lead to parity between supply and demand, stability and predictive value, respectively. In particular, without continuity, small errors in the observation of parameters of an economy may lead to entirely different predicted equilibria.

Although mathematical economists studied very extensively questions of existence and uniqueness for several concrete and realistic market models, the question of continuity was studied only in very abstract settings. Megiddo and Vazirani [10] attempted to rectify this situation by starting with perhaps the simplest market model – the linear case of Fisher’s model. They showed that the mapping giving the unique vector of equilibrium prices is continuous and the correspondence giving the set of equilibrium allocations, is upper hemicontinuous. In this paper, we determine continuity properties of equilibrium prices, allocation and utilities for the models given in the Abstract; our results are in the following table.

	Equilibrium Price	Equilibrium Allocation	Equilibrium Utility
Fisher-PL	No	No	No
Fisher-SC	Yes	Yes	Yes
AD-L	No	Yes	Yes
EG	Yes	not defined	Yes

---

\* Supported by NSF Grant CCF-0728640.

[10] crucially used the Eisenberg-Gale convex program [5] for proving their results, hence their proofs were steeped in polyhedral combinatorics. For the Arrow-Debreu model with linear utilities, we use a convex program due to Jain [6]. For the remaining three cases, such convex programs are not known, and instead, we use the combinatorial structure of equilibria for proving our theorems. The groundwork for discovering such structure was done in [4] in the context of obtaining a polynomial time algorithm for computing the equilibrium the linear case of Fisher’s model. Each of the remaining three cases is a generalization of this case and the relevant structure is also a generalization of that for the Fisher’s linear case. For the first case in the table, we use structure found by [12] (and used for proving that equilibrium prices and allocations are rational numbers). For the second case, we use structure found by [11] (and used for obtaining a polynomial time algorithm for computing the equilibrium). We refer the reader to [4,6,11,12,7] for definitions of the market models studied.

## 2 Eisenberg-Gale Markets

Eisenberg-Gale market (EG, in brief) is an abstract market model. The equilibrium utilities can be captured as the optimal solution of the following Eisenberg-Gale convex program [5]:

$$\begin{aligned} \max \quad & \sum_{i=1}^n m_i \log u_i \\ \text{s.t.} \quad & \sum_{i=1}^n a_{ij} u_i \leq b_j, \quad \forall j \in [n'] \\ & u_i \geq 0, \quad \forall i \in [n] \end{aligned}$$

Here,  $m_i, a_{ij}, b_j$  are constants.

By the KKT conditions,  $\mathbf{u} = (u_i)_{i \in [n]}$  is the optimal solution iff there exists some  $\mathbf{p} = (p_j)_{j \in [n']}$  such that the following hold: (1) For  $\forall j \in [n']$ ,  $p_j \geq 0$  and  $p_j > 0 \Rightarrow \sum_{i=1}^n a_{ij} u_i = b_j$ ; (2) For  $\forall i \in [n]$ ,  $\frac{m_i}{u_i} = \sum_{j=1}^{n'} a_{ij} p_j$ . We call an optimal dual solution  $\mathbf{p} = (p_1, \dots, p_{n'})$  an equilibrium price.

We introduce a matrix  $A = (a_{ij})$  and vectors  $\mathbf{b} = (b_j)$ ,  $\mathbf{m} = (m_i)$  to specify the coefficients in the Eisenberg-Gale convex program.

**Theorem 1.** *The equilibrium utility  $\mathbf{u}(\mathbf{m}, A, \mathbf{b})$  in the given EG market is a continuous function.*

*Proof.* Assume that  $\mathbf{m}^k, A^k, \mathbf{b}^k$  converge to  $\mathbf{m}^0, A^0, \mathbf{b}^0$  respectively. For any  $k \geq 0$ , let  $\mathcal{P}^k$  be the feasible polytope corresponding to  $A^k, \mathbf{b}^k$  and let  $\mathbf{u}^k$  be the corresponding optimal solution. We may assume  $\{\mathbf{u}^k\}$  is convergent. Suppose  $\mathbf{u}^k \rightarrow \mathbf{u}'$ , we want to show  $\mathbf{u}' = \mathbf{u}^0$ .

Let  $f(\mathbf{m}, \mathbf{u}) = \sum m_i \log u_i$  be the objective function of EG convex program. Since it is strictly convex, an EG convex program has a unique optimal solution.

By our assumption,  $\mathbf{u}^0$  is the optimal solution corresponding to  $\mathbf{m}^0, A^0, \mathbf{b}^0$ , hence if we can show: (i)  $\mathbf{u}'$  is feasible, i.e.  $\mathbf{u}' \in \mathcal{P}^0$  and (ii)  $f(\mathbf{m}^0, \mathbf{u}') \geq f(\mathbf{m}^0, \mathbf{u}^0)$ , then we are done.

In the above, (i) can be shown easily: From the condition  $A^k \cdot \mathbf{u}^k \leq \mathbf{b}^k, \mathbf{u}^k \geq 0$ , taking the limit, we have  $A^0 \cdot \mathbf{u}' \leq \mathbf{b}^0, \mathbf{u}' \geq 0$ , i.e.  $\mathbf{u}' \in \mathcal{P}^0$ .

Now we prove (ii). If for some  $i, u_i^0 = 0$ , then we have  $f(\mathbf{m}^0, \mathbf{u}^0) = -\infty$  and (ii) holds trivially. So we may assume that  $u_i^0 \geq 2\delta_1, \forall i \in [n]$ , for some constant  $\delta_1 > 0$ . For each  $k$ , let  $dist(\mathbf{u}^0, \mathcal{P}^k)$  be the distance between  $\mathbf{u}^0$  and  $\mathcal{P}^k$ . Let  $\mathbf{w}^k$  be a point in  $\mathcal{P}^k$  which is closest to  $\mathbf{u}^0$ . Since  $\mathcal{P}^k \rightarrow \mathcal{P}^0$  and  $\mathbf{u}^0 \in \mathcal{P}^0$ , we know that  $\mathbf{w}^k \rightarrow \mathbf{u}^0$ . Now, we can assume that  $\{\mathbf{w}^k\}$  and  $\mathbf{u}^0$  are in a closed region  $\mathcal{S}$  without the origin. (Indeed, we can take  $\mathcal{S}$  to be  $\{x \mid \delta_1 \leq \|x\| \leq \|u^0\| + \delta_1\}$ )

By considering the derivative of  $f$ , it's easy to see,  $f$  is *Lipchitz continuous* with respect to the second coordinate in a region without the origin, i.e. there exists a constant  $L > 0$  such that:  $f(\mathbf{m}^k, \mathbf{u}^0) - f(\mathbf{m}^k, \mathbf{w}^k) \leq L\|\mathbf{u}^0 - \mathbf{w}^k\| = L \cdot dist(\mathbf{u}^0, \mathcal{P}^k)$

Thus  $f(\mathbf{m}^k, \mathbf{u}^0) \leq f(\mathbf{m}^k, \mathbf{w}^k) + L \cdot dist(\mathbf{u}^0, \mathcal{P}^k) \leq f(\mathbf{m}^k, \mathbf{u}^k) + L \cdot dist(\mathbf{u}^0, \mathcal{P}^k)$  where the last inequality holds because  $\mathbf{u}^k$  maximizes  $f(\mathbf{m}^k, \mathbf{u})$  in  $\mathcal{P}^k$  and  $\mathbf{w}^k \in \mathcal{P}^k$ .

Let  $k \rightarrow +\infty$ , we have  $f(\mathbf{m}^0, \mathbf{u}^0) \leq f(\mathbf{m}^0, \mathbf{u}')$ .

KKT condition gives us a connection between equilibrium utility and equilibrium price. Since we have shown that the equilibrium utility is continuous, by KKT, we can show:

**Theorem 2.** *The equilibrium price  $\mathbf{p}(\mathbf{m}, A, \mathbf{b})$  in the given EG market is a upper hemicontinuous set-valued function.*

### 3 Arrow-Debreu Market with Linear Utilities

In this model, agent  $i$ 's utility as a function of allocation  $x$  is linear:  $U_i(x) = \sum_j u_{ij} x_{ij}$ . Given the utility functions, we write the  $u_{ij}$ 's in a matrix  $U = (u_{ij})$ . We use  $X(U)$  and  $T(U)$  to denote the equilibrium allocation and utility. Note that a multiple of an equilibrium price in **AD-L** is again an equilibrium price, so we adopt the notion of *normalized equilibrium price* which is defined to be an equilibrium price whose  $l_1$  norm is 1. We use  $P(U)$  to denote the set-valued function of normalized equilibrium prices.

We start by giving a convex program due to Jain [6] that captures equilibrium allocations for this model. For any given utilities, construct a directed graph  $G$  with  $n$  vertices to represent the  $n$  agents. Draw an edge from  $i$  to  $j$  if  $u_{ij} > 0$ . Define  $w(i, j) = \frac{\sum_{1 \leq t \leq n} u_{it} x_{it}}{u_{ij}}$ . Then consider the following convex program:

$$\forall j : \sum_i x_{ij} = 1,$$

$$\forall i, j : x_{ij} \geq 0,$$

$$\text{For every cycle } C \text{ of } G : \prod_{(i,j) \in C} w(i, j) \geq 1.$$



**Theorem 3.** ([6]) *An allocation  $\mathbf{x}$  is an equilibrium allocation if and only if  $\mathbf{x}$  is a feasible solution of the above convex program.*

Now we prove that equilibrium allocations and utilities are upper hemicontinuous set-valued functions. The idea is to show that a small perturbation won't affect the feasible region of the convex program much.

**Theorem 4.** *In AD-L, equilibrium allocation and utility are upper hemicontinuous.*

*Proof.* Suppose  $\{U^k\}$  is a sequence of utilities and  $U^k \rightarrow U^0$ . Suppose  $\{\mathbf{x}^k\}$  is a sequence of allocations such that  $\mathbf{x}^k \in X(U^k)$  and  $\mathbf{x}^k \rightarrow \mathbf{x}^0$ . We want to show  $\mathbf{x}^0 \in X(U^0)$ .

Because  $\mathbf{x}^k \in X(U^k)$ , by Theorem 3, for every  $k$ ,  $\mathbf{x}^k$  is a feasible solution to the convex program. Let  $G^k$  be the graph formed by all the edges  $(i, j)$  where  $u_{ij}^k > 0$ . Then we have:

$$\forall j : \sum_i x_{ij}^k = 1, \tag{1}$$

$$\forall i, j : x_{ij}^k \geq 0, \tag{2}$$

$$\text{For every cycle } C \text{ of } G^k : \prod_{(i,j) \in C} w^k(i, j) \geq 1. \tag{3}$$

Now consider the graph  $G^0$ . Because  $U^k \rightarrow U^0$ , when  $k$  is large enough, for every  $i, j$  such that  $u_{ij}^0 > 0$ , we have  $u_{ij}^k > 0$ . Therefore if an edge  $(i, j)$  is presented in  $G^0$ , it is also presented in  $G^k$ . This implies if a cycle  $C$  is presented in  $G^0$ , it is also presented in  $G^k$  when  $k$  is large enough. Therefore  $\mathbf{x}^k$  satisfies a weaker condition than (3): *For every cycle  $C$  of  $G^0 : \prod_{(i,j) \in C} w^k(i, j) \geq 1$ .*

Now in the above inequalities, let  $k \rightarrow \infty$ , we have  $\mathbf{x}^0 \in X(U^0)$ , hence  $X(U)$  is upper hemicontinuous.

Since the equilibrium utility is a linear function of the equilibrium allocation, we have that  $T(U)$  is also upper hemicontinuous.

At last, we show that the equilibrium price is not upper hemicontinuous by the following example: There are two agents, each of them has one good. Suppose  $\{U^k\}_{k=1}^\infty$  is a sequence of utilities where  $U^k = \begin{pmatrix} 1 - \frac{1}{k} & 0 \\ 0 & 1 - \frac{1}{k} \end{pmatrix}$ . Let  $p^k = (\frac{1}{k}, 1 - \frac{1}{k})$ . It's easy to see that  $p^k \in P(U^k)$  and  $X^k = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$  is the corresponding equilibrium allocation.

Since  $U^k \rightarrow U^0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$  but  $p^k \rightarrow (0, 1) \notin P(U^0)$ , we conclude that  $P(U)$  is not upper hemicontinuous.

## 4 Fisher’s Market Model

### 4.1 Linear Utilities

In **Fisher-L** model, for each  $i, j$ ,  $U_{ij}(x_{ij}) = u_{ij}x_{ij}$  is a linear function. Each buyer  $i$  has an initial endowment of money  $e_i$ . The amount of each good is assumed to be 1 unit. We use  $\mathbf{e}$  to denote  $(e_1, \dots, e_n)$  and  $\mathbf{U}$  to denote the matrix  $(u_{ij})$ .

Given a price  $\mathbf{p}$ , we can decide whether  $\mathbf{p}$  is an equilibrium price by looking at the following network  $N(\mathbf{e}, \mathbf{U}, \mathbf{p})$  defined in [4]: Consider a bipartite graph with bipartition  $B$  and  $G$ , vertices in  $B$  and  $G$  represent buyers and goods respectively. Under the price  $\mathbf{p}$ , compute the optimal bang-per-buck  $\alpha_i$  of buyer  $i$ . Connect good  $j$  to buyer  $i$  if good  $j$  is desired by buyer  $i$ , i.e.  $u_{ij}/p_j = \alpha_i$ . Set the capacity of the edge to be  $\infty$ . Then add a source  $s$  and a sink  $t$  into the graph. Connect the source to each good and connect each buyer to the sink. For each  $j \in G$ , set the capacity of the edge  $(s, j)$  to be  $p_j$  and for each buyer  $i \in B$ , set the capacity of the edge  $(i, t)$  to be  $e_i$ .

The following theorem reveals the connection between the network and the market equilibrium:

**Theorem 5.** ([4])  $\mathbf{p}$  is the equilibrium price if and only if both  $(s, G \cup B \cup t)$  and  $(s \cup G \cup B, t)$  are min cuts in  $N(\mathbf{e}, \mathbf{U}, \mathbf{p})$ . Moreover, when both cuts are min cuts, let  $\mathbf{f}$  be a maximum flow, then the equilibrium allocation is given by  $x_{ij} = f_{(j,i)}/p_j$  for any directed edge  $(j, i)$ .

There is a simple way to test whether  $(s, G \cup B \cup t)$  is a min cut or not. For any subset  $S$  of goods, let  $\Gamma(S)$  be the neighbor of  $S$ , which consists of the buyers who desire some goods in  $S$ , i.e.  $\Gamma(S) = \{i \in B : \exists j \in S, s.t. (j, i) \in N(\mathbf{e}, \mathbf{U}, \mathbf{p})\}$ . The set  $S$  is called *tight*, if  $\sum_{j \in S} p_j = \sum_{i \in \Gamma(S)} e_i$ . It is called *overtight* if  $\sum_{j \in S} p_j > \sum_{i \in \Gamma(S)} e_i$ .

**Theorem 6.** ([4])  $(s, G \cup B \cup t)$  is a min cut if and only if no subset of  $G$  is *overtight*.

To show continuity of the equilibrium price, we only need to show that for any  $\epsilon > 0$ , when a small enough perturbation is made to  $\mathbf{U}$  and  $\mathbf{e}$ , the change in the equilibrium price can be bounded by  $\epsilon$ .

We can view a perturbation in the following way: the perturbation occurs on  $u_{ij}$ ’s and  $e_i$ ’s one by one. Therefore, it suffices to prove the following two lemmas:

**Lemma 1.** For any  $\epsilon > 0$  and for any buyer  $k$ , if  $|e_k - e'_k|$  is small enough, then  $|\mathbf{p} - \mathbf{p}'| < \epsilon$  where  $\mathbf{p}$  and  $\mathbf{p}'$  are the equilibrium prices corresponding to  $e_k$  and  $e'_k$  respectively.

**Lemma 2.** For any  $\epsilon > 0$  and for any buyer  $k$  and good  $l$ , if  $|u_{kl} - u'_{kl}|$  is small enough, then  $|\mathbf{p} - \mathbf{p}'| < \epsilon$  where  $\mathbf{p}$  and  $\mathbf{p}'$  are the equilibrium prices corresponding to  $u_{kl}$  and  $u'_{kl}$  respectively.

In the above,  $\|\mathbf{p} - \mathbf{p}'\|$  denote the  $l_1$  norm of the vector  $\mathbf{p} - \mathbf{p}'$ .

*Proof (Proof of Lemma 7).* Consider buyer  $k$ . Suppose her money has a small perturbation  $\delta$ , i.e.  $e'_k = e_k(1 + \delta)$ . We assume  $\delta > 0$ . In the case when  $\delta < 0$ , similar argument applies. Let  $\mathbf{e}'$  be the money vector  $(e_1, \dots, e_{k-1}, e'_k, e_{k+1}, \dots, e_n)$ .

Consider the network  $N'$  by increasing the capacity of the edge  $(k, t)$  to  $e'_k$ . In  $N'$ ,  $(s \cup G \cup B, t)$  is not a min cut. Let  $C$  be a *maximal min cut* in  $N'$ , which is defined to be the min cut with maximum number of vertices in the  $s$ -side. This min cut has the form  $(s \cup G_2 \cup B_2, G_1 \cup B_1 \cup t)$ , where  $B_1, B_2 \subseteq B$  and  $G_1, G_2 \subseteq G$ . The following are some observations: (1)  $k$  must be in  $B_1$ ; (2) there is no edge from  $G_2$  to  $B_1$ ; (3) we can drop edges from  $G_1$  to  $B_2$ ; (4) no subset of  $G_1$  is tight.

Now for all  $j \in G_1$ , define  $p'_j = p_j(1 + \epsilon')$  where  $\epsilon'$  satisfies the equation  $e_k \delta = \epsilon' \sum_{j \in G_1} p_j$ . For all  $j \in G_2$ , let  $p'_j = p_j$ . By this, we get a new price vector  $\mathbf{p}'$ . By our choice of  $\epsilon'$ , we have that in the new network, both  $(s, G \cup B \cup t)$  and  $(s \cup G \cup B, t)$  are min cuts, hence  $\mathbf{p}'$  is the equilibrium price corresponds to  $\mathbf{e}'$ . Thus for any  $\epsilon > 0$ , choose  $\delta$  small enough, we can make sure  $\|\mathbf{p}' - \mathbf{p}\| < \epsilon$ .

*Proof (Proof of Lemma 8).* Suppose for buyer  $k$  and good  $l$ ,  $u'_{kl} = u_{kl}(1 + \delta)$ . Again, we assume  $\delta > 0$  and the case when  $\delta < 0$  is similar. Let  $\mathbf{p}$  be the equilibrium price before the perturbation. Let  $\mathbf{U}'$  be the new utility matrix. We may assume good  $l$  is desired by buyer  $k$  under price  $\mathbf{p}$ .

In this case, when we increase  $u_{kl}$  to  $u'_{kl}$ , under the price  $\mathbf{p}$ , buyer  $k$  desires only good  $l$ . We construct a network  $N'$  by dropping all the edges of the form  $(j, k), j \neq l$  from  $N(\mathbf{e}, \mathbf{U}, \mathbf{p})$ . Let  $C = (s \cup G_2 \cup B_2, G_1 \cup B_1 \cup t)$  be a maximum min cut in  $N'$ . All the four observations in the proof of claim 4 hold in this case. In addition, good  $l$  must be in  $G_1$  and any good  $j$  such that  $(j, k)$  is presented in  $N(\mathbf{e}, \mathbf{U}, \mathbf{p})$  must be in  $G_2$ .

We construct a new price  $\mathbf{p}'$  in the following way: For each  $j \in G_1$ , let  $p'_j = p_j(1 + \epsilon_1)$ ; For each  $j \in G_2$ , let  $p'_j = p_j(1 - \epsilon_2)$ . Here  $\epsilon_1, \epsilon_2$  are solutions to the system

$$\epsilon_1 \sum_{j \in G_1} p_j = \epsilon_2 \sum_{j \in G_2} p_j$$

$$\frac{1 + \delta}{1 + \epsilon_1} = \frac{1}{1 - \epsilon_2}.$$

By this, in the new network, all the lost edges come back and when  $\delta$  is small enough, no subset of  $G$  is overtight, hence  $\mathbf{p}'$  is the new equilibrium price. At last, for any  $\epsilon > 0$ , when  $\delta$  is small enough, we can make sure  $\|\mathbf{p}' - \mathbf{p}\| < \epsilon$ .

Hence, we have given an alternative proof of the following:

**Theorem 7.** ([10]) *In Fisher-L, the equilibrium price is continuous.*

### 4.2 Spending Constraint and Piecewise-Linear, Concave Utilities

For the spending constraint model, a similar network (see [11]) as the one given in the last section can be set up to decide whether a given price is an equilibrium price or not, therefore our combinatorial proof in the last section generalizes naturally to **Fisher-SC**. Since the idea is similar, we omit the proof in this extended abstract.

**Theorem 8.** *In Fisher-SC, the equilibrium price is continuous, the equilibrium allocation and utility are upper hemicontinuous.*


For piecewise linear utility model, even though it is a natural market model, surprisingly, we have examples to show the following negative results:

**Theorem 9.** *In Fisher-PL, the equilibrium price, allocation and utility are not upper hemicontinuous.*

*Proof.* The following example shows that the equilibrium price, which is a set-valued function, is not upper hemicontinuous. Suppose in the market, there is only one buyer with money  $m$  and there are two goods, each of 1 unit. Suppose  $U_1, U_2$  are her utility functions for good 1 and good 2:  $U_1(x) = x$  if  $0 \leq x \leq 1$  and  $U_1(x) = 1$  for all  $x > 1$ ;  $U_2(x) = x$  for all  $x$ .

Suppose the price for good 1 is  $p_1$  and the price for good 2 is  $p_2$ . Then  $p = (p_1, p_2)$  is an equilibrium price if and only if  $p_1 > 0, p_2 > 0, \frac{1}{p_1} \geq \frac{1}{p_2}$  and  $p_1 + p_2 = m$ . Now we take  $p^k = (\frac{1}{k}, m - \frac{1}{k})$ . For each  $k, p^k$  is an equilibrium price. However,  $p^k \rightarrow (0, m)$  which is not an equilibrium price. Thus the set of equilibrium prices is not upper hemicontinuous.

For equilibrium allocation and utility, consider the following example. There are two buyers and three goods in the market. Suppose we have  $\{(m^k, U^k)\}_{k \geq 0}$  where  $m^k = (1, 1 + \frac{1}{k})$  and the picture of  $U^k$  is below, The slopes of the line segments are indicated in the picture.



It is easy to see  $p^k = (\frac{1}{k}, 1, 1)$  and

$$x^k = \begin{pmatrix} 1 - \frac{1}{2} - \frac{1}{k} & \frac{1}{2} \\ 0 & \frac{1}{2} + \frac{1}{k} \end{pmatrix}$$

are equilibrium price and allocation. However,

$$x^k \rightarrow x^0 = \begin{pmatrix} 1 - \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{1}{2} \end{pmatrix},$$

which is not an equilibrium allocation for  $(m^0, U^0)$ . Therefore the equilibrium allocation is not upper hemicontinuous. Similar for equilibrium utility.

**Acknowledgement.** We would like to thank Gagan Goel for many valuable suggestions.

## References

1. Arrow, K.J., Debreu, G.: Existence of an equilibrium for a 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 (2000)
3. Debreu, G.: *Mathematical Economics: Twenty papers of Gerard Debreu*. Cambridge University Press, Cambridge (1986)
4. Devanur, N., Papadimitriou, C.H., Saberi, A., Vazirani, V.V.: Market equilibrium via a primal-dual-type algorithm. *Journal of ACM* 55(5) (2008)
5. Eisenberg, E., Gale, D.: Consensus of subjective probabilities: the Pari-Mutuel method. *The Annals of Mathematical Statistics* 30, 165–168 (1959)
6. Jain, K.: A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. *SIAM J. Comput.* 37(1), 303–318 (2007)
7. Jain, K., Vazirani, V.: Eisenberg-Gale Markets: Algorithms and Game-Theoretic Properties. In: *Proceedings of 38th IEEE Annual Symposium on Theory of Computing* (2006)
8. Kelly, F.: Charging and rate control for elastic traffic. *European Transactions on Telecommunication* 8, 33–37 (1997)
9. Kelly, F., Vazirani, V.: Rate control as a market equilibrium. Unpublished note (2002)
10. Megiddo, N., Vazirani, V.V.: Continuity properties of equilibrium prices and allocations in linear fisher markets. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 362–367. Springer, Heidelberg (2007)
11. Vazirani, V.V.: Spending constraint utilities, with applications to the Adwords market (2006) (submitted)
12. Vazirani, V.V., Yannakakis, M.: Market Equilibrium under Separable, Piecewise-Linear, Concave Utilities (submitted, 2009)

# Route Distribution Incentives

Joud Khoury<sup>1</sup>, Chaouki T. Abdallah<sup>1</sup>, Kate Krause<sup>2</sup>, and Jorge Crichigno<sup>1</sup>

<sup>1</sup> ECE Department, MSC01 1100,

1 University of New Mexico, Albuquerque NM 87131  
{jkhoury, chaouki, jcrichigno}@ece.unm.edu

<sup>2</sup> Economics Department, University of New Mexico  
1915 Roma NE/Economics Bldg., Albuquerque NM 87131  
kkrause@unm.edu

**Abstract.** We present an incentive model for route distribution in the context of path vector routing protocols and focus on the Border Gateway Protocol (BGP). We model BGP route distribution and computation using a game in which a BGP speaker advertises its prefix to its direct neighbors promising them a reward for further distributing the route deeper into the network. The result of this cascaded route distribution is an advertised prefix and hence reachability of the BGP speaker. We first study the convergence of BGP protocol dynamics to a unique outcome tree in the defined game. We then study the existence of equilibria in the *full information* game considering competition dynamics focusing on the simplest two classes of graphs: 1) the line (and the tree) graphs which involve no competition, and 2) the ring graph which involves competition.

## 1 Introduction

The Border Gateway Protocol (BGP) [12] is a policy-based path vector protocol and is the de-facto protocol for Internet interdomain routing. BGP is intrinsically about distributing route information about destinations (which are IP prefixes) to establish paths in the network. Path discovery, or simply discovery hereafter, starting with some destination prefix is the outcome of route distribution and route computation. Accounting for and sharing the cost of discovery is an interesting problem and its absence from current path discovery schemes has led to critical economic and scalability concerns. As an example, the BGP control plane functionality is oblivious to cost. A node (BGP speaker) that advertises a provider-independent prefix (identifier) does not pay for the cost of being discoverable. Such a cost, which may be large given that the prefix is maintained at every node in the Default Free Zone (DFZ), is paid by the rest of the network. For example, Herrin [5] has preliminarily analyzed the non-trivial cost of maintaining a BGP route. Such incentive mismatch in the current BGP workings is further exacerbated by provider-independent addressing, multi-homing, and traffic engineering practices [11]. The fact that the number of BGP prefixes in the global routing table (or RIB) is constantly increasing at a rate of roughly 100,000 entries every 2 years and is expected to reach a total of 388,000 entries in 2011 [6], has motivated us to devise a model that accounts for distribution incentives in BGP.

A large body of work has focused on choosing the right incentives given that Autonomous Systems (AS) are self-interested, utility-maximizing agents. Most previous

work has ignored the control plane incentives [1] (route advertisement) and has instead focused on the forwarding plane incentives (e.g. transit costs). One possible explanation is based on the fact that a node has an incentive to distribute routes to destinations since the node will get paid for transiting traffic to these destinations, and hence route distribution is ignored as it becomes an artifact of the transit process. We argue that this assumption is not economically viable by considering the arrival of a new customer (BGP speaker). While the servicing edge provider makes money from transiting the new customer’s traffic to the customer, the middle providers do not necessarily make money while still incurring the cost to maintain and distribute the customer’s route information. In this work, we separate the control plane incentives (incentives to distribute route information) from the forwarding plane incentives (incentives to forward packets) and use game theory to model a BGP distribution game. The main problem we are interested in is how to allow BGP prefix information to be distributed globally while aligning the incentives of all the participating agents.

**Model and Results.** We synthesize many of the ideas and results from [2,4,8,9] into a coherent model for studying BGP route distribution incentives. A destination  $d$  is willing to invest some initial amount of money  $r_d$  to get its route information to be globally distributed. Since  $d$  may only advertise its prefix to its direct neighbors,  $d$  must incentivize them to further distribute the route. The neighbors then incentivize their neighbors, and so on. While this work takes BGP as the motivating application, we are interested in the general setting of distributing a good to a set of agents. In this paper, we define a *BGP distribution game* by building upon the general model for studying BGP devised by Griffin et. al in [4]. We assume *full information* since our main goal is to study the existence of equilibria rather than how to reach the equilibrium. Studying the equilibria for arbitrary graph structures is difficult given the complexity of the strategic dependencies and the competition dynamics. Since we are not aware of general existence results that apply to our game, we initially focus on two simple graphs: 1) the line (and the tree) graphs which involve no competition, and 2) the ring graph which involves competition. Our results are detailed in section 3 and more fully in [7].

**Related work.** The Simple Path Vector Protocol (SPVP) formalism [4] develops sufficient conditions for the outcome of a path vector protocol to be stable. A respective game-theoretic model was developed by Levin [9] to capture these conditions and incentives in a game theoretic setting. Feigenbaum et. al study incentive issues in BGP by considering least cost path (LCP) policies [1] and more general policies [2]. Our model is fundamentally different from [1] (and other works based in mechanism design) in that the prices are strategic, the incentive structure is different, and we do not assume the existence of a central “designer” (or bank) that allocates payments to the players but is rather completely distributed as in real markets. The bank assumption is limiting in a distributed setting, and an important question posed in [2] is whether the bank can be replaced by direct payments by the nodes. Li et. al [10] study an incentive model for query relaying in peer-to-peer (p2p) networks based on rewards, upon which Kleinberg

<sup>1</sup> In this paper, we use the term “control plan” to refer only to route prefix advertisements (not route updates) as we assume that the network structure is static.

et. al [8] build to model a more general class of trees. In [8], Kleinberg and Raghavan allude to a similar version of our distribution game in the context of query incentive networks. They pose the general question of whether an equilibrium exists for general Directed Acyclic Graphs (DAGs) in the query propagation game. Both of these probabilistic models do not account for competition. While we borrow the basic idea, we address the different problem of route distribution rather than information seeking.

## 2 The General Game

Borrowing notation from [29], we consider a graph  $G = (V, E)$  where  $V$  is a set of  $n$  nodes (alternatively termed players, or agents) each identified by a unique index  $i = \{1, \dots, n\}$ , and a destination  $d$ , and  $E$  is the set of edges or links. Without loss of generality (WLOG), we study the BGP discovery/route distribution problem for some fixed destination AS with prefix  $d$  (as in [429]). The model is extendable to all possible destinations (BGP speakers) by noticing that route distribution and computation are performed independently per prefix. The destination  $d$  is referred to as the *advertiser* and the set of players in the network are termed *seekers*. Seekers may be distributors who participate in distributing  $d$ 's route information to other seeker nodes or consumers who simply consume the route. For each seeker node  $j$ , Let  $P(j)$  be the set of all routes to  $d$  that are known to  $j$  through advertisements,  $P(j) \subseteq \mathcal{P}(j)$ , the latter being the set of all simple routes from  $j$ . The empty route  $\phi \in \mathcal{P}(j)$ . Denote by  $R_j \in P(j)$  a simple route from  $j$  to the destination  $d$  with  $R_j = \phi$  when no route exists at  $j$ , and let  $(k, j)R_j$  be the route formed by concatenating link  $(k, j)$  with  $R_j$ , where  $(k, j) \in E$ . Denote by  $B(i)$  the set of direct neighbors of node  $i$  and let  $next(R_i)$  be the next hop node on the route  $R_i$  from  $i$  to  $d$ . Finally, define node  $j$  to be an *upstream* node relative to node  $i$  when  $j \in R_i$ . The opposite holds for a *downstream* node. The general distribution game is as follows: destination  $d$  first exports its prefix (identifier) information to its neighbors promising them a reward  $r_d \in \mathbb{Z}^+$  which directly depends on  $d$ 's utility of being discoverable. A node  $i$ , a player, in turn receives offers from its neighbors where each neighbor  $j$ 's offer takes the form of a reward  $r_{ji}$ . We use  $r_{next(R_i)}$  to refer to the reward that the upstream parent from  $i$  on  $R_i$  offers to  $i$ .

**Strategy Space:** Given a set of advertised routes  $P(i)$  where each route  $R_i \in P(i)$  is associated with a promised reward  $r_{next(R_i)} \in \mathbb{Z}^+$ , a *pure strategy*  $s_i \in S_i$  of an autonomous node  $i$  comprises two decisions:

First, after receiving offers from neighboring nodes, pick a single “best” route  $R_i \in P(i)$  (where “best” is defined shortly in Theorem 1);

Second, pick a reward vector  $r_i = [r_{ij}]_j$  promising a reward  $r_{ij}$  to each candidate neighbor  $j \in B(i)$  that it has not received a competing offer from (i.e., such that  $r_{ji} < r_{ij}$  where  $r_{ji} = 0$  means that  $i$  did not receive an offer from  $j$ ). Then export the route and reward to the respective candidate neighbors. The distribution process repeats up to some depth that is directly dependent on the initial investment  $r_d$  as well as on the strategies of the players.



**Cost:** The cost of participation is local to the node and includes for example the cost associated with the effort spent in maintaining the route information. We assume that every player  $i$  incurs a cost of participation  $c_i$  and for simplicity we take  $c_i = c = 1$ .

**Utility:** A strategy profile  $\mathbf{s} = (s_1, \dots, s_n)$  and a reward  $r_d$  define an outcome of the game<sup>2</sup>. Every outcome determines a set of paths to destination  $d$  given by  $O_d = (R_1, \dots, R_n)$ . A utility function  $u_i(\mathbf{s})$  for player  $i$  associates every outcome with a real value in  $\mathbb{R}$ . We use the notation  $s_{-i}$  to refer to the strategy profile of all players excluding  $i$ . A simple class of utility functions we experiment with rewards a node linearly based on the number of sales that the node makes. This model incentivizes distribution and potentially requires a large initial investment from  $d$ . More clearly, define  $N_i(\mathbf{s}) = \{j \in V \setminus \{i\} | i \in R_j\}$  to be the set of nodes that pick their best route to  $d$  going through  $i$  (nodes downstream of  $i$ ) and let  $\delta_i(\mathbf{s}) = |N_i(\mathbf{s})|$ . Let the utility of a node  $i$  from an outcome or strategy profile  $\mathbf{s}$  be:

$$u_i(\mathbf{s}) = (r_{next(R_i)} - c_i) + \sum_{\{j | i = next(R_j)\}} (r_{next(R_i)} - r_{ij})(\delta_j(\mathbf{s}) + 1) \quad (1)$$

The first term  $(r_{next(R_i)} - c_i)$  of (1) is incurred by every participating node and is the one unit of reward from the upstream parent on the chosen best path minus the local cost. Based on the fixed cost assumption, we often drop this first term when comparing player payoffs from different strategies since the term is always positive when  $c = 1$ . The second term of (1) (the summation) is incurred only by distributors and is the total profit made by  $i$  where  $(r_{next(R_i)} - r_{ij})(\delta_j(\mathbf{s}) + 1)$  is  $i$ 's profit from the sale to neighbor  $j$  (which depends on  $\delta_j$ ). A rational selfish node will always try to maximize its utility by picking  $s_i = (R_i, [r_{ij}]_j)$ . There is an inherent tradeoff between  $(r_{next(R_i)} - r_{ij})$  and  $(\delta_j(\mathbf{s}))$  s.t.  $i = next(R_j)$  when trying to maximize the utility in Equation (1) in the face of competition as shall become clear later. A higher promised reward  $r_{ij}$  allows the node to compete (and possibly increase  $\delta_j$ ) but cuts the profit margin. Finally, we implicitly assume that the destination node  $d$  gets a constant marginal utility of  $r_d$  for each distinct player that maintains a route to  $d$  - the marginal utility of being discoverable by any seeker - and declares  $r_d$  truthfully to its neighbors i.e.,  $r_d$  is not strategic.

**Assumptions:** We take the following simplifying assumptions to keep our model tractable:

1. the advertiser  $d$  does not differentiate among the different players (ASes).
2. the advertised rewards are integers and are strictly decreasing with depth i.e.  $r_{ij} \in \mathbb{Z}^+$  and  $r_{ij} < r_{next(R_i)}, \forall i, j$  and let 1 unit be the cost of distribution.
3. finally, our choice of the utility function isolates a class of policies which we refer to as the Highest Reward Path (HRP). We assume for the scope of this work that transit costs are extraneous to the model.

**Convergence under HRP.** Before proceeding with the game model, we first prove the following theorem which results in the Highest Reward Path (HRP) policy. All proofs may be found in the full version of this paper [7].

<sup>2</sup> We abuse notation hereafter and we refer to the outcome with simply the strategy profile  $\mathbf{s}$  where it should be clear from context that an outcome is defined by the tuple  $\langle \mathbf{s}, r_d \rangle$ .

**Theorem 1.** *In order to maximize its utility, node  $i$  must always pick the route  $R_i$  with the highest promised reward i.e. such that  $r_{next(R_i)} \geq r_{next(R_l)}, \forall R_l \in P(i)$ .*

Theorem (I) implies that a player could perform her two actions sequentially, by first choosing the highest reward route  $R_i$ , then deciding on the reward vector  $r_{ij}$  to export to its neighbors. Thus, we shall represent player  $i$ 's strategy hereafter simply with the rewards vector  $[r_{ij}]$  and it should be clear that player  $i$  will always pick the “best” route to be the route with the highest promised reward. When the rewards are equal however, we assume that a node breaks ties consistently. Given the asynchronous nature of BGP, we ask the question of whether the BGP protocol dynamics converge to a unique outcome tree  $T_d$  under some strategy profile  $s$  [4]. From Theorem (I), it may be shown that the BGP outcome converges under any strategy profile  $s$ , including the equilibrium (see [7] for proof). This result allows us to focus on the existence of equilibria.

## 2.1 The Static Multi-stage Game with Fixed Schedule


We restrict the analysis of equilibria to the simple line and ring graphs. In order to apply the correct solution concept, we fix the *schedule* of play (i.e. who plays when?) based on the inherent order of play in the model. We resort to the *multi-stage game with observed actions* [3] where stages in our game have no temporal semantics. Rather, stages identify the network positions which have strategic significance due to the strictly decreasing rewards assumption. Formally, and using notation from [3], each player  $i$  plays only once at stage  $k > 0$  where  $k$  is the distance from  $i$  to  $d$  in number of hops. At every other stage, the player plays the “do nothing” action. The game starts at stage 1 after  $d$  declares  $r_d$ . Players at the same stage play simultaneously, and we denote by  $a^k = (a_1^k, \dots, a_n^k)$  the set of player actions at stage  $k$ , the stage- $k$  action profile. Further, denote by  $h^{k+1} = (r_d, a^1, \dots, a^k)$ , the *history* at the end of stage  $k$  which is simply the initial reward  $r_d$  concatenated with the sequence of actions at all previous stages. We let  $h^1 = (r_d)$ . Finally,  $h^{k+1} \in H^{k+1}$  the latter being the set of all possible stage- $k$  histories. When the game has a finite number of stages, say  $K + 1$ , then a terminal history  $h^{K+1}$  is equivalent to an outcome of the game (which is a tree  $T_d$ ) and the set of all outcomes is  $H^{K+1}$ . The pure-strategy of player  $i$  who plays at stage  $k > 0$  is a function of the history and is given by  $s_i : H^k \rightarrow \mathbb{R}^{m_i}$  where  $m_i$  is the number of direct neighbors of player  $i$  that are at stage  $k + 1$  (implicitly, a player at stage  $k$  observes the full history  $h^k$  before playing). We resort to the multi-stage model (the fixed schedule) on our simple graphs to eliminate the synchronization problems inherent in the BGP protocol and to focus instead on the existence of equilibria. By restricting the analysis to the fixed schedule, we do not miss any equilibria (see [7]). The key concept here is that it is the *information sets* [3] that matter rather than the time of play i.e. since all the nodes at distance 1 from  $d$  observe  $r_d$  before playing, all these nodes belong to the same information set whether they play at the same time or at different time instants.

Starting with  $r_d$  (which is  $h^1$ ), it is clear how the game produces actions at every later stage based on the player strategies resulting in a terminal action profile or outcome. Hence, given  $r_d$ , an outcome in  $H^{K+1}$  may be associated with every strategy profile  $s$  and so the definition of Nash equilibrium remains unchanged (see [3] for definitions of *Nash equilibrium*, *proper subgame*, and *subgame perfection*). In our game, each

stage begins a new subgame which restricts the full game to a particular history. For example, a history  $h^k$  begins a subgame  $G(h^k)$  such that the histories in the subgame are restricted to  $h^{k+1} = (h^k, a^k)$ ,  $h^{k+2} = (h^k, a^k, a^{k+1})$ , and so on. Hereafter, the general notion of equilibrium we use is the Nash equilibrium and we shall make it clear when we generalize to subgame perfect equilibria. We are only interested in pure-strategy equilibria [3] and in studying the existence question as the incentive  $r_d$  varies.

### 3 Equilibria on the Line Graph, the Tree, and the Ring Graph

In the general game model defined thus far, the tie-breaking preferences of the players is a defining property of the game, and every outcome (including the equilibrium) depends on the initial reward/utility  $r_d$  of the advertiser. In the same spirit as [8] we inductively construct the equilibrium for the line graph of Figure 1(a) given the utility function of Equation (1). We present the result for the line which may be directly extended to trees. Before proceeding with the construction, notice that for the line,  $m_i = 1$  for all players except the leaf player since each of those players has a single downstream neighbor. In addition,  $\delta_i(s) = \delta_j(s) + 1, \forall i, j$  where  $j$  is  $i$ 's child ( $\delta_i = 0$  when  $i$  is a leaf). We shall refer to both the player and the stage using the same index since our intention should be clear from the context. For example, the child of player  $i$  is  $i + 1$  and its parent is  $i - 1$  where player  $i$  is the player at stage  $i$ . Additionally, we simply represent the history  $h^{k+1} = (r_k)$  for  $k > 0$  where  $r_k$  is the reward promised by player  $k$  (player  $k$ 's action). The strategy of player  $k$  is therefore  $s_k(h^k) = s_k(r_{k-1})$  which is a singleton (instead of a vector) since  $m_i = 1$  (for completeness, let  $r_0 = r_d$ ). This is a *perfect information* game [3] since a single player moves at each stage and has complete information about the actions of all players at previous stages. Backward induction may be used to construct the subgame-perfect equilibrium. We construct the equilibrium strategy  $s^*$  inductively as follows: first, for all players  $i$ , let  $s_i^*(x) = 0$  when  $x \leq c$  (where  $c$  is assumed to be 1). Then assume that  $s_i^*(x)$  is defined for all  $x < r$  and for all  $i$ . Obviously, with this information, every player  $i$  may compute  $\delta_i(x, s_{-i}^*)$  for all  $x < r$ .



**Fig. 1.** (a) Line graph: a player's index is the stage at which the player plays;  $d$  advertises at stage 0;  $K = n$ ; (b) Ring graph with even number of players: (i) 2-stage game, (ii) 3-stage game, and general (iii)  $K$ -stage game

This is simply due to the fact that  $\delta_i$  depends on the downstream players from  $i$  who must play an action or reward strictly less than  $r$ . Finally, for all players  $i$  we let  $s_i^*(r) = \arg \max_x (r - x) \delta_i(x, s_{-i}^*)$  where  $x < r$ .

**Theorem 2.** *The strategy profile  $s^*$  is a subgame-perfect equilibrium.*

The proof may be directly extended to the tree since each player in the tree has a single upstream parent as well and backward induction follows in the same way. On the tree, the strategies of the players that play simultaneously at each stage are also independent.

**Competition: the ring.** We present next a negative result for the ring graph. In a ring, each player has a degree = 2 and  $m_i = 1$  for all players except the leaf player. We consider rings with an even number of nodes due to the direct competition dynamics. Figure 1(b) shows the 2-, the 3-, and general  $K$ -stage versions of the game. In the multi-stage game, after observing  $r_d$ , players 1 and 2 play simultaneously at stage 1 promising rewards  $r_1$  and  $r_2$  respectively to their downstream children, and so on. We refer to the players at stage  $j$  using ids  $2j - 1$  and  $2j$  where the stage of a player  $i$ , denoted as  $l(i)$ , may be computed from the id as  $l(i) = \lceil \frac{i}{2} \rceil$ . For the rest of the discussion, we assume WLOG that the player at stage  $K$  (with id  $2K - 1$ ) breaks ties by picking the route through the left parent  $2K - 3$ . For the 2-stage game in Figure 1(b)(i), it is easy to show that an equilibrium always exists in which  $s_1^*(r_d) = s_2^*(r_d) = (r_d - 1)$  when  $r_d > 1$  and 0 otherwise. This means that player 3 enjoys the benefits of *perfect competition* due to the Bertrand-style competition [3] between players 1 and 2. The equilibrium in this game is independent of player 3’s preference for breaking ties. We now present the following negative result,

**Claim 1.** *The 3-stage game induced on the ring (of Figure 1(b)(ii)) does not have a subgame-perfect equilibrium. Particularly, there exists a class of subgames for  $h^1 = r_d > 5$  for which there is no Nash equilibrium.*

The value  $r_d > 5$  signifies the breaking point of equilibrium or the reward at which player 2, when maximizing her utility  $(r_d - r_2) \delta_2$ , will always oscillate between competing for 5 (by playing large  $r_2$ ) or not (by playing small  $r_2$ ). This negative result for the game induced on the 3-stage ring may be directly extended to the general game for the  $K$ -stage ring by observing that a class of subgames  $G(h^{K-2})$  of the general  $K$ -stage game are identical to the 3-stage game. While the full game does not always have an equilibrium when  $K > 2$  stages, we shall show next that there always exists an equilibrium for a special subgame.

**Growth of Incentives, and a Special Subgame.** We next answer the following question: Find the minimum incentive  $r_d^*$ , as a function of the depth of the network  $K$  (equivalently the number of stages in the multi-stage game), such that there exists an equilibrium outcome for the subgame  $G(r_d^*)$  that is a spanning tree. We seek to compute the function  $f$  such that  $r_d^* = f(K)$ . First, we present a result for the line, before extending it to the ring. On the line,  $K$  is simply the number of players i.e.  $K = n$ , and  $f(K)$  grows exponentially with the depth  $K$  as follows:

**Lemma 1.** *On the line graph, we have  $f(0) = 0$ ,  $f(1) = 1$ ,  $f(2) = 2$ , and  $\forall k > 2$ ,  $f(k) = (k - 1)f(k - 1) - (k - 2)f(k - 2)$ .*

We now revisit the the  $K$ -stage game of Figure 1(b)(iii) on the ring and we focus on a specific subgame which is the restriction of the full game to  $h_1 = r_d^* = f(K)$ , and we denote this subgame by  $G(r_d^*)$ . Consider the following strategy profile  $\mathbf{s}^*$  for the subgame: players at stage  $j$  play  $s_{2j-1}^*(h^j) = f(K-j)$ , and  $s_{2j}^*(h^j) = f(K-j-1)$ ,  $\forall 1 \leq j \leq K-1$ , and let  $s_{2K-1}^*(h^K) = 0$ .

**Theorem 3.** *The profile  $\mathbf{s}^*$  is a Nash equilibrium for the subgame  $G(r_d^*)$  on the  $K$ -stage ring,  $\forall K > 2$ .*

This result may be interpreted as follows: if the advertiser were to play strategically assuming she has a marginal utility of at least  $r_d^*$  and is aiming for a spanning tree (global discoverability), then  $r_d^* = f(K)$  will be her Nash strategy in the game induced on the  $K$ -stage ring,  $\forall K > 2$  (given  $\mathbf{s}^*$ ). We can now extend the growth result of Lemma 1 to the ring denoting by  $f_r(K)$  the growth function for the ring.

**Corollary 1.** *On the ring graph, we have  $f_r(k) = f(k)$  as given by Lemma 1.*

In this paper, we have studied the equilibria existence question for a simple class of graphs. Many questions remain to be answered including extending the results to general network structures (and to the Internet *small-world* connectivity graph), relaxing the fixed cost assumption, quantifying how hard is it to find the equilibria, and devising mechanisms to get to them. All these questions are part of our ongoing work [7].

## References

1. Feigenbaum, J., Papadimitriou, C., Sami, R., Shenker, S.: A bgp-based mechanism for lowest-cost routing. *Distrib. Comput.* 18(1), 61–72 (2005)
2. Feigenbaum, J., Ramachandran, V., Schapira, M.: Incentive-compatible interdomain routing. In: *EC 2006*, pp. 130–139. ACM, New York (2006)
3. Fudenberg, D., Tirole, J.: *Game Theory*. MIT Press, Cambridge (1991)
4. Griffin, T.G., Shepherd, F.B., Wilfong, G.: Policy disputes in path-vector protocols. In: *ICNP 1999*, p. 21. IEEE Computer Society, Washington (1999)
5. Herrin, W.: What does a bgp route cost (2008), <http://bill.herrin.us/network/bggpcost.html>
6. Huston, G.: Bgp in 2008 (2008), <http://www.potaroo.net/ispcol/2009-03/bgp2008.html>
7. Khoury, J., Abdallah, C.T., Krause, K., Crichigno, J.: Route distribution incentives, <http://arxiv.org/abs/0909.3558v1>
8. Kleinberg, J., Raghavan, P.: Query incentive networks. In: *FOCS 2005*, pp. 132–141. IEEE Computer Society, Washington (2005)
9. Levin, H., Schapira, M., Zohar, A.: Interdomain routing and games. In: *STOC 2008*, pp. 57–66. ACM, New York (2008)
10. Li, C., Yu, B., Sycara, K.: An incentive mechanism for message relaying in unstructured peer-to-peer systems. In: *AAMAS 2007*, pp. 1–8. ACM, New York (2007)
11. Meyer, D., Zhang, L., Fall, K.: Report from the iab workshop on routing and addressing. Internet RFC 4984 (September 2007)
12. Rekhter, Y., Li, T., Hares, S.: RFC 4271: A border gateway protocol 4, bgp-4 (2006)

# Wiretapping a Hidden Network<sup>\*</sup>

Haris Aziz<sup>2</sup>, Oded Lachish<sup>1</sup>, Mike Paterson<sup>1</sup>, and Rahul Savani<sup>3</sup>

<sup>1</sup> Department of Computer Science, University of Warwick, CV4 7AL Coventry, UK  
{oded,msp}@dcs.warwick.ac.uk

<sup>2</sup> Institut für Informatik, Universität München, 80538 München, Germany  
aziz@tcs.ifi.lmu.de

<sup>3</sup> Department of Computer Science, University of Liverpool, L69 3BX Liverpool, UK  
rahul.savani@liverpool.ac.uk

**Abstract.** We consider the problem of maximizing the probability of hitting a strategically chosen hidden *virtual network* by placing a wiretap on a single link of a communication network. This can be seen as a two-player win-lose (zero-sum) game that we call the *wiretap game*. The *value* of this game is the greatest probability that the wiretapper can secure for hitting the virtual network. The value is shown to be equal the reciprocal of the *strength* of the underlying graph. We provide a polynomial-time algorithm that finds a linear-sized description of the maxmin-polytope, and a characterization of its extreme points. It also provides a succinct representation of all equilibrium strategies of the wiretapper that minimize the number of pure best responses of the hider. Among these strategies, we efficiently compute the *unique* strategy that maximizes the least punishment that the hider incurs for playing a pure strategy that is not a best response. Finally, we show that this unique strategy is the nucleolus of the recently studied simple cooperative *spanning connectivity game*.

**Keywords:** Network security, nucleolus, wiretapping, zero-sum game.

## 1 Introduction

We consider the problem of maximizing the probability of hitting a strategically chosen hidden *virtual network* by placing a wiretap on a single link of a communication network, represented by an undirected, unweighted graph. This can be seen as a two-player win-lose (zero-sum) game that we call the *wiretap game*. A pure strategy of the wiretapper is an edge to tap, and of his opponent, the *hider*, a choice of virtual network, a *connected spanning subgraph*. The *wiretapper* wins, with payoff one, when he picks an edge in the network chosen by the hider, and loses, with payoff zero, otherwise. Thus, the *value* of this game is the greatest probability that the wiretapper can secure for hitting the hidden network. He

---

<sup>\*</sup> This research was supported in part by EPSRC projects EP/D067170/1, EP/G064679/1, and by the Centre for Discrete Mathematics and its Applications (DIMAP), EPSRC award EP/D063191/1.

does this by playing a maxmin strategy, which is a probability distribution on the edges. The value also equals the smallest probability that the hider can secure, which she does by playing a minmax strategy, which is a probability distribution on connected spanning subgraphs.

**Our results.** The value is shown to be equal to the *strength* of the underlying graph [6]. We obtain in polynomial time a linear number of simple two-variable inequalities that define the maxmin-polytope, and a characterization of its extreme points. In contrast, the natural description of the maxmin-polytope is as the solutions to a linear program with exponentially many constraints. This allows us to efficiently find all equilibrium strategies of the wiretapper that minimize the number of pure best responses of the hider, and, among these strategies, the *unique* strategy that maximizes the least punishment that the hider incurs for playing a pure strategy that is not a best response. This special maxmin strategy corresponds to the nucleolus of the *spanning connectivity game*, a simple cooperative game [1].

**Related work.** The strength of an unweighted graph, which has a central role in our work, is also called the edge-toughness, and relates to the classical work of Nash-Williams [8] and Tutte [12]. Cunningham [4] generalized the concept of strength to edge-weighted graphs and proposed a strongly polynomial-time algorithm to compute it. Computing the strength of a graph is a special type of ratio optimization in the field of submodular function minimization [5]. Cunningham used the strength of a graph to address two different one-player optimization problems: the optimal attack and reinforcement of a network. The prime-partition we use is a truncated version of the principal-partition, first introduced by Narayanan [7] and Tomizawa [11]. The principal-partition was used in an extension of Cunningham's work to an online setting [9].

The nucleolus of the spanning connectivity game can be seen as a special maxmin strategy in the wiretap game. The connection between the nucleolus of a cooperative game and equilibrium strategies in a zero-sum game has been investigated before in a general context [10]. However, in many cases the nucleolus is hard to compute. Our positive results for the spanning connectivity game are in contrast to the negative results presented in [1], where it is shown that the problems of computing the Shapley values and Banzhaf values are #P-complete for the spanning connectivity game.

## 2 The Wiretap Game

The strategic form of the wiretap game is defined implicitly by the graph  $G = (V, E)$ . The pure strategies of the wiretapper are the edges  $E$  and the pure strategies of the hider are the set of connected spanning subgraphs  $\mathcal{S}$ . An element of  $\mathcal{S}$  is a set of edges, with a typical element denoted by  $S$ . The wiretapper receives payoff one if the edge he chooses is part of the spanning subgraph chosen by the hider, and receives payoff zero otherwise. Thus, the value of the game is



the probability that the wiretapper can secure for wiretapping the connected spanning subgraph chosen by the hider.

Let  $\Delta(A)$  be the set of mixed strategies (probability distributions) on a finite set  $A$ . By the well-known minmax theorem for finite zero-sum games, the wiretap game  $\Gamma(G)$  has a unique *value*, defined by

$$val(\Gamma) = \max_{x \in \Delta(E)} \min_{S \in \mathcal{S}} \sum_{e \in S} x_e = \min_{y \in \Delta(\mathcal{S})} \max_{e \in E} \sum_{\{S \in \mathcal{S}: e \in S\}} y_S . \tag{1}$$

The equilibrium or *maxmin* strategies of the wiretapper are the solutions  $\{x \in \Delta(E) \mid \sum_{e \in S} x_e \geq val(\Gamma) \text{ for all } S \in \mathcal{S}\}$  to the following linear program, which has the optimal value  $val(\Gamma)$ .

$$\begin{aligned} & \max z \\ & \text{s.t. } \sum_{e \in S} x_e \geq z \text{ for all } S \in \mathcal{S} , \\ & \quad x \in \Delta(E) . \end{aligned} \tag{2}$$

Playing any maxmin strategy guarantees the wiretapper a probability of successful wiretapping of at least  $val(\Gamma)$ . The equilibrium or *minmax* strategies of the hider are  $\{y \in \Delta(\mathcal{S}) \mid \sum_{\{S \in \mathcal{S}: e \in S\}} y_S \leq val(\Gamma) \text{ for all } e \in E\}$ . Playing any minmax strategy guarantees the hider to suffer a probability of successful wiretapping of no more than  $val(\Gamma)$ . The following simple observation shows the importance of minimum connected spanning graphs in the analysis of the wiretap game. For a mixed strategy  $x \in \Delta(E)$  and pure strategy  $S \in \mathcal{S}$ , the resulting probability of a successful wiretap is  $\sum_{e \in S} x_e$ . We denote by  $G^x$  the edge-weighted graph comprising the graph  $G$  with edge weights  $x(e)$  for all  $e \in E$ . Let  $w^*(x)$  be the weight of a minimum connected spanning graph of  $G^x$ .

**Fact 1.** *The set of pure best responses of the hider against the mixed strategy  $x \in \Delta(E)$  is*

$$\{S \in \mathcal{S} \mid \sum_{e \in S} x_e = w^*(x)\} .$$

We could define the wiretap game by only allowing the hider to pick spanning trees, however, our definition with connected spanning subgraphs allows a clean connection to the spanning connectivity game.

### 3 Overview of Results

In this section, we present our results. Proofs of the results appear in [2]. We start with the basic notations and definitions. From here on we fix a connected graph  $G = (V, E)$ . Unless mentioned explicitly otherwise, any implicit reference to a graph is to  $G$  and  $\alpha$  is an edge-distribution, which is a probability distribution on the edges  $E$ . For ease, we often refer to the weighted graph  $G^\alpha$  simply by  $\alpha$ , where this usage is unambiguous. For a subgraph  $H$  of  $G$ , we denote by  $\alpha(H)$  the sum  $\sum_{e \in E(H)} \alpha(e)$ , where  $E(H)$  is the edge set of  $H$ . We refer to equilibrium strategies of the wiretapper as maxmin-edge-distributions.



**Definition 1.** For every edge-distribution  $\alpha$ , we denote its distinct weights by  $x_1^\alpha > \dots > x_m^\alpha \geq 0$  and define  $\mathcal{E}(\alpha) = \{E_1^\alpha, \dots, E_m^\alpha\}$  such that  $E_i^\alpha = \{e \in E \mid \alpha(e) = x_i^\alpha\}$  for  $i = 1, \dots, m$ .

Our initial goal is to characterize those partitions  $\mathcal{E}(\alpha)$  that can arise from maxmin-edge-distributions  $\alpha$ . We start with the following simple setting. Assume that the wiretapper is restricted to choosing a strategy  $\alpha$  such that  $|\mathcal{E}(\alpha)| = 2$ , and  $x_2^\alpha = 0$ . Thus, the wiretapper’s only freedom is the choice of the set  $E_1^\alpha$ . What is his best possible choice? By Fact [II](#), a best response against  $\alpha$  is a minimum connected spanning subgraph  $H$  of  $\alpha$ . So the wiretapper should choose  $E_1^\alpha$  so as to maximize  $\alpha(H)$ . How can such an  $E_1^\alpha$  be found? To answer, we relate the weight of a minimum connected spanning subgraph  $H$  of  $\alpha$  to  $E_1^\alpha$ .

To determine  $\alpha(H)$ , we may assume about  $H$  that for every connected component  $C$  of  $(V, E \setminus E_1^\alpha)$  we have  $E(H) \cap E(C) = E(C)$ , since  $\alpha(e) = 0$  for every  $e \in E(C)$ . We can also assume that  $|E_1^\alpha \cap E(H)|$  is the number of connected components in  $(V, E \setminus E_1^\alpha)$  minus 1, since this is the minimum number of edges in  $E(H)$  that a connected spanning subgraph may have. To formalize this we use the following notation.

**Definition 2.** Let  $E' \subseteq E$ . We set  $C_G(E')$ , to be the number of connected components in the graph  $G \setminus E'$ , where  $G \setminus E'$  is a shorthand for  $(V, E \setminus E')$ . If  $E' = \emptyset$  we just write  $C_G$ .

Using the above notation, a connected spanning subgraph  $H$  is a minimum connected spanning subgraph of  $\alpha$  if  $|H \cap E_1^\alpha| = C_G(E_1^\alpha) - C_G = C_G(E_1^\alpha) - 1$ . Now we can compute  $\alpha(H)$ . By definition,  $x_1^\alpha = \frac{1}{|E_1^\alpha|}$  and  $x_2^\alpha = 0$  and therefore

$$\alpha(H) = \frac{C_G(E_1^\alpha) - C_G}{|E_1^\alpha|}.$$

We call this ratio that determines  $\alpha(H)$  the cut-rate of  $E_1^\alpha$ . Note that it uniquely determines the weight of a minimum connected spanning subgraph of  $\alpha$ .

**Definition 3.** Let  $E' \subseteq E$ . The cut-rate of  $E'$  in  $G$  is denoted by  $cr_G(E')$  and defined as follows.

$$cr_G(E') := \begin{cases} \frac{C_G(E') - C_G}{|E'|} & \text{if } |V| > 1 \text{ and } |E'| > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

We write  $cr(E')$ , unless we make a point of referring to a different graph.

Thus, when  $|\mathcal{E}(\alpha)| = 2$  and  $x_2^\alpha = 0$ , a best choice of  $E_1^\alpha$  is one for which  $cr(E_1^\alpha)$  is maximum. Since  $E$  is finite, an  $E_1^\alpha$  that maximizes  $cr(E_1^\alpha)$  exists.

**Definition 4.** The cut-rate of  $G$  is defined as  $opt := \max_{E' \subseteq E} cr(E')$ .

By *opt*, we always refer to the cut-rate of the graph  $G$ . In case we refer to the cut-rate of some other graph, we add the name of the graph as a subscript.

The value  $opt$  is a well known and studied attribute of a graph. It is equal to the reciprocal of the strength of a graph, as defined by Gusfield [6] and named by Cunningham [4]. There exists a combinatorial algorithm for computing the strength, and hence  $opt$ , that runs in time polynomial in the *size* of the graph, by which we always mean  $|V| + |E|$ .

We generalize the above technique to the case that  $\alpha$  is not restricted.

**Definition 5.** For  $\ell = 1, \dots, |\mathcal{E}(\alpha)|$  we set

$$cr_\ell^\alpha = \frac{C_G(\cup_{i=1}^\ell E_i^\alpha) - C_G(\cup_{i=1}^{\ell-1} E_i^\alpha)}{|E_\ell^\alpha|}.$$

**Proposition 1.** Let  $H$  be a minimum connected spanning subgraph of  $\alpha$ . Then  $|E(H) \cap E_\ell^\alpha| = |E_\ell^\alpha| cr_\ell^\alpha$  for every  $\ell$  such that  $x_\ell^\alpha > 0$ .

Using Proposition 1 we can relate the weight of a minimum connected spanning subgraph of  $\alpha$  to the sets of  $\mathcal{E}(\alpha)$ . This relationship also characterizes the maxmin-edge-distributions, which are the edge-distributions whose minimum connected spanning subgraph weight is the maximum possible.

**Theorem 1.** Let  $H$  be a minimum connected spanning subgraph of  $\alpha$  and  $m = |\mathcal{E}(\alpha)|$ . Then  $\alpha(H) \leq opt$  and we have  $\alpha(H) = opt$  if and only if

1.  $cr_\ell^\alpha = opt$  for  $\ell = 1, \dots, m - 1$ , and
2. if  $cr_m^\alpha \neq opt$  then  $x_m^\alpha = 0$ .

An immediate implication of Theorem 1 is that  $opt$  is an upper bound on the value the wiretapper can achieve. This also follows from the well-known fact that the fractional packing number of spanning trees of a graph is equal to the strength of a graph, which in turn follows from the theorems of Nash-Williams [8] and Tutte [12] on the integral packing number (see also [3]). Since we have already seen that indeed the wiretapper can achieve  $opt$  by distributing all probability mass equally over an edge set that has cut-rate  $opt$ , we get the following.

**Corollary 1.** The value of the wiretap game is  $opt$ .

We know what the value of the game is and we know a characterization of the  $\mathcal{E}(\alpha)$ 's for maxmin-edge-distributions  $\alpha$ . Yet this characterization does not give us a simple way to find maxmin-edge-distributions. Resolving this is our next goal.

**Definition 6.** Let  $\mathcal{E}_1, \mathcal{E}_2$  be partitions of  $E$ . Then  $\mathcal{E}_1$  refines  $\mathcal{E}_2$  if for every set  $E' \in \mathcal{E}_1$  there exists a set  $E'' \in \mathcal{E}_2$  such that  $E' \subseteq E''$ .

Thus, there exists a partition of  $E$  that is equal to  $\mathcal{E}(\beta)$  for some maxmin-edge-distribution  $\beta$  and refines  $\mathcal{E}(\gamma)$  for every maxmin-edge-distribution  $\gamma$ . We call such a partition the *prime-partition*. It is unique since there can not be different partitions that refine each other.

**Definition 7.** The prime-partition  $\mathcal{P}$  is the unique partition that is equal to  $\mathcal{E}(\beta)$  for some maxmin-edge-distribution  $\beta$  and refines  $\mathcal{E}(\gamma)$  for every maxmin-edge-distribution  $\gamma$ .

**Theorem 2.** The prime-partition exists and can be computed in time polynomial in the size of  $G$ .

The prime-partition  $\mathcal{P}$  reveals a lot about the structure of the maxmin-edge-distributions. Yet by itself  $\mathcal{P}$  does not give us a simple means for generating maxmin-edge-distributions. Using the algorithm for finding  $\mathcal{P}$  one can show that, depending on  $G$ , there may be a unique element in  $\mathcal{P}$  whose edges are assigned 0 by every maxmin-edge-distribution.

**Lemma 1.**  $cr_G(E) \neq opt$  if and only if there exists a unique set  $D \in \mathcal{P}$  such that for every maxmin-edge-distribution  $\alpha$  and  $e \in D$  we have  $\alpha(e) = 0$ . If  $D$  exists then it can be found in time polynomial in the size of  $G$ .

From here on we shall always refer to the set  $D$  in Lemma 1 as the *degenerate set*. For convenience, if  $D$  does not exist then we shall treat both  $\{D\}$  and  $D$  as the empty set. We use the prime-partition to define a special subset of the minimum connected spanning subgraphs that we call the *omni-connected-spanning-subgraphs*.

**Definition 8.** A connected spanning subgraph  $H$  is an omni-connected-spanning-subgraph if for every  $P \in \mathcal{P} \setminus \{D\}$  we have  $|E(H) \cap P| = |P| \cdot opt$ .

**Proposition 2.** There exists an omni-connected-spanning-subgraph.

The omni-connected-spanning-subgraphs are the set of the hider's pure strategies that are best responses against every maxmin-edge-distribution.

**Proposition 3.** For every edge-distribution  $\alpha$  such that  $\mathcal{P}$  refines  $\mathcal{E}(\alpha)$  and  $\alpha(e) = 0$  for every  $e \in D$  and omni-connected-spanning-subgraph  $H$ , we have  $\alpha(H) = opt$ .

The importance of omni-connected-spanning-subgraphs stems from the following scenario. Assume that  $\mathcal{P}$  refines  $\mathcal{E}(\alpha)$  and  $\alpha(e) = 0$  for every  $e \in D$ , and let  $H$  be an omni-connected-spanning-subgraph. By Proposition 3, we know that  $\alpha(H) = opt$ . Suppose we can remove from  $H$  an edge from  $E(H) \cap P$ , where  $P$  is a nondegenerate element of  $\mathcal{P}$ , and add a new edge from another set  $P' \setminus E(H)$  in order to get a new connected spanning subgraph. Assume  $\alpha$  assigns to the edge removed strictly more weight than it assigns to the edge added. Then the new connected spanning subgraph has weight strictly less than  $\alpha(H)$  and hence strictly less than  $opt$ , since  $\alpha(H) = opt$  by Proposition 3. Consequently,  $\alpha$  is not a maxmin-edge-distribution and we can conclude that any edge-distribution  $\beta$  that assigns to each edge in  $P$  strictly more weight than to the edges in  $P'$  is not a maxmin-edge-distribution. This intuition is captured by the following definition, which leads to the characterization of maxmin-edge-distributions in Theorem 3.

**Definition 9.** Let  $P, P' \in \mathcal{P} \setminus \{D\}$  be distinct. Then  $P$  leads to  $P'$  if and only if there exists an *omni-connected-spanning-subgraph*  $H$  with  $e \in P \setminus E(H)$  and  $e' \in P' \cap E(H)$  such that  $(H \setminus \{e'\}) \cup \{e\}$  is a *connected spanning subgraph*.

**Definition 10.** Let  $P, P' \in \mathcal{P} \setminus \{D\}$  be distinct. We say that  $P$  is a *parent* of  $P'$  (conversely  $P'$  a *child* of  $P$ ) if  $P$  leads to  $P'$  and there is no  $P'' \in \mathcal{P}$  such that  $P$  leads to  $P''$  and  $P''$  leads to  $P'$ . We refer to the relation as the *parent-child relation* and denote it by  $\mathcal{O}$ .

**Definition 11.** An *edge-distribution*  $\alpha$  agrees with  $\mathcal{O}$  if  $\mathcal{P}$  refines  $\mathcal{E}(\alpha)$  and for every  $P \in \mathcal{P} \setminus \{D\}$  that is a parent of  $P' \in \mathcal{P} \setminus \{D\}$  and  $e \in P, e' \in P'$  we have  $\alpha(e) \geq \alpha(e')$ , and for every  $e \in D$  we have  $\alpha(e) = 0$ .

**Theorem 3.** An *edge-distribution*  $\alpha$  is a *maxmin-edge-distribution* if and only if it agrees with  $\mathcal{O}$ .

Theorem 3 defines a linear inequality for each parent and child in the relation  $\mathcal{O}$ . Along with the inequalities that define a probability distribution on edges, this gives a small number of two-variable inequalities describing the maxmin-polytope. In [2], we characterize the extreme points of the maxmin-polytope.

**Theorem 4.** The *parent-child relation*  $\mathcal{O}$  can be computed in time polynomial in the size of  $G$ .

The wiretapper will in general have a choice of infinitely many maxmin-edge-distributions. To choose a maxmin-edge-distribution, it is natural to consider refinements of the Nash equilibrium property that are beneficial to the wiretapper if the hider does not play optimally. First we show how to minimize the number of pure best responses of the hider. To do this, we use the relation  $\mathcal{O}$  to characterize a special type of maxmin-edge-distribution which achieves this. We call this a *prime-edge-distribution*. The prime-edge-distributions are characterized by the following lemma.

**Definition 12.** A *maxmin-edge-distribution*  $\alpha$  is a *prime-edge-distribution* if the number of the hider’s pure best responses against it is the minimum possible.

**Lemma 2.** An *edge-distribution*  $\gamma$  is a *prime-edge-distribution* if and only if  $\gamma(e) > 0$  for every  $e \in E \setminus D$ , and for every  $P, P' \in \mathcal{P} \setminus \{D\}$  such that  $P$  is a parent of  $P'$  and every  $e \in P, e' \in P'$ , we have  $\gamma(e') > \gamma(e)$ .

Using this characterization one can easily check whether  $\alpha$  is a prime-edge-distribution and one can also easily construct a prime-edge-distribution.

We have already seen how to minimize the number of pure best responses of the hider, by playing a prime-edge-distribution. We now show how to uniquely maximize the weight of a pure second-best response by choosing between prime-edge-distributions. This maximizes the least punishment that the hider will incur for picking a non-optimal pure strategy.

Against a prime-edge-distribution, the candidates for pure second-best responses are those connected spanning subgraphs that differ from omni-connected-spanning-subgraphs in at most two edges. For each parent and child we have at

least one of these second-best responses. A second-best response either is a best response with one extra edge, or it differs from a best response in two edges, where it has one less edge in a child of  $\mathcal{O}$  and one more in the child's parent.

We are only interested in the case that  $opt < 1$ , since the graph has  $opt = 1$  if and only if it contains a bridge, in which case the value of the game is 1 and the hider does not have a second-best response. So we assume that  $opt < 1$ .

Intuitively, to maximize the weight of a second-best response, we want to minimize the number of distinct weights. The minimum number of distinct positive weights we can achieve for a prime-edge-distribution is equal to the number of elements in the longest chain in the parent-child relation. This motivates the following definition.

**Definition 13.** We define  $\mathcal{L}_1, \mathcal{L}_2, \dots$  inductively as follows. The set  $\mathcal{L}_1$  is all the sinks of  $\mathcal{O}$  excluding  $D$ . For  $j = 2, \dots$ , we have that  $\mathcal{L}_j$  is the set of all the sinks when all elements of  $\{D\} \cup (\cup_{i=1, \dots, j-1} \mathcal{L}_i)$  have been removed from  $\mathcal{O}$ .

Note that  $\mathcal{O}$  is defined only over nondegenerate elements of  $\mathcal{P}$  and hence the degenerate set is not contained in any of  $\mathcal{L}_1, \mathcal{L}_2, \dots$ .

The following theorem shows that there is a unique prime-edge-distribution that maximizes the difference between the payoff of a best and second-best response. This unique prime-edge-distribution turns out to be the nucleolus of the spanning connectivity game, as explained in [2]. For convenience, we refer to this strategy as the nucleolus.

**Theorem 5.** Let  $L_i = \cup_{E' \in \mathcal{L}_i} E'$  for  $i = 1, \dots, t$ . Let

$$\kappa = \frac{1}{\sum_{i=1}^t i \cdot |L_i|}.$$

The nucleolus  $\nu$  has  $\nu(e) = i \cdot \kappa$  for every  $i \in \{1, \dots, t\}$  and  $e \in L_i$  and  $\nu(e) = 0$  otherwise.

## References

1. Aziz, H., Lachish, O., Paterson, M., Savani, R.: Power indices in spanning connectivity games. In: AAIM: Algorithmic Aspects in Information and Management, pp. 55–67 (2009)
2. Aziz, H., Lachish, O., Paterson, M., Savani, R.: Wiretapping a hidden network. Technical report (2009), <http://arxiv.org/abs/0909.5293>
3. Chakrabarty, D., Mehta, A., Vazirani, V.V.: Design is as easy as optimization. In: ICALP: International Colloquium on Automata, Languages and Programming, pp. 477–488 (2006)
4. Cunningham, W.H.: Optimal attack and reinforcement of a network. J. ACM 32(3), 549–561 (1985)
5. Fujishige, S.: Submodular Functions and Optimization. Annals of Discrete Mathematics, vol. 58. Elsevier, Amsterdam (2005)
6. Gusfield, D.: Connectivity and edge-disjoint spanning trees. Inf. Process. Lett. 16(2), 87–89 (1983)

7. Narayanan, H.: Theory of matroids and network analysis. PhD thesis, IIT, Bombay (1974)
8. Nash-Williams, C.S.A.: Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.* 36, 445–450 (1961)
9. Patkar, S., Narayanan, H.: Fast on-line/off-line algorithms for optimal reinforcement of a network and its connections with principal partition. In: *FST TCS: Foundations of Software Technology and Theoretical Computer Science*, pp. 94–105. Springer, Heidelberg (2000)
10. Potters, J.A.M., Tijs, S.H.: The nucleolus of a matrix game and other nucleoli. *Math. Oper. Res.* 17(1), 164–174 (1992)
11. Tomizawa, N.: Strongly irreducible matroids and principal partition of a matroid into strongly irreducible minors. *Electron. and Commun.* 59(A), 1–10 (1976)
12. Tutte, W.T.: On the problem of decomposing a graph into  $n$  connected factors. *J. London Math. Soc.* 36, 221–230 (1961)

# Refining the Cost of Cheap Labor in Set System Auctions

Ning Chen, Edith Elkind, and Nick Gravin

Division of Mathematical Sciences, Nanyang Technological University, Singapore

**Abstract.** In *set system auctions*, a single buyer needs to purchase services from multiple competing providers, and the set of providers has a combinatorial structure; a popular example is provided by shortest path auctions [17]. In [3] it has been observed that if such an auction is conducted using first-price rules, then, counterintuitively, the buyer's payment may go down if some of the sellers are prohibited from participating in the auction. This reduction in payments has been termed "the cost of cheap labor". In this paper, we demonstrate that the buyer can attain further savings by setting lower bounds on sellers' bids. Our model is a refinement of the original model of [3]: indeed, the latter can be obtained from the former by requiring these lower bounds to take values in  $\{0, +\infty\}$ . We provide upper and lower bounds on the reduction in the buyer's payments in our model for various set systems, such as minimum spanning tree auctions, bipartite matching auctions, single path and  $k$ -path auctions, vertex cover auctions, and dominating set auctions. In particular, we illustrate the power of the new model by showing that for vertex cover auctions, in our model the buyer's savings can be linear, whereas in the original model of [3] no savings can be achieved.

## 1 Introduction

Combinatorial procurement auctions, or *set system auctions*, play an important role in electronic commerce [13]. In such auctions, a buyer (center) needs to purchase products or services from a number of competing sellers, and the subsets of sellers that satisfy the buyer's requirements can be characterized combinatorially. A well-known example is provided by path auctions [17], where the buyer's aim is to obtain a path in a network whose edges are owned by selfish agents; other examples include minimum spanning tree auctions [14], bipartite matching auctions [3], and vertex cover auctions [2]. An important research goal in this setting is the minimization of the buyer's total payment. While most of the work on this topic focuses on dominant-strategy incentive compatible mechanisms (e.g., [14, 9, 7, 5]), the properties of Nash equilibria of first-price auctions have recently received a lot of attention as well [8, 6, 3].

An interesting — and, perhaps, counterintuitive — property of set system auctions is that the buyer can lower her total payment by prohibiting some of the agents from participating in the auction. In other words, reducing competition in the market can benefit the buyer. This has been observed for VCG mechanisms by Elkind [5] in the context of path auctions (see also [4]). Later, Chen and Karlin [3] discovered that this can also happen in first-price auctions for a variety of set systems. They labeled this

phenomenon “the cost of cheap labor”, and provided tight bounds on the cost of cheap labor in several set systems.

Prohibiting an agent from participating in an auction can be interpreted as requiring him to raise his bid to  $+\infty$ . The goal of this paper is to explore a more general approach, namely, allowing the center to place arbitrary lower bounds on all sellers’ bids, in a manner reminiscent of using reserve prices in combinatorial auctions. Clearly, this technique is more flexible than simply deleting agents, and hence the resulting savings, which we term “the refined cost of cheap labor”, may be even higher than the cost of cheap labor, as defined in [3]. In this paper, we study the benefits of this approach by quantifying the refined cost of cheap labor for a number of well-known set systems.

We start by providing general upper and lower bounds on the refined cost of cheap labor for arbitrary set systems (Section 3). We then consider several classes of set systems for which we can show that the refined cost of cheap labor and the cost of cheap labor coincide. These include matroids and (single) paths considered in [3], as well as a richer set system not considered in [3], namely,  $k$ -paths. For  $k$ -path auctions, we significantly extend the techniques of [3] to provide tight bounds on the (refined) cost of cheap labor.

We then move on to vertex cover set systems. In these set systems, deleting an agent creates a monopoly, and hence the cost of cheap labor is exactly 1. On the other hand, artificially inflating the agents’ bids may prove to be very profitable for the buyer: we show that there exist vertex cover auctions for which the refined cost of cheap labor is linear in the number of agents, matching the general upper bound of Section 3. Finally, we consider set systems that are based on dominating sets and perfect bipartite matchings. For such set systems, we show that both the cost of cheap labor and the refined cost of cheap labor can be quite large, and also that these two quantities can differ by a large factor. These set systems illustrate that setting lower bounds on the sellers’ bids is a very powerful — yet simple and practically applicable — technique. Thus, we believe that the refined cost of cheap labor is an important characteristic of a set system auction, which deserves further study.

## 2 Preliminaries

A *set system* is a pair  $(E, \mathcal{F})$ , where  $E$  is the *ground set* and  $\mathcal{F} \subseteq 2^E$  is a collection of *feasible* subsets of  $E$ . Throughout the paper, we only consider set systems with  $|E| < +\infty$  and set  $n = |E|$ . The set  $\mathcal{F}$  can be listed explicitly, or defined combinatorially. In this paper, we consider the following set systems:

- *spanning trees*: the set  $E$  is the set of all edges of a given graph  $G$  and  $\mathcal{F}$  consists of all sets  $S \subseteq E$  that contain a spanning tree. This is a special case of a more general *matroid* set system [12], in which the set  $E$  is the ground set of a given matroid  $M$ , and the set  $\mathcal{F}$  is the collection of all subsets of  $2^E$  that contain a base of  $M$ .
- *perfect bipartite matchings*: the set  $E$  is the set of all edges of a given bipartite graph  $G$  and  $\mathcal{F}$  consists of all sets  $S \subseteq E$  that contain a perfect bipartite matching.
- *$k$ -paths*: the set  $E$  is the set of all edges of a given network  $G$  with a source  $s$  and a sink  $t$ , and  $\mathcal{F}$  consists of all sets  $S \subseteq E$  that contain  $k$  edge-disjoint  $s$ - $t$  paths.



- *vertex covers*: the set  $E$  is the set of all vertices of a given graph  $G$ , and  $\mathcal{F}$  consists of all sets  $S \subseteq E$  that contain a vertex cover of  $G$ .
- *dominating sets*: the set  $E$  is the set of all vertices of a given graph  $G$ , and  $\mathcal{F}$  consists of all sets  $S \subseteq E$  that contain a dominating set of  $G$ , i.e., for each vertex  $v \notin S$ , there is  $u \in S$  such that there is an edge between  $u$  and  $v$ .

Observe that all set systems listed above are *upwards closed*, i.e.,  $S \in \mathcal{F}$  implies  $S' \in \mathcal{F}$  for any  $S' \supseteq S$ . A set system is said to be *monopoly-free* if  $\bigcap_{S \in \mathcal{F}} S = \emptyset$ . Throughout this paper, we restrict ourselves to upwards closed, monopoly-free set systems.

In a *set system auction* for a set system  $(E, \mathcal{F})$ , each  $e \in E$  is owned by a selfish agent, and there exists a center (auctioneer) who wants to purchase a feasible solution, i.e., an element of  $\mathcal{F}$ . Each agent  $e \in E$  has a cost  $c_e \geq 0$ , which is incurred if this element is used in the solution purchased by the center. We will refer to a triple  $(E, \mathcal{F}, \mathbf{c})$ , where  $\mathbf{c} = (c_e)_{e \in E}$  as a *market*. For any subset  $S \subseteq E$ , we write  $c(S)$  to denote  $\sum_{e \in S} c_e$ .

Throughout the paper, we assume that the sale is conducted by means of a first-price auction: each agent  $e$  announces his *bid*  $b_e$ , indicating how much he wants to be paid for the use of his element, the auctioneer selects the cheapest feasible set breaking ties in an arbitrary (but deterministic) way, and all agents in the winning set are paid their bid. Thus, the payoff of a winning agent  $e$  with bid  $b_e$  is  $b_e - c_e$ , whereas the payoff of any losing agent is 0. The agents are selfish, i.e., they aim to maximize their payoff. Therefore, we are interested in *Nash equilibria (NE)* of such auctions, i.e., vectors of bids  $\mathbf{b} = (b_e)_{e \in E}$  such that no agent  $e$  can increase his payoff by bidding  $b'_e \neq b_e$  as long as all other agents bid according to  $\mathbf{b}$ . We restrict ourselves to equilibria in which no agent bids below their cost, i.e.,  $b_e \geq c_e$  for all  $e \in E$ .

Unfortunately, as shown in [8], for some markets and some tie-breaking rules, first-price auctions may have no NE in pure strategies. However, they do have  $\varepsilon$ -*Nash equilibria* in pure strategies for any  $\varepsilon > 0$ , i.e., a bid vector such that no agent can unilaterally change his bid to increase the payoff by more than  $\varepsilon$ . Moreover, for any market  $(E, \mathcal{F}, \mathbf{c})$ , there exists a tie-breaking rule (e.g., one that favors the feasible set with the smallest cost) that ensures the existence of a pure NE. Thus, in what follows, we will ignore the issues of existence of pure NE, and use the term “Nash equilibrium” to refer to a bid vector that is a pure NE of a first-price auction for a given market under *some* tie-breaking rule, or, equivalently, can be obtained as a limit of  $\varepsilon$ -NE for that market as  $\varepsilon \rightarrow 0$ .

Following the approach of [3], we will focus on NE of set system auctions that are *buyer-optimal*, i.e., minimize the center’s total payment. For a given market  $(E, \mathcal{F}, \mathbf{c})$ , we denote the center’s total payment in such a NE with the smallest total payment by  $\nu(E, \mathcal{F}, \mathbf{c})$ .

This quantity is similar to—but different from—the quantity  $\nu_0$  that is used in [9] as a benchmark to measure the frugality of dominant-strategy set system auctions. Indeed, the latter can be interpreted as the minimal total payment in a buyer-optimal *efficient* NE of a first-price auction (i.e., a NE in which the winning set  $S$  satisfies  $S \in \operatorname{argmin}_{S \in \mathcal{F}} c(S)$ ), in which, in addition, all losing agents bid their cost.

### 3 Refined Cost of Cheap Labor: General Bounds

In this section, we introduce our new measure of the cost of cheap labor, which we will call the *refined cost of cheap labor*, and compare it to the notion of the cheap labor cost introduced in [3].

The following definition of cheap labor cost is adapted from [3].

**Definition 1.** Given a market  $(E, \mathcal{F}, \mathbf{c})$ , its cheap labor cost  $\delta_1(E, \mathcal{F}, \mathbf{c})$  is defined as follows:

$$\delta_1(E, \mathcal{F}, \mathbf{c}) = \max_{S \subseteq E} \frac{\nu(E, \mathcal{F}, \mathbf{c})}{\nu(S, \mathcal{F}[S], \mathbf{c}[S])},$$

where  $\mathcal{F}[S] = \{S' \in \mathcal{F} \mid S' \subseteq S\}$ , and  $\mathbf{c}[S] = (c_e)_{e \in S}$ . The cheap labor cost of a set system  $(E, \mathcal{F})$  is defined as  $\delta_1(E, \mathcal{F}) = \sup_{\mathbf{c}} \delta_1(E, \mathcal{F}, \mathbf{c})$ .

Informally,  $\delta_1(E, \mathcal{F})$  measures how much the center can save by removing some of the agents from the system. Alternatively, the center's actions can be interpreted as setting the costs of some agents to  $+\infty$  (or some appropriately large number). The notion of *refined cheap labor cost*, which we will now introduce, allows the center more flexibility, permitting him to raise the cost of any agent  $e \in E$  to any value between its cost  $c_e$  and  $+\infty$ .

**Definition 2.** Given a market  $(E, \mathcal{F}, \mathbf{c})$ , its refined cheap labor cost  $\delta_2(E, \mathcal{F}, \mathbf{c})$  is defined as follows:

$$\delta_2(E, \mathcal{F}, \mathbf{c}) = \sup_{\mathbf{c}' \succeq \mathbf{c}} \frac{\nu(E, \mathcal{F}, \mathbf{c}')}{\nu(E, \mathcal{F}, \mathbf{c}')},$$

where  $\mathbf{c}' \succeq \mathbf{c}$  means that  $c'_e \geq c_e$  for all  $e \in E$ . The refined cheap labor cost of a set system  $(E, \mathcal{F})$  is defined as  $\delta_2(E, \mathcal{F}) = \sup_{\mathbf{c}} \delta_2(E, \mathcal{F}, \mathbf{c})$ .

As argued above, Definition 1 can be obtained from Definition 2 by requiring that  $c'_e \in \{c_e, +\infty\}$  for all  $e \in E$ . The following theorem provides some simple bounds on  $\delta_1$  and  $\delta_2$ .

**Theorem 1.** Fix a market  $(E, \mathcal{F}, \mathbf{c})$ , and let  $S$  be a cheapest feasible solution in  $\mathcal{F}$  with respect to  $\mathbf{c}$ . Then the following inequalities hold:

$$1 \leq \delta_1(E, \mathcal{F}, \mathbf{c}) \leq \delta_2(E, \mathcal{F}, \mathbf{c}) \leq |S|.$$

In what follows, we present upper and lower bounds on  $\delta_2$  for specific set systems.

### 4 Spanning Trees and Other Matroids

For any spanning tree set system, artificially inflating the agents' costs cannot lower the center's payments, i.e.,  $\delta_1 = \delta_2 = 1$  (where  $\delta_1 = 1$  is shown in [3]). In fact, this result holds for the more general case of matroid set systems. We refer the readers to [12] for a formal definition of a matroid.

**Theorem 2.** For any matroid market  $\mathcal{M} = (E, \mathcal{F}, \mathbf{c})$  we have  $\delta_2(E, \mathcal{F}, \mathbf{c}) = 1$ .

## 5 Paths and $k$ -Paths

Throughout this section, for a given network  $G = (V, E)$  with a source  $s$  and a sink  $t$ , we denote by  $\mathcal{F}_k$  the collection of sets of edges that contain  $k$  edge-disjoint paths from  $s$  to  $t$ .

For  $k$ -paths set systems, it turns out that the optimal cost reduction can be achieved by simply deleting edges in  $E$ , i.e.,  $\delta_1 = \delta_2$ . Furthermore,  $\delta_2 = \delta_1 \leq k + 1$  for any network, and this bound is tight, i.e. for any  $k$  there is a  $k$ -path set system  $(E, \mathcal{F}_k)$  with  $\delta_1(E, \mathcal{F}_k) = \delta_2(E, \mathcal{F}_k) = k + 1$ . This generalizes the result of [3], which proves this claim for  $k = 1$ .

**Theorem 3.** *For any network  $G = (V, E)$  with a source  $s$  and a sink  $t$ , and any cost vector  $\mathbf{c}$ , we have  $\delta_2(E, \mathcal{F}_k, \mathbf{c}) = \delta_1(E, \mathcal{F}_k, \mathbf{c})$ .*

We also give a tight bound on the cost of cheap labor (and hence, by Theorem 3, a tight bound on the refined cost of cheap labor) in any  $k$ -paths set system.

**Theorem 4.** *For any network  $G = (V, E)$  with a source  $s$  and a sink  $t$ , and any cost vector  $\mathbf{c}$ , we have  $\delta_1(E, \mathcal{F}_k, \mathbf{c}) \leq k + 1$ , and this bound is tight.*

## 6 Vertex Covers

In this section, we consider vertex cover auctions. In these auctions, as well as in the auctions considered in Section 7, the sellers are the vertices. Therefore, in these two sections we depart from the standard graph-theoretic notation, and use  $E$  to denote the set of vertices of a graph  $G$ , and  $H$  to denote the set of edges of  $G$ . Also, we denote by  $\mathcal{F}$  the collection of all sets of vertices that contain a vertex cover (respectively, a dominating set) for  $G$ .

The vertex cover set systems demonstrate that  $\delta_1$  and  $\delta_2$  can be very different: for any such set system  $\delta_1 = 1$ , whereas  $\delta_2$  can be linear in  $|E|$ .

**Proposition 1.** *For any graph  $G = (E, H)$  and any costs  $\mathbf{c}$ , we have  $\delta_1(E, \mathcal{F}, \mathbf{c}) = 1$ .*

In contrast, we will now show that there is a graph  $G = (E, H)$  with  $|E| = n$  such that the corresponding set system  $(E, \mathcal{F})$  satisfies  $\delta_2(E, \mathcal{F}) = \Omega(n)$ .

**Proposition 2.** *There exists a graph  $G = (E, H)$  and a cost vector  $\mathbf{c}$  that satisfy  $\delta_2(E, \mathcal{F}, \mathbf{c}) \geq \frac{n-3}{2}$ , where  $n = |E|$ .*

*Proof.* Consider a graph  $G$  obtained from complete graph  $K_{n-2}$  by adding two new vertices  $u$  and  $u'$  and connecting them to two adjacent vertices  $v$  and  $v'$  of  $K_{n-2}$ , respectively (see Fig. 1). In addition, consider a cost vector  $\mathbf{c}$  given by  $c_v = c_{v'} = 1$ , and  $c_e = 0$  for  $e \neq v, v'$ .

For the cost vector  $\mathbf{c}$ , it can be seen that the buyer-optimal NE  $\mathbf{b}$  is  $b_u = b_{u'} = 0$ ,  $b_e = 1$  for  $e \neq u, u'$ . Thus,  $\nu(E, \mathcal{F}, \mathbf{c}) = n - 3$ . On the other hand, consider a cost vector  $\mathbf{c}' \succeq \mathbf{c}$  given by  $c'_v = c'_{v'} = c'_u = c'_{u'} = 1$  and  $c_e = 0$  for  $e \neq v, v', u, u'$ . It is easy to see that for this cost vector, the buyer-optimal NE  $\mathbf{b}'$  satisfies  $\mathbf{b}' = \mathbf{c}'$  and the winning set consists of all vertices of  $K_{n-2}$ . Hence,  $\nu(E, \mathcal{F}, \mathbf{c}') = 2$ , and we have  $\delta_2(E, \mathcal{F}, \mathbf{c}) \geq \frac{n-3}{2}$ .




Fig. 1. Vertex Cover




Fig. 2. Perfect Bipartite Matching

### 7 Dominating Sets

For dominating sets, note that deleting an agent that corresponds to a vertex  $e$  is not equivalent to deleting the vertex  $e$  itself from the graph:  $e$  still needs to be dominated, even though it cannot be a member of a feasible set.

For dominating sets,  $\delta_1$  does not necessarily equal  $\delta_2$ . Furthermore,  $\delta_1$  and  $\delta_2$  can be as large as  $\Omega(\sqrt{n})$ . We will now present two examples to illustrate this. Both examples are obtained by a modification of the construction used in the last section.

**Definition 3.** Given a complete graph  $K_n$ ,  $n \geq 3$ , let  $K'_n$  be the graph obtained from  $K_n$  by replacing each of its edges  $(v_i, v_j)$  by a pair of edges  $(v_i, w_{ij}), (w_{ij}, v_j)$ . Define  $W = \{w_{ij}\}_{i,j \in \{1, \dots, n\}}$  and  $V = \{v_i\}_{i \in \{1, \dots, n\}}$ .

**Proposition 3.** There exists a graph  $G = (E, H)$  and a cost vector  $\mathbf{c}$  that satisfy  $\delta_1(E, \mathcal{F}, \mathbf{c}) = 1$  and  $\delta_2(E, \mathcal{F}, \mathbf{c}) = \Omega(\sqrt{n})$ .

The graph  $G$  is constructed from  $K'_n$  by selecting two adjacent vertices  $v, v' \in V$  and adding three new vertices  $t, u, u'$  and  $n + 2$  new edges  $(u, v), (u', v'), (t, v)_{v \in V}$  (see Fig. 3 [left]). For cost vector  $\mathbf{c}$ , we set  $c_e = n^2$  for  $e \in W$ ,  $c_v = c_{v'} = 1$ ,  $c_u = c_{u'} = c_t = 0$ , and  $c_e = 0$  for  $e \in V \setminus \{v, v'\}$ .

**Proposition 4.** There is a graph  $G = (E, H)$  such that  $\delta_1(E, \mathcal{F}) = \Omega(\sqrt{n})$  and  $\delta_2(E, \mathcal{F}) = \Omega(\sqrt{n})$ .

The graph  $G$  is constructed from  $K'_n$  by selecting a vertex  $v \in V$  and adding three new vertices  $t, u, u'$  and  $n + 3$  new edges  $(u, v), (u', v), (u, u'), (t, v)_{v \in V}$  (see Fig. 3 [right]). For cost vector  $\mathbf{c}$ , we set  $c_e = n^2$  for  $e \in W$ ,  $c_v = 1$ ,  $c_u = c_{u'} = c_t = 0$ , and  $c_e = 0$  for  $e \in V \setminus \{v\}$ .

### 8 Perfect Bipartite Matchings

Perfect bipartite matching systems have a similar flavor to dominating set systems— $\delta_2$  can be very different from  $\delta_1$ , and both of them can be very large. For perfect matching




Fig. 3. Dominating Set (graph  $G$  [left] and  $G'$  [right] with  $n = 5$ )

in bipartite graphs, [3] shows that there is a graph  $G$  such that the corresponding set system satisfies  $\delta_1(E, \mathcal{F}) = \Omega(n)$ . As any bipartite matching in a graph with  $n$  edges has size  $O(n)$ , by Theorem 1 we have the following claim.

**Proposition 5.** *There is a graph  $G = (V, E)$  such that  $\delta_1(E, \mathcal{F}) = \Theta(n)$  and  $\delta_2(E, \mathcal{F}) = \Theta(n)$ , where  $n = |E|$ .*

Proposition 5 shows that in the worst case  $\delta_1$  and  $\delta_2$  coincide. However, they can also differ by a linear factor.

**Proposition 6.** *There is a graph  $G = (V, E)$  such that  $\delta_1(E, \mathcal{F}) = 1$  and  $\delta_2(E, \mathcal{F}) = \Omega(n)$ .*

*Proof.* Consider the graph shown in Fig. 2. For any cost vector  $\mathbf{c}$ , since we cannot delete any edge without creating a monopoly, we have  $\delta_1(E, \mathcal{F}, \mathbf{c}) = 1$ .

On the other hand, to see that  $\delta_2(E, \mathcal{F}) = \Omega(n)$ , consider a cost vector  $\mathbf{c}$  where  $c_{(u_i, u)} = 1$  for  $i = 3, \dots, n$ , and  $c_e = 0$  for any other edge  $e \in E$ . In any buyer-optimal Nash equilibrium  $\mathbf{b}$ , we have to set  $b_{(u_i, v_i)} = 1$  for  $i = 3, \dots, n$ , which implies that  $\nu(E, \mathcal{F}, \mathbf{c}) = n - 2$ . Consider another cost vector  $\mathbf{c}' \succeq \mathbf{c}$ , where  $c_{(u_i, u)} = 1$  for  $i = 1, \dots, n$  and  $c_e = 0$  for any other edge  $e \in E$ . It can be seen that  $\nu(E, \mathcal{F}, \mathbf{c}') = 1$ , and thus  $\delta_2(E, \mathcal{F}, \mathbf{c}) \geq n - 2$ .

### 9 Conclusions and Future Work

We have introduced the notion of refined cost of cheap labor for set system auctions, and analyzed it for several classes of set systems. A number of questions suggest themselves for further study. First, in this paper we largely ignored computational issues related to our problem, such as, e.g., computing the refined cost of cheap labor for a

given set system, or identifying an optimal or close-to-optimal modified cost vector  $\mathbf{c}'$ . We believe that this is a fruitful topic that deserves to be investigated further. Another promising research direction is bounding the ratio between  $\nu$  and  $\nu_0$ , i.e., the additional cost of requiring the winning set to be optimal with respect to the true costs; this quantity can be seen as “the cost of efficiency”. In particular, it would be interesting to see if the latter can be bounded in terms of the (refined) cost of cheap labor.

## Acknowledgements

This work was supported by a NAP grant, NRF grant NRF-RF2009-08, and A-STAR SINGA Graduate Scholarship (Singapore).

## References

1. Archer, A., Tardos, É.: Frugal Path Mechanisms. In: SODA 2002, pp. 991–999 (2002)
2. Calinescu, G.: Bounding the Payment of Approximate Truthful Mechanisms. In: Fleischer, R., Trippen, G. (eds.) ISAAC 2004. LNCS, vol. 3341, pp. 221–233. Springer, Heidelberg (2004)
3. Chen, N., Karlin, A.R.: Cheap Labor Can Be Expensive. In: SODA 2007, pp. 707–715 (2007)
4. Conitzer, V., Sandholm, T.: Failures of the VCG Mechanism in Combinatorial Auctions and Exchanges. In: AAMAS 2006, pp. 521–528 (2006)
5. Elkind, E.: True Costs of Cheap Labor are Hard to Measure: Edge Deletion and VCG Payments in Graphs. In: EC 2005, pp. 108–116 (2005)
6. Elkind, E., Goldberg, L.A., Goldberg, P.W.: Frugality Ratios and Improved Truthful Mechanisms for Vertex Cover. In: EC 2007, pp. 336–345 (2007)
7. Elkind, E., Sahai, A., Steiglitz, K.: Frugality in Path Auctions. In: SODA 2004, pp. 701–709 (2004)
8. Immorlica, N., Karger, D., Nikolova, E., Sami, R.: First-Price Path Auctions. In: EC 2005, pp. 203–212 (2005)
9. Karlin, A.R., Kempe, D., Tamir, T.: Beyond VCG: Frugality of Truthful Mechanisms. In: FOCS 2005, pp. 615–626 (2005)
10. Krishna, V.: Auction Theory. Academic Press, London (2002)
11. Nisan, N., Ronen, A.: Algorithmic Mechanism Design. In: STOC 1999, pp. 129–140 (1999)
12. Oxley, J.: Matroid Theory. Oxford University Press, Oxford (1992)
13. Sandholm, T.: Expressive Commerce and Its Application to Sourcing: How We Conducted \$35 Billion of Generalized Combinatorial Auctions. In: ICEC 2007, pp. 349–350 (2007)
14. Talwar, K.: The Price of Truth: Frugality in Truthful Mechanisms. In: Alt, H., Habib, M. (eds.) STACS 2003. LNCS, vol. 2607, pp. 608–619. Springer, Heidelberg (2003)

# On the Power of Mediators<sup>\*</sup>

Josep Díaz<sup>1</sup>, Dieter Mitsche<sup>1</sup>, Navin Rustagi<sup>1</sup>, and Jared Saia<sup>2</sup>

<sup>1</sup> Departament Llenguatges i Sistemes, Universitat Politècnica de Catalunya  
{diaz,dmitsche}@lsi.upc.edu

<sup>2</sup> Department of Computer Science, University of New Mexico  
{rustagi,saia}@cs.unm.edu

**Abstract.** We consider a problem at the intersection of distributed computing and game theory, namely: Is it possible to achieve the “windfall of malice” even without the actual presence of malicious players? Our answer to this question is “Yes and No”. Our positive result is that for the virus inoculation game, it is possible to achieve the windfall of malice by use of a mediator. Our negative result is that for symmetric congestion games that are known to have a windfall of malice, it is not possible to design a mediator that achieves this windfall. In proving these two results, we develop novel techniques for mediator design that we believe will be helpful for creating non-trivial mediators to improve social welfare in a large class of games.

## 1 Introduction

Recent results show that malicious players in a game may, counter-intuitively, improve social welfare [7,4,9]. For example, in [7] it is showed that for a virus inoculation game, the existence of malicious players will actually lead to better social welfare for the remaining players than if such malicious players are absent. This improvement in the social welfare with malicious players has been referred to as the “windfall of malice” [4]. The existence of the windfall of malice for some games leads to an intriguing question: Can we achieve the windfall of malice even without the actual presence of malicious players?

We show that the answer to the previous question is sometimes “Yes”. How do we achieve the beneficial impact of malicious players without their actual presence? Our approach is to use a mediator. Informally, a mediator is a trusted third party that suggests actions to each player. The players retain free will and can ignore the mediator’s suggestions. The mediator proposes actions privately to each player, but the algorithm the mediator uses to decide what to propose is public knowledge. In this paper, we introduce a general technique for designing mediators that is inspired by careful study of the “windfall of malice” effect. In our approach, the mediator makes a random choice of one of two possible

---

<sup>\*</sup> This research was partially supported by: for the 1st. 2nd authors, the FP7-ICT-21527 (FRONTS) and the TIN2007-66523 (FORMALISM), for the 3rd. and 4th. authors by NSF CAREER Award 0644058, NSF CCR-0313160, and an AFOSR MURI grant.

configurations, where a configuration is just a set of proposed actions for each player. The first configuration is optimal: the mediator proposes a set of actions that achieves the social optimum (or very close to it). The second configuration is “fear inducing”: the mediator proposes a set of actions that leads to catastrophic failure for those players who do not heed the mediators advice. The purpose of the second configuration is to ensure that the players follow the advice of the mediator when the optimal configuration is chosen. Thus, the random choice of which configuration is chosen must be hidden from the players. We show the applicability of our technique by using it to design a mediator for the virus inoculation game from [7] that achieves a social welfare that is asymptotically optimal.

We also show the limits of our technique by proving an impossibility result that shows that for a large class of games, no mediator will improve the social welfare over the best Nash equilibrium. In particular, this impossibility result holds for the congestion games that in [4] is shown to have a windfall of malice.

**Related Work.** The concept of a mediator is closely related to that of a correlated equilibrium, which was introduced by Aumann in [3]. In particular, if a mediator proposes actions to the players such that it is in the best interest of each player to follow the mediators proposal, then the mediator is said to implement a correlated equilibrium. There are several recent results on correlated equilibrium and mediators. The authors in [8] give polynomial time algorithms that can optimize over correlated equilibria, via a LP approach, for a large class of multiplayer games that are “succinctly representable”. Christodoulou et al. [6] study the price of anarchy and stability in congestion games where each edge has a linear cost function with positive coefficients. They show that in such a setting, the price of anarchy for pure equilibrium is almost the same as the price of anarchy of correlated equilibrium. Balcan et al. [5], describe techniques for moving from a high cost Nash equilibrium to a low cost Nash equilibrium via a “public service advertising campaign”. They show that in many games, even if not all players follow instructions, it is possible to ensure such a move. While their result does not explicitly consider mediators, it is similar in flavor to ours in the sense that an outside third party is acting to improve social welfare. Recent work by Abraham et al. [1] presents distributed algorithms that enable a group of players to implement a mediator, entirely through point-to-point communication, even when there is a constant fraction of adversarial players.

**Basic definitions and notation.** A *correlated equilibrium* is a probability distribution over strategy vectors that ensures that no player has incentive to deviate. We define a *configuration* for a given game to be a vector of pure strategies for that game, one for each player. We define a *mediator* for a game to be a probability distribution  $\mathcal{D}(\mathcal{C})$  over a finite set of different configurations  $\mathcal{C}$ . The set of configurations  $\mathcal{C}$  and the distribution  $\mathcal{D}(\mathcal{C})$  are known to all players. However, the actual configuration chosen is unknown, and the advice the mediator gives to a particular player based on the chosen configuration is known only to



that player. We say that a mediator is *valid* if all players are incentivized to follow its advice. In this case, the mediator implements a correlated equilibrium. From a distributed computing viewpoint, the major difference between a correlated equilibrium and a Nash equilibrium is that in a correlated equilibrium, players share a global coin, but in a Nash equilibrium, players only have access to private coins.

Throughout this paper, we will only consider mediators that treat all players equally, i.e., once having decided (by a random experiment according to  $\mathcal{D}(\mathcal{C})$ ) which is the configuration the mediator is choosing from, all players have the same probability to be proposed a particular strategy. Also, throughout the paper we assume that the number of strategic players,  $n$ , is very large (tending to infinity). Finally, we will use the notation  $a(n) \sim b(n)$  if  $a(n) = b(n)(1 \pm o(1))$ . We also use the notation  $[n] = \{1, \dots, n\}$ .

## 2 Virus Inoculation Game

We now describe the *virus inoculation game* from [2,7]. There are  $n$  players, each corresponding to a node in a square grid  $G$ . Each player has two choices: either to inoculate itself (at a cost of 1) or to do nothing and risk infection (which costs  $L$ ). After the decision of the nodes to inoculate or not, one node selected uniformly at random is infected with a virus. A node  $v$  that chooses not to inoculate gets infected by the virus if either the virus starts at  $v$  or the virus starts at another node  $v'$  and there is a path of not inoculated nodes connecting  $v$  and  $v'$ .

The *attack graph*  $G_a$  is the graph induced on  $G$  by the set of all nodes that do not inoculate. Aspnes et al. [2] proved that in a pure Nash equilibrium every component of the attack graph has size  $n/L$ . The social welfare achieved in such an equilibrium is thus  $\Theta(n)$ . Following Moscibroda et al. [7], we will focus on outcomes of the game on the grid. It is proved there that the minimum social welfare on the grid is  $\Theta(n^{2/3}L^{1/3})$ , which occurs when the components in  $G_a$  are of size  $(n/L)^{2/3}$ . This implies that the cost of anarchy for this game is large when  $L$  is large. However, Moscibroda et al. show that the existence of enough Byzantine players, who can never be trusted to inoculate, ensures that the social welfare of any Nash equilibrium is slightly better than  $\Theta(n)$ .

Based on the result from [7], we observe that the main problem in this game is that the individual players do not have enough fear of being infected. In particular, they are unable to achieve the optimal social welfare because they form connected components in  $G_a$  that are too large. Thus, we design a mediator that randomly chooses between two configurations (see Figure 1). The first configuration is optimal: all components in  $G_a$  are of size  $(n/L)^{2/3}$ . The second configuration is “fear inducing”: any node that does not inoculate in this configuration has probability about 1/2 of being infected. The only purpose of the second configuration is to ensure that the selfish players follow the advice of the mediator when the optimal configuration is chosen.

We now formally describe the mediator for this game.<sup>1</sup> The mediator will choose randomly between one of the following two configurations  $C_1$  and  $C_2$ .

**Configuration  $C_1$ :** The mediator proposes a pattern of inoculation ensuring that 1) each component in  $G_a$  is of size no more than  $(\frac{n}{L})^{2/3}$ ; 2) each node is advised to inoculate with equal probability; and 3) the probability that a fixed node is advised to inoculate is at most  $2(L/n)^{1/3}$ . It does this as follows.

1. The mediator chooses a random integer  $x$  uniformly in  $[0, (n/L)^{1/3} - 1]$ .
2. For every node  $v$  in row  $r$  and column  $c$ , if one of the following two conditions hold, the mediator proposes  $v$  to inoculate: 1)  $r \equiv x \pmod{(n/L)^{1/3}}$ ; or 2)  $c \equiv x \pmod{(n/L)^{1/3}}$ . Otherwise the mediator tells  $v$  not to inoculate.

**Configuration  $C_2$ :** The mediator proposes a pattern of inoculation such that 1) all nodes that do not inoculate are in one giant component in  $G_a$ ; 2) each node has equal probability of being chosen to inoculate; and 3) the probability that a fixed node is advised to inoculate is  $\frac{1}{2} - \frac{1}{2\sqrt{n}}$ . The mediator accomplishes this in the following manner:

1. The mediator flips a coin. If it comes up heads, it proposes that all nodes in even columns do not inoculate. If it comes up tails, it proposes that all nodes in odd columns do not inoculate.
2. The mediator chooses a random integer,  $x$ , uniformly in  $[1, \sqrt{n}]$ . For each of the columns that have not already been told not to inoculate, the mediator proposes that each node in that column inoculate except for the  $x$ -th node in that column.




Fig. 1. The  $\sqrt{n} \times \sqrt{n}$  grid with configurations  $C_1, C_2$

For these two configurations  $C_1$  and  $C_2$  we now define the probability distribution  $\mathcal{D}(\{C_1, C_2\})$  with  $p_1 = (1 - cL^{-2/3}n^{-1/3})$  and  $p_2 = cL^{-2/3}n^{-1/3}$ , where  $c > 0$  can be chosen to be any small constant satisfying  $c > 4L/(L - 2)$ . The following result shows that  $\mathcal{D}(\{C_1, C_2\})$  is asymptotically optimal.

<sup>1</sup> For ease of analysis, we assume that both  $\sqrt{n}$  and  $(\frac{n}{L})^{1/3}$  are integers. Also,  $\sqrt{n}$  should be an integer multiple of  $(\frac{n}{L})^{1/3}$  (this assumption can be removed easily without effecting our asymptotic results).

**Theorem 1.**  $\mathcal{D}(\{C_1, C_2\})$  is a mediator with social welfare  $\Theta(n^{2/3}L^{1/3})$ .

*Proof.* Define by  $\mathcal{E}_I^j$  the event that the mediator advises player  $j$  to inoculate and define by  $\mathcal{E}_{\bar{I}}^j$  the event that the mediator advises player  $j$  not to inoculate. Since all players are to be treated equally by the mediator, we omit  $j$ . Let  $\mathcal{E}_A$  be the event that a node gets infected by the virus, and denote by  $\mathcal{C}_A$  the infection cost. Denote by  $\mathcal{C}_I$  the cost of inoculation. We need to show that  $\mathcal{D}(\{C_1, C_2\})$  yields a mediator, i.e. that  $\mathbf{E}[\mathcal{C}_A|\mathcal{E}_I] \geq \mathbf{E}[\mathcal{C}_I|\mathcal{E}_I] = 1$  and  $\mathbf{E}[\mathcal{C}_A|\mathcal{E}_{\bar{I}}] \leq \mathbf{E}[\mathcal{C}_I|\mathcal{E}_{\bar{I}}] = 1$ , which is equivalent to showing that: (1)  $\Pr(\mathcal{E}_A|\mathcal{E}_I) \geq 1/L$  and (2)  $\Pr(\mathcal{E}_A|\mathcal{E}_{\bar{I}}) \leq 1/L$ .

Let  $\mathcal{E}_i, i = 1, 2$ , be the event that  $C_i, i = 1, 2$  is chosen. To prove (1), observe that

$$\Pr(\mathcal{E}_1|\mathcal{E}_I) = \Pr(\mathcal{E}_1, \mathcal{E}_I)/\Pr(\mathcal{E}_I) \sim \frac{p_1(2(L/n)^{1/3})}{p_1(2(L/n)^{1/3}) + p_2(1/2 - 1/(2\sqrt{n}))},$$

and similarly for  $\Pr(\mathcal{E}_2|\mathcal{E}_I)$ . Now, plugging in the values of  $p_1, p_2$  and using that  $L \in o(n)$  we get<sup>2</sup>

$$\begin{aligned} \Pr(\mathcal{E}_A|\mathcal{E}_I) &= \Pr(\mathcal{E}_A, \mathcal{E}_1|\mathcal{E}_I) + \Pr(\mathcal{E}_A, \mathcal{E}_2|\mathcal{E}_I) \\ &= \Pr(\mathcal{E}_A|\mathcal{E}_1, \mathcal{E}_I)\Pr(\mathcal{E}_1|\mathcal{E}_I) + \Pr(\mathcal{E}_A|\mathcal{E}_2, \mathcal{E}_I)\Pr(\mathcal{E}_2|\mathcal{E}_I) \\ &\sim \frac{2}{L^{2/3}n^{1/3}}\Pr(\mathcal{E}_1|\mathcal{E}_I) + \frac{1}{2}\Pr(\mathcal{E}_2|\mathcal{E}_I) \sim \frac{c}{2c + 4L}, \end{aligned}$$

which is greater than  $1/L$  for  $c > (4L)/(L - 2)$ . Reasoning in a similar way we get,

$$\Pr(\mathcal{E}_A|\mathcal{E}_{\bar{I}}) \sim \frac{c + 2}{2L^{2/3}n^{1/3}},$$

which is smaller than  $1/L$  since  $L \in o(n)$ , so we proved (2). To compute the social cost for this mediator, let  $\mathcal{I}_1$  ( $\bar{\mathcal{I}}_1$ ) be the set of nodes that inoculate (respectively do not inoculate) in  $C_1$ , and let  $\mathcal{I}_2$  ( $\bar{\mathcal{I}}_2$ ) be the set of nodes that inoculate (respectively do not inoculate) in  $C_2$ . Then the social cost for the mediator can be written as

$$p_1(|\mathcal{I}_1| + \sum_{v \in \bar{\mathcal{I}}_1} L\Pr(\mathcal{E}_A|\mathcal{E}_1, \mathcal{E}_{\bar{I}})) + p_2(|\mathcal{I}_2| + \sum_{v \in \bar{\mathcal{I}}_2} L\Pr(\mathcal{E}_A|\mathcal{E}_2, \mathcal{E}_{\bar{I}})) \sim \Theta(n^{2/3}L^{1/3}).$$

### 3 Impossibility Result

In light of the results in the previous section, a natural question is: Is it possible to design a mediator that will always improve the social welfare in any game for which there is a windfall of malice? Unfortunately, the answer to this question is “No”, as we show in this section. In particular, we show that the congestion games which Babaioff, Kleinberg and Papadimitriou [4] have proven have a windfall of malice effect do not admit a mediator that is able to improve the social welfare. In fact, we prove a stronger impossibility result, showing that

<sup>2</sup> If  $L = \theta(n)$ , then any pure Nash equilibrium is trivially asymptotically optimal.

for any non-atomic, symmetric congestion game where the cost of a path never decreases as a function of the flow through that path (of which class of games, the examples in [4] are special instances), no mediator can improve the social optimum.

A non-atomic, symmetric *congestion game* (henceforth, simply a congestion game) is a specified by a set of  $n \rightarrow \infty$  players; a set of  $E$  facilities (or edges);  $A \subset 2^E$  actions (or paths); and finally, for each facility  $e$  a cost function  $f_e$  associated with that facility. A pure strategy profile  $\mathcal{A} = (A_1, \dots, A_n)$  is a vector of actions, one for each player. The cost of player  $i$  for action profile  $\mathcal{A}$  is given by  $F_i(\mathcal{A}) = \sum_{e \in A_i} f_e(x_e(\mathcal{A}))$  where  $x_e(\mathcal{A})$  is the fraction of players using  $e$  in  $\mathcal{A}$ . As in [4], we assume that the game is *non-atomic*: since  $n \rightarrow \infty$  the contribution of a single player to the flow over a facility is negligible; and *symmetric*: all players have the same cost functions.

For an action  $a$  and a flow  $x \in [0, 1]$ , let  $\mathcal{F}_h(a, x)$  be the maximum possible cost of following action  $a$  when the total fraction of players following this action is  $x$ , where the maximum is taken over all ways that the remaining flow of  $1 - x$  can be distributed over other actions. Similarly, let  $\mathcal{F}_\ell(a, x)$  be the *minimum* cost of following action  $a$  when the total fraction of players following this action is  $x$ .

The following theorem says that for congestion games where the cost function of every action is non-decreasing in the fraction of players performing that action, the coordination between the agents in order to establish a correlated equilibrium will not decrease the social cost.

**Theorem 2.** *Consider a non-atomic, anonymous congestion game. If for all  $a \in A$  and  $0 \leq x \leq x' \leq 1$ ,  $\mathcal{F}_h(a, x) \leq \mathcal{F}_\ell(a, x')$  then the smallest social cost achieved by a correlated equilibrium is no less than the smallest social cost achieved by a Nash equilibrium.*

We next give a high level sketch of how we prove this theorem. We will fix a non-atomic, anonymous congestion game  $G$  with  $q$  actions,  $a_1, \dots, a_q$ , and  $n$  players. We define a *configuration*,  $C$ , for such a game to be a partitioning of the set of players across the  $q$  actions. We note that the number of possible configurations is finite; in particular,  $q^n$ . We next fix a mediator,  $M$ , for this game. We assume the mediator uses  $\ell$  different configurations  $C_1, \dots, C_\ell$ ; that  $0 \leq x_{i,j} \leq 1$  is the fraction of the players in configuration  $C_j$  assigned to action  $a_i$ ; and that  $c_{i,j} \in \mathbb{R}$  is the cost in configuration  $C_j$  for action  $a_i$ . We further assume that for all  $j \in [\ell]$ ,  $p_j$  is the probability with which the mediator  $M$  chooses  $C_j$ .

For any two actions  $a, a'$  we define the *a posteriori cost* of  $a$  given  $a'$  as the expected cost for a player of performing action  $a$  when action  $a'$  is suggested by the mediator  $M$ ; formally,  $\text{POST}(a, a') = \mathbf{E}[C_a | \mathcal{E}_{a'}]$ , where  $C_a$  is a random variable (over the configuration chosen by the mediator) and  $\mathcal{E}_{a'}$  is the event that action  $a'$  is recommended by the mediator. We define the *a priori cost* of action  $a$  as the cost of a player completely ignoring what the mediator suggests and always performing action  $a$ ; formally,  $\text{PRI}(a) := \sum_{j=1}^{\ell} p_j c_{i,j}$ .

The sketch behind our proof for this theorem is as follows. First, we show in Lemma 1 that for all actions  $a$ , if the cost of  $a$  is non-decreasing in the flow

through  $a$ , then  $\text{POST}(a, a) \geq \text{PRI}(a)$ . This is done by repeated decompositions of terms in summations for the a priori and posterior costs. Next, let  $Y$  be the cost of a player following the advice of the mediator, and let  $X$  be the cost of the player if she ignores the advice of the mediator and always chooses the action  $a$  that minimized  $\text{PRI}(a)$ . In Lemma 2 we show that  $E(Y) \leq E(X)$ . This lemma is shown by summing up inequality constraints on the mediator. Finally, we use these two lemmas to show the main theorem by showing that if Lemma 1 holds, then  $E(Y) > E(X)$ . The main technical challenge is the fact that we must show that  $E(Y) > E(X)$  even though Lemma 1 does not necessarily give a strict inequality. We address this problem by a subtle case analysis in the proof of the main theorem, and by augmenting Lemma 1 to show that in some cases, the inequality it implies is strict.

We now sketch the main points and defer the full proof to the full version of our paper. Observe that the condition for all  $a \in A$  and  $0 \leq x \leq x' \leq 1$ ,  $\mathcal{F}_h(a, x) \leq \mathcal{F}_\ell(a, x')$  implies that for all  $i \in [m], \forall j, k \in [\ell]$  we have that  $x_{ij} \leq x_{ik}$  implies  $c_{ij} \leq c_{ik}$ , and so the conditions of the following lemma are satisfied.

**Lemma 1.** *Given  $\ell \geq 2$  configurations  $C_1, \dots, C_\ell$ , with corresponding probabilities  $p_r > 0, r \in [\ell]$ . If for  $i \in [m], \forall j, k \in [\ell]$  we have that  $x_{ij} \leq x_{ik}$  implies  $c_{ij} \leq c_{ik}$ , then  $\text{POST}(a_i, a_i) \geq \text{PRI}(a_i)$ . Moreover, if for any  $i \in [q]$ , not all  $c_{ij}, j \in [\ell]$  are the same, then  $\text{POST}(a_i, a_i) > \text{PRI}(a_i)$ .*

Define by  $a_{pri} := \text{argmin}_a \text{PRI}(a)$ . Given a mediator over a fixed set of configurations, let  $X$  be the random variable denoting the cost of an arbitrary player when he decides to use action  $a_{pri}$ , i.e.,  $\mathbf{E}[X] = \sum_{j=1}^\ell p_j c_{a_{pri}, j}$ . Let  $Y$  be a random giving the cost of a player that follows the advice of the mediator, i.e.,  $\mathbf{E}[Y] = \sum_{i=1}^m \text{POST}(a_i, a_i) \mathbf{Pr}(\mathcal{E}_i) = \sum_{i=1}^m \sum_{j=1}^\ell p_j x_{ij} c_{ij}$ . In the following lemma we give the relationship between  $Y$  and  $X$ . The proof is based on summing the  $q-1$  inequalities resulting from constraints for a correlated equilibrium, and again it is deferred to the full version.

**Lemma 2.** *For any mediator we have  $\mathbf{E}[Y] \leq \mathbf{E}[X]$ .*

To prove the main theorem, denote by  $a_{post} := \text{argmin}_s \text{POST}(s, s)$  the action with minimum a posteriori cost. We will consider two cases.

Case 1: Not all actions have the same a posteriori cost. Using Lemma 2.2:

$$\mathbf{E}[Y] > \text{POST}(a_{post}, a_{post}) \geq \text{PRI}(a_{pri}) = \mathbf{E}[X].$$

Case 2: All actions have the same a posteriori cost. Assume it is not true that there is an action that does not have equal costs in each configuration. Then the cost of each action is the same in every configuration, and so any particular configuration must be a Nash equilibrium that achieves social cost equal to the social cost of the correlated equilibrium. Thus, we let  $a_x$  be some action that does not have the same cost in all configurations. Then using Lemma 3.2:

$$\mathbf{E}[Y] = \text{POST}(a_x, a_x) \geq \text{PRI}(a_{pri}) = \mathbf{E}[X].$$

In both cases we have  $\mathbf{E}[Y] > \mathbf{E}[X]$ . This however contradicts Lemma 2, hence there can not exist a correlated equilibrium achieving social cost less than the optimal Nash equilibrium.

## 4 Conclusion

We have shown that a mediator can improve the social welfare in some strategic games with a positive windfall of malice. Several open questions remain including the following. First, can we determine necessary and sufficient conditions for a game to allow a mediator that improves social welfare over the best Nash? In particular, can we find such conditions for general congestion games? What about arbitrary anonymous games? Second, for games where each player can choose among  $k$  actions, can we say how many configurations are needed by any mediator? Preliminary work in this direction shows that for 2 actions, sometimes more than 2 configurations are needed. Finally, can we use approaches similar to those in this paper for designing mediators for multi-round games? We have already made some preliminary progress in this direction for multi-round games where the number of rounds is determined by a geometric random variable.

## References

1. Abraham, I., Dolev, D., Gonen, R., Halper, J.: Distributed Computing Meets Game Theory: Robust mechanisms for rational secret sharing and multiparty computation. In: Principles of Distributed Computing, PODC (2006)
2. Aspnes, J., Chang, K., Yampolskiy, A.: Inoculation strategies for victims of viruses and the sum-of-squares partition problem. In: Symposium on Discrete Algorithms, SODA (2005)
3. Aumann, R.: Subjectivity and correlation in randomized games. *Mathematical Economics* 1, 67–96 (1974)
4. Babaioff, M., Kleinberg, R., Papadimitriou, C.H.: Congestion games with malicious players. In: ACM Conference on Electronic Commerce (2007)
5. Balcan, M., Blum, A., Mansour, Y.: Improved equilibria via public service advertising. In: ACM Symposium on Discrete Algorithms, SODA (2009)
6. 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)
7. Moscibroda, T., Schmid, S., Wattenhofer, R.: When selfish meets evil: Byzantine players in a virus inoculation game. In: Principles of Distributed Computing, PODC (2006)
8. Papadimitriou, C., Roughgarden, T.: Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)* 5(14) (2008)
9. Roth, A.: The price of malice in linear congestion games. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 118–125. Springer, Heidelberg (2008)

# Strong Nash Equilibria in Games with the Lexicographical Improvement Property

Tobias Harks\*, Max Klimm, and Rolf H. Möhring

Institut für Mathematik, Technische Universität Berlin  
{harks,klimm,moehring}@math.tu-berlin.de

**Abstract.** We provide an axiomatic framework for the the well studied lexicographical improvement property and derive new results on the existence of strong Nash equilibria for a very general class of congestion games with bottleneck objectives. This includes extensions of classical load-based models, routing games with splittable demands, scheduling games with malleable jobs, and more.

## 1 Introduction

A main criticism about Nash equilibria is the fact that they do not consider deviations of coalitions. To cope with the issue of coordination, we adopt the solution concept of a *strong equilibrium* (SNE for short) proposed in [2]. In a SNE, no coalition (of any size) can deviate and strictly improve the private cost of each of its members. Even though a SNE may rarely exist, it forms a very robust and appealing stability concept.

One of the most successful approaches in establishing existence of PNE (as well as SNE) is the potential function approach initiated in [12]: one defines a real-valued function  $P$  on the set of strategy profiles of the game and shows that every improving move of a coalition strictly reduces the value of  $P$ . Given that the set of strategy profiles is finite, every sequence of improving moves reaches a SNE. In particular, the global minimum of  $P$  is a SNE. For most games, however, it is hard to prove or disprove the existence of such a potential function.

In this paper, we introduce vector-valued potential functions and say that a game has the lexicographic improvement property (LIP) if there is a vector-valued function that lexicographically decreases for every improving move. A game has the  $\pi$ -LIP if the vector of private cost itself constitutes such a vector-valued potential function.

The main contribution of this paper is twofold. We first study desirable properties of arbitrary finite and infinite games having the LIP and  $\pi$ -LIP, respectively. These properties concern the existence of SNE, efficiency and fairness of SNE, and computability of SNE. Secondly, we identify an important class of games that we term bottleneck congestion games for which we can actually prove the  $\pi$ -LIP and, hence, prove that these games possess SNE with the above desirable properties.

Before we outline our results in more detail, we briefly explain the importance of bottleneck objectives in congestion games with respect to real-world applications. Referring to previous work by Keshav, it has been pointed out in [5] that the performance of

---

\* Research supported by the Federal Ministry of Education and Research (BMBF grant 03MOPAI1).

a communication network is closely related to the performance of its bottleneck (most congested) link. This behavior is also stressed by [3], who investigated PNE in routing games with bottleneck objectives. Similar observations are reported in [14], where the applicability of selfish routing models to realistic models of the Internet is investigated.

**Our Results.** We characterize games having the LIP by means of the existence of a generalized strong ordinal potential function. The proof of this characterization is constructive, that is, given a game  $G$  having the LIP for a function  $\phi$ , we explicitly construct a generalized strong ordinal potential  $P$ . We then investigate games having the  $\pi$ -LIP with respect to efficiency and fairness of SNE. Our characterization implies that there are SNE satisfying various efficiency and fairness properties, e.g., bounds on the prices of stability and anarchy, Pareto optimality, and min-max fairness.

We establish that bottleneck congestion games possess the  $\pi$ -LIP and thus possess SNE with the above mentioned properties. Moreover, our characterization of games having the LIP implies that bottleneck congestion games possess the strong finite improvement property.

In contrast to most congestion games considered so far, we require only that the cost functions on the facilities satisfy three properties: "non-negativity", "independence of irrelevant choices", and "monotonicity". Roughly speaking, the second and third conditions assume that the cost of a facility solely depends on the *set* of players using the respective facility and that this cost decreases if some players leave this facility. Thus, this framework extends classical *load-based* models in which the cost of a facility depends on the number or total weight of players using the respective facility.

We then study the LIP in *infinite* games, that is, games with infinite strategy spaces that can be described by compact subsets of  $\mathbb{R}^p$ ,  $p \in \mathbb{N}$ . Infinite games need not admit a strong potential function even if the LIP is satisfied. We prove, however, that continuity of  $\phi$  in the definition of LIP is sufficient for the existence of SNE. Our existence proof (which can be found in [9]) is constructive, that is, we outline an algorithm whose output is a SNE.

We then introduce *infinite bottleneck congestion games*. An infinite bottleneck congestion game arises from a bottleneck congestion game  $G$  by allowing players to fractionally distribute a certain *demand* over the pure strategies of  $G$ . We prove that these games have the  $\pi$ -LIP provided that the cost functions on the facilities are non-negative and non-decreasing. It turns out, however, that the function  $\pi$  may be discontinuous on the strategy space (even if the cost functions on the facilities are continuous). Thus, the existence of SNE does not immediately follow. We solve this difficulty by generalizing the LIP. As a consequence, we obtain for the first time the existence of SNE for infinite bottleneck congestion games with non-decreasing and continuous cost functions.

In the final section, we show that the methods presented here also apply to a more general framework.

Because of lack of space, the results of this work are presented without proofs. For the proofs, examples and further explanatory notes, we refer to [9].

**Related Work.** Congestion games were introduced in [16] and further studied in [12]. The existence of SNE in congestion games with monotone increasing cost functions has been studied in [10]. They showed that SNE need not exist in such games and gave



a structural characterization of the strategy space for symmetric (and quasi-symmetric) congestion games that admit SNE. Based on the work in [12], they also introduced the concept of a *strong potential function*, that is, a function on the set of strategy profiles that decreases for every profitable deviation of a coalition. The existence of (correlated) SNE in congestion games with non-increasing cost functions is explored in [18].

Several authors studied the existence and efficiency (price of anarchy and stability) of PNE and SNE in various specific classes of congestion games. Referring to an observation of Mehlhorn, a lexicographic argument was used first in [8] in order to establish the existence of PNE in singleton congestion games. Similar arguments have been applied to job scheduling games in [6] and to more general job scheduling games, where the processing time of a machine may depend on the set of jobs scheduled on the respective machine in [7]. Moreover, this line of argumentation has been used in [1] to prove even the existence of SNE in scheduling games on unrelated machines. They further studied differences between PNE and SNE and derived bounds on the (strong) price of anarchy and stability, respectively.

Bottleneck congestion games with network structure have been considered by Banner and Orda [3]. They studied existence of PNE in the unsplittable flow and in the splittable flow setting, respectively. They observed that standard techniques (such as Kakutani's fixed-point theorem) for proving existence of PNE do not apply to bottleneck routing games as the private cost functions may be discontinuous. They proved existence of PNE (but not SNE) by showing that bottleneck games are better reply secure, quasi-convex, and compact. Under these conditions, they recall Reny's existence theorem [15] for better reply secure games with possibly discontinuous private cost functions. In contrast, we show the existence of SNE with direct and constructive methods.

## 2 Preliminaries

We consider strategic games  $G = (N, X, \pi)$ , where  $N = \{1, \dots, n\}$  is the non-empty and finite set of players,  $X = \prod_{i \in N} X_i$  is the non-empty strategy space, and  $\pi : X \rightarrow \mathbb{R}_+^n$  is the combined *private cost* function that assigns a private cost vector  $\pi(x)$  to each strategy profile  $x \in X$ . These games are cost minimization games and we assume additionally that the private cost functions are non-negative. A strategic game is called *finite* if  $X$  is finite. We use standard game theory notation; for a coalition  $S \subseteq N$  we denote by  $-S$  its complement and by  $X_S = \prod_{i \in S} X_i$  we denote the set of strategy profiles of players in  $S$ .

A pair  $(x, (y_S, x_{-S})) \in X \times X$  is called an *improving move* if  $\pi_i(x_S, x_{-S}) - \pi_i(y_S, x_{-S}) > 0$  for all  $i \in S$ . We denote by  $I(S)$  the set of improving moves of coalition  $S$  and we set  $I := \bigcup_{S \subseteq N} I(S)$ . We call a sequence of strategy profiles  $\gamma = (x^0, x^1, \dots)$  an *improvement path* if every tuple  $(x^k, x^{k+1}) \in I$ . A strategy profile  $x$  is a *strong Nash equilibrium* (SNE) if  $(x, (y_S, x_{-S})) \notin I$  for all  $\emptyset \neq S \subseteq N$  and  $y_S \in X_S$ .

In recent years, much attention has been devoted to games admitting the finite improvement property (FIP), that is, each path of single-handed (one player) deviations is finite. Equivalently, we say that  $G$  has the strong finite improvement property (SFIP) if every improvement path is finite. A necessary and sufficient condition for the SFIP is the existence of a generalized *strong* ordinal potential function, that is, a function  $P : X \rightarrow \mathbb{R}$  such that  $P(x) - P(y) > 0$  for all  $(x, y) \in I$ .

It is known that both the SFIP and the existence of a generalized strong ordinal potential are hard to prove or disprove for a particular game. We define a class of games that we call *games with the Lexicographical Improvement Property (LIP)* and show that such games possess a generalized strong ordinal potential.

**Definition 1 (Sorted lexicographical order).** Let  $a, b \in \mathbb{R}_+^q$  and denote by  $\tilde{a}, \tilde{b} \in \mathbb{R}_+^q$  be the sorted vectors derived from  $a, b$  by permuting the entries in non-increasing order. Then,  $a$  is strictly sorted lexicographically smaller than  $b$  (written  $a < b$ ) if there exists an index  $m$  such that  $\tilde{a}_i = \tilde{b}_i$  for all  $i < m$ , and  $\tilde{a}_m < \tilde{b}_m$ .

**Definition 2 (Lexicographical improvement property,  $\pi$ -LIP).** A finite strategic game  $G = (N, X, \pi)$  possesses the lexicographical improvement property (LIP) if there exist  $q \in \mathbb{N}$  and a function  $\phi : X \rightarrow \mathbb{R}_+^q$  such that  $\phi(x) > \phi(y)$  for all  $(x, y) \in I$ .  $G$  has the  $\pi$ -LIP if  $G$  has the LIP for  $\phi = \pi$ .

The function  $\phi$  is a generalized strong ordinal potential if  $q = 1$ . Taking the  $M$ -norm of  $\phi$  for a sufficiently large  $M$  it is easy to verify that the LIP is equivalent to the existence of a generalized strong ordinal potential, regardless of  $q$ .

**Theorem 1.** Let  $G = (N, X, \pi)$  be a finite strategic game. Then,  $G$  has the LIP if and only if there exists  $\phi : X \rightarrow \mathbb{R}_+^q$ ,  $q \in \mathbb{N}$ , and  $M \in \mathbb{N}$  such that  $P(x) = \sum_{i=1}^q \phi_i(x)^M$  is a generalized strong ordinal potential function for  $G$ .

### 3 Properties of SNE in Games with the $\pi$ -LIP

As the existence of SNE in games with the LIP is guaranteed, it is natural to ask which properties these SNE may satisfy. In recent years, several notions of efficiency have been discussed in the literature, see [11]. We here cover the price of stability, Pareto optimality and min-max-fairness.

**Price of Stability.** We study the efficiency of SNE with respect to the optimum of a predefined social cost function. Given a game  $G = (N, X, \pi)$  and a social cost function  $C : X \rightarrow \mathbb{R}_+$ , whose minimum is attained in a strategy profile  $y \in X$ , let  $X^{\text{SNE}} \subseteq X$  denote the set of strong Nash equilibria. Then, the strong price of stability for  $G$  with respect to  $C$  is defined as  $\inf_{x \in X^{\text{SNE}}} C(x)/C(y)$ . We consider the following natural social cost functions: the sum-objective or  $L_1$ -norm defined as  $L_1(x) = \sum_{i \in N} \pi_i(x)$ , the  $L_p$ -objective or  $L_p$ -norm,  $p \in \mathbb{N}$ , defined as  $L_p(x) = (\sum_{i \in N} \pi_i(x)^p)^{1/p}$ , and the min-max objective or  $L_\infty$ -norm defined as  $L_\infty(x) = \max_{i \in N} \{\pi_i(x)\}$ .

**Proposition 1.** Let  $G$  be a strategic game with the  $\pi$ -LIP. Then, the strong price of stability w.r.t.  $L_\infty$  is 1, and for any  $p \in \mathbb{R}$ , the strong price of stability w.r.t.  $L_p$  is smaller than  $n$ .

**Pareto Optimality.** Pareto optimality is one of the fundamental concepts studied in economics, see Osborne and Rubinstein [13]. For a strategic game  $G = (N, X, \pi)$ , a strategy profile  $x$  is called *weakly Pareto efficient* if there is no  $y \in X$  such that  $\pi_i(y) < \pi_i(x)$  for all  $i \in N$ . A strategy profile  $x$  is *strictly Pareto efficient* if there is no  $y \in X$  such that  $\pi_i(y) \leq \pi_i(x)$  for all  $i \in N$ , where at least one inequality is strict.

So strictly Pareto efficient strategy profiles are those strategy profiles for which every improvement of a coalition of players is to the expense of at least one player outside the coalition. Pareto optimality has also been studied in the context of congestion games, see Chien and Sinclair [4] and Holzman and Law-Yone [10].

**Proposition 2.** *Let  $G$  be a finite strategic game with the  $\pi$ -LIP. Then, there is a SNE that is strictly Pareto optimal.*

**Min-Max-Fairness.** We next define the notion of min-max fairness, which is a central topic in resource allocation in communication networks, see Srikant [19] for an overview and pointers to the large body of research in this area. While strict Pareto efficiency requires that there is no improvement to the expense of anyone, the notion of min-max-fairness is stricter. Here, it is only required that there is no improvement at the cost of someone who receives already higher costs (while an improvement that increases the cost of a player with smaller original cost is allowed). It is easy to see that every min-max-fair strategy profile constitutes a strict Pareto optimum, but the converse need not hold. A strategy profile  $x$  is called *min-max fair* if for any other strategy profile  $y$  with  $\pi_i(y) < \pi_i(x)$  for some  $i \in N$ , there exists  $j \in N$  such that  $\pi_j(x) \geq \pi_j(y)$  and  $\pi_j(y) > \pi_j(x)$ .

**Proposition 3.** *Let  $G$  be a finite strategic game with the  $\pi$ -LIP. Then, there is a SNE that is min-max fair.*

## 4 Bottleneck Congestion Games

We now present a rich class of games satisfying the  $\pi$ -LIP. We call these games *bottleneck congestion games*. They are natural generalizations of variants of congestion games. In contrast to standard congestion games, we focus on *makespan-objectives*, that is, the cost of a player when using a set of facilities only depends on the highest cost of these facilities.

**Definition 3 (Congestion model).** *A tuple  $\mathcal{M} = (N, F, X, (c_f)_{f \in F})$  is called a congestion model if  $N = \{1, \dots, n\}$  is a non-empty, finite set of players,  $F = \{1, \dots, m\}$  is a non-empty, finite set of facilities, and  $X = \prod_{i \in N} X_i$  is the set of strategies. For each player  $i \in N$ , her collection of pure strategies  $X_i$  is a non-empty, finite set of subsets of  $F$ . Given a strategy profile  $x$ , we define  $N_f(x) = \{i \in N : f \in x_i\}$  for all  $f \in F$ . Every facility  $f \in F$  has a cost function  $c_f : \prod_{i \in N} X_i \rightarrow \mathbb{R}_+$  satisfying*

**Non-negativity:**  $c_f(x) \geq 0$  for all  $x \in X$

**Independence of Irrelevant Choices:**  $c_f(x) = c_f(y)$  for all  $x, y \in X$  with  $N_f(x) = N_f(y)$

**Monotonicity:**  $c_f(x) \leq c_f(y)$  for all  $x, y \in X$  with  $N_f(x) \subseteq N_f(y)$ .

Bottleneck congestion games generalize congestion games in the definition of the cost functions on the facilities. For bottleneck congestion games, the only requirements are that the cost  $c_f(x)$  on facility  $f$  for strategy profile  $x$  only depends on the set of players using  $f$  in their strategy profile and that costs are increasing with larger sets.

**Definition 4 (Bottleneck congestion game).** Let  $\mathcal{M}$  be a congestion model. The corresponding bottleneck congestion game is the strategic game  $G(\mathcal{M}) = (N, X, \pi)$  in which  $\pi$  is defined as  $\pi = \times_{i \in N} \pi_i$  and  $\pi_i(x) = \max_{f \in x_i} c_f(x)$ .

We are now ready to state our main result concerning bottleneck congestion games, providing a large class of games that satisfies the  $\pi$ -LIP.

**Theorem 2.** Let  $G(\mathcal{M})$  be a bottleneck congestion game with allocation model  $\mathcal{M}$ . Then,  $G$  fulfills the LIP for the functions  $\phi : X \rightarrow \mathbb{R}_+^n$  and  $\psi : X \rightarrow \mathbb{R}_+^{mn}$  defined as

$$\phi_i(x) = \pi_i(x) \quad \text{for all } i \in N, \quad \psi_{i,f}(x) = \begin{cases} c_f(x) & \text{if } f \in x_i \\ 0 & \text{else} \end{cases} \quad \text{for all } i \in N, f \in F.$$

As a corollary of Theorem 2 we obtain that each bottleneck congestion game has the  $\pi$ -LIP and hence possesses the SFIP. In addition, the results on the price of stability, Pareto optimality and min-max-fairness apply.

The class of bottleneck congestion games contains some special cases of particular interest. For example, *scheduling games* on unrelated machines can be seen as bottleneck congestion games in which all strategies are singletons. As the private costs received from each machine may depend on the *set* of players on that machine we cover certain *interference games* as well, see [9] for details. Another special case of bottleneck congestion games are bottleneck routing games. To the best of our knowledge, this work establishes for the first time the existence of the FIP and SFIP in such games. In the full version of this paper [9], we identify some special cases of bottleneck routing games in which the SNE can be computed in polynomial time.

## 5 Infinite Bottleneck Congestion Games

From  $\mathcal{M}$  we derive an *infinite congestion model*  $\mathcal{IM} = (N, F, X, d, \Delta, (c_f)_{f \in F})$ , where  $d \in \mathbb{R}_+^n$ ,  $\Delta = \Delta_1 \times \dots \times \Delta_n$ , and  $\Delta_i = \{\xi_i = (\xi_{i1}, \dots, \xi_{in_i}) : \sum_{j=1}^{n_i} \xi_{ij} = d_i, \xi_{ij} \geq 0, j = 1, \dots, n_i\}$ . The strategy profile  $\xi_i = (\xi_{i1}, \dots, \xi_{in_i})$  of player  $i$  can be interpreted as a distribution of non-negative *intensities* over the elements in  $X_i$  satisfying  $\sum_{j=1}^{n_i} \xi_{ij} = d_i$  for  $d_i \in \mathbb{R}_+, i \in N$ . Clearly,  $\Delta_i$  is a compact subset of  $\mathbb{R}_+^{n_i}$  for all  $i \in N$ . For a profile  $\xi = (\xi_1, \dots, \xi_n)$ , we define the set of used facilities of player  $i$  as  $F_i(\xi) = \{f \in F : \text{there exists } j \in \{1, \dots, n_i\} \text{ with } f \in x_{ij} \text{ and } \xi_{ij} > 0\}$ . We define the *load* of player  $i$  on  $f$  under profile  $\xi$  by  $\xi_i^f = \sum_{x_{ij} \in X_i : f \in x_{ij}} \xi_{ij}, i \in N, f \in F$ . In contrast to finite bottleneck congestion games, we assume that cost functions  $c_f : X \rightarrow \mathbb{R}_+$  only depend on the *total load* defined as  $\ell_f(\xi) = \sum_{i \in N} \xi_i^f$ , and are continuous and non-decreasing.

**Definition 5 (Infinite bottleneck congestion game).** Let  $\mathcal{IM} = (N, F, X, d, \Delta, (c_f)_{f \in F})$  be an infinite congestion model derived from  $\mathcal{M}$ . The corresponding infinite bottleneck congestion game is the strategic infinite game  $G(\mathcal{IM}) = (N, \Delta, \pi)$ , where  $\pi$  is defined as  $\pi = \times_{i \in N} \pi_i$  and  $\pi_i(\xi) = \max_{f \in F_i(\xi)} c_f(\ell_f(\xi))$ .

Examples of such games are bottleneck routing games with splittable demands.

**Theorem 3.** Let  $G(\mathcal{IM}) = (N, \Delta, \pi)$  be an infinite bottleneck congestion game. Then,  $G(\mathcal{IM})$  has the LIP for the functions  $\phi : \Delta \rightarrow \mathbb{R}_+^n$  and  $\psi : X \rightarrow \mathbb{R}_+^{mn}$  defined as

$$\phi_i(\xi) = \pi_i(\xi), \quad \text{for all } i \in N, \quad \psi_{i,f}(\xi) = \begin{cases} c_f(\ell_f(\xi)), & \text{if } f \in F_i(\xi) \\ 0, & \text{else} \end{cases} \quad \text{for all } i \in N, f \in F.$$

We remark that for infinite games, the LIP does not imply the existence of a strong potential function. However, if the LIP is satisfied for a continuous function and the strategy space is compact, we can prove the existence of a SNE, see [9] for details. Unfortunately, a simple example reveals that the functions used in Theorem 3 for proving the LIP are not necessarily continuous. In order to obtain a the LIP for a continuous function, we generalize the notion of lexicographical ordering to ordered sets that are different from  $(\mathbb{R}, \leq)$ . To this end, consider a totally ordered set  $(\mathcal{A}, \leq_{\mathcal{A}})$ . Similar to Definition 1, we introduce a lexicographical order on  $\mathcal{A}$ -valued vectors. For two vectors  $a, b \in \mathcal{A}^q$ , let  $\tilde{a}$  and  $\tilde{b}$  be two vectors that arise from  $a$  and  $b$  by ordering them w.r.t  $\leq_{\mathcal{A}}$  in non-increasing order. We say that  $a$  is  $\mathcal{A}$ -lexicographically smaller than  $b$ , written  $a <_{\mathcal{A}} b$  if there is  $m \in \{1, \dots, q\}$  such that  $\tilde{a}_i =_{\mathcal{A}} \tilde{b}_i$  for all  $i < m$  and  $\tilde{a}_m <_{\mathcal{A}} \tilde{b}_m$ . A game satisfies the  $\mathcal{A}$ -LIP if there are  $q \in \mathbb{N}$  and a function  $\phi : X \rightarrow \mathcal{A}^q$  such that  $\phi(x) >_{\mathcal{A}} \phi(y)$  for all  $(x, y) \in I$ .

The following theorem establishes the  $\mathcal{A}$ -LIP for infinite bottleneck congestion games, where  $(\mathcal{A}, \leq_{\mathcal{A}}) = (\mathbb{R}^2, \leq_{lex})$  and  $\leq_{lex}$  denotes the ordinary lexicographical order (that does not involve any sorting of the entries) on  $\mathbb{R}^2$ , that is,  $(a_1, a_2) <_{lex} (b_1, b_2)$  if either  $a_1 < b_1$  or  $(a_1 = b_1 \text{ and } b_2 < b_2)$ .

**Theorem 4.** *Let  $(\mathcal{A}, \leq_{\mathcal{A}}) = (\mathbb{R}^2, \leq_{lex})$  and let  $G(IM) = (N, \Delta, \pi)$  be an infinite bottleneck congestion game. Then,  $G(IM)$  has the  $\mathcal{A}$ -LIP for  $\phi : \Delta \rightarrow \mathcal{A}^m$  defined as  $\phi_f(\xi) = (c_f(\ell_f(\xi)), \ell_f(\xi))$  for all  $f \in F$ .*

The function  $\phi$  in Theorem 4 is continuous. Hence, we obtain for the first time the existence of SNE for a variety of games such as scheduling games with malleable jobs, bottleneck routing games with splittable demands, etc. Note that, compared with the proof in [3], our result gives also an alternative and constructive proof for the existence of PNE in bottleneck routing games with splittable demands.

## 6 Extensions

We present three extensions that (as we feel) are the most interesting ones.

Regarding  $\alpha$ -approximate strong Nash equilibria it is possible to soften the conditions on the cost functions that establish the existence of an equilibrium point. In fact, we can show that every infinite bottleneck congestion game with bounded cost functions possesses an  $\alpha$ -approximate SNE for every  $\alpha > 0$ .

A natural generalization of bottleneck congestion games can be obtained by assuming that players are *heterogeneous* with respect to the cost of the most expensive facility, that is, they attach different *values* to the cost of the most expensive facility.

Assuming that higher costs on facilities are associated with higher private costs we can show that these games possess the LIP (though not the  $\pi$ -LIP).

## Acknowledgments

We thank Michal Feldman for giving an inspiring talk in Berlin about SNE.

## References

1. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. *Games and Economic Behavior* 65(2), 289–317 (2009)
2. Aumann, R.: Acceptable points in general cooperative n-person games. In: *Contributions to the Theory of Games*, vol. 4 (1959)
3. Banner, R., Orda, A.: Bottleneck routing games in communication networks. *IEEE Journal on Selected Areas in Communications* 25(6), 1173–1179 (2007)
4. Chien, S., Sinclair, A.: Strong and pareto price of anarchy in congestion games. In: Albers, S., et al. (eds.) *ICALP 2009, Part I. LNCS*, vol. 5555, pp. 279–291. Springer, Heidelberg (2009)
5. Cole, R., Dodis, Y., Roughgarden, T.: Bottleneck links, variable demand, and the tragedy of the commons. In: *Proc. of the 17th SODA*, pp. 668–677 (2006)
6. Even-Dar, E., Kesselman, A., Mansour, Y.: Convergence time to Nash equilibrium in load balancing. *ACM Trans. Algorithms* 3(3), 32 (2007)
7. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure Nash equilibria. In: *Proc. of the 36th STOC*, pp. 604–612 (2004)
8. Fotakis, D., Kontogiannis, S., Koutsoupias, E., Mavronicolas, M., 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)
9. Harks, T., Klimm, M., Möhring, R.H.: Strong Nash equilibria in games with the lexicographical improvement property. Technical Report 17, COGA, TU Berlin (2009), <http://www.math.tu-berlin.de/coga/publications/techreports/2009/Report-017-2009.xhtml>
10. Holzman, R., Law-Yone, N.: Strong equilibrium in congestion games. *Games and Economic Behavior* 21(1-2), 85–101 (1997)
11. 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)
12. Monderer, D., Shapley, L.S.: Potential games. *Games and Economic Behavior* 14(1), 124–143 (1996)
13. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. The MIT Press, Cambridge (1994)
14. Qiu, L., Yang, Y.R., Zhang, Y., Shenker, S.: On selfish routing in internet-like environments. *IEEE/ACM Trans. Netw.* 14(4), 725–738 (2006)
15. Reny, P.J.: On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica* 67(5), 1029–1056 (1999)
16. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2(1), 65–67 (1973)
17. Rozenfeld, O.: Strong equilibrium in congestion games. Master’s thesis, Israel Institut of Technology (2007)
18. Rozenfeld, O., Tennenholtz, M.: Strong and correlated strong equilibria in monotone congestion games. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) *WINE 2006. LNCS*, vol. 4286, pp. 74–86. Springer, Heidelberg (2006)
19. Srikant, R.: *The Mathematics of Internet Congestion Control*. Birkhäuser, Boston (2003)

# A Note on Strictly Competitive Games

Ilan Adler<sup>1</sup>, Constantinos Daskalakis<sup>2</sup>, and Christos H. Papadimitriou<sup>3</sup>

<sup>1</sup> IEOR, UC Berkeley  
adler@ieor.berkeley.edu  
<sup>2</sup> CSAIL, MIT  
costis@csail.mit.edu  
<sup>3</sup> CS, UC Berkeley  
christos@cs.berkeley.edu

**Abstract.** Strictly competitive games are a class of 2-player games often quoted in the literature to be a proper generalization of zero-sum games. Other times it is claimed, e.g. by Aumann, that strictly competitive games are only payoff transformations of zero-sum games. But to the best of our knowledge there is no proof of such claim. We shed light to this point of confusion in the literature, showing that any strictly competitive game is indeed a payoff transformation of a zero sum-game; in fact, an affine transformation. We offer two proofs of this fact, one combinatorial and one algebraic.

## 1 Introduction

A two-person game is *strictly competitive* [1] if it has the following property: if both players change their mixed strategies, then either there is no change in the expected payoffs, or one of the two expected payoffs increases and the other decreases. That is, all pairs of mixed strategies are Pareto optimal. Mathematically, a game  $(A, -B)$  is strictly competitive if for any two pairs of mixed strategies  $(x, y)$  and  $(x', y')$ ,  $x^T A y - x'^T A y'$  and  $x^T B y - x'^T B y'$  have the same sign [1]. Obviously, these games generalize zero-sum games (the case  $A = B$ ). The question is, how much more general than zero-sum games is this class?

There is much confusion in the literature about this question. Aumann writes “Recall that a strictly competitive game is defined as a two-person game in which if one outcome is preferred to another by one player, the preference is reversed for the other. Since randomized strategies are admitted, this condition applies also to mixed outcomes (probability mixtures of pure outcomes). From this it may be seen that a two-person game is strictly competitive if and only if, for an appropriate choice of utility functions, the utility payoffs of the players sum to zero in each square of the matrix.”

Notice that “appropriate choice” is not defined, and no proof, or outline, is given. Aumann’s insight above is mirrored elsewhere in the literature, e.g. in the textbooks, [3, 4], also without proof. Elsewhere, in lieu of proof a rather

---

<sup>1</sup> For our purposes the *sign* function takes on three values, +, −, and 0.



straightforward weaker fact is pointed out: Let  $a_1, \dots, a_n$  be  $n$  events and let  $u(a_i), v(a_i)$  be the utilities of players 1 and 2 respectively. Now suppose that for any pair of probability vectors  $p, q$  we have  $\sum p_i u(a_i) \geq \sum q_i u(a_i)$  iff  $\sum p_i v(a_i) \leq \sum q_i v(a_i)$ . Then it is easy to see that there exist affine transformations of  $u, v$ , call them  $u', v'$ , such that  $v'(a_i) = -u'(a_i)$  for all  $i$ . Aumann's assertion is stronger: Its hypothesis is that the inequalities hold for all distributions *that are products of mixed strategies*. To increase the confusion, in [5], strictly competitive games are defined with  $x, y, x', y'$  above restricted to pure strategies—this is a proper generalization, albeit of no interest. And elsewhere, strictly competitive games are treated as a proper generalization of zero-sum games.

In this note we prove Aumann's assertion. In fact, we give two very different proofs, one combinatorial and one algebraic (both are quite nontrivial). Let  $A$  and  $B$  be  $m \times n$  real matrices. By  $\Delta_n$  we denote all distributions (mixed strategies) over  $[n]$ . We say that matrix  $B$  is an *affine variant* of matrix  $A$  if for some  $\lambda > 0$  and unrestricted  $\mu$ ,  $B = \lambda \cdot A + \mu \cdot U$ , where  $U$  is  $m \times n$  all-ones matrix. Our main result is the following:

**Theorem 1.** *If for all  $x, x' \in \Delta_m$  and  $y, y' \in \Delta_n$ ,  $x^T A y - x'^T A y'$  and  $x^T B y - x'^T B y'$  have the same sign, then  $B$  is an affine variant of  $A$ .*

Note that the converse is trivial.

## 2 A Combinatorial Proof

Consider the strictly competitive game  $(A, -B)$  with at least two pure strategies for each player. Let

$$a_{max} = \max_{ij} A_{ij}, \quad a_{min} = \min_{ij} A_{ij}$$

and

$$b_{max} = \max_{ij} B_{ij}, \quad b_{min} = \min_{ij} B_{ij}.$$

**Lemma 1.** *For all  $i, j$ :*

$$a_{ij} = a_{max} \iff b_{ij} = b_{max}; \tag{1}$$

$$a_{ij} = a_{min} \iff b_{ij} = b_{min}. \tag{2}$$

*Proof.* We only show the first assertion. The other assertion can be shown similarly. Suppose there exist  $i, j$  such that  $a_{ij} = a_{max}$ , but  $b_{ij} < b_{max}$ . Let then  $k, \ell$  be such that  $b_{k\ell} = b_{max}$ . If  $x, x'$  are the pure strategies  $i, k$  and  $y, y'$  the pure strategies  $j, \ell$ , then the pairs of strategies  $(x, y)$  and  $(x', y')$  violate the condition of strict competitiveness. ■

**Corollary 1.**  $a_{max} = a_{min} \iff b_{max} = b_{min}$ .

If  $a_{max} = a_{min}$  and  $b_{max} = b_{min}$ , then clearly  $B$  is an affine variant of  $A$ . If  $a_{max} > a_{min}$  and  $b_{max} > b_{min}$ , we define the following affine variants of the matrices  $A$  and  $B$ .



$$A' = \frac{1}{a_{max} - a_{min}}[A - a_{min}U],$$

$$B' = \frac{1}{b_{max} - b_{min}}[B - b_{min}U].$$

Observe that all entries of  $A'$ ,  $B'$  are in  $[0, 1]$ ; in particular, both the value 0 and the value 1 appear as entries in both  $A'$  and  $B'$ . Moreover,  $(A', -B')$  is a strictly competitive game. We show the following.

**Lemma 2.**  $A' = B'$ .

*Proof.* Suppose that  $A' \neq B'$ . By Lemma 1 and by rearranging the rows and columns of  $A'$  and  $B'$ , we can assume without loss of generality that  $A'_{11} = B'_{11} = 1$  and either  $A'_{22} = B'_{22} = 0$  (case 1) or  $A'_{12} = B'_{12} = 0$  (case 2). Let  $D = B' - A'$  and let  $|D_{rs}| = \max_{ij} |D_{ij}|$ . For  $0 \leq p \leq 1$ , let  $x(p)$ ,  $y(p)$  be probability vectors whose non-zero elements are:

- Case 1:  $x_1(p) = y_1(p) = p, \quad x_2(p) = y_2(p) = 1 - p;$
- Case 2:  $x_1(p) = 1, \quad y_1(p) = p, \quad y_2(p) = 1 - p.$

Since  $x(p)^T A' y(p) = 0$  for  $p = 0$  and  $x(p)^T A' y(p) = 1$  for  $p = 1$ , there exists  $\bar{p}$  such that  $x(\bar{p})^T A' y(\bar{p}) = A'_{rs}$ . Assuming  $D_{rs} \neq 0$ , we have  $0 < \bar{p} < 1$ . Since the game is strictly competitive, we have that  $x(\bar{p})^T B' y(\bar{p}) = B'_{rs}$ . If this weren't the case, then by taking  $x'$  to be the pure strategy  $r$  and  $y'$  the pure strategy  $s$  we would obtain a contradiction to the strict competitiveness of the game by considering the pairs of mixed strategies  $(x(\bar{p}), y(\bar{p}))$  and  $(x', y')$ .

Given the above we have  $x(\bar{p})^T (A' + D) y(\bar{p}) = B'_{rs}$ , which implies that  $x(\bar{p})^T D y(\bar{p}) = B'_{rs} - A'_{rs} = D_{rs}$ . Noting that  $D_{11} = 0$ ,  $x_1(\bar{p}) \cdot y_1(\bar{p}) > 0$ ,  $|D_{rs}| = \max_{ij} |D_{ij}|$ , and that  $x(\bar{p})^T D y(\bar{p})$  is a weighted average of the elements in  $D$ , we can't have  $x(\bar{p})^T D y(\bar{p}) = D_{rs}$ . Thus  $D = 0$ , implying  $A' = B'$ . ■

Since  $A'$ ,  $B'$  are affine variants of  $A$ ,  $B$ , this completes the proof of Theorem 1.

### 3 An Algebraic Proof

For any matrix  $A$  we consider the polynomial  $p_A(z) = x^T A y - x'^T A y'$ , where by  $z$  we denote the vector of variables  $x, y, x', y'$ . The hypothesis then states that  $p_A(z)$  and  $p_B(z)$  always have the same sign.

First note that, as polynomials,  $p_A$  and  $p_B$  are irreducible (there is no way to factor them without getting extra terms involving both primed and unprimed variables). Consider now the polynomial  $p_{A+B}$ , and consider a  $z^*$  such that  $p_{A+B}(z^*) = 0$ . It is easy to see that such a  $z^*$  exists. We claim that also  $p_A(z^*) = 0$  — otherwise,  $p_A(z^*)$  and  $p_B(z^*) = p_{A+B}(z^*) - p_A(z^*)$  would have opposite signs.

We conclude that the roots of the irreducible polynomial  $p_{A+B}(z)$  are a subset of the roots of the irreducible polynomial  $p_A(z)$ . It follows from Hilbert's Nullstellensatz [2] that  $p_A(z)$  is a multiple of  $p_{A+B}(z)$  (where we used that  $p_A$  is

irreducible)<sup>2</sup> since  $p_A$  is irreducible, a constant multiple. Therefore,  $p_A(z)$  and  $p_B(z)$  are multiples of one another, and thus positive multiples.

Now, it is easy to see that  $p_C(z) = p_D(z)$  iff  $C$  and  $D$  differ by a multiple of  $U$  (that is, the multiples of  $U$  comprise the kernel of the homomorphism from matrices  $A$  to polynomials  $p_A$ ). We conclude that  $B$  is an affine variant of  $A$ , completing the proof.

## References

1. Aumann, R.J.: Game Theory. In: Eatwell, J., Milgate, M., Newman, P. (eds.) *The New Palgrave: A Dictionary of Economics*, pp. 460–482. Macmillan & Co., London (1987)
2. Becker, T., Weispfenning, V.: *Gröbner Bases: A Computational Approach to Commutative Algebra*. Springer, New York (1993)
3. Luce, R.D., Raiffa, H.: *Games and Decisions: Introduction and Critical Survey*. Wiley, New York (1957)
4. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press, Cambridge (1994)
5. Raiffa, H.: Arbitration Schemes for Generalized Two-Person Games. In: Kuhn, H.W., Tucker, A.W. (eds.) *Contributions to the Theory of Games*, vol. 2, pp. 361–388 (1953)

---

<sup>2</sup> Strictly speaking, the application of the Nullstellensatz requires that there is a non-singular point for the two polynomials with neighborhoods isomorphic to balls, but this is easy, if a little technical, to see—any point in the interior will do.

# The Efficiency of Fair Division<sup>\*</sup>

Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos,  
and Maria Kyropoulou

Research Academic Computer Technology Institute and  
Department of Computer Engineering and Informatics  
University of Patras, 26500 Rio, Greece

**Abstract.** We study the impact of fairness on the efficiency of allocations. We consider three different notions of fairness, namely proportionality, envy-freeness, and equitability for allocations of divisible and indivisible goods and chores. We present a series of results on the price of fairness under the three different notions that quantify the efficiency loss in fair allocations compared to optimal ones. Most of our bounds are either exact or tight within constant factors. Our study is of an optimistic nature and aims to identify the potential of fairness in allocations.

## 1 Introduction

Fair division (or fair allocation) dates back to the ancient times and has found applications such as border settlement in international disputes, greenhouse gas emissions reduction, allocation of mineral riches in the ocean bed, inheritance, divorces, etc. In the era of the Internet, it appears regularly in distributed resource allocation and cost sharing in communication networks.

We consider allocation problems in which a set of *goods* or *chores* has to be allocated among several players. Fairness is an apparent desirable property in these situations and means that each player gets a *fair share*. Depending on what the term “fair share” means, different notions of fairness can be defined. An orthogonal issue is *efficiency* that refers to the total happiness of the players. An important notion that captures the minimum efficiency requirement from an allocation is that of *Pareto-efficiency*; an allocation is Pareto-efficient if there is no other allocation that is strictly better for at least one player and is at least as good for all the others.

*Model and problem statement.* We consider two different allocation scenarios, depending on whether the items to be allocated are goods or chores. In both cases, we distinguish between *divisible* and *indivisible* items.

The problem of allocating divisible goods is better known as *cake-cutting*. In instances of cake-cutting, the term *cake* is used as a synonym of the whole set of goods to be allocated. Each player has a *utility function* on each piece

---

<sup>\*</sup> This work is partially supported by the European Union under IST FET Integrated Project FP6-015964 AEOLUS and Cost Action IC0602 “Algorithmic Decision Theory”, and by a “Caratheodory” basic research grant from the University of Patras.

of the cake corresponding to the happiness of the player if she is allocated the particular piece; this function is non-negative and additive. We assume that the utility of each player for the whole cake is 1. Divisibility means that the cake can be cut in arbitrarily small pieces which can then be allocated to the players. In instances with indivisible goods, the utility function of a player is defined over sets of items; again, utilities are non-negative and additive and the utility of each player for the whole set of items is 1. Each item cannot be cut in pieces and has to be allocated as a whole to some player. Given an allocation, the utility of a player is simply the sum of her utilities over the (pieces of) items she receives. An allocation with  $n$  players is proportional if the utility of each player is at least  $1/n$ . It is envy-free if the utility of a player is not smaller than the utility she would have when exchanging the (pieces of) items she gets with the items of any other player. It is equitable if the utilities of all players are equal. An allocation is *optimal* if it maximizes the total utility of all players, i.e., each (piece of) item is allocated to the player that values it the most (ties are broken arbitrarily).

In instances with divisible chores, each player has a *disutility function* for each piece of the cake which denotes the regret of the player when she is allocated the particular piece. Again, the disutility functions are non-negative and additive and the disutility of a player for the whole cake is 1. The case of indivisible chores is defined accordingly; indivisibility implies that an item cannot be cut into pieces and has to be allocated as a whole to some player. Given an allocation, the disutility of a player is simply the sum of her disutilities over the (pieces of) items she receives. An allocation with  $n$  players is proportional if the disutility of each player is at most  $1/n$ . It is envy-free if the disutility of a player is not larger than the disutility she would have when exchanging the (pieces of) items she gets with the items of any other player. It is equitable if the disutilities of all players are equal. An allocation is optimal if it minimizes the total disutility of all players, i.e., each (piece of) item is allocated to the player that values it the least (ties are broken arbitrarily).

Note that envy-freeness implies proportionality. Furthermore, instances with divisible items always have proportional, envy-free, or equitable allocations. It is not hard to see that this is not always the case for instances with indivisible items. Furthermore, there are instances in which no optimal allocation is fair.

Models similar to ours have been considered in the literature; the focus has been on the design of protocols for achieving proportionality [3,4,8], envy-freeness [3,6,7], and equitability [3] or on the design of approximation algorithms in settings where fulfilling the fairness objective exactly is impossible [2,5]. However, the related literature seems to have neglected the issue of efficiency. Although several attempts have been made to characterize fair division protocols in terms of Pareto-efficiency [3], the corresponding results are almost always negative. Most of the existing protocols do not even provide Pareto-efficient solutions and this seems to be due to the limited amount of information they use for the utility functions of the players. Recall that in the case of divisible goods and chores, complete information about the utility or disutility functions of the players may

not be compactly representable. Furthermore, Pareto-efficiency is rather unsatisfactory, since it may imply that an allocation is far from optimal.

Instead, in the current paper we are interested in quantifying the decrease of efficiency due to fairness (*price of fairness*). Our study has an optimistic nature and aims to identify the potential of fairness in allocations. We believe that such a study is well-motivated since the knowledge of tight bounds on the price of fairness may detect whether a fair allocation can be improved. In many settings, complete information about the utility functions of the players is available (e.g., in a divorce) and computing an efficient and fair allocation may not be infeasible. Fair allocations can be thought of as counterparts of equilibria in strategic games; hence, our work is similar in spirit to the line of research that studies the price of stability in games [11].

In order to capture the price of fairness, we define the price of proportionality, envy-freeness, and equitability. Given an instance  $I$  for the allocation of goods, its price of proportionality (resp., envy-freeness, resp., equitability) is defined as the ratio of the total utility of the players in the optimal allocation for  $I$  over the total utility of the players in the best proportional (resp., envy-free, resp., equitable) allocation for  $I$ . Similarly, if  $I$  is an instance for the allocation of chores, its price of proportionality (resp., envy-freeness, resp., equitability) is defined as the ratio of the total disutility of the players in the best proportional (resp., envy-free, resp., equitable) allocation for  $I$  over the total disutility of the players in the optimal allocation for  $I$ . The price of proportionality (resp., envy-freeness, resp., equitability) of a class  $\mathcal{I}$  of instances is then the maximum price of proportionality (resp., envy-freeness, resp., equitability) over all instances of  $\mathcal{I}$ . The classes of instances considered in this paper are defined by the number of players, the type of items (goods or chores), and their divisibility property (divisible or indivisible). We remark that, in order for the price of proportionality, envy-freeness, and equitability to be well-defined, in the case of indivisible items, we assume that the class of instances contains only those ones for which proportional, envy-free, and equitable allocations, respectively, do exist.

*Overview of results.* In this paper we provide upper and lower bounds on the price of proportionality, envy-freeness, and equitability in fair division with divisible and indivisible goods and chores. Our work reveals an almost complete picture. In all subcases except the price of envy-freeness with divisible goods and chores, our bounds are either exact or tight within a small constant factor.

Table 1 summarizes our results. For divisible goods, the price of proportionality is very close to 1 (i.e.,  $8 - 4\sqrt{3} \approx 1.072$ ) for two players and  $\Theta(\sqrt{n})$  in general. The price of equitability is slightly worse for two players (i.e.,  $9/8$ ) and  $\Theta(n)$  in general. Our lower bound for the price of proportionality implies the same lower bound for the price of envy-freeness; while a very simple upper bound of  $n - 1/2$  completes the picture for divisible goods. For indivisible goods, we present an exact bound of  $n - 1 + 1/n$  on the price of proportionality while we show that the price of envy-freeness is  $\Theta(n)$  in this case. Although our upper bounds follow by very simple arguments, the lower bounds use quite involved constructions. The price of equitability is proven to be finite only for the case of two players. These

**Table 1.** Summary of our results (lower and upper bounds)

	LB	UB	$n = 2$	LB	UB	$n = 2$
Price of	Divisible goods			Indivisible goods		
Proportionality	$\Omega(\sqrt{n})$	$O(\sqrt{n})$	$8 - 4\sqrt{3}$	$n - 1 + 1/n$	$n - 1 + 1/n$	$3/2$
Envy-freeness	$\Omega(\sqrt{n})$	$n - 1/2$		$\frac{3n+7}{9} - O(1/n)$	$n - 1/2$	
Equitability	$\frac{(n+1)^2}{4n}$	$n$	$9/8$	$\infty$	$\infty$	$2$
	Divisible chores			Indivisible chores		
Proportionality	$\frac{(n+1)^2}{4n}$	$n$	$9/8$	$n$	$n$	$2$
Envy-freeness	$\frac{(n+1)^2}{4n}$	$\infty$		$\infty$	$\infty$	
Equitability	$n$	$n$	$2$	$\infty$	$\infty$	$\infty$

results are presented in Section 2. For divisible chores, the price of proportionality is  $9/8$  for two players and  $\Theta(n)$  in general while the price of equitability is exactly  $n$ . For indivisible chores, we present an exact bound of  $n$  on the price of proportionality while both the price of envy-freeness and the price of equitability are infinite. These last results imply that in the case of indivisible chores, envy-freeness and equitability are usually incompatible with efficiency. These results are presented in Section 3. Due to lack of space, many proofs have been omitted.

## 2 Fair Division with Goods

In this section, we focus on fair division and goods. We begin by presenting our results for the case of divisible goods.

**Theorem 1.** *For  $n$  players and divisible goods, the price of proportionality is  $\Theta(\sqrt{n})$ .*

*Proof.* Consider an instance with  $n$  players and let  $\mathcal{O}$  denote the optimal allocation and  $OPT$  be the total utility of  $\mathcal{O}$ . We partition the set of players into two sets, namely  $L$  and  $S$ , so that if a player obtains utility at least  $1/\sqrt{n}$  in  $\mathcal{O}$ , then she belongs to  $L$ , otherwise she belongs to  $S$ . Clearly,  $OPT < |L| + |S|/\sqrt{n}$ . We now describe how to obtain a proportional allocation  $\mathcal{A}$ ; we distinguish between two cases depending on  $|L|$ .

We first consider the case  $|L| \geq \sqrt{n}$ ; hence,  $|S| \leq n - \sqrt{n}$ . Then, for any negligibly small item that is allocated to a player  $i \in L$  in  $\mathcal{O}$ , we allocate to  $i$  a fraction of  $\sqrt{n}/n$  of the item, while we allocate to each player  $i \in S$  a fraction of  $\frac{n-\sqrt{n}}{|S|} \geq 1/n$ . Furthermore, for any negligibly small item that is allocated to a player  $i \in S$  in  $\mathcal{O}$ , we allocate to each player  $i \in S$  a fraction of  $1/|S| > 1/n$ . In this way, all players obtain a utility of at least  $1/n$ , while all items are fully allocated; hence,  $\mathcal{A}$  is proportional. For every player  $i \in L$ , her utility in  $\mathcal{A}$  is exactly  $1/\sqrt{n}$  times her utility in  $\mathcal{O}$ , while every player  $i \in S$  obtained a utility strictly less than  $1/\sqrt{n}$  in  $\mathcal{O}$  and obtains utility at least  $1/n$  in  $\mathcal{A}$ . So, we conclude that the total utility in  $\mathcal{A}$  is at least  $1/\sqrt{n}$  times the optimal total utility.

Otherwise, let  $|L| < \sqrt{n}$ . Since  $OPT < |L| + |S|/\sqrt{n}$ , we obtain that  $OPT < 2\sqrt{n} - 1$ , while the total utility of any proportional allocation is at least 1. Hence, in both cases we obtain that the price of proportionality is  $O(\sqrt{n})$ . We continue by presenting a lower bound of  $\Omega(\sqrt{n})$ .

Consider the following instance with  $n$  players and  $m < n$  items. Player  $i$ , for  $i = 1, \dots, m$ , has utility 1 for item  $i$  and 0 for any other item, while player  $i$ , for  $i = m + 1, \dots, n$ , has utility  $1/m$  for any item. In the optimal allocation, item  $i$ , for  $i = 1, \dots, m$ , is allocated to player  $i$ , and the total utility is  $m$ . Consider any proportional allocation and let  $x$  be the sum of the fractions of the items that are allocated to the last  $n - m$  players. The total utility of these players is  $x/m$ . Clearly,  $x \geq m(n - m)/n$ , otherwise some of them would obtain a utility less than  $1/n$  and the allocation would not be proportional. The first  $m$  players are allocated the remaining fraction of  $m - x$  of the items and their total utility is at most  $m - x$ . The total utility of all players is  $m - x + x/m \leq \frac{m^2 + n - m}{n}$ . We conclude that the price of proportionality is at least  $\frac{mn}{m^2 + n - m}$  which becomes more than  $\sqrt{n}/2$  by setting  $n = m^2$ .  $\square$

For the price of equitability, we can show that when the number of players is large, equitability may provably lead to less efficient allocations.

**Theorem 2.** *For  $n$  players and divisible goods, the price of equitability is at most  $n$  and at least  $\frac{(n+1)^2}{4n}$ .*

Since every envy-free allocation is also proportional, the lower bound on the price of proportionality also holds for envy-freeness. Interestingly, in the case of two players, there always exist almost optimal proportional or equitable allocations. Recall that in this case proportionality and envy-freeness are equivalent.

**Theorem 3.** *For two players and divisible goods, the price of proportionality (or envy-freeness) is  $8 - 4\sqrt{3} \approx 1.072$ , and the price of equitability is  $9/8$ .*

*Proof.* We only present the result for proportionality here. Consider an optimal allocation  $\mathcal{O}$  and a proportional allocation  $\mathcal{E}$  that maximizes the total utility of the players. We partition the cake into four parts  $A, B, C$ , and  $D$ :  $A$  is the part of the cake which is allocated to player 1 in both  $\mathcal{O}$  and  $\mathcal{E}$ ,  $B$  is the part which is allocated to player 2 in both  $\mathcal{O}$  and  $\mathcal{E}$ ,  $C$  is the part which is allocated to player 1 in  $\mathcal{O}$  and to player 2 in  $\mathcal{E}$ , and  $D$  is the part of the cake which is allocated to player 1 in  $\mathcal{E}$  and to player 2 in  $\mathcal{O}$ . In the following, we use the notation  $u_i(X)$  to denote the utility of player  $i$  for part  $X$  of the cake.

Since  $\mathcal{O}$  maximizes the total utility, we have  $u_1(A) \geq u_2(A)$ ,  $u_1(B) \leq u_2(B)$ ,  $u_1(C) \geq u_2(C)$ , and  $u_1(D) \leq u_2(D)$ . First observe that if  $u_1(C) = u_2(C)$  and  $u_1(D) = u_2(D)$ , then  $\mathcal{E}$  has the same total utility with  $\mathcal{O}$ . So, in the following we assume that this is not the case.

We consider the case  $u_1(C) > u_2(C)$ ; the other case is symmetric. In this case, we also have that  $u_1(D) = u_2(D) = 0$ . Assume otherwise that  $u_2(D) > 0$ . Then, there must be a subpart  $X$  of  $C$  for which player 1 has utility  $x$  and player 2 has utility at most  $x \cdot u_2(C)/u_1(C)$  and a subpart  $Y$  of  $D$  for which player 2 has

utility  $x$  and player 1 has utility at least  $x \cdot u_1(D)/u_2(D)$ . Then, the allocation in which player 1 gets parts  $A, X$ , and  $D - Y$  and player 2 gets parts  $B, C - X$ , and  $Y$  is proportional and has larger utility than  $\mathcal{E}$ .

Now, we claim that  $u_2(A) = 1/2$ . Clearly, since  $\mathcal{E}$  is proportional, the utility of player 2 in  $\mathcal{E}$  is at least  $1/2$ , i.e.,  $u_2(B) + u_2(C) \geq 1/2$ . Since the utilities of player 2 sum up to 1 over the whole cake, we also have that  $u_2(A) \leq 1/2$ . If it were  $u_2(A) < 1/2$ , then we would have  $u_2(B) + u_2(C) > 1/2$ . Then, there would exist a subpart  $X$  of  $C$  for which player 2 has utility  $x$  for some  $x \leq 1/2 - u_2(A)$  and player 1 has utility larger than  $x$ . By allocating  $X$  to player 1 instead of player 2, we would obtain another proportional allocation with larger total utility.

Also, it holds that  $u_2(A)/u_1(A) \leq u_2(C)/u_1(C)$ . Otherwise, there would exist a subpart  $X$  of  $C$  for which player 1 has utility  $x$  and player 2 has utility  $u_2(X)$  at most  $x \cdot u_2(C)/u_1(C)$  and a subpart  $Y$  of  $A$  for which player 1 has utility  $x$  and player 2 has utility  $u_2(Y)$  at least  $x \cdot u_2(A)/u_1(A) > x \cdot u_2(C)/u_1(C) \geq u_2(X)$ . By allocating the subpart  $X$  to player 1 and subpart  $Y$  to player 2, we would obtain another proportional allocation with larger total utility.

By the discussion above, we have  $u_2(C) \geq \frac{u_1(C)}{2u_1(A)}$ . We are now ready to bound the ratio of the total utility of  $\mathcal{O}$  over the total utility of  $\mathcal{E}$  which will give us the desired bound. We obtain that the price of proportionality is

$$\begin{aligned} \frac{u_1(A) + u_2(B) + u_1(C)}{u_1(A) + u_2(B) + u_2(C)} &= \frac{u_1(A) + 1/2 - u_2(C) + u_1(C)}{u_1(A) + 1/2} \\ &\leq \frac{u_1(A) + 1/2 + u_1(C) \left(1 - \frac{1}{2u_1(A)}\right)}{u_1(A) + 1/2} \\ &\leq \frac{u_1(A) + 1/2 + (1 - u_1(A)) \left(1 - \frac{1}{2u_1(A)}\right)}{u_1(A) + 1/2} \end{aligned}$$

where the last inequality follows since  $u_1(A) \geq u_2(A) = 1/2$  and  $u_1(C) \leq 1 - u_1(A)$ . The last expression is maximized to  $8 - 4\sqrt{3}$  for  $u_1(A) = \frac{1+\sqrt{3}}{4}$  and the upper bound follows.

In order to prove the lower bound, it suffices to consider a cake consisting of two parts  $A$  and  $B$ . Player 1 has utility  $u_1(A) = 1$  and  $u_1(B) = 0$  and player 2 has utility  $u_2(A) = \sqrt{3} - 1$  and  $u_2(B) = 2 - \sqrt{3}$ . □

Moreover, it is easy to show an upper bound of  $n - 1/2$  for the price of envy-freeness for both divisible and indivisible goods.

We next present our results that hold explicitly for indivisible goods; these results are either exact or tight within a constant factor.

**Theorem 4.** *For  $n$  players and indivisible goods, the price of proportionality is  $n - 1 + 1/n$ .*

*Proof.* We begin by proving the upper bound. Consider an instance and a corresponding optimal allocation. If this allocation is proportional, then the price of proportionality is 1; assume otherwise. In any proportional allocation, each



player has utility at least  $1/n$  on the pieces of the cake she receives and the total utility is at least 1. Since the optimal allocation is not proportional, some player has utility less than  $1/n$  and the total utility in the optimal allocation is at most  $n - 1 + 1/n$ .

We now present the lower bound. Consider the following instance with  $n$  players and  $2n - 1$  items. Let  $0 < \epsilon < 1/n$ . For  $i = 1, \dots, n - 1$ , player  $i$  has utility  $\epsilon$  for item  $i$ , utility  $1 - 1/n$  for item  $i + 1$ , utility  $1/n - \epsilon$  for item  $n + i$  and utility 0 for all other items. The last player has utility  $1/n - \epsilon$  for items  $1, 2, \dots, n - 1$ , utility  $1/n + (n - 1)\epsilon$  for item  $n$ , and utility 0 for all other items.

We argue that the only proportional allocation assigns items  $i$  and  $n + i$  to player  $i$  for  $i = 1, \dots, n - 1$ , and item  $n$  to player  $n$ . To see that, notice that each player must be allocated at least one of the first  $n$  items, regardless of what other items she obtains, in order to be proportional. Since there are  $n$  players, each of them must be allocated exactly one of the first  $n$  items. Now, consider player  $n$ . It is obvious that she must be allocated item  $n$ , since she has utility strictly less than  $1/n$  for any other item. The only available items (with positive utility) left for player  $n - 1$  are items  $n - 1$  and  $2n - 1$ , and it is easy to see that both of them must be allocated to her. Using the same reasoning for players  $n - 2, n - 3, \dots, 1$ , we conclude that the only proportional allocation is the aforementioned one, which has total utility  $1 + (n - 1)\epsilon$ .

Now, the total utility of the optimal allocation is lower-bounded by the total utility of the allocation where player  $i$  gets items  $i + 1$  and  $n + i$ , for  $i = 1, \dots, n - 1$ , and player  $n$  gets the first item. The total utility obtained by this allocation is  $(1 - 1/n + 1/n - \epsilon)(n - 1) + \frac{1}{n} - \epsilon = n - 1 + 1/n - n\epsilon$ . By selecting  $\epsilon$  to be arbitrarily small, the theorem follows.  $\square$

The above lower bound construction uses instances with no envy-free allocation and, hence, the lower bound on the price of proportionality does not extend to envy-freeness. We have a slightly weaker lower bound on the price of envy-freeness for indivisible goods which uses a more involved construction.

**Theorem 5.** *For  $n$  players and indivisible goods, the price of envy-freeness is at least  $\frac{3n+7}{9} - O(1/n)$ .*

Unfortunately, equitability may lead to arbitrarily inefficient allocations of indivisible goods when the number of players is at least 3.

**Theorem 6.** *For  $n$  players and indivisible goods, the price of equitability is 2 for  $n = 2$  and infinite for  $n > 2$ .*

### 3 Fair Division with Chores

Our next theorem considers divisible chores.

**Theorem 7.** *For  $n$  players and divisible chores, the price of proportionality is at most  $n$  and at least  $\frac{(n+1)^2}{4n}$ , and the price of equitability is  $n$ .*

Since every envy-free allocation is also proportional, the lower bound on the price of proportionality also holds for envy-freeness. We also have a matching upper bound of  $9/8$  for proportionality (or envy-freeness) in the case  $n = 2$ .

Finally, we consider the case of indivisible chores. Although the price of proportionality is bounded, the price of envy-freeness and equitability is infinite.

**Theorem 8.** *For  $n$  players and indivisible chores, the price of proportionality is  $n$ , whereas the price of envy-freeness (for  $n \geq 3$ ) and equitability (for  $n \geq 2$ ) is infinite.*

*Proof.* Due to lack of space, we only present the case of envy-freeness. Consider the following instance with  $n$  players and  $2n$  items. Let  $\epsilon < 1/(2n)$ . For  $i = 1, \dots, n-2$ , player  $i$  has disutility  $1/n$  for the first  $n$  items and disutility 0 for every other item. Player  $n-1$  has disutility 0 for the first  $n-1$  items, disutility  $\epsilon$  for item  $n$ , disutility  $1/n$  for items  $n+1, \dots, 2n-1$  and disutility  $1/n - \epsilon$  for item  $2n$ . Finally, player  $n$  has disutility 0 for the first  $n-1$  items, disutility  $1/(2n)$  for items  $n$  and  $2n$ , and disutility  $1/n$  for items  $n+1, \dots, 2n-1$ .

Clearly, the optimal allocation has total disutility  $\epsilon$  and is obtained by allocating items  $n+1, \dots, 2n$  to players  $1, \dots, n-2$ , item  $n$  to player  $n-1$ , and items  $1, \dots, n-1$  either to player  $n-1$ , or to player  $n$ . In each case, player  $n-1$  envies player  $n$ . Furthermore, the allocation in which player  $i$ , for  $i = 1, \dots, n$  is allocated items  $i$  and  $i+n$  is envy-free. The remark that concludes this proof is that there cannot exist an envy-free allocation having negligible disutility (i.e., less than  $1/(2n)$ ).  $\square$

## 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. *SIAM Journal of Computing* 38(4), 1602–1623 (2008)
2. Bansal, N., Sviridenko, M.: The Santa Claus problem. In: *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC 2006)*, pp. 31–40 (2006)
3. Brams, S.J., Taylor, A.D.: *Fair division: From cake-cutting to dispute resolution*. Cambridge University Press, Cambridge (1996)
4. Edmonds, J., Pruhs, K.: Cake-cutting really is not a piece of cake. In: *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2006)*, pp. 271–278 (2006)
5. Kumar, A., Kleinberg, J.: Fairness measures for resource allocation. In: *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS 2000)*, pp. 75–85 (2000)
6. Lipton, R., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: *Proceedings of the 5th ACM Conference on Electronic Commerce (EC 2004)*, pp. 125–131 (2004)
7. Procaccia, A.: Thou shalt covet thy neighbor's cake. In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009* (2009)
8. Woeginger, G.J., Sgall, J.: On the complexity of cake-cutting. *Discrete Optimization* 4, 213–220 (2007)

# Sequential Bidding in the Bailey-Cavallo Mechanism

Krzysztof R. Apt<sup>1,2</sup> and Evangelos Markakis<sup>3</sup>

<sup>1</sup> CWI, Science Park 123, 1098 XG Amsterdam

<sup>2</sup> Institute of Logic, Language and Computation, University of Amsterdam  
The Netherlands

<sup>3</sup> Athens University of Economics and Business  
Department of Informatics, Athens, Greece

**Abstract.** We are interested in mechanisms that maximize social welfare. In [2] this problem was studied for multi-unit auctions and for public project problems, and in each case social welfare undominated mechanisms were identified. One way to improve upon these optimality results is by allowing the players to move sequentially. With this in mind, we study here a sequential version of the Bailey-Cavallo mechanism, a natural mechanism that was proved to be welfare undominated in the simultaneous setting by [2]. Because of the absence of dominant strategies in the sequential setting, we focus on a weaker concept of an optimal strategy. We proceed by introducing natural optimal strategies and show that among all optimal strategies, the one we introduce generates maximal social welfare. Finally, we show that the proposed strategies form a safety level equilibrium and within the class of optimal strategies they also form a Pareto optimal ex-post equilibrium [1].

## 1 Introduction

In many resource allocation problems a group of agents would like to determine who among them values a given object the most. A natural way to approach this problem is by viewing it as a single unit auction. Such an auction is traditionally used as a means of determining by a seller to which bidder and for which price the object is to be sold. The absence of a seller however changes the perspective and leads to different considerations since in our setting, the payments that the agents need to make flow out of the system (are “burned”). Instead of maximizing the revenue of the seller we are thus interested in maximizing the final social welfare.

This has led to the problem of finding mechanisms that are optimal in the sense that no other feasible, efficient and incentive compatible mechanism generates a larger social welfare. Recently, in [2] this problem was studied for two domains: multi-unit auctions with unit demand bidders and the public project problem of [8]. For the first domain a class of optimal mechanisms (which includes the Bailey-Cavallo mechanism) was identified, while for the second one

---

<sup>1</sup> A full version of this work along with all the missing proofs is available at <http://pages.cs.aueb.gr/~markakis/research/pubs.html>

it was proved that the pivotal mechanism is optimal. Other related aspects and objectives have also been recently studied in a series of works on redistribution and money-burning mechanisms, see among others [9,12,10,13,16,7].

We continue this line of research by relaxing the assumption of simultaneity and allowing the players to move sequentially. This setup has been recently studied in [3] for the public project problem and here we consider such a modified setting for the case of single unit auctions. We call it *sequential bidding* as the concept of a “sequential auction” usually refers to a sequence of auctions, see, e.g. [14, chapter 15].

Hence we assume that there is a single object for sale and the players announce their bids sequentially in a fixed order. In contrast to the open cry auctions each player announces his bid *exactly once*. Once all bids have been announced, a mechanism is used to allocate the object to the highest bidder and determine the payments. Such a sequential setting can be very natural in many decision making or coordination problems without a central authority.

## 1.1 Results

We study here a sequential version of the Bailey-Cavallo mechanism of [5] and [6], as being a simplest, natural and most intuitive mechanism in the class of OEL mechanisms [11]. Our main results start in Section 4, where we first show that in a large class of sequential Groves auctions no *dominant strategies* exist. Therefore we settle on a weaker concept, that of an *optimal strategy*. An optimal strategy is a natural relaxation of the notion of dominant strategy, which also captures precisely the way a “prudent” player would play (see Lemma 1).

We proceed in Section 5 with proposing optimal strategies that differ from truth telling in the Bailey-Cavallo mechanism. We show that the proposed strategies yield maximal social welfare among all possible vectors of optimal strategies. Finally in Section 6 we further clarify the nature of the proposed strategies by studying what type of equilibrium they form. First we point that they do not form an ex-post equilibrium, a concept criticized in [4] and [1], where an alternative notion of a *safety-level equilibrium* was introduced for pre-Bayesian games. This concept captures the idea of an equilibrium in the case when each player is “prudent”. We prove that the proposed strategies form a safety-level equilibrium. We also show that our strategies form a Pareto optimal ex-post equilibrium within the class of optimal strategies.

## 2 Preliminaries

Assume that there is a finite 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*  $\Theta_i$  and an (*initial*) *utility function*  $v_i : D \times \Theta_i \rightarrow \mathbb{R}$ . Let  $\Theta := \Theta_1 \times \dots \times \Theta_n$ . A *decision rule* is a function  $f : \Theta \rightarrow D$ . A mechanism is given by a pair of functions  $(f, t)$ , where  $f$  is the decision rule and  $t = (t_1, \dots, t_n)$  is the tax function that determines the players’ payments. We assume that the (*final*) *utility function* for player

$i$  is a function  $u_i$  defined by  $u_i(d, t_1, \dots, t_n, \theta_i) := v_i(d, \theta_i) + t_i$ . Thus, when the true type of player  $i$  is  $\theta_i$  and his announced type is  $\theta'_i$ , his final utility under the mechanism  $(f, t)$  is:

$$u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i) = v_i(f(\theta'_i, \theta_{-i}), \theta_i) + t_i(\theta'_i, \theta_{-i}),$$

Given a sequence  $a := (a_1, \dots, a_j)$  of reals we denote the least  $l$  such that  $a_l = \max_{k \in \{1, \dots, j\}} a_k$  by  $\operatorname{argmax} a$ . A **single item sealed bid auction**, is modelled by choosing  $D = \{1, \dots, n\}$ , each  $\Theta_i$  to be the set of non-negative reals and  $f(\theta) := \operatorname{argmax} \theta$ . Hence the object is sold to the highest bidder and in the case of a tie we allocate the object to the player with the lowest index.<sup>2</sup>

By a **Groves auction** we mean a Groves mechanism for an auction setting (for details on Groves mechanisms see [15]). Below, given a sequence  $\theta$  of reals we denote by  $\theta^*$  its reordering in descending order. Then  $\theta_k^*$  is the  $k$ th largest element in  $\theta$ . For example, for  $\theta = (1, 5, 0, 3, 2)$  we have  $(\theta_{-2})_2^* = 2$  since  $\theta_{-2} = (1, 0, 3, 2)$ . The **Vickrey auction** is the pivotal mechanism for an auction (also referred to as the VCG mechanism). In it the winner pays the second highest bid.

The **Bailey-Cavallo** mechanism, in short **BC auction**, was originally proposed in [5] and [6]. To define it note that each Groves mechanism is uniquely determined by its **redistribution function**  $r := (r_1, \dots, r_n)$ . Given the redistribution function  $r$ , the tax for player  $i$  is defined by  $t_i(\theta) := t_i^p(\theta) + r_i(\theta_{-i})$ , where  $t_i^p$  is the tax of player  $i$  in the Vickrey auction. So we can think of a Groves auction as first running the pivotal mechanism and then redistributing some amount of the pivotal taxes.

The BC auction is a Groves mechanism defined by using the following redistribution function  $r := (r_1, \dots, r_n)$  (assuming that  $n \geq 3$ ):

$$r_i(\theta_{-i}) := \frac{(\theta_{-i})_2^*}{n}$$

It can be seen that the BC auction always yields at least as high social welfare as the pivotal mechanism. Note also that the aggregate tax is 0 when the second and third highest bids coincide.

### 3 Sequential Mechanisms

We are interested in sequential mechanisms, where players announce their types according to a fixed order, say,  $1, 2, \dots, n$ . Each player  $i$  *observes* the actions announced by players  $1, \dots, i-1$  and uses this information to decide which action to select. Thus a **strategy** of player  $i$  is now a function  $s_i : \Theta_1 \times \dots \times \Theta_{i-1} \times \Theta_i \rightarrow \Theta_i$ . Then if the vector of types that the players have is  $\theta$  and the vector of strategies that they decide to follow is  $s(\cdot) := (s_1(\cdot), \dots, s_n(\cdot))$ , the resulting vector of the selected actions will be denoted by  $[s(\cdot), \theta]$ , where  $[s(\cdot), \theta]$  is defined inductively by  $[s(\cdot), \theta]_1 := s_1(\theta_1)$  and  $[s(\cdot), \theta]_{i+1} := s_{i+1}([s(\cdot), \theta]_1, \dots, [s(\cdot), \theta]_i, \theta_{i+1})$ .

<sup>2</sup> If we make a different assumption on breaking ties, some of our proofs need to be adjusted, but similar results hold.

Given  $\theta \in \Theta$  and  $i \in \{1, \dots, n\}$  we denote the sequence  $\theta_{i+1}, \dots, \theta_n$  by  $\theta_{>i}$  and the sequence  $\Theta_{i+1}, \dots, \Theta_n$  by  $\Theta_{>i}$ , and similarly with  $\theta_{\leq i}$  and  $\Theta_{\leq i}$ .

A strategy  $s_i(\cdot)$  of player  $i$  is called **dominant** if for all  $\theta \in \Theta$ , all strategies  $s'_i(\cdot)$  of player  $i$  and all vectors  $s_{-i}(\cdot)$  of strategies of players  $j \neq i$

$$u_i((f, t)((s_i(\cdot), s_{-i}(\cdot)), \theta)), \theta_i) \geq u_i((f, t)((s'_i(\cdot), s_{-i}(\cdot)), \theta)), \theta_i),$$

We call a joint strategy  $s(\cdot) = (s_1(\cdot), \dots, s_n(\cdot))$

- an **ex-post equilibrium** if for all  $i \in \{1, \dots, n\}$ , all strategies  $s'_i(\cdot)$  of player  $i$  and all joint types  $\theta \in \Theta$

$$u_i((f, t)((s_i(\cdot), s_{-i}(\cdot)), \theta)), \theta_i) \geq u_i((f, t)((s'_i(\cdot), s_{-i}(\cdot)), \theta)), \theta_i),$$

- a **safety-level equilibrium** if for all  $i \in \{1, \dots, n\}$ , all strategies  $s'_i(\cdot)$  of player  $i$  and all  $\theta_{\leq i} \in \Theta_{\leq i}$

$$\min_{\theta_{>i} \in \Theta_{>i}} u_i((f, t)((s_i(\cdot), s_{-i}(\cdot)), \theta)), \theta_i) \geq \min_{\theta_{>i} \in \Theta_{>i}} u_i((f, t)((s'_i(\cdot), s_{-i}(\cdot)), \theta)), \theta_i).$$

Intuitively, given the types  $\theta_{\leq i} \in \Theta_{\leq i}$  of players  $1, \dots, i$  and the vector  $s(\cdot)$  of strategies used by the players, the quantity  $\min_{\theta_{>i} \in \Theta_{>i}} u_i((f, t)((s(\cdot), \theta)), \theta_i)$  is the minimum payoff that player  $i$  can guarantee to himself.

## 4 Sequential Groves Auctions

In Groves auctions truth telling is a dominant strategy. In the case of sequential Groves auctions the situation changes as for a wide class, which includes sequential BC auctions no dominant strategies exist (except for the last player).

**Theorem 1.** *Consider a sequential Groves auction. Suppose that for player  $i \in \{1, \dots, n - 1\}$ , the redistribution function  $r_i$  is such that there exists  $z > 0$  such that  $r_i(0, 0, \dots, z, 0, \dots, 0) \neq r_i(0, \dots, 0) + z$  (in the first term  $z$  is in the  $i$ th argument of  $r_i$ ). Then no dominant strategy exists for player  $i$ .*

In light of this negative result, we would like to identify strategies that players could choose. We therefore focus on a concept that formalizes the idea that the players are “prudent” in the sense that they want to avoid the *winner’s curse* by winning the item at a too high price. Such a player  $i$  could argue as follows: if his actual type is no more than the currently highest bid among players  $1, \dots, i - 1$ , then he can safely bid up to the currently highest bid. On the other hand, if his actual type is higher than the currently highest bid among players  $1, \dots, i - 1$ , then he should bid truthfully (overbidding can result in a winner’s curse and underbidding can result in losing). Lemma □ below shows that the above intuition is captured by the following definition.

**Definition 1.** *We call a strategy  $s_i(\cdot)$  of player  $i$  **optimal** if for all  $\theta \in \Theta$  and all  $\theta'_i \in \Theta_i$*

$$u_i((f, t)(s_i(\theta_1, \dots, \theta_i), \theta_{-i}), \theta_i) \geq u_i((f, t)(\theta'_i, \theta_{-i}), \theta_i).$$

By choosing truth telling as the strategies of players  $j \neq i$  we see that each dominant strategy is optimal. For player  $n$  the concepts of dominant and optimal strategies coincide.

Definition 1 is a natural relaxation of the notion of dominant strategy as it calls for optimality w.r.t. a restricted subset of the other players' strategies. Call a strategy of player  $j$  *memoryless* if it does not depend on the types of players  $1, \dots, j - 1$ . Then a strategy  $s_i(\cdot)$  of player  $i$  is optimal if for all  $\theta \in \Theta$  it yields a best response to all joint strategies of players  $j \neq i$  in which the strategies of players  $i + 1, \dots, n$  are memoryless. In particular, an optimal strategy is a best response to the truth telling by players  $j \neq i$ .

The following lemma provides the announced characterization of optimal strategies. For any  $i$ , define  $\bar{\theta}_i := \max_{j \in \{1, \dots, i-1\}} \theta_j$ . We stipulate here and elsewhere that for  $i = 1$  we have  $\bar{\theta}_1 = -1$  so that for  $i = 1$  we have  $\theta_i > \bar{\theta}_i$ .

**Lemma 1.** *In each sequential Groves auction a strategy  $s_i(\cdot)$  is optimal for player  $i$  if and only if the following holds for all  $\theta_1, \dots, \theta_i$ :*

- (i) *Suppose  $\theta_i > \bar{\theta}_i$  and  $i < n$ . Then  $s_i(\theta_1, \dots, \theta_i) = \theta_i$ .*
- (ii) *Suppose  $\theta_i > \bar{\theta}_i$  and  $i = n$ . Then  $s_i(\theta_1, \dots, \theta_i) > \bar{\theta}_i$ .*
- (iii) *Suppose  $\theta_i \leq \bar{\theta}_i$  and  $i < n$ . Then  $s_i(\theta_1, \dots, \theta_i) \leq \bar{\theta}_i$ .*
- (iv) *Suppose  $\theta_i < \bar{\theta}_i$  and  $i = n$ . Then  $s_i(\theta_1, \dots, \theta_i) \leq \bar{\theta}_i$ .*

Note that no conclusion is drawn when  $\theta_n = \max_{j \in \{1, \dots, n-1\}} \theta_j$ . Player  $n$  can place then an arbitrary bid.

The following simple observation, see [3], provides us with a sufficient condition for checking whether a strategy is optimal in a sequential Groves mechanism.

**Lemma 2.** *Consider a Groves mechanism  $(f, t)$ . Suppose that  $s_i(\cdot)$  is a strategy for player  $i$  such that for all  $\theta \in \Theta$ ,  $f(s_i(\theta_1, \dots, \theta_i), \theta_{-i}) = f(\theta)$ . Then  $s_i(\cdot)$  is optimal in the sequential version of  $(f, t)$ .*

In particular, truth telling is an optimal strategy.

## 5 Sequential BC Auctions

As explained in the Introduction the BC mechanism cannot be improved upon in the simultaneous case, as shown in [2]. As we shall see here, the final social welfare can be improved in the sequential setting by appropriate optimal strategies that deviate from truth telling.

Theorem 1 applies for the BC auction, therefore no dominant strategies exist. We will thus focus on the notion of an optimal strategy. As implied by Lemma 2 many natural optimal strategies exist. In the sequel we will focus on the following optimal strategy which is tailored towards welfare maximization as we exhibit later on:

$$s_i(\theta_1, \dots, \theta_i) := \begin{cases} \theta_i & \text{if } \theta_i > \max_{j \in \{1, \dots, i-1\}} \theta_j \\ (\theta_1, \dots, \theta_{i-1})_1^* & \text{if } \theta_i \leq \max_{j \in \{1, \dots, i-1\}} \theta_j \\ & \text{and } i \leq n - 1 \\ (\theta_1, \dots, \theta_{i-1})_2^* & \text{otherwise} \end{cases} \tag{1}$$

According to strategy  $s_i(\cdot)$  if player  $i$  cannot be a winner when bidding truthfully he submits a bid that equals the highest current bid if  $i < n$  or the second highest current bid if  $i = n$ . Note that  $s_i(\cdot)$  is indeed optimal in the sequential BC auction, since Lemma 2 applies.

We now exhibit that within the universe of optimal strategies, if all players follow  $s_i(\cdot)$ , maximal social welfare is generated. Given  $\theta$  and a vector of strategies  $s(\cdot)$ , define the final social welfare of a sequential mechanism  $(f, t)$  as:

$$SW(\theta, s(\cdot)) = \sum_{i=1}^n u_i((f, t)([s(\cdot), \theta]), \theta_i) = \sum_{i=1}^n v_i(f([s(\cdot), \theta]), \theta_i) + \sum_{i=1}^n t_i([s(\cdot), \theta]).$$

**Theorem 2.** *In the sequential BC auction for all  $\theta \in \Theta$  and all vectors  $s'(\cdot)$  of optimal players' strategies,*

$$SW(\theta, s(\cdot)) \geq SW(\theta, s'(\cdot))$$

where  $s(\cdot)$  is the vector of strategies  $s_i(\cdot)$  defined in 1.

The maximal final social welfare of the sequential BC auction under  $s(\cdot)$  is always greater than or equal to the final social welfare achieved in a BC auction when players bid truthfully.

## 6 Implementation in Safety-Level Equilibrium

In this section we clarify the status of the strategies studied in Section 5 by analyzing what type of equilibrium they form. The notion of an ex-post equilibrium is somewhat problematic, since in pre-Bayesian games (the games we study here are a special class of such games) it has a different status than Nash equilibrium in strategic games. Indeed, as explained in 11, there exist pre-Bayesian games with finite sets of types and actions in which no ex-post equilibrium in mixed strategies exists.

The vector of strategies  $s_i(\cdot)$  defined in 10 is not an ex-post equilibrium in the sequential BC auction. Indeed, take three players and  $\theta = (1, 2, 5)$ . Then for player 1 it is advantageous to deviate from  $s_1(\cdot)$  strategy and submit, say 4. This way player 2 submits 4 and player's 1 final utility becomes  $4/3$  instead of  $2/3$ .

We believe that an appropriate equilibrium concept for the (sequential) pre-Bayesian games is the safety-level equilibrium introduced by 4 and 11 and defined in Section 3. In the case of sequential mechanisms it captures a cautious approach by focusing on each player's guaranteed payoff in view of his lack of any information about the types of the players who bid after him. We have the following result.

**Theorem 3.** *The vector of strategies  $s_i(\cdot)$  defined in 11 is a safety-level equilibrium in the sequential BC auction.*

One natural question is whether one can extend our Theorem 2 to show that our proposed vector of strategies in 10 generates maximal social welfare among all safety-level equilibria. The answer to this turns out to be negative as illustrated by the next example:



*Example 1.* Consider truth telling as the strategy for players  $1, \dots, n - 2$  and  $n$ . For any  $i$ , define  $\hat{\theta}_i := \max_{j \in \{1, \dots, i-1\}} [s(\cdot), \theta]_j$ . For player  $n - 1$  define the strategy:

$$s'_{n-1}(\theta_1, \dots, \theta_{n-1}) = \begin{cases} \hat{\theta}_{n-1} + \epsilon & \text{if } \theta_{n-1} > \hat{\theta}_{n-1}, \\ \theta_{n-1} & \text{otherwise.} \end{cases}$$

where  $\epsilon$  is a positive number in the interval  $(\hat{\theta}_{n-1}, \theta_{n-1})$ . This vector of strategies forms a safety-level equilibrium (we omit the proof here due to lack of space). Consider now the vector  $\theta = (0, 0, \dots, 1, 15, 16)$ . The sum of taxes under the set of strategies we have defined will be  $\frac{2\epsilon}{n}$ . On the other hand, under the vector  $s(\cdot)$  defined in (1), the sum of the taxes is  $\frac{2}{n}(15 - 1)$ .  $\square$

The set of safety-level equilibria is quite large. The above example illustrates that we can construct many other safety-level equilibria, by slight deviations from the truth telling strategy. In fact, there are even equilibria in which some players overbid and yet for some type vectors they generate higher social welfare than our proposed strategies. These equilibria, however, may be unlikely to form by prudent players and Theorem 2 guarantees that among equilibria where all players are prudent our proposed strategies generate maximal welfare.

Finally, if we assume that players select only optimal strategies, then we could consider an ex-post equilibrium in the universe of optimal strategies. We have then the following positive result.

**Theorem 4.** *If we allow only deviations to optimal strategies, then in the sequential BC auction, the vector of strategies  $s_i(\cdot)$  defined in (1) is an ex-post equilibrium that is also Pareto optimal.*

## 7 Final Remarks

This paper and our previous recent work, [3], forms part of a larger research endeavour in which we seek to improve the social welfare by considering sequential versions of commonly used incentive compatible mechanisms. The main conclusion of [3] and of this work is that in the sequential version of single-item auctions and public project problems there exist optimal strategies that deviate from truth telling and can increase the social welfare. Further, the vector of these strategies generates the maximal social welfare among the vectors of all optimal strategies. Here, we also showed that the vector of the introduced strategies forms a safety-level equilibrium.

We would like to undertake a similar study of the sequential version of the incentive compatible mechanism proposed in [17], concerned with purchasing a shortest path in a network.

## Acknowledgements

We thank Hervé Moulin and Moshe Tennenholtz for helpful discussions on the subject of this paper.

## References

1. Aghassi, M., Bertsimas, D.: Robust game theory. *Mathematical Programming* 107, 231–273 (2006)
2. Apt, K.R., Conitzer, V., Guo, M., Markakis, V.: Welfare undominated Groves mechanisms. In: Papadimitriou, C., Zhang, S. (eds.) *WINE 2008*. LNCS, vol. 5385, pp. 426–437. Springer, Heidelberg (2008)
3. Apt, K.R., Estévez-Fernández, A.: Sequential pivotal mechanisms for public project problems, *Computing Research Repository, CoRR* (2008), <http://arxiv.org/abs/0810.1383>; To appear in *SAGT 2009: 2nd Symposium on Algorithmic Game Theory*
4. Ashlagi, I., Monderer, D., Tennenholtz, M.: Resource selection games with unknown number of players. In: *AAMAS 2006: Proceedings 5th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, pp. 819–825. ACM Press, New York (2006)
5. Bailey, M.J.: The demand revealing process: To distribute the surplus. *Public Choice* 91(2), 107–126 (1997)
6. Cavallo, R.: Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In: *AAMAS 2006: Proceedings of the 5th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, pp. 882–889. ACM Press, New York (2006)
7. Cavallo, R.: Efficiency and redistribution in dynamic mechanism design. In: *EC 2008: ACM Conference on Electronic Commerce*, pp. 220–229 (2008)
8. Clarke, E.: Multipart pricing of public goods. *Public Choice* 11, 17–33 (1971)
9. 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 Press, New York (2007)
10. Guo, M., Conitzer, V.: Better redistribution with inefficient allocation in multi-unit auctions with unit demand. In: *EC 2008: Proceedings of the 8th ACM Conference on Electronic Commerce*, pp. 210–219 (2008)
11. Guo, M., Conitzer, V.: Optimal-in-expectation redistribution mechanisms. In: *AAMAS 2008: Proceedings of the 7th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems*, pp. 1047–1054 (2008)
12. Guo, M., Conitzer, V.: Undominated VCG redistribution mechanisms. In: *AAMAS 2008: Proc. of 7th Int. Conf. on Autonomous Agents and Multi Agent Systems* (2008)
13. Hartline, J., Roughgarden, T.: Optimal mechanism design and money burning. In: *STOC 2008: ACM Symposium on Theory of Computing*, pp. 75–84 (2008)
14. Krishna, V.: *Auction Theory*, 3rd edn. Academic Press, New York (2002)
15. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Microeconomic Theory*. Oxford University Press, Oxford (1995)
16. Moulin, H.: Efficient, strategy-proof and almost budget-balanced assignment, Working Paper (March 2007)
17. Nisan, N., Ronen, A.: Algorithmic mechanism design (extended abstract). In: *STOC*, pp. 129–140 (1999)

# Subsidized Prediction Markets for Risk Averse Traders<sup>\*</sup>

Stanko Dimitrov<sup>1</sup>, Rahul Sami<sup>2</sup>, and Marina Epelman<sup>1</sup>

<sup>1</sup> Department of Industrial and Operations Engineering

<sup>2</sup> School of Information

University of Michigan, Ann Arbor, MI 48109

{sdimitro, rsami, mepelman}@umich.edu

**Abstract.** In this paper we study the design and characterization of prediction markets in the presence of traders with unknown risk-aversion. We formulate a series of desirable properties for any “market-like” forecasting mechanism. We present a randomized mechanism that satisfies all these properties while guaranteeing that it is myopically optimal for each trader to trade honestly, regardless of her degree of risk aversion. We observe, however, that the mechanism has an undesirable side effect: the traders’ expected reward, normalized against the inherent value of their private information, decreases exponentially with the number of traders. We prove that this is unavoidable: any mechanism that is myopically strategyproof for traders of all risk types, while also satisfying other natural properties of “market-like” mechanisms, must sometimes result in a player getting an exponentially small normalized expected reward.

## 1 Introduction and Related Work

Prediction markets are markets designed and deployed to aggregate information about future events by having agents with private beliefs trade in these markets. One market format that is gaining in popularity is the *market scoring rule* [8]. A market scoring rule is a market mechanism with an automated market maker that guarantees liquidity, effectively subsidizing the market to incentivize trade.

Hanson [8] has shown that, for risk-neutral agents who are myopic, it is optimal for each to reveal their true beliefs on the traded event. This results leaves two questions. The first question, partially addressed by Chen *et al.* [6], is: *What happens when agents take into account future payoffs?*

In this paper, we tackle the second question: *What happens when agents are not risk-neutral?* In practice, most people are better modeled as being risk-averse in their decision making. Therefore, we model traders as expected-utility maximizers with an arbitrary weakly monotone and concave utility function. Current prediction market mechanisms, like the Market Scoring Rule or the Dynamic Pari-mutuel Market [11], do not always give appropriate incentives to risk-averse traders. For example, a sufficiently risk-averse informed trader, who knows that an event will occur with 80% probability

---

<sup>\*</sup> This work was supported by the National Science Foundation under award CCF-0728768. The proofs have been omitted from this short paper; the extended version can be found at <http://www.umich.edu/~rsami/papers/RiskAverseMarkets.pdf> [7]

even though it is currently priced at 50%, may not want to push up the price in a market because of the 20% chance of making a loss. This suggests that current subsidized prediction markets may converge to a non-truthful price in a sequential equilibrium.

If traders have known risk aversion, the scoring rules could be adjusted, retaining the original incentive properties. In this paper, we focus on the setting where traders have *unknown* risk aversion, and study whether it is possible to modify the market mechanism to guarantee myopic honesty and preserve other desirable properties. We first list a set of properties that any prediction market-like mechanism must satisfy: (1) myopic strategyproofness; (2) sequential trade, giving traders the opportunity to update beliefs; (3) a variant of *sybilproofness*, capturing the idea that trading under multiple identities does not yield any direct advantage; and (4) the expected subsidy should be bounded.

We propose one mechanism that satisfies all of these properties in the presence of traders with unknown risk-averse preferences. The key building block of our result is a sweepstakes technique, developed by Allen [2]. Unfortunately, the proposed mechanism reduces the expected reward exponentially with the number of agents.

We then establish that exponentially decreasing rewards are unavoidable for any mechanism satisfying all the properties listed above. To exclude trivial examples of decreasing rewards, we normalize all rewards by a measure of the intrinsic informativeness of a trader's private information. We show that exponential decrease in the *normalized* expected reward is necessary for any mechanism that satisfies the properties we propose in the presence of arbitrarily risk-averse agents.

## 1.1 Related Work

In this section we discuss some of the previous work in prediction markets and scoring rules. Hanson [8] introduced the concept of a market scoring rule, a form of subsidized prediction market, and proved a myopic strategyproofness property for risk-neutral traders, as well as a bound on the total subsidy. Pennock [11] introduced another mechanism, the dynamic pari-mutuel market, for a subsidized prediction market. Both these mechanisms introduce some of the properties in section 2. However, both mechanisms assume risk neutrality of the traders, which is not assumed in this paper. Lambert *et al.* [9] introduce a class of self-financed wagering mechanism along with the properties such mechanisms must satisfy. The authors assume risk neutral traders and an absence of subsidy in the mechanism. Chen and Pennock [5], and later generalized by Agrawal *et al.* [1], consider a risk-averse market maker in a subsidized market and show that the market maker has bounded subsidy in most forms of risk aversion. However, unlike our paper, the incentive consequence of risk-averse *traders* is not addressed.

Several prediction market mechanisms are extensions of proper scoring rules. The notion of scoring rules was introduced by Brier [3], in the form of the quadratic scoring rule (which is proper), to measure the accuracy of weather forecasters. *Proper* scoring rules provide a way to reward forecasters such that honest reports are made. Most of the early work on scoring rules assumed that forecasters were risk neutral.

There has been some research on addressing risk aversion in scoring rules. Winkler and Murphy [12] showed that, if forecasters have a *known* risk type, scoring rules can be transformed to recapture the honest reporting property. One approach to handling

forecasters with unknown risk type, as proposed by Chen *et al.* [4] and Offerman *et al.* [10], is to figure out every participant's risk type by asking them a series of questions, and then calibrate their future reports using this data. This mechanism may work for a prediction market mechanism if the group of traders can be pre-screened. However, this may not be the case, and ideally we would like to have an "online" mechanism that can handle traders regardless of their risk type without any calibration. Allen [2] proposed one such "online" scoring rule for forecasters with arbitrary risk type. Our mechanism is based on Allen's result, and we discuss this idea in section 3.

## 2 Model, Notation, and Definitions

In this section and section 4 we consider a class of mechanisms defined by a set of properties satisfied by most subsidized prediction markets described in literature. We do not claim that the outlined properties are sufficient to completely characterize the space of prediction market mechanisms; rather, they identify a class of broad market-like mechanisms. We first describe our basic model of the information and interaction setting in which the mechanism operates, and then list the properties that the mechanisms we study must satisfy.

We consider a class of mechanisms designed to aggregate information from a set of *agents* (or *traders*) in order to forecast the outcome of a future event  $\omega$ . Each agent  $i$  receives a private information signal,  $s_i$ , relevant to the outcome of the event; we assume  $s_i$  is binary, as is  $\omega$ .

A central feature of the market-like mechanisms we consider is that the agents express a predicted probability of the event in the mechanism, through a sequence of public trades or *reports*. Other agents can update their beliefs based on the observed history of reports. We use  $r_k \in [0, 1]$  to denote the  $k$ th report made in the market, and let  $\mu_k = (r_1, \dots, r_{k-1})$  denote the history up to the start of the  $k$ th trade.  $r_k$  can thus depend on  $\mu_k$  as well as any private information available to the trader making the report. We let  $n$  denote the total number of trades in the market.

We assume the identity of the agents making the reports cannot be verified, and the total number of agents participating is unknown. As each agent's signal is static, there is no need for any agent to trade more than once. Therefore, we will treat each report as if were made by separate traders and is natural for a market setting. However, an agent may masquerade as multiple agents, which is a consideration of sybilproofness.

Once the true outcome of the event is realized,  $\omega = 1$  if the event occurs and  $\omega = 0$  otherwise, the mechanism determines the reward for every agent. The reward for agent  $i$ ,  $\rho(r_i, \mu_i, n, \omega)$ , is a function of the agent's report, market state, the total number of agents participating in the mechanism, and the event outcome. We allow the mechanism to randomize the distribution of the rewards, and we propose one such mechanism in section 3. We assume that the reward does not depend on the value of any reports made in the future. This is a nontrivial technical assumption that enables us to simplify the analysis of agents' myopic strategies, as agents can make decisions based on their current beliefs about the outcome, without forming beliefs about future agents' signals and strategies. This assumption is satisfied by most securities markets as well as market scoring rule markets, but not necessarily true for pari-mutuel markets.

Every agent  $i$  values the distributed reward according to her value function  $V_i(\cdot)$ , where  $V_i(\cdot)$  is a weakly monotone increasing concave function. We make the normalizing assumption that  $V_i(0) = 0$ . In order to make her report, an agent maximizes her expected reward, with respect to her true belief  $p_i$ , over the outcome of the event and any randomization of the mechanism over the rewards, written as  $E_{\omega \sim p_i} V_i(\rho(r, \mu_i, n, \omega))$ . Though there may be other sources of uncertainty in the mechanism, we do not consider them in our model.

We identify the properties mechanisms should satisfy by examining literature in prediction market design and other wagering mechanisms. Hanson [8], in introducing the market scoring rule, had a subsidized prediction market be *myopically strategy proof* and have *bounded market subsidy*, both defined below. The same properties also hold in the dynamic pari-mutuel market introduced by Pennock [11]. As both of the mechanisms were subsidized, both had *guaranteed liquidity* by having a market maker that is always willing to trade with an agent. Prediction markets provide *anonymity*, i.e., the reward given due to a report is independent of who made the report. Finally, prediction markets are *sybilproof*, meaning that an agent reporting once with some information is no better off reporting twice in the market with the exact same information. Though anonymity and sybilproofness were not explicitly stated by Hanson or Pennock, they still hold in their proposed mechanisms and were explicitly defined by Lambert et al. [9]. We use a relaxed version of sybilproofness by requiring agents to be no better off reporting twice, as opposed to having the same payoff as presented by Lambert et al.. Using the notation established above, we formally define the desired properties:

**P1: Myopically Strategyproof:** If an agent making trade  $i$  has true belief  $p_i$ , and trades only once in a market, her reported belief will be her true belief. Mathematically,

$$\forall n, i \in \{1..n\}, \mu_i \ p_i = \operatorname{argmax}_{r \in [0,1]} E_{\omega \sim p_i} V_i(\rho(r, \mu_i, n, \omega)). \tag{1}$$

Further, we also require  $\max_{r \in [0,1]} E_{\omega \sim p_i} V_i(\rho(r, \mu_i, n, \omega)) \geq 0$  so that myopic strategyproofness includes a standard individual rationality condition.

**P2: Sybilproofness:** An agent is no worse off reporting once honestly than making any two consecutive reports  $r^{(1)}, r^{(2)}$  with the same information. Mathematically,

$$\forall n, i \in \{1..n\}, \mu_i \ E_{\omega \sim p_i} V_i(\rho(p_i, \mu_i, n, \omega)) \geq E_{\omega \sim p_i} V_i(\rho(r^{(1)}, \mu_i, n + 1, \omega) + \rho(r^{(2)}, \mu_{i+1}, n + 1, \omega)). \tag{2}$$

**P3: Bounded Subsidy:** The expected subsidy the market maker needs to invest into the market is bounded by a value  $\beta$ :

$$\forall n, i \in \{1..n\}, \mu_i, r_i \ \sum_{\text{all players } i} E_{\omega} \rho(r_i, \mu_i, n, \omega) < \beta$$

To summarize, we define the class of market-like mechanisms to be all mechanisms that are *anonymous*, *guarantee liquidity*, *myopically strategy proof*, *sybilproof*, and have *bound market subsidy*.

Before we introduce our results, we must introduce the concepts of information structure, report informativeness, and normalized expected reward.

**Information Structure:** We define an *information structure* to consist of a set of possible signal realizations for each trader, and the posterior probability of events given a subset of signal realizations (equivalently, the joint probability of signal realizations and the true outcomes).

**Informativeness:** For a given information structure, we define *informativeness* of an agent  $k$ , given a history  $\mu_k$ , as the expected reduction in forecasting error, as measured by the reduction in entropy of the event, after conditioning on  $k$ 's signal.

**Normalized Expected Reward:** The informativeness and the reward of each agent may deviate. Therefore, in order to compare the reward an agent receives from a report, we define the *normalized expected reward* as the ratio of the expected reward to the informativeness of the report given the history up to that point.

### 3 Proposed Mechanism

In this section we review the work presented by Allen [2] and then present one mechanism that satisfies the properties outlined in section 2.

Allen shows that an agent with a monotone value function  $V(\cdot)$  with unknown risk preference will set her report,  $\hat{p}$ , to her true belief,  $p$ , on an outcome  $\omega$  if she participates in a sweepstakes. According to the sweepstakes, the agent will receive a reward of 1 with probability  $q(\hat{p}) = 1 - (1 - \hat{p})^2$  if the event occurs and probability  $\hat{q}(\hat{p}) = 1 - \hat{p}^2$  if the event does not occur. Allen's result follows from the fact that the expected value is linear in probabilities and the value function is monotonically increasing.

Now consider the following serial sweepstakes that is a derivative of Allen's result:

1. An agent observes the previous agents' reports, and plays an individual sweepstake as defined by Allen with sweepstake functions described above.
2. The outcome of the event is observed.
3. If there are  $n$  reports in the mechanism, then each player reporting  $\hat{p}$  wins a reward of 1 with probability  $q(\hat{p}) = \frac{1}{4^n}(1 - (1 - \hat{p})^2)$  if the event occurs and  $\hat{q}(\hat{p}) = \frac{1}{4^n}(1 - \hat{p}^2)$  if the event does not occur.

The mechanism above possesses all of the properties outlined in section 2; however, the mechanism distributes rewards that are exponentially decreasing with the number of agents. Moreover, if all the reports are equally informative, then the normalized expected reward also decreases exponentially with the number of reports.

### 4 Impossibility Result

**Theorem 1.** *If an anonymous, guaranteed liquidity mechanism satisfies properties P1–P3, then, there is a family of information structures  $I^{(n)}$ , each parameterized by a number  $n$  of agents, such that, even if all agents perform perfect Bayesian updating according to the structure  $I^{(n)}$  and report their posteriors honestly, the minimum normalized expected reward of an agent must decrease exponentially with  $n$ .*



We start by showing that if agents with arbitrary risk-averse preferences are to participate in a mechanism in our class, all rewards must be non-negative. We then observe that the informativeness of a report is a constant multiple of the square of the differences between the posteriors after every report, so long as the posteriors are bounded in  $[0.5 - c, 0.5 + c]$ , for  $c \leq 0.2$ . We then show that for any two sequential reports made under two different posterior beliefs, the expected reward from the reports under the first posterior is a constant multiple of the expected reward under the second posterior, so long as the two posterior beliefs are bounded within  $[0.5 - c, 0.5 + c]$ ,  $c \leq 0.2$ .

The theorem proof follows by inductively building a family of information structures starting with the structure of two agents both making symmetric reports. In the base structure  $I^{(2)}$ , two agents start with a common prior of 0.5 on the event and receive a binary signal. If the agents report honestly, each agent will change the posterior report by  $\pm \frac{c}{2}$ . This results in information structure  $I^{(2)}$  with bounded posteriors between  $[0.5 - c, 0.5 + c]$ . We now consider a sybil attack in this setting, where one of the agents is making two reports under the same priors. To make this consideration we construct  $I^{(3)}$ . This construction is dependent on the mechanism: If the expected reward of the first agent is larger than 4 times the expected reward of the second, we construct a structure  $I^{(3')}$ , otherwise  $I^{(3'')}$ .

In either case, we split one of the reports into two, such that the histories up to the split report are consistent with  $I^{(2)}$ . By the sybilproofness property, we know that the sum of the expected rewards of the split reports is no larger than the original. Even accounting for the reduction in informativeness, we show that there is a report with normalized expected reward of at most  $\gamma$  times that of the split report in  $I^{(2)}$ , where  $\gamma = \frac{0.5+c}{0.5-c} \cdot \frac{80}{81} < 1$  for a suitable  $c$ .

From  $I^{(3)}$  we construct  $I^{(4)}$  in a similar manner. Iteratively applying this procedure we show that there exists at least one report in  $I^{(n)}$  with normalized expected reward that is exponentially smaller than one of the two reports in  $I^{(2)}$ .

## 5 Conclusion

In this paper we present one mechanism that satisfies properties in section 2 that allows agents with arbitrary risk-averse value function to participate. However, this mechanism requires that the normalized expected reward exponentially decrease with the number of agents. We then show that as long as the risk aversion structure of the agents is not known, for any mechanism in the class of interest that allows agents with arbitrary risk-averse value functions to participate, the normalized expected reward must decrease exponentially with the number of agents.

## References

- [1] Agrawal, S., Delage, E., Peters, M., Wang, Z., Ye, Y.: A unified framework for dynamic parimutuel information market design. In: EC 2009: Proceedings of the tenth ACM conference on Electronic commerce, pp. 255–264. ACM, New York (2009)
- [2] Allen, F.: Discovering personal probabilities when utility functions are unknown. *Management Science* 33(4), 542–544 (1987)



- [3] Brier, G.W.: Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78(1), 1–3 (1950)
- [4] Chen, K.-Y., Fine, L.R., Huberman, B.A.: Forecasting uncertain events with small groups. In: *Proceedings of the Third ACM Conference on Electronic Commerce (EC 2001)*, Tampa, Florida, pp. 58–64 (2001)
- [5] Chen, Y., Pennock, D.M.: A utility framework for bounded-loss market makers. In: *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI 2007)*, pp. 49–56 (2007)
- [6] Chen, Y., Dimitrov, S., Sami, R., Reeves, D., Pennock, D., Hanson, R., Fortnow, L., Gonen, R.: Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*
- [7] Dimitrov, S., Sami, R., Epelman, M.: Subsidized prediction markets for risk averse traders (2009) (manuscript), <http://www.umich.edu/~rsami/papers/RiskAverseMarkets.pdf>
- [8] Hanson, R.: Combinatorial information market design. *Information Systems Frontiers* 5(1), 107–119 (2003)
- [9] Lambert, N.S., Langford, J., Wortman, J., Chen, Y., Reeves, D., Shoham, Y., Pennock, D.M.: Self-financed wagering mechanisms for forecasting. In: *EC 2008: Proceedings of the 9th ACM conference on Electronic commerce*, pp. 170–179. ACM, New York (2008)
- [10] Offerman, T., Sonnemans, J., van de Kuilen, G., Wakker, P.P.: A truth-serum for non-bayesians: Correcting proper scoring rules for risk attitudes. Working paper (June 2008), <http://www1.fee.uva.nl/creed/pdf/files/propscr30jun08.pdf>
- [11] Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: *EC 2004: Proceedings of the 5th ACM conference on Electronic commerce*, pp. 170–179. ACM Press, New York (2004)
- [12] Winkler, R.L., Murphy, A.H.: Nonlinear utility and the probability score. *Journal of Applied Meteorology* 9, 143–148 (1970)

# Envy, Multi Envy, and Revenue Maximization

Amos Fiat and Amiram Wingarten

School of Computer Science, Tel Aviv University

**Abstract.** We study the envy free pricing problem faced by a seller who wishes to maximize revenue by setting prices for bundles of items. If there is an unlimited supply of items and agents are single minded then we show that finding the revenue maximizing envy free allocation/pricing can be solved in polynomial time by reducing it to an instance of weighted independent set on a perfect graph.

We define an allocation/pricing as *multi envy free* if no agent wishes to replace her allocation with the union of the allocations of some set of other agents and her price with the sum of their prices. We show that it is *coNP*-hard to decide if a given allocation/pricing is multi envy free. We also show that revenue maximization multi envy free allocation/pricing is *APX* hard.

Furthermore, we give efficient algorithms and hardness results for various variants of the highway problem.

## 1 Introduction

We consider the combinatorial auction setting where there are several different items for sale, not all items are identical, and agents have valuations for subsets of items. We allow the seller to have identical copies of an item. We distinguish between the case of limited supply (e.g., physical goods) and that of unlimited supply (e.g., digital goods). Agents have known valuations for subsets of items. We assume free disposal, i.e., the valuation of a superset is  $\geq$  the valuation of a subset. Let  $S$  be a set of items, agent  $i$  has valuation  $v_i(S)$  for set  $S$ . The valuation functions,  $v_i$ , are public knowledge. Ergo, we are not concerned with issues of truthfulness or incentive compatible bidding. Our concern here is to maximize revenue in an envy free manner.

Our goal is to determine prices for sets of items while (approximately) maximizing revenue. The output of the mechanism is a payment function  $p$  that assigns prices to sets of items and an allocation  $a$ . Although there are exponentially many such sets, we will only consider payments functions that have a concise representation. For a set of items  $S$  let  $p(S)$  be the payment required for set  $S$ . Let  $a_i$  be the set assigned to agent  $i$ .

In general, every agent  $i$  has valuation function  $v_i$  defined over every subset of items.

Given a payment function  $p$ , and a set of valuation functions  $v_i$ , let  $z_i = \max_S (v_i(S) - p(S))$ , and let  $\mathcal{S}_i$  to be a collection of sets such that  $S \in \mathcal{S}_i$  if and only if  $v_i(S) - p(S) = z_i$ .

We now distinguish between two notions of envy freeness.

**Definition 1.** We say that  $(a, p)$  is envy free if

- If  $z_i > 0$  then  $a_i \in \mathcal{S}_i$ .
- If  $z_i = 0$  then either  $a_i \in \mathcal{S}_i$  or  $a_i = \emptyset$ .
- If  $z_i < 0$  then  $a_i = \emptyset$ .

**Definition 2.** A pricing  $p$  is monotone if for each subset  $S$  and for each collection of subsets  $C$  such that  $S \subseteq \bigcup_{T \in C} T$  the following inequality holds:  $p(S) \leq \sum_{T \in C} p(T)$ .

**Definition 3.** An allocation/pricing  $(a, p)$  is multi envy free if it is envy free and its pricing is monotone.

Clearly, multi envy-freeness is a more demanding requirement than envy-freeness, so any allocation/pricing that is multi envy-free is also envy-free. In item pricing, a price is set for every item, identical copies of an item are priced the same, and the price of a set is the sum of the individual item prices. In subset pricing one may assign a sets of items prices that cannot be consistently expressed as a sum of the item prices comprising the set. E.g., discounts for volume would not generally be consistent with item pricing.

In the unlimited supply setting, item pricing is always multi envy-free (and hence also envy-free). With a limited supply of items, achieving item-pricing [multi] envy-freeness is not automatic. Circumstances may arise where some agent has a valuation less than the price of some set of items she is interested in, but there is insufficient supply of these items. An envy-free solution must avoid such scenarios. Even so, for limited or unlimited supply, item pricing is envy-free if and only if item pricing is multi envy-free (this follows from the monotonicity of the item pricing).

For subset pricing, it does not necessarily follow that every allocation/pricing that is envy-free is also multi envy-free.

Although the definitions above are valid in general, we are interested in single minded bidders, and more specifically in a special case of single minded bidders called the highway problem ([7,11]) where items are ordered and agents bid for consecutive interval of items.

## 2 Our Results

Unfortunately, due to lack of space we have omitted all proofs and constructions with one exception. In this extended abstract we show that for unlimited supply, and single minded bidders, finding the envy free allocation/pricing that maximizes revenue is polynomial time. Missing proofs can be found in the full version of this paper [5].

In the full version of this paper we show gaps in revenue between the item pricing (where envy freeness and multi envy freeness are equivalent), multi envy freeness, and envy freeness. These results are summarized in Table 1. These gaps are for single minded bidders. In all cases (single minded bidders or not)

**Table 1.** Revenue gaps for single minded bidders ( $n$  items,  $m$  agents)

	Lower Bound	[Multi] Envy-free item pricing	Multi envy-free subset pricing	Envy-free subset pricing	Type of Instance
Limited	#1	$H_n$ $H_m$	$n$ $m$		Highway
Unlimited	#2		1 1	$H_n$ $H_m$	Single Minded
Unlimited	#3	1 1	$H_n$ $\log \log m$		Highway

$$\begin{aligned}
 \text{Revenue}([\text{Multi}] \text{ EF item pricing}) &\leq \text{Revenue}(\text{Multi EF subset pricing}) \\
 &\leq \text{Revenue}(\text{EF subset pricing}) \\
 &\leq \text{Social Welfare}.
 \end{aligned}$$

Clearly, if a lower bound holds for unlimited supply it also holds for (big enough) limited supply.

All of our lower bound constructions are for single minded bidders, for single minded bidders with unlimited supply the bounds are almost tight as from [7] it follows directly:

$$(\text{Social welfare})/(\text{Envy-free item pricing}) \leq H_m + H_n.$$

In the full version of this paper, we also show the following:

1. The decision problem of whether an allocation/pricing is multi envy free is *coNP*-hard.
2. Finding an allocation/pricing that is multi envy free and maximizes the revenue is *APX*-hard.
3. For the the highway problem, if all capacities are  $O(1)$  then the (exact) revenue maximizing *envy free* allocation/pricing can be computed in polynomial time. *I.e.*, the problem is fixed parameter tractable with respect to the capacity.
4. For the highway problem with  $O(1)$  capacities, we give a FPTAS for revenue maximization on the more difficult Multi envy-free setting.

### 3 Related Work

Much of the work on envy free revenue maximization is on item pricing rather than on subset pricing. Guruswami et al. [7] give an  $O(\log m + \log n)$ -approximation for the general single minded problem, where  $n$  is the number of items and  $m$  is the number of agents. This result was extended by Balcan et al. [2] to an  $O(\log m + \log n)$ -approximation for arbitrary valuations and unlimited supply using single fixed pricing which is basically pricing all bundles with the same price. Demaine et al. [4] show that the general item pricing problem with unlimited availability of items is hard to approximate within a (semi-)logarithmic factor.

## 4 Notation and Definitions

The capacity of an item is the number of (identical) copies of the item available for sale. The supply can be *unlimited supply* or *limited supply*. In a *limited supply* seller is allowed to sell up to some fixed amount of copies of each item. In the *unlimited supply* setting, there is no limit as to how many units of an item can be sold.

We consider *single-minded* bidders, where each agent has a valuation for a bundle of items,  $S_i$ , and has valuation 0 for all sets  $S$  that are not supersets of  $S_i$ . The valuation function for  $i$ ,  $v_i$ , has a succinct representation as  $(S_i, v_i)$  where  $v_i = v_i(S_i)$ . For every  $S$  such that  $S_i \subseteq S$ ,  $v_i(S) = v_i(S_i) = v_i$ , for all other sets  $S'$ ,  $v_i(S') = 0$ . Without loss of generality, an allocation/pricing must either have  $a_i = \emptyset$  and  $p_i(a_i) = 0$  or  $a_i = S_i$  and  $p_i(a_i) \leq v_i(S_i)$ .

For single minded bidders, the definition of envy free can be simplified:

**Observation 4.** *For single minded bidders an allocation/pricing is envy free if and only if*

1. *For any two agents  $i$  and  $j$  with non-empty allocations, if  $S_i \subseteq S_j$ , it must be that  $p(S_i) \leq p(S_j)$*
2. *For any agent  $i$  that receives nothing and any other agent  $j$  that receives  $S_j$ , where  $S_i \subseteq S_j$ , it must be that  $v(S_i) \leq p(S_j)$*

## 5 Polynomial Time Envy-Free Revenue Maximization (Unlimited Supply, Single Minded Bidders)

For the unlimited supply setting we show that:

**Theorem 5.** *For single minded bidders the revenue maximizing envy free allocation/pricing with unlimited supply can be computed in polynomial time.*

Allocating a bundle at price  $p$  means that any bundle that is a superset and has valuation  $< p$  must not be allocated. We transform the problem into a perfect graph  $H$  and then compute the revenue maximizing allocation/pricing by computing a maximal independent set on  $H$  (which can be done in polytime for perfect graph). A similar construction was used by Chen et al. [3] for envy free pricing with metric substitutability.

### 5.1 Construction of Graph $H$

For each  $i \in \{1, \dots, m\}$ , define

$$A(i) = \{1 \leq j \leq m \mid S_i \subseteq S_j \text{ and } v_j < v_i\}.$$

Given price  $p$  for agent  $i$ , all requests in  $A(i)$  with valuation  $< p$  cannot be allocated. For each agent  $i$ , define an ordering  $\pi_i$  on  $A(i)$  in non-decreasing order of valuation. I.e., for each pair  $j, k$  such that  $1 \leq j \leq k \leq n_i$ , where

$n_i = |A(i)|$ , the valuations must be ordered,  $v_{\pi_i(j)} \leq v_{\pi_i(k)}$  (ties are broken arbitrarily).

We construct an undirected vertex-weighted graph  $H$  as follows. For each agent  $i$  we associate  $n_i + 1$  weighted vertices. These vertices constitute the  $T(i)$  component, the vertices of which are  $\{i_1, i_2, \dots, i_{n_i+1}\}$ . The set of all  $H$  vertices is defined as  $\bigcup_{i \in V} T_i$ . The weight of each vertex in  $T(i)$  is given as follows:  $w(i_1) = v_{\pi_i(1)}$ ,  $w(i_2) = v_{\pi_i(2)} - v_{\pi_i(1)}$ ,  $\dots$ ,  $w(i_n) = v_{\pi_i(n_i)} - v_{\pi_i(n_i-1)}$ , and  $w(i_{n_i+1}) = v_i - v_{\pi_i(n_i)}$ .

By definition, all weights are non-negative and  $\sum_{j \in T(i)} w(j) = v_i$ . Some vertices in  $T(i)$  have edges to vertices of  $T(j)$  such that  $j \in A(i)$  (connecting a vertex  $i_k \in T(i)$  to a component  $T(j)$  means that there are edges from  $i_k$  to all vertices of  $T(j)$ ). More specifically, Vertex  $i_{n_i+1}$  is connected to all components  $T(j)$  for  $j \in A(i)$ . Vertex  $i_k$ ,  $1 \leq k \leq n_i$ , is connected to all components  $T(\pi_i(j))$  for all  $1 \leq j < k$ . *E.g.*, vertex  $i_3$  is connected to components  $T(\pi_i(1))$  and  $T(\pi_i(2))$  (See Figure [□](#)). It is easy to see that for any  $a < b$ ,  $i_b$  is connected to each component that  $i_a$  is connected to (and possibly to additional components).

**Lemma 6.** *The value of the maximum weighted independent set on  $H$  is equal to the revenue obtained from the optimal envy free pricing of our problem.*

*Proof.* We prove that a maximal revenue envy-free allocation/pricing can be translated into an independent set in  $H$  and that a maximal independent set in  $H$  can be translated into a revenue maximizing envy free allocation/pricing.

(envy free  $\Rightarrow$  independent set) It is easy to see that the price  $p_i$  is equal to one of the valuations  $v_j$  for  $j$  such that  $j \in A(i)$  or  $j = i$  (otherwise the prices can be increased). We will choose all vertices of  $i_k \in T(i)$  such that  $k \leq \pi(j)$ . By construction, the sum of weights is  $v_j = p_i$ .


By the envy free pricing we have that  $\forall i, j : S_i \subset S_j \Rightarrow p_i \leq p_j$ . If this does not result in an independent set, then there exist two vertices,  $i_k$  and  $j_m$ , such that  $j \in A(i)$  and there is an edge between them. But the price paid by agent  $i$  ( $\geq \sum_{1 \leq t < k} w(i_t)$ ) is greater than  $j$ 's valuation (by construction of  $H$ ) so  $p_i > p_j$ , a contradiction.

(independent set  $\Rightarrow$  envy free) By construction, a node  $i_k$  in component  $T(i)$  has an edge to all neighbors of  $i_m$ ,  $1 \leq m < k$ . Therefore, vertices in a maximal independent set from  $T(i)$  are of the form  $\{i_k | k \leq i^{max}\}$  for some  $1 \leq i^{max} \leq n_i$ . We transform the independent set into a pricing as follows:

- If none of  $T(i)$ 's vertices were chosen to the independent set then agent  $i$  receives nothing.
- If vertices  $\{i_k | k \leq i^{max}\} \subset T(i)$  were chosen to the independent set then agent  $i$  receives  $S_i$  at price  $\sum_{k \leq i^{max}} w(i_k)$

Assume that the pricing is not envy-free and we have requests  $i, j$  such that  $S_i \subset S_j$  and  $p_i > p_j$ :

- If  $p_j = v_j$  then  $j \in A(i)$ ,  $p_i > v_j$  which implies that  $i_{i^{max}}$  has an edge to all vertices of  $T(j)$ . But both  $i_{i^{max}}$  and  $j_1$  were chosen to the independent set, a contradiction.



**Fig. 1.** Pricing problems to independent set problems: Single minded agents and their bundles/valuations, each agent would like to buy a set of products (the black balls) at a price less than her valuation (the numbers in the bundle requests are valuations).  $A$  is represented as a dependency graph (top right), there is an edge from vertex  $i$  to  $A(i)$ , (note there is no edge between 2 and 7 since  $2 \leq 7$ ). The graph  $H$  appears in the center, the edges between  $T(i)$  components are described at the bottom.

- If  $p_j < v_j$  then let  $a$  be the minimum index in  $1, \dots, n_j$  such that vertex  $j_a$  was not chosen to the independent set. This is only possible because  $j_a$  has edges to some  $T(j')$ , at least one of whose vertices,  $j'_1$ , was chosen. Now,  $v_{j'} < p_j$  which implies that  $v_{j'} < p_i$ , thus  $j' \in A(i)$ . As  $i_{i^{max}}$  has an edge to all vertices of  $T(j')$ , we derive a contradiction to both  $i_{i^{max}}$  and  $j'_1$  being in the independent set.

**Lemma 7.** *H is a comparability graph.*

*Proof.* Orient an edge from a node  $v$  of  $T(i)$  to a node  $w$  of  $T(j)$  from  $v$  to  $w$  if  $j \in A(i)$ . It is not hard to show that this orientation guarantees transitivity which implies that the graph is a comparability graph.

A graph is said to be perfect if the chromatic number (the least number of colors needed to color the graph) of every induced subgraph equals the clique number of that subgraph. Lemma 7 implies:

**Corollary 8.** *H is a perfect graph.*

Maximal weighted independent set can be solved in polynomial time on perfect graphs [6]. By Lemma 6 and Corollary 8 we conclude that finding the optimal envy free allocation/pricing in the general single minded setting can be done in polynomial time. This completes the proof of Theorem 5.

## References

1. Balcan, M.-F., Blum, A.: Approximation algorithms and online mechanisms for item pricing. In: EC 2006: Proceedings of the 7th ACM conference on Electronic commerce, pp. 29–35 (2006)
2. Balcan, M.-F., Blum, A., Mansour, Y.: Item pricing for revenue maximization. In: EC 2008: Proceedings of the 9th ACM conference on Electronic commerce, pp. 50–59. ACM, New York (2008)
3. Chen, N., Ghosh, A., Vassilvitskii, S.: Optimal envy-free pricing with metric substitutability. In: EC 2008: Proceedings of the 9th ACM conference on Electronic commerce, pp. 60–69. ACM, New York (2008)
4. Demaine, E.D., Hajiaghayi, M.T., Feige, U., Salavatipour, M.R.: Combination can be hard: approximability of the unique coverage problem. In: SODA 2006: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm, pp. 162–171. ACM, New York (2006)
5. Fiat, A., Wingarten, A.: Envy, multi envy, and revenue maximization (2009), <http://arxiv.org/abs/0909.4569>
6. Grötschel, M., Lovász, L., Schrijver, A.: The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* 1, 169–197 (1981)
7. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: SODA 2005: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms, pp. 1164–1173 (2005)



# Nudging Mechanisms for Technology Adoption

Jacomo Corbo<sup>1</sup> and Yevgeniy Vorobeychik<sup>2</sup>

<sup>1</sup> The Wharton School of Business,  
University of Pennsylvania, Philadelphia PA 19104, USA  
jacomoc@wharton.upenn.edu

<sup>2</sup> School of Engineering and Applied Sciences,  
University of Pennsylvania, Philadelphia PA 19104, USA  
yev@seas.upenn.edu

**Abstract.** We study the adoption dynamics of two competing technologies and the efficacy of viral pricing strategies for driving adoption. Our model considers two incompatible technologies of differing quality and a market in which user valuations are heterogeneous and subject to network effects. We provide partial characterization results about the structure and robustness of equilibria and give conditions under which the higher quality technology purveyor can make significant inroads into the competitor's market share. We then show that myopic best-response dynamics in our setting are monotonic and convergent, and propose two pricing mechanisms that use this insight to help the entrant technology seller tip the market in its favor. Comparable implementations of both mechanisms reveals that the non-discriminatory strategy, based on a calculated public price subsidy, is less costly and just as effective as a discriminatory policy.

## 1 Introduction

We study the pattern and dynamics of adoption of two technologies of differing qualities, where consumers' valuations are heterogeneous and subject to network effects (*positive externalities*) from the installed base. Such positive externalities occur in a variety of settings: for example, the value of an Internet chat client to a user increases as more of his friends use the same chat client service. Network externalities have an important impact on buyer decisions and can be leveraged to sell a technology more effectively. In this paper, we study the efficacy of simple (discriminatory and non-discriminatory) pricing mechanisms to generate adoption cascades.

Throughout, we take the perspective of the higher quality (entrant) technology seller and study the market through the lens of incomplete information about actual user valuations. As such, we present probabilistic characterizations of equilibrium market shares under exogenously specified price and quality differences between two technologies, which are suggestive of the likelihood of successful market entry of the higher quality technology. We observe, for example, that substantial network effects can often stand in the way of significant

initial market penetration, even when the entrant offers a substantial quality improvement.

Even so, our characterization of best response dynamics offers a way for an entrant with the better product to tip the market considerably in its favor. We study two such pricing strategies, which we call “nudging” mechanisms, that leverage the monotonicity of best response dynamics and focus on enticing a small group of non-adopters to make the switch, with the hope that this fosters an adoption avalanche. In the full paper, we report that a non-discriminatory nudging mechanism based on a public price subsidy is as effective in growing market share as a discriminatory alternative that “seeds” (*provides a targeted subsidy to*) a sufficient number of non-adopters, all the while costing less. The full paper also addresses the general problem of profit maximization under uncertainty in our setting for both discriminatory and non-discriminatory pricing regimes. There, we show that the problem is NP-Hard in general, but observe that the simple nudging approach can serve as effective heuristics.

The role of network externalities on equilibrium adoption of standards and technologies has received much interest recently. Indeed, an active line of research in economics, mathematical sociology, and marketing is concerned with studying how behaviors, decisions, and trends propagate through a population. These types of diffusion processes naturally suggest a game [7,10], which in turn suggests the use of incentive mechanisms to maximize some specified objective, like profit [4] or influence [9,3,2]. Furthermore, our focus on pricing mechanisms for a seller is more closely related to the work by [4] on marketing in social networks.

## 2 The Model

Suppose there are two competing technologies (1 and 2), each with a fixed and commonly known one-dimensional measure of quality  $q_1$  and  $q_2$  and prices  $p_1$  and  $p_2$  respectively. We assume that  $0 \leq q_k, p_k \leq 1$ , doing so primarily for convenience of exposition as most of our results easily generalize. In what follows, differences between qualities and prices will be most germane to our analysis, in which case we denote  $\Delta q = q_2 - q_1 > 0$  and  $\Delta p = p_1 - p_2$ . We assume throughout that  $q_2 > q_1$  or  $\Delta q > 0$  (where the inequality is strict).

There are  $N$  players that are interested in adopting one of these two technologies; moreover, we require that they adopt one of the two technologies. Each player  $i$  has a parameter  $\theta_i \in [0, 1]$  which determines his strength of preference for a higher quality technology.  $\theta_i$  are distributed i.i.d. according to some distribution  $F()$  with density  $f()$ , which we assume to be strictly positive on  $[0, 1]$  and zero everywhere else. Thus,  $F()$  is non-atomic and strictly increasing on  $[0, 1]$ . The utility that player  $i$  receives from choosing technology  $k \in \{1, 2\}$  is given by:

$$u_i^k = \theta_i q_k + v(x_k) - p_k,$$

where  $x_k$  is the proportion of players currently choosing technology  $k$  and  $v : [0, 1] \rightarrow [0, 1]$  is a continuous and strictly increasing function representing

network effects on player preferences over the two technologies. Assume that  $v(0) = 0$  and  $v(1) = 1$ . If  $x_1 = x$  and  $x_2 = 1 - x$ , define  $\Delta v(x) = v(x) - v(1 - x)$ .

Observe that the particular realizations of  $\theta_i$  for all players define a static game between them. The equilibrium outcomes can be viewed as the fixed points of player best-response dynamics, in which case we need not assume that player types are public knowledge. We now define equilibria (fixed points) formally.

Let  $a_i \in \{1, 2\}$  be a technology choice of player  $i$ , which we subscript by a time (iteration)  $t$  where appropriate.

**Definition 1.** *A profile of technology choices  $a$  of all players and corresponding adoption proportions  $x_1 = x$  and  $x_2 = 1 - x$  constitute an equilibrium if for every  $i \in I$  with  $u_i^1(x) > u_i^2(x)$ ,  $a_i = 1$  and for every  $i$  with  $u_i^2(x) > u_i^1(x)$ ,  $a_i = 2$ .*

While we consider iterative best response and its fixed points in the “underlying” complete information game, we take the perspective of a technology seller and assume that consumer types  $\theta_i$  are unknown, although we have a prior distribution  $F()$  over these. We view the adoption outcomes and dynamics through the lens of this distribution.

We now make several very natural observations about the monotonicity of player preferences, which, while only of limited interest in their own right, will serve as building blocks for other results below. The first lemma states that, given a fixed market share of each technology, if  $\theta$  is a marginal type (a player with this type is indifferent between the two technologies), then a player  $j$  with  $\theta_j > \theta$  prefers technology 2, whereas a player with  $\theta_j < \theta$  prefers technology 1. The second lemma states that increasing the market share of technology 1 will not induce any player who currently prefers technology 1 to switch to technology 2 (the converse is true as well).

**Lemma 1.** *Let  $\theta_i$  be the type of player  $i$ ,  $u_i^k$  the corresponding utility of player  $i$  given fixed market shares  $x_1$  and  $x_2$ , and suppose that  $u_i^1 \geq u_i^2$ . Then  $\theta_j < \theta_i$  for some player  $j$  implies that  $u_j^1 > u_j^2$ . Similarly, if  $u_i^1 \leq u_i^2$ , then  $\theta_j > \theta_i$  for some player  $j$  implies  $u_j^1 < u_j^2$ .*

*Proof.* For the first case, note that if  $u_i^1 \geq u_i^2$ , then  $\theta_i q_1 + v(x_1) - p_1 \geq \theta_i q_2 + v(1 - x_1) - p_2$ , or, equivalently,  $v(x_1) - p_1 \geq \theta_i \Delta q + v(1 - x_1) - p_2 > \theta_j \Delta q + v(1 - x_1) - p_2$ . The latter inequality then implies that  $\theta_j q_1 + v(x_1) - p_1 > \theta_j q_2 + v(1 - x_1) - p_2$ . The second case is proved analogously.

**Lemma 2.** *Fix a player  $i$  with type  $\theta_i$  and suppose that  $x_1 \leq x'_1$  and  $u_i^1(x_1) \geq u_i^2(x_1)$ . Then  $u_i^1(x'_1) \geq u_i^2(x'_1)$ . Conversely, if  $x_1 \geq x'_1$  and  $u_i^1(x_1) \leq u_i^2(x_1)$ , then  $u_i^1(x'_1) \leq u_i^2(x'_1)$ .*

*Proof.*  $u_i^1(x_1) \geq u_i^2(x_1)$  means that  $\theta_i q_1 + v(x_1) - p_1 \geq \theta_i q_2 + v(1 - x_1) - p_2$ .  $x_1 \leq x'_1$  implies that  $v(x_1) \leq v(x_1)'$  and  $v(1 - x_1) \geq v(1 - x'_1)$ . Hence  $\theta_i q_1 + v(x'_1) - p_1 \geq \theta_i q_2 + v(1 - x'_1) - p_2$ . A symmetric argument proves the converse.

A useful result that follows directly from these lemmas is that every equilibrium corresponds to some  $\bar{\theta}$  which separates all player types into those who prefer

technology 1 and those who prefer technology 2, that is, all players with  $\theta_i \leq \bar{\theta}$  choose (and prefer) technology 1 and all players with  $\theta_i > \bar{\theta}$  choose (and prefer) technology 2. Without loss of generality and to simplify the analysis, we assume from now on that all indifferent players select technology 1.

**Theorem 1.** *Let  $a$  be an equilibrium profile with corresponding  $x_1 = x$  and  $x_2 = 1 - x$ . Then there exists  $\bar{\theta}$  such that  $\theta_i \leq \bar{\theta}$  if and only if  $a_i = 1$ .*

*Proof.* Let  $\theta$  be such that given  $x$ , the player with preference  $\theta$  is indifferent between the two technologies. Then by Lemma 1, any  $i$  with  $\theta_i < \theta$  strictly prefers technology 1, whereas any  $i$  with  $\theta_i > \theta$  strictly prefers technology 2. Letting  $\bar{\theta} = \theta$  yields the result.

We now establish the monotonicity and convergence of best-response dynamics, crucial properties for the nudging mechanisms in the next section. Let  $n_t$  denote the number of players in iteration  $t$  that adopt technology 1. Let  $x_t = \frac{n_t}{N}$  be the fraction of all players adopting technology 1 and let  $1 - x_t$  be the fraction of players adopting technology 2. The dynamics will be characterized by the change in the numbers (and, consequently, proportions) of players adopting technology 1 as iterations progress. Let  $a_{i,t}$  be the adoption strategy of player  $i$  at time  $t$ . Thus,  $a_{i,t} = 1$  means that player  $i$  chooses technology 1 at time  $t$ .

**Definition 2.** *Let  $a_t$  be a vector of player strategies at time  $t$ ,  $n_t$  the corresponding number of adopters of technology 1,  $x_t$  the corresponding proportion of technology 1 adopters.  $a_{t+1}$  is a best response to  $a_t$  when for all  $i$ ,  $\theta_i q_1 + v(x_t) - p_1 > \theta_i q_2 + v(1 - x_t) - p_2 \Rightarrow a_{i,t+1} = 1$  and  $\theta_i q_1 + v(x_t) - p_1 < \theta_i q_2 + v(1 - x_t) - p_2 \Rightarrow a_{i,t+1} = 0$ .*

**Lemma 3.** *Suppose that at time  $t$  players make some arbitrary choices  $a_t$  which result in some  $x_t$ . Let  $a_{t+1}$  be a best response to  $a_t$ . Then there is  $\bar{\theta}_{t+1}$  such that  $\forall \theta_i \leq \bar{\theta}_{t+1}, a_{i,t+1} = 1$  and  $\forall \theta_i > \bar{\theta}_{t+1}, a_{i,t+1} = 2$ .*

*Proof.* Sort players by  $\theta_i$  and let the indices  $i$  of players correspond to this ordering, with  $\theta_1$  being the smallest type. Let  $\theta_j$  be the largest type with  $u_j^1 \geq u_j^2$ . By Lemma 1, all players with  $j' < j$  strictly prefer 1, and since  $j$  is the largest such type, all players with  $j' > j$  strictly prefer 2. Setting  $\bar{\theta} = \theta_j$  completes the proof.

By Lemma 3, we can assume without loss of generality that there exists such  $\bar{\theta}_0$  which separates the players into those who prefer technology 1 and those who prefer 2. We will therefore assume that for every  $t$  there is some player  $\theta_j = \bar{\theta}_t$  with  $\bar{\theta}_t$  as in Lemma 3.

**Lemma 4.** *Let  $\theta_j = \bar{\theta}_t$  and suppose that  $\bar{\theta}_t$  is not an equilibrium. Then either (a)  $u_j^1 < u_j^2$  or (b)  $u_{j+1}^1 > u_{j+1}^2$ .*

*Proof.* Suppose that neither (a) nor (b) hold, but  $\bar{\theta}_t$  is not an equilibrium. That means that  $u_j^1 \geq u_j^2$  and  $u_{j+1}^2 \geq u_{j+1}^1$ . By Lemma 1, this implies that for all  $j' \leq j$ ,  $u_{j'}^1 \geq u_{j'}^2$ , and for all  $j' > j$ ,  $u_{j'}^2 \geq u_{j'}^1$ , which means that  $\bar{\theta}_t$  is an equilibrium, a contradiction.

We now define the best-response dynamic.

**Definition 3.** A sequence  $\{a_t\}_t$  is a best response (BR) dynamics if  $a_{t+1}$  is a best response to  $a_t$  for all  $t$ .

**Definition 4.**  $\{x_t\}_t$  is monotone if  $x_1 \leq x_0 \Rightarrow x_{t+1} \leq x_t \quad \forall t$  and  $x_1 \geq x_0 \Rightarrow x_{t+1} \geq x_t \quad \forall t$ .

The following result establishes that best response dynamics is monotone in the above sense, and from this it will follow directly that best response dynamics converges to an equilibrium from any starting point.

**Theorem 2.** If  $\{x_t\}_t$  is generated by best response dynamics, then it is monotone for  $t \geq 1$ . Furthermore,  $x_{t+1} = x_t$  if and only if  $\bar{\theta}_t$  is an equilibrium.

*Proof.* Note that by Lemma 3, for all  $t > 0$ , there is  $\bar{\theta}_t$  which separates player types linearly based on preferences. Thus, let  $t \geq 1$ . Now, assume that condition (a) in Lemma 4 holds at time  $t$ . Then  $\Delta n_t < 0$  since by Lemma 1 we are guaranteed that all players adopting 2 at time  $t$  will still prefer it at time  $t + 1$ , and (a) means that, additionally, the marginal type prefers it also. If  $\Delta n_t < 0$ , then clearly  $x_{t+1} < x_t$ . By Lemma 3, there again exists  $\bar{\theta}_{t+1}$  which separates player types linearly based on their preferences. Then, if  $x_{t+1}$  (and corresponding  $\bar{\theta}_{t+1}$ ) is an equilibrium, it remains constant forever after. Suppose  $x_{t+1}$  is not an equilibrium. By Lemma 2 and Lemma 4 it follows that condition (a) obtains again (i.e., it cannot be that (b) holds by Lemma 2), and, hence (by the same argument)  $\Delta n_{t+1} < 0$ . By induction, then,  $\Delta n_{t'} \leq 0$  for all  $t' \geq t$ , and the sequence is stationary if and only if the equilibrium is reached. If we assume that condition (b) holds in Lemma 4, the result can be proved by a symmetric argument.

**Corollary 1.** Best response dynamics converges from any starting point.

The monotonicity and convergence of best response dynamics suggest a simple class of simple mechanisms to affect substantial changes in the market shares of the two technologies. We explore these presently.

### 3 Market Penetration with “Nudging”

The idea of inducing general consensus in networks by influencing a set of key players is certainly not new. It has appeared in the context of interdependent security games 5 as a minimum critical coalition (i.e., the minimal set of players to seed that is sufficient to incentivize the rest to invest in security), as well as in more abstract treatments of maximizing network influence, such as 9, where a set of network nodes are seeded so as to achieve as great a number of additional adopters as possible. More significantly, we make use of another weapon of influence which is now available to us: price. Specifically, we analyze the effect of lowering the market price of the entrant technology on its adoption under incomplete information.

### 3.1 Nudging via Seeding

Our first mechanism for effecting a change is via seeding a subset of players who are technology 1 adopters under the current equilibrium. We would like to know how many players we have to seed in order to have an additional voluntary adoption which (the hope is) will start an adoption avalanche. To simplify notation below, we define  $f(x) = F\left(\frac{\Delta v(x) - \Delta p}{\Delta q}\right)$ . Additionally, we let  $x' = x - k/N$ , where  $x$  is the current equilibrium adoption of technology 1 and  $k$  is the number that we will seed with technology 2. Thus,  $x'$  is the resulting proportion of technology 1 adopters after  $k$  of the initial adopters are seeded with the higher quality technology.

**Theorem 3.** *Given an equilibrium  $x$ , at least  $l = xyN$  (with  $y$  a desired proportion) of  $xN$  technology 1 adopters will prefer technology 2 with probability at least  $1 - z$  if  $k$  or more technology 1 adopters are randomly seeded, where  $k$  solves*

$$\sum_{j=l}^{xN} \binom{xN}{j} \left[ \frac{f(x) - F\left(\frac{\Delta v(x') - \Delta p}{\Delta q}\right)}{f(x)} \right]^j \left[ \frac{F\left(\frac{\Delta v(x') - \Delta p}{\Delta q}\right)}{f(x)} \right]^{xN-j} = 1 - z.$$

*Proof.* The probability that exactly  $l$  of  $K$  players will adopt is binomial with parameters  $p$  and  $K$ , where  $p = \left[ \frac{f(x) - F\left(\frac{\Delta v(x') - \Delta p}{\Delta q}\right)}{f(x)} \right]$ , which is just the probability that  $\theta_i \geq \frac{\Delta v(x') - \Delta p}{\Delta q}$  (i.e., that  $i$  now adopts) given  $\theta_i < f(x)$  (i.e., that  $i$  hadn't yet adopted). The result follows since it is the probability of a union of mutually exclusive events.

An especially clean and convenient form of nudging would attempt to influence only one additional player to adopt, and hope that an adoption avalanche follows a tiny “nudge”.

**Theorem 4.** *Given an equilibrium  $x$ , at least 1 player out of  $xN$  technology 1 adopters will prefer technology 2 with probability at least  $1 - z$  if  $k$  or more technology 1 adopters are randomly seeded, where*

$$k \geq N \left[ x - \Delta v^{-1} \left( \Delta q F^{-1} \left( f(x) z^{\frac{1}{xN}} \right) + \Delta p \right) \right]. \tag{1}$$

*Proof.* Since probability that at least one player adopts is the same as 1 less the probability that no one adopts technology 2, and since  $\theta_i$  are i.i.d. for all current non-adopters of the better technology, we need  $\left[ \frac{F\left(\frac{\Delta v(x') - \Delta p}{\Delta q}\right)}{f(x)} \right]^{xN} \leq z$ . Solving first for  $x'$  and then for  $k$  yields the desired result.

Letting  $k(x)$  be the minimal number of players that we need to seed to incentivize a switch by a single non-adopter (as given in Theorem 4), we now consider what

happens when the number of players  $N$  grows large while the proportion of non-adopters  $x$  remains fixed. Note that for any  $z > 0$ , if  $N$  is sufficiently large,  $z^{\frac{1}{xN}}$  tends to 1. Thus,  $\Delta q F^{-1}\left(f(x)z^{\frac{1}{xN}}\right)$  becomes  $\Delta q F^{-1}(f(x)) = (\Delta v(x) - \Delta p)$ . Consequently,  $k(x) \rightarrow N(x - \Delta v^{-1}(v(x))) = 0$ . This is very much congruent with our expectations: when the number of players is very large, even after a very small change to the current state we are quite likely to find at least one player whose preferences now flip towards the higher quality technology (another way of saying this is that the probability of finding a marginal user becomes very large when  $N$  grows).

A related question is how  $k(x)$  changes as  $x$  falls when  $N$  becomes large. We address this question in the special case of  $F(\theta) = \theta$  (uniform distribution on player types) and  $v(x) = x$  (linear network effects). In this case,

$$k(x) = N \left[ x - \frac{\Delta q(2x - (1 + \Delta p))z^{\frac{1}{xN}} + (1 + \Delta p)}{2} \right].$$

Differentiating  $k(x)$  with respect to  $x$  we get

$$\frac{dk(x)}{dx} = N \left[ 1 - \frac{1}{2} \left( 2z^{\frac{1}{xN}} - (2x - (1 + \Delta p))z^{\frac{1}{xN}} \frac{\log(z)}{x^2 N} \right) \right].$$

As  $N$  grows large,  $z^{\frac{1}{xN}}$  tends to 1 and  $\frac{\log(z)}{x^2 N}$  to 0, and thus  $\frac{dk(x)}{dx} \rightarrow N(1 - 1) = 0$ . We thus observe that when  $N$  is very large,  $k(x)$  is essentially insensitive to the changes in adoption proportions. Both this observation and the tendency of  $k(x)$  towards zero in general as the number of technology users grows offer a very optimistic perspective on what previously may have seemed as a rather intractable problem of breaking incumbent network externalities: at relatively little cost, via a series of minor nudges to the market share, the higher quality entrant can establish a firm foothold in the market.

### 3.2 Nudging via Posted Prices

In the last section we presented a characterization of the impact that seeding  $k$  players has on adoption preferences of  $l$  others who currently prefer the incumbent technology. We now proceed towards a similar characterization of nudging that uses the price of the entrant technology as an incentive mechanism.

Let  $p_2$  be the initial price of technology 2, and  $p'_2$ , the new (discounted) price. Let  $\Delta p' = p_1 - p'_2$ .

**Theorem 5.** *Given an equilibrium  $x$ , at least  $l = xyN$  (with  $y$  a desired proportion) of  $xN$  technology 1 adopters will prefer technology 2 with probability at least  $1 - z$  under  $\Delta p'$ , where  $\Delta p'$  solves*

$$\sum_{j=l}^{xN} \binom{xN}{j} \left[ \frac{f(x) - F\left(\frac{\Delta v(x) - \Delta p'}{\Delta q}\right)}{f(x)} \right]^j \left[ \frac{F\left(\frac{\Delta v(x) - \Delta p'}{\Delta q}\right)}{f(x)} \right]^{xN-j} = 1 - z.$$

Again, if we restrict ourselves to targeting just one player, we obtain the following characterization.

**Theorem 6.** *Given an equilibrium  $x$ , at least 1 player out of  $xN$  technology 1 adopters will prefer technology 2 with probability at least  $1 - z$  under  $p'_2$  if*

$$p'_2 \leq \Delta q F^{-1} \left( f(x) z^{\frac{1}{xN}} \right) - \Delta v(x) + p_1.$$

We now let  $p'_2(x)$  be the largest new price that satisfies the inequality in Theorem 6. Performing a similar limiting analysis as  $N$  grows large while  $x$  remains fixed, we can observe that  $p'_2(x) \rightarrow p_2$ , that is, only a tiny drop in price can effect adoption by at least one (marginal) user. Similarly, differentiating  $p'_2(x)$  with respect to  $x$  and letting  $N$  go to infinity (under the assumption of  $F(\theta) = \theta$  and  $v(x) = x$ ) gives us  $\frac{dp'_2(x)}{dx} \rightarrow 0$ . Thus,  $p'_2(x)$  is nearly insensitive to the changes in adoption proportions when the user population is very large. The news, again, is optimistic: no matter how widespread the adoption of the incumbent technology is, as long as the number of users is large, a sequence of small price changes can allow the entrant to make inroads into the market.

While the above nudging mechanisms are described in the context of adoption growth, they apply to a number of problems. In the full paper, we look at profit maximization using discriminatory and non-discriminatory nudging mechanisms.

## References

1. Auriol, E., Benaim, M.: Standardization in decentralized economies. *American Economic Review* 90(3), 550–570 (2000)
2. Blume, L.E.: The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 387–424 (1993)
3. Chang, M.-H.: Discovery and diffusion of knowledge in an endogenous social network. *Modeling, Computing, and Mastering Complexity 2003 01*, Society for Computational Economics (2003)
4. Hartline, J., Mirrokni, V.S., Sundararajan, M.: Optimal marketing strategies over social networks. In: *Seventeenth International World Wide Web Conference*, pp. 189–198 (2008)
5. Heal, G.M., Kunreuther, H.: Interdependent security: A general model. Working paper (2004)
6. Immorlica, N., Kleinberg, J., Mahdian, M., Wexler, T.: The role of compatibility in the diffusion of technologies in social networks. In: *ACM Conference on Electronic Commerce*, pp. 75–83 (2007)
7. Jackson, M.O., Yariv, L.: Diffusion of behavior and equilibrium properties in network games. *American Economic Review* 97(2), 92–98 (2007)
8. Jin, Y., Sen, S., Guerin, R., Hosanagar, K., Zhang, Z.-L.: Dynamics of competition between incumbent and emerging network technologies. In: *Third International Workshop on Economics of Networked Systems*, pp. 49–54 (2008)
9. Kempe, D., Kleinberg, J., Tardos, É.: Maximizing the spread of influence through a social network. In: *KDD 2003: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146. ACM, New York (2003)
10. Morris, S.: Contagion. *Review of Economic Studies* 67(1), 57–78 (2000)
11. Sundararajan, A.: Local Network Effects and Complex Network Structure. *The B.E. Journal of Theoretical Economics* 7(1), Article 46 (Contributions)



# Mechanism Design for Complexity-Constrained Bidders

Ravi Kumar<sup>1</sup>, Mohammad Mahdian<sup>1</sup>, and Amin Sayedi<sup>2,\*</sup>

<sup>1</sup> Yahoo! Research, 701 Sunnyvale, CA 94089

{ravikumar, mahdian}@yahoo-inc.com

<sup>2</sup> Carnegie Mellon University, Pittsburgh, PA 15213

ssayedir@cmu.edu

**Abstract.** A well-known result due to Vickery gives a mechanism for selling a number of goods to interested buyers in a way that achieves the maximum social welfare. In practice, a problem with this mechanism is that it requires the buyers to specify a large number of values. In this paper we study the problem of designing optimal mechanisms subject to constraints on the complexity of the bidding language in a setting where buyers have additive valuations for a large set of goods. This setting is motivated by sponsored search auctions, where the valuations of the advertisers are more or less additive, and the number of keywords that are up for sale is huge. We give a complete solution for this problem when the valuations of the buyers are drawn from simple classes of prior distributions. For a more realistic class of priors, we show that a mechanism akin to the broad match mechanism currently in use provides a reasonable bicriteria approximation.

## 1 Introduction

Consider the following setting: there are  $m$  buyers who are interested in buying  $n$  goods, each with a unit supply. Each buyer has a value for each good, and her valuation for a bundle of goods is simply the sum of her valuations for each good in the bundle.

This is perhaps the simplest model for selling multiple non-identical goods, as the buyers' valuations are assumed to be additive and not combinatorial. From a mechanism design perspective, designing optimal (i.e., social welfare maximizing) auctions for this setting is trivial: simply run an independent second-price auction for each good. Each buyer will have the incentive to bid her true value for each good, and each good will be allocated to the buyer who has the maximum value for it.

The problem with this simple mechanism is that each bidder has to provide  $n$  values, one for each type of good, as her bid. This is especially problematic in applications such as sponsored search auctions, where the number of different types of goods is quite large or possibly infinite. This motivates the following problem, which is the main subject of this paper: *what is the maximum social welfare that can be achieved with a mechanism that is restricted to ask each bidder for only a small amount of information?*

---

\* Part of this work was done while the author was at Yahoo! Research.

<sup>1</sup> The valuations of bidders (i.e., advertisers) in sponsored search auctions are quite close to being additive. The only non-additive valuations that sponsored search systems allow bidders to specify are those involving budget constraints. However, binding budget constraints seem to be rare. Furthermore, since in sponsored search auctions advertisers can bid on any user query, the number of goods that are available for sale is essentially infinite.

To make this more precise, we denote the value of a bidder  $i$  for a good  $j$  by  $v_{ij}$ , and assume a prior on the  $n$ -tuple  $(v_{ij})_j$  of values, i.e., we assume these values are picked from a given joint distribution. This prior is supposed to capture the information on how the valuation of a bidder on different goods are correlated. The algorithm is allowed to query each bidder  $i$  for a fixed number  $k$  of  $v_{ij}$ 's, and based on the responses it gets, it allocates each good to one bidder. We evaluate the algorithm based on the social welfare it achieves, i.e., the sum of the valuations of the bidders for the goods they receive. Our objective is to design an allocation algorithm that achieves the maximum expected social welfare (where the expectation is over the draw of the valuations from the prior) among all algorithms that ask each bidder for at most  $k$  of her values.

Designing the optimal algorithm in the above model is often a complicated task, and depends on the type of prior we assume on the valuations. The main results of this paper include exact and approximate solutions of this problem for a few simple yet important classes of distributions, and a proof that for a more realistic class of distributions, clustering-based bidding languages — akin to the concept of *broad match* or *advanced match* that is currently in use — provide a reasonable bicriteria approximation for the optimal algorithm. Although the major technical contribution of the paper is on the problem of designing the optimal or approximately optimal allocation algorithms, we will also observe that these algorithms, combined with suitable payment schemes, result in mechanisms with good incentive properties.

**Related work.** Ronen [4] and Ronen and Saberi [5] studied the design of revenue-optimal mechanisms when the mechanism designer has communication complexity constraints in accessing the distribution of bidders' values. Mechanism design for single-item auctions under a constraint on the number of bits that each bidder can send to the auctioneer was studied by Blumrosen, Nisan, and Segal [1]. Earlier, Nisan and Segal [3] studied the communication complexity of maximizing social welfare in combinatorial auctions. The main difference between this line of work and ours is that in their models the constraint is on the *communication complexity* of the bidders, whereas our model focuses on *query complexity*, if each bidder is thought of as an oracle that can be queried for their value for each good. This is because in our motivating application (sponsored search auctions), the costly operation for an advertiser is to compute their value for a keyword, and not transmitting information to the auctioneer.

The setting of additive utilities for a large number of goods (motivated by sponsored search auctions) was studied by Mahdian and Wang [2]. Their focus is on a specific class of bidding languages called *clustering-based* bidding languages that are similar to the broad match scheme used in sponsored search systems. We address a more basic question: finding the most socially efficient algorithm without imposing any constraint other than the complexity constraint on the bidding language.

## 2 Model

Assume there are  $m$  buyers numbered 1 through  $m$ , and  $n$  different goods  $1, \dots, n$  that are offered for sale ( $n \gg m$ ). Without loss of generality, we can normalize the supply of each good to one. Buyer  $i$  has a non-negative real value  $v_{ij}$  for good  $j$ . We assume

that the valuations of buyers are additive, i.e., the value that buyer  $i$  has for a set  $S$  of goods is simply  $\sum_{j \in S} v_{ij}$ .

We assume a prior on the values  $v_{ij}$ . For simplicity, we assume that the valuation of different bidders are independently and identically distributed, i.e., there is a joint distribution  $\mathcal{D}$  over the set of all  $n$ -tuples of non-negative numbers, and for each buyer  $i$ , the tuple of valuations  $(v_{ij})_j$  is drawn independently from  $\mathcal{D}$ .

Given the distribution  $\mathcal{D}$  and an integer  $k$ , the problem is to design an algorithm that asks each buyer for at most  $k$  of their values (i.e., asks buyer  $i$  for  $v_{ij}$  for at most  $k$  different values of  $j$  to be chosen by the mechanism), and among all such algorithms achieves the maximal social welfare. This problem can be studied both in an adaptive and in a non-adaptive framework. We focus on the *non-adaptive* variant, i.e., the algorithm asks all the questions at once and then receives the answers. The adaptive version (where the algorithm asks the questions one by one and can use the earlier answers to decide what question to ask next), although theoretically intriguing, is of less practical value in sponsored search systems<sup>2</sup>.

Perhaps the simplest possible case is when the distribution  $\mathcal{D}$  is a product distribution, i.e., all  $v_{ij}$ 's are drawn independently. Even in this case, finding the optimal algorithm is non-trivial. However, we show (Section 3) that a mechanism based on *spreading* the questions among different goods is optimal. A more interesting class of distributions correspond to the case where each buyer  $i$  has a one-dimensional *type*  $t_i$  drawn from a distribution  $\mathcal{D}^*$ , and the values of buyer  $i$  on various goods are independent *conditioned on the type*  $t_i$  (Section 4). Finally, we study a model where the set of goods are partitioned into a number of clusters, and the valuation of each buyer on each cluster followed the model of buyers with types described above, while the values are independent for distinct clusters (Section 5).

**Incentive properties.** Most of this paper is devoted to the problem of designing the optimal allocation algorithm subject to the number of queries this algorithm can make. Ideally, we would like to match such an algorithm with a suitable payment scheme to turn it into an incentive-compatible mechanism, i.e., a mechanism where bidders are better off answering the questions truthfully. However, in our setting, we need to be careful about the notion of incentive compatibility, for the following reason. If we assume that the bidder knows her value for all goods, then for any mechanism that infers something from the values of queried goods about the unknown values, there is some chance that the queried goods under- or over-represent the values of other goods. In such cases the bidder might have an incentive to bid untruthfully to “correct” the mistake of the allocation algorithm. This intuitive argument can be made precise to show that with the strict notion of incentive compatibility in dominant strategies, essentially no non-trivial learning can be done.

However, it is possible (details omitted) to get around this problem by weakening the incentive requirement to one of the following:

<sup>2</sup> This is mainly because the algorithm cannot interleave the questions asked from two different advertisers in any way it wants, due to timing constraints. However, a hybrid between the adaptive and non-adaptive models, where the algorithm can use the answers provided by an advertiser to decide its next questions of the same (but not other) advertisers might be feasible in practice. We will comment on this in Section 6.

(1) If the bidder does not know her value for goods about which she is not queried, then her best strategy is to answer the questions truthfully; this can be thought of as an *ex ante* notion of incentive compatibility.

(2) As the number of questions that can be asked of each bidder grows, the incentive to deviate from truthfulness quickly approaches zero.

### 3 Independent Valuations

In this section we consider the case where the prior distribution  $\mathcal{D}$  is a product distribution, i.e., each value  $v_{ij}$  is picked independently from a distribution. This is a simple case since there is no *learning* involved: the answers that the mechanism receives on one good cannot help with the allocation of other goods. However, the problem is still non-trivial as it involves optimally distributing the queries among different goods.

For simplicity of exposition, we further assume that all  $v_{ij}$ 's are independently and *identically* distributed according to a distribution with cdf  $F$  and pdf  $f$ . The assumption that the distributions for different goods are identical is not necessary, but will simplify the statement of our results. The following lemma gives the optimal allocation of the goods, given the questions that the mechanism asks and their answers.

**Lemma 1.** *Suppose that values of bidders for a good come from a distribution with cdf  $F(\cdot)$  and pdf  $f(\cdot)$  and expectation  $\mu$ . If the auctioneer knows the values of  $i < m$  bidders in a set  $S$  for this good, then the expected welfare  $W(i)$  of allocating this good is maximized when it is given to the bidder in  $S$  with the maximum known value  $v$  if  $v \geq \mu$  or to an arbitrary bidder not in  $S$  if  $v < \mu$ . Furthermore,  $W(i) = M - \int_{\mu}^M F(x)^i dx$ , where  $M$  is the upper bound for the valuation (i.e.,  $F(M) = 1$ ).*

Using this, we show that the expected welfare that the mechanism gets from a good is a concave function of the number of queries it makes about the value of that good.

**Lemma 2.** *The function  $W(i) = M - \int_{\mu}^M F(x)^i dx$  is a concave function of  $i$ , i.e.,  $W(i) - W(i+1) \geq W(i+1) - W(i+2)$  for every  $i$ .*

**Theorem 1.** *If all values  $v_{ij}$  are drawn iid from a distribution  $F$  with expectation  $\mu$ , then any mechanism of the following form is optimal: ask each bidder for their value for  $k$  goods in such a way that each good is asked about  $\lfloor km/n \rfloor$  or  $\lceil km/n \rceil$  times, and then allocate each good as prescribed by Lemma 1*

If getting an approximation to the optimal social welfare is enough, we show that when  $km \leq n$ , a simple mechanism that allocates all the goods arbitrarily without asking any questions about the values extracts at least half of the social welfare of the optimal mechanism. In other words, when  $n$  is large, a mechanism with limited knowledge about values cannot do much better than a random mechanism.

**Lemma 3.** *If  $n \geq mk$ , and all the values  $v_{ij}$  are independent, a mechanism that allocates the good randomly to the bidders extracts welfare at least  $\text{OPT}/2$ , where  $\text{OPT}$  is the welfare of optimal mechanism subject to the same constraints. Furthermore, the bound is tight.*

The assumption  $n \geq km$  in this lemma (also in Theorem 3) is indeed necessary.

## 4 Bidders with Types

In reality, buyers’ valuations for different goods are not independent. For example, if a buyer has high valuation for a good, she is more likely to have high valuation for other related goods as well. To capture this, we consider the following class of distributions: each buyer  $i$  has a one-dimensional type  $t_i$  drawn from a distribution  $\mathcal{D}^*$  over non-negative numbers. Conditioned on the type  $t_i$ , the values of this buyer for the goods are independently and identically distributed according to some distribution  $\mathcal{D}(t_i)$ . This is a good model for settings where all the goods are related. In this case, the buyer  $i$ ’s *base value* for the goods is determined by the type  $t_i$ , and her value for an individual good depends on her base value as well as another factor that is independently distributed. E.g., the value of the buyer for a good can be equal to her base value plus an iid noise. Another important example in the context of sponsored search is when the base value of the buyer (the advertiser) is her *value per conversion*, and her value for an individual good (keyword) is her value per conversion times a keyword-specific *conversion rate* [\[6\]](#).

As in Section [\[3\]](#), we give an exact optimal allocation algorithm and a 2-approximation, in the case that the number of goods  $n$  is larger than  $mk$ . We start with a lemma that gives an upper bound on the social welfare of the optimal algorithm. Before stating the lemma, we need to define a few random variables  $a_{i,j}$ ’s,  $\mu_i$ ’s,  $\hat{\mu}_i$ ’s, and  $\hat{\mu}$ :

**Definition 1.** Consider the following random experiment: for every  $i = 1, \dots, m$ , generate a random number  $t_i$  according to  $\mathcal{D}^*$ , and then for every  $j = 1, \dots, k$ , generate  $a_{i,j}$  iid according to  $\mathcal{D}(t_i)$ . Let the random variable  $\mu_i$  denote the expected value of a random variable distributed according to  $\mathcal{D}(t)$  where  $t$  is drawn from  $\mathcal{D}^*$ , conditioned on  $k$  samples of this distribution (generated with the same  $t$ ) being  $a_{i,1}, \dots, a_{i,k}$ . Finally, let  $\hat{\mu} = \max_{j=1, \dots, m} \{\mu_j\}$  and  $\hat{\mu}_i = \max_{j \neq i} \{\mu_j\}$ . Note that the random variables  $\mu_i$ ’s,  $\hat{\mu}_i$ ’s, and  $\hat{\mu}$  are functions of the random variables  $a_{i,j}$ ’s.

**Lemma 4.** Assume  $n \geq mk$ . Then the social welfare of any algorithm that asks at most  $k$  questions from each buyer is at most  $E[\sum_{i=1}^m \sum_{j=1}^k \max(a_{i,j}, \hat{\mu}_i) + (n - km)\hat{\mu}]$ .

**Theorem 2.** Assume  $n \geq mk$ . The following algorithm is optimal: query each buyer for the values of  $k$  goods in such a way that no good is queried more than once, compute the values of  $\mu_i$ ’s as in Definition [\[1\]](#) and allocate each good  $j$  either to the bidder who queried about  $j$ , or to the one who has the maximum value of  $\mu_i$ , whichever is larger.

Since the algorithm given in the above theorem might be too complicated for implementation, or impractical as it treats different buyers asymmetrically, we also present a simple and natural algorithm that is a 2-approximation to the optimal welfare.

---

<sup>3</sup> One might argue that the conversion rate of an advertiser for different keywords are not independent. This is in fact true, however, in sponsored search systems it is common to assume that conversion rates are separable, i.e., the conversion rate of an advertiser for a keyword is the product of an advertiser-specific conversion rate and a keyword-specific one. Such a system can be captured by our model by letting the type  $t_i$  be the value per conversion times the advertiser-specific conversion rate.

**Theorem 3.** *If  $n \geq k(m + 1)$ , the algorithm  $\mathcal{A}$  that asks all the buyers about their value for the first  $k$  goods, allocates each of these goods to the buyer with the highest value, and allocates all other goods to the buyer with the highest value of  $\mu_i$  (as in Definition 7) gets a welfare that is at least half of the welfare of the optimal algorithm in expectation. Moreover, the bound is tight.*

*Proof.* Consider an optimal algorithm OPT. It is clear that the welfare that  $\mathcal{A}$  receives from the first  $k$  goods is at least as large as the welfare that OPT gets on those goods. For the remaining goods, it is easy to adapt the proof of Lemma 4 to show that the welfare that OPT gets from goods  $k + 1, \dots, n$  is at most  $E[\sum_{i=1}^m \sum_{j=1}^k \max(a_{i,j}, \hat{\mu}_i) + (n - k - km)\hat{\mu}]$ , where  $a_{i,j}$ 's,  $\mu_i$ 's,  $\hat{\mu}_i$ 's, and  $\hat{\mu}$  are as in Definition 7. This is at most

$$E[\sum_{i=1}^m \sum_{j=1}^k (a_{i,j} + \hat{\mu}) + (n - k - km)\hat{\mu}]. \quad (1)$$

We claim that  $E[a_{i,j}] = E[\mu_i]$ . This can be proved using the following experiment: draw  $t$  from  $\mathcal{D}^*$  and then  $k + 1$  numbers  $a_{i,1}, \dots, a_{i,k+1}$  from  $\mathcal{D}(t)$ . Clearly, we have  $E[a_{i,j}] = E[a_{i,k+1}]$ . On the other hand, fixing the values of  $a_{i,1}, \dots, a_{i,k}$ , the value of  $a_{i,k+1}$  has a distribution with mean  $\mu_i$ . Therefore, taking the expectation over values of  $a_{i,1}, \dots, a_{i,k}$ , we have  $E[a_{i,k+1}] = E[\mu_i]$ . Therefore,  $E[a_{i,j}] = E[\mu_i] \leq E[\hat{\mu}]$ .

Using this inequality, the expression in (1) is at most  $(n - k - km + 2km)E[\hat{\mu}] \leq 2(n - k)E[\hat{\mu}]$ , where the latter inequality follows from the assumption  $n \geq k(m + 1)$ .

On the other hand, we analyze the expected welfare that algorithm  $\mathcal{A}$  receives from goods  $k + 1, \dots, n$  as follows: the answers that  $\mathcal{A}$  gets on queries that it makes on the first  $k$  goods are distributed as the  $a_{i,j}$ 's of Definition 7. Therefore, fixing the values of  $a_{i,j}$ 's, the expected welfare that  $\mathcal{A}$  gets on any of the goods  $k + 1, \dots, n$  is precisely  $\hat{\mu}$ . Thus, the total expected welfare of  $\mathcal{A}$  on these goods is  $(n - k)E[\hat{\mu}]$ .

To sum things up, the total welfare of  $\mathcal{A}$  on the first  $k$  goods is at least that of OPT on these goods, and its total welfare on the other  $n - k$  goods is at least half that of OPT. Therefore,  $\mathcal{A}$  is a 2-approximation to OPT.

For the tightness, suppose there is only one type. Value of a bidder for a good is 0 with probability  $1 - \epsilon$  and is 1 with probability  $\epsilon$ . Also, let  $n = k(m + 1)$ . Since all bidders are of the same type,  $\mathcal{A}$  allocates all goods  $k + 1, \dots, n$  to an arbitrary bidder for expected welfare of  $(n - k)\epsilon = m k \epsilon$ . Therefore, the total welfare of this algorithm is at most  $k + m k \epsilon$ . On the other hand, OPT asks one question about each of the first  $m k$  goods, and gets an expected welfare of  $\epsilon + (1 - \epsilon)\epsilon$  on each such good. Therefore, the total welfare of OPT is at least  $2m k \epsilon - m k \epsilon^2$ . Taking  $\epsilon = 1/\sqrt{m}$  and letting  $m$  grow, the welfare of  $\mathcal{A}$  will tend to half that of OPT.  $\square$

The proof of the above theorem also implies the following, which will be useful in Section 5.

**Lemma 5.** *Assume  $n \geq km$ . An algorithm that knows the types of all buyers and allocates all goods to the buyer that has the maximum expected value (given the type) without asking any question gets a welfare of at least half of the optimal algorithm.*

**Non-uniform supplies.** We can extend these results to a more general model in which every good  $j$  has some supply  $s_j$ . In other words, value of buyer  $i$  for good  $j$  is  $s_j v_{ij}$ . By



slightly modifying algorithm  $\mathcal{A}$  in Theorem 3 and replacing the condition  $n \geq (m+1)k$  by  $\sum_{j=1}^n s_j \geq (m+1) \sum_{j=1}^k s_j$ , we can show that algorithm  $\mathcal{A}$  can be adopted for this more general model. More formally, assume without loss of generality that  $s_1 \geq \dots \geq s_n$ . If algorithm  $\mathcal{A}$  in Theorem 3 asks all the buyers about the  $k$  goods with largest amount of supply, namely  $s_1, \dots, s_k$ , it gets a welfare that is at least half of the welfare of the optimal algorithm in expectation.

## 5 A Cluster Model

The model studied in Section 4 captures situations where all goods are related, and therefore buyers' valuations for these goods are correlated by their *types*. Our final model captures the more realistic setting where goods can be partitioned into a number  $c$  of *clusters*, with each cluster containing a number of related goods. Different clusters are unrelated, and each buyer has a type for each cluster. The valuations of a buyer for the goods within a cluster follows the model from Section 4 with the buyer's type for the cluster. This is a good model for applications such as sponsored search, where the goods (keywords) can be partitioned based on their *topic*, with keywords within a topic being related and keywords from different topics being essentially independent.

Formally, there are  $c$  disjoint clusters, with the  $j$ th cluster containing  $n_j$  goods. Buyer  $i$  has a type  $t_{ij}$  for cluster  $j$ , drawn iid according to a distribution  $\mathcal{D}^*$ .<sup>4</sup> Given these types, the value of this buyer for each good in cluster  $j$  is picked iid according to a distribution  $\mathcal{D}(t_{ij})$ . We denote the expected value of this distribution by  $\mu_{ij}$ . Note that  $\mu_{ij}$  is a function of  $t_{ij}$  and is therefore a random variable.

The problem of designing the optimal algorithm boils down to deciding how the queries of each buyer should be allocated across different clusters. In one extreme, we might want to ask many queries about the same cluster to get a better estimate of the values of the buyer for that cluster, and forgo other clusters. In the other extreme, we might want to spread the queries evenly across different clusters, to get a rough estimate of the values of the buyer for all clusters. Interestingly, there are cases where each of these strategies outperforms the other by an arbitrary factor. This suggests that finding the optimal algorithm in the cluster model is a difficult problem, as the optimal allocation of the queries can be highly dependent on the nature of underlying distribution.

On the positive side, we can show that if we are allowed to ask more queries than OPT, enough to ask a logarithmic number of queries for each cluster, then we can get a good estimate of  $\mu_{ij}$ 's for each cluster and therefore the machinery from Section 4 gives us a good approximation to the optimal welfare. This result (Theorem 4) is essentially a bicriteria approximation of the optimal algorithm. To prove this result, we start with the following lemma (proof omitted), which shows that an algorithm that knows all  $\mu_{ij}$ 's can guarantee welfare of at least half of the optimal algorithm, if all clusters are large.

**Lemma 6.** *Suppose we are given all expected values  $\mu_{ij}$ , and for every  $j$ ,  $n_j \geq mk$ . An algorithm  $\mathcal{A}$  that allocates all goods of each cluster to the buyer with highest expected value gets welfare of at least  $\text{OPT}/2$  where OPT is the optimal expected welfare.*

<sup>4</sup> It is not hard to see that our result holds in the more general case where the distribution of types for different clusters are different (but still independent).

Let  $L$  denote the maximum, over the choice of the type  $t$ , of the ratio between the maximum value from the distribution  $\mathcal{D}(t)$  and the expected value of  $\mathcal{D}(t)$ ; this parameter captures how difficult it is to estimate the mean of the distribution by sampling.

**Theorem 4.** *Assume  $n_j \geq mk, \forall j$ . There is an algorithm  $\mathcal{A}$  that asks  $O(cL^2\epsilon^{-2} \log m)$  queries, and achieves at least a  $\frac{0.9}{2(1+\epsilon)^2}$  fraction of the welfare of the optimal algorithm.*

## 6 Conclusions

We studied the problem of designing the optimal allocation algorithms for an auction setting where buyers have additive values, but the number of different types of goods is large and we are limited in the number of queries we can make to each buyer. We believe this is a promising direction for research, as in many realistic situations, we are faced with valuations that lie in a high-dimensional space, even though the distribution that these valuations come from can allow for “learning” the valuation given only a small-dimensional sample. The problem we studied is essentially about the interaction between this learning aspect, and the optimization aspect of queries. In other words, in our optimization, we must take into account not only the additional information that a query gives us, but also the amount of additional welfare it can lead to.

There are many open problems left for future research. For example, our results for the model of typed buyers (Section 4) cannot handle the case with non-identical distributions, except for the case of non-uniform supplies. Extending the results to such a setting seems difficult, as it requires ways to capture the information contents of each distribution. Also, it would be nice to get rid of the  $n \geq mk$  assumption in our results. In particular, in Section 4, it might be possible to prove that the algorithm that spreads the questions almost evenly across the goods is optimal. Finally, it would be interesting to solve the adaptive variant of the problem, or at least the more practical “hybrid” between adaptive and non-adaptive variants where the questions the mechanism asks can depend on the previous answers of the same buyer but not the other buyers.

## References

1. Blumrosen, L., Nisan, N., Segal, I.: Auctions with severely bounded communication. *JAIR* 28, 233–266 (2007)
2. Mahdian, M., Wang, G.: Clustering-based bidding languages for sponsored search. In: *Proc. 17th ESA*, pp. 167–178 (2009)
3. Nisan, N., Segal, I.: The communication requirements of efficient allocations and supporting prices. *JET* 129(1), 192–224 (2006)
4. Ronen, A.: On approximating optimal auctions. In: *Proc. 3rd EC*, pp. 11–17 (2001)
5. Ronen, A., Saberi, A.: On the hardness of optimal auctions. In: *Proc. 43rd FOCS*, pp. 396–405 (2002)



# Priority Right Auction for Komi Setting (Extended Abstract)

Xiaotie Deng<sup>1</sup> and Qi Qi<sup>2</sup>

<sup>1</sup> Department of Computer Science,  
City University of Hong Kong, Hong Kong SAR, P.R. China  
deng@cs.cityu.edu.hk

<sup>2</sup> Department of Management Science and Engineering,  
Stanford University, Stanford, California, USA  
kaylaqi@stanford.edu

**Abstract.** We develop an auction system that determines a fair number of komi for holding black stones in a GO game between two players of equal strength, and hence the right to the first move. It is modeled as a priority right pricing problem that demands for budget-balanced and egalitarian conditions, where a negative utility is associated with the losers. We establish results involved with the incentive compatible properties for this problem under both deterministic and randomized protocols.

## 1 Introduction

GO is a competitive territorial game<sup>1</sup> between two players on a grid of 19 lines by 19 lines. The players take turns to place their stones at the intersections of the lines, the black stone holder (BLACK) first, and alternating with the white stone holder (WHITE). The goal of a player is to occupy more grid points than the opponent.

Being the first to move, BLACK has an advantage. For two players of equal strength, if WHITE has  $w$  grid points, BLACK is required to occupy more than  $b = w + k$  grid points to win for some  $k > 0$ , a number referred to as *komi* by the Japanese. Despite of a long history, the GO game does not have a unified set of rules yet, one of which is how to decide the value for komi. It is a problem of allocating the priority right to one of the two players, and determining a price that the winner pays to the loser.

Priority right is an important phenomenon in human activities. This issue arises in a competitive environment where the winner of the priority right may have an advantage and the loser(s) may suffer disadvantages such as home team advantage in sports.

For the game of GO, a fixed komi is adopted, but there is a disagreement on what it should be. For example, AGA (American GO Association) sets 6.5 as

---

<sup>1</sup> Originated from China, stretching back some four thousand years ago. The legend goes that it was invented by Emperor YAO for his son Dan Zhu.

the current standard for komi. Chinese rule for today's komi is equivalent to 7.5. As a further complication, komi is dependent on strengths and preferences of individual players who have their own opinions, corresponding to their different types of strengths in terms of attack and defense.

Auction has naturally emerged as an alternative solution to determine the compensation point for players with different private values for the komi. Two types of solutions have already been appeared in GO competitions over the Internet: the cut-&-choose method and the English auction [1].

However, the above straight forward applications of typical pricing models are not fully fair in the GO game. With the cut-&-choose method, the cutter sets a komi that equalizes its probability of winning in either black or white. The chooser, however, will be able to take an advantage by choosing one with higher probability to win. The English auction on the other hand, favors the player with a large private komi who pays its opponent's (smaller) komi for the first move.

Therefore, to be completely fair in setting the komi between two different players, we would need to design a system that decides on a komi which results in the same utility to each. In addition, modeling the priority right as an indivisible good for the two players and komi as money, the problem requires for a budget-balanced condition.

Motivated by the application, we set to study the problem of finding a strategy-proof auction protocol that is budget-balanced and

- *We derive a relationship for allocation and pricing in our model for strategy-proof protocols that observe egalitarian and budget-balanced conditions.* As a corollary, it follows that no deterministic incentive compatible protocol exists under those conditions.
- *We design a randomized auction protocol that is both egalitarian and budget-balanced, and that, at the same time, truth-telling strategy is a (weakly) best response for everyone.*
- *We prove that the randomized protocol is unique in the sense that there is no other randomized protocol that would satisfy the above three conditions.* In this randomized protocol, no player would gain by misreporting his true private value. However, it would not be worse off either. By the uniqueness result, there is no randomized strategy-proof protocol that is simultaneously budget-balanced and egalitarian.
- *We extend all our results to other related models in multi-player environment.* We establish similar results as in the two-player model for cases where many players are present, both the protocol and the uniqueness result.

We should present those results in the subsequent sections.

## 2 Model and Notations

To solve the problem of competing for a priority right in games (called PRP for short) such as GO, we consider an auction system to price it and allocate

it. The priority right can be considered as an item for sale in the auction. We first consider the case of two risk neutral players,  $i = \{1, 2\}$ . Each has a private value,  $v_1$  and  $v_2$ , for the priority right. The protocol assigns the item to one of them depends on their bids  $b_1$  and  $b_2$ . If  $i$  is the winner of the right, it gains  $v_i$ , while the loser  $j$  gains  $-v_j$ .

Let  $o_i(b_1, b_2) = 1$  if the item is assigned to player  $i$ ; else  $o_i(b_1, b_2) = -1$ . It is required that the item is assigned to one player only:  $o_1(\cdot) + o_2(\cdot) = 0$ . At the same time, Player  $i$  will be charged  $t_i(b_1, b_2)$ , which is positive for the winner and negative for the loser, i.e., which will be paid to the loser as a compensation. With those notations, we formulate the payoffs of player  $i$ 's,  $i = 1, 2$ , by the usual quasi-linear utility functions.

**Definition 1.** *Quasi-linear utility functions: The utility function of player  $i$ ,  $i = 1, 2$ , is*

$$u_i(b_1, b_2) = o_i(b_1, b_2) \times v_i - t_i(b_1, b_2).$$

As the payment is between the winner and the loser, we have the following budget requirement.

**Definition 2.** *Budget-Balancedness: A protocol for priority assignment is budget-balanced if the total charge to the two players are zero, i.e.,*

$$t_1(b_1, b_2) + t_2(b_1, b_2) = 0.$$

In the context of the GO game, a completely fair komi setting would even out the differences of the two players' private values, which leads us to the egalitarian condition. Pazner and Schmeidler [11] proposed a general concept of egalitarian-equivalence, which equalizes two agents if their allocations are equivalent to some reference bundle consisting of goods and money. What we need here is a simple monetary-only reference bundle.

**Definition 3.** *Egalitarian: A protocol for priority assignment is egalitarian if the protocol leads to equal utilities for both players:  $u_1(b_1, b_2) = u_2(b_1, b_2)$ .*

Vickrey laid down the fundamental principle in auction design to provide bidders an incentive to bid their true value [12]. There are some minor differences in the definitions of the concept, with different names, in the literatures [2,3,8,9,10]. We should focus on the following two related concepts: one standard incentive compatible concept and another weaker concept of incentive compatibility, call equal-or-dominant incentive compatibility. For our application, a subtle variation in their definitions results in a big difference.

**Definition 4.** *Dominant Incentive Compatibility (or strategy-proof or truthful): A protocol for priority assignment is Dominant Incentive Compatibility if under the protocol, truth-revelation is everyone's dominant-strategy:  $\forall i, b_{-i}, v'_i$ ,*

$$u_i(v_i, b_{-i}) \geq u_i(v'_i, b_{-i}) \text{ (with at least one strict inequality),}$$

where  $b_{-i}$  is the bidding profile of all other bidders except bidder  $i$ .

**Definition 5.** *Equal-or-dominant Incentive Compatibility (EIC): A protocol for priority assignment is EIC if under the protocol, truth-revelation is always a (weakly) best response for every player no matter what are other players' strategies:  $\forall i, b_{-i}, v'_i$ ,*

$$u_i(v_i, b_{-i}) \geq u_i(v'_i, b_{-i}).$$

Informally, in a Dominant Incentive Compatible protocol, every player has incentive to bid truthfully, while an EIC protocol only requires that no player has incentive not to bid truthfully.

Finally, we also list some definitions we will refer to later during the discussion, though they are not our main focuses.

**Definition 6.** *Truthful Nash Equilibrium Protocol: A protocol has truthful Nash equilibrium if revealing the true private value is a Nash equilibrium for the players:*

**Definition 7.** *Envy-free: A protocol is envy-free if none of the players would be better off by exchange the solution with another player in the same game.*

**Definition 8.** *Efficiency(or allocation efficiency): A protocol is efficient if its solution always maximizes the total value over all agents.*

### 3 On Incentive Compatibility

We first characterize the price structure in our model under the conditions of egalitarian and budget-balancedness. As an immediate corollary, the characterization derives the impossibility result for deterministic incentive compatible protocols satisfying egalitarian and budget-balanced conditions for our model. The non-existence deterministic protocol result drives us to study randomized protocols.

#### 3.1 Characterization of Budget-Balancedness and Egalitarian Property

In this subsection, we characterize protocols satisfying both budget-balanced and egalitarian conditions.

**Lemma 1.** *An incentive compatible protocol for two-player PRP is egalitarian and budget-balanced if and only if*

$$t_i(b_1, b_2) = o_i(b_1, b_2) \frac{b_1 + b_2}{2}.$$

**Theorem 1.** *There does not exist a deterministic protocol for two-player PRP which is incentive compatible, budget-balanced and egalitarian.*

### 3.2 On Randomized Incentive Compatible Protocols

We consider a randomized protocol for two-player PRP as follows:

#### Protocol 1 (EQUAL\_CUT)

1. Each player submits its own bid:  $b_i$ ,  $i = 1, 2$ .
2. The system sets price to be  $\frac{b_1+b_2}{2}$ , and randomly selects a winner for the priority right.
3. Winner yields to the opponent a compensation of  $\frac{b_1+b_2}{2}$ .

**Theorem 2.** *EQUAL\_CUT is an EIC protocol for the players, and satisfies budget-balanced and ex ante egalitarian conditions.*

**Corollary 1.** *EQUAL\_CUT is ex post egalitarian under truth-telling.*

The result presents a simple protocol satisfies budget-balanced and egalitarian conditions, and is EIC. Examining it carefully, bidding truthfully is always an optimal strategy, but it is not (weakly) dominant in that there is no strategy which is dominated by bidding the truth. Similar protocol was known previously to Morgan [7] for a different problem of partnership dissolution under the common value model. It leaves open the question whether there is a randomized strategy-proof protocol.

Next, we will show a uniqueness theorem that EQUAL\_CUT is in essence the only candidate for the egalitarian, budget-balanced dominant IC protocol.

**Proposition 1.** *Any protocol that satisfies dominant incentive compatible, budget balanced and egalitarian, will choose the winner with equal probability and set the payment as  $\frac{b_1+b_2}{2}$  no matter who is the winner.*

It follows immediately,

**Corollary 2.** *EQUAL\_CUT is the only protocol that simultaneously satisfies EIC, budget-balanced and egalitarian conditions.*

Furthermore, Corollary [2] implies the following conditions.

1. No randomized EIC protocol could be simultaneously budget-balanced, egalitarian and envy-free;
2. No randomized EIC protocol could be simultaneously budget-balanced, egalitarian and efficient.

However, EQUAL\_CUT doesn't satisfy the dominant condition. It follows that we don't have a (randomized) dominant incentive compatible protocol for our problem.

**Theorem 3.** *There is no randomized protocol that simultaneously satisfies dominant incentive compatibility, egalitarian and budget-balanced conditions.*

Note that under the risk-averse condition, EQUAL\_CUT is a Nash implementation of the properties of budget-balancedness and egalitarian. In the game of go, players only care the difference of the utilities. In EQUAL\_CUT protocol, if player 1 bids truthfully, player 2's truthful bidding results in  $u_2 - u_1 = 0$  for both outcomes of the random pick. Any other bids of player 2 would result in a non-zero utility for both of the two outcomes though the expected utility is zero. Therefore, any risk-averse player would bid truthfully if another does so. The Nash implementation theory was developed especially for such problem when impossibility is present obstructing social goal of certain properties [46]. Our results can be a positive one in such cases.

### 4 Generalization

In this section, we consider one priority right for sale to  $n$  players. In an alternative formulation, we may view this problem as having three types of goods: one type of a single item representing the priority right (PR), and another of  $(n-1)$  items representing no such right (no-PR), as well as a numeraire good (money) initially at a total zero amount. The players prefer PR but dislike no-PR. Each player's valuations of the  $n$  goods are  $((n-1)v_i, -v_i, -v_i, \dots, -v_i)$ , i.e.,  $(n-1)v_i$  for the PR and  $-v_i$  for each of no-PRs. The problem is to decide how to assign each player one of the  $n$  goods and a certain amount of money such that each player has the same utility.

More formally, we have  $n$  players  $i = \{1, 2, \dots, n\}$ . Each has a private value  $v_i$ . The protocol assigns priority right to one of them depending on their bidding profile  $\mathbf{b} = (b_1, b_2, \dots, b_n)$ . Let  $o_i(\mathbf{b}) = n - 1$  for player  $i$  who wins the priority right; and  $o_i(\mathbf{b}) = -1$ , otherwise. The single winner condition implies:  $\sum_i o_i(\mathbf{b}) = 0$ . Denote by  $t_i(\mathbf{b})$  the amount charged to player  $i$ ,  $i = 1, 2, \dots, n$  by the auction system when the bid vector is  $\mathbf{b}$ . The utility function of player  $i$ ,  $i = 1, 2, \dots, n$ , will be:

$$u_i(v_i, \mathbf{b}) = o_i(\mathbf{b})v_i - t_i(\mathbf{b}).$$

We define the allocation and pricing for the above model as the  $n$ -player priority right pricing problem ( $n$ -PRP for short). We first give the formal definition for budget-balancedness and egalitarian for  $n$  player case.

**Definition 9.** *Budget-Balancedness: A protocol for priority assignment is budget-balanced if the total charge to all players are zero, i.e.,*

$$\sum_{i=1}^n t_i(\mathbf{b}) = 0.$$

**Definition 10.** *Egalitarian: A protocol for priority assignment is egalitarian if the protocol leads to equal utilities for all players:  $u_1(\mathbf{b}) = u_2(\mathbf{b}) = \dots = u_n(\mathbf{b})$ .*

Next we characterize incentive compatible protocols which are egalitarian and budget-balanced.

**Lemma 2.** *An incentive compatible protocol for n-PRP is egalitarian and budget-balanced if and only if the winner  $i$  pays*

$$t_i = (n - 2)b_i + \frac{1}{n} \sum_{j=1}^n b_j$$

and the loser  $j \neq i$  receives

$$-t_j = b_j + b_i - \frac{1}{n} \sum_{k=1}^n b_k.$$

Denote  $w_i(b_1, b_2, \dots, b_n) \geq 0$  as the probability that player  $i$  wins the auction under our protocol (as a function of the bidding profiles), we have  $\sum_{i=1}^n w_i(b_1, b_2, \dots, b_n) = 1$ . Note that  $w_i(b_1, b_2, \dots, b_n)$  is dependent on the particular randomized protocol. We immediately have the following corollary.

**Corollary 3.** *Under the conditions of Lemma 2, the expected utility of player  $i$  can be written as*

$$U_i = w_i((n - 1)v_i - (n - 2)b_i - \frac{1}{n} \sum_{j=1}^n b_j) + \sum_{j \neq i} w_j(-v_i + b_i + b_j - \frac{1}{n} \sum_{k=1}^n b_k).$$

We obtain a similar uniqueness result as the two-player case.

**Theorem 4.** *There is no randomized dominant incentive compatible protocol but a unique EIC protocol for the n-player PRP, which is budget-balanced and egalitarian.*

## 5 Discussion

Our study opens up a possibility of placing the study of priority right under the auction model. Further applications are possible with priority right auction protocols such as home court advantages in pairwise competitions. Other applications include games of many players where the dealer has an advantage over other players.

Further more, such types of advantages may not be the same for different players in the group to become the dealer. The disadvantages of losers may also be different with respect to different leaders. Those situations are in wide existence in human activities such as chairmanship in elected committees, choice of reserve currencies, etc. In this sense, our model falls under a general framework of endogenous valuations such as discussed in [5].

In addition, auction may not be the only possible protocol for determining winners of priority rights. The particular setting of GO allows us to develop a tight bound for the incentive compatibility paradigm. This may not always be possible for the general framework. Instead, our negative results reveal the difficulty in this paradigm and may suggest further effort to be made using different methodologies.

**Acknowledgement.** Work reported in this work is sponsored by a matching research grant from Comet Electronics (HK) Ltd and Research Grants Council of Hong Kong SAR (Proj. No. 9220046).

## References

1. <http://senseis.xmp.net/?KomiBidding>
2. Archer, A., Tardös, É.: Truthful Mechanisms for One-Parameter Agents. In: Proceedings of 42nd Annual Symposium on Foundations of Computer Science, Las Vegas, Nevada, USA, pp. 482–491 (2001)
3. Goldberg, A., Hartline, J., Karlin, A., Saks, M., Wright, A.: Competitive Auctions. *Games and Economic Behavior* 55(2), 242–269 (2006)
4. Jackson, M.O.: A Crash Course in Implementation Theory. *Social Choice and Welfare* 18, 655–708 (2001)
5. Jehiel, P., Moldovanu, B.: Strategic Non-participation. *Rand Journal of Economics* 27, 84–98 (1996)
6. Maskin, E., Sjostrom, T.: Implementation theory. In: Arrow, K.J., Sen, A.K., Suzumura, K. (eds.) *Handbook of Social Choice and Welfare*, 1st edn., vol. 1, ch. 5. Elsevier, Amsterdam (2002)
7. Morgan, J.: Dissolving a partnership (un)fairly. *Economic Theory* 23, 909–923 (2004)
8. Myerson, R.B.: Optimal Auction Design. *Mathematics of Operations Research* 6(1), 58–73 (1981)
9. Nisan, N., Ronen, A.: Algorithmic Mechanism Design (Extended Abstract). In: Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, Atlanta, Georgia, USA, pp. 129–140. ACM, New York (1999)
10. Parkes, D.C.: Computational Mechanism Design. In: Lecture notes of Tutorials at 10th Conf. on Theoretical Aspects of Rationality and Knowledge (TARK 2005), Institute of Mathematical Sciences, University of Singapore (2008) (to appear)
11. Pazner, E., Schmeidler, D.: Egalitarian-equivalent allocations: a new concept of economic equity. *Q. J. Econ.* 92, 671–687 (1978)
12. Vickrey, W.: Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance* 16, 8–37 (1961)



# Randomized Online Algorithms for the Buyback Problem

B.V. Ashwinkumar and Robert Kleinberg\*

Cornell University, Ithaca, NY  
{ashwin85,rdk}@cs.cornell.edu

**Abstract.** In the matroid buyback problem, an algorithm observes a sequence of bids and must decide whether to accept each bid at the moment it arrives, subject to a matroid constraint on the set of accepted bids. Decisions to reject bids are irrevocable, whereas decisions to accept bids may be canceled at a cost which is a fixed fraction of the bid value. We present a new randomized algorithm for this problem, and we prove matching upper and lower bounds to establish that the competitive ratio of this algorithm, against an oblivious adversary, is the best possible. We also observe that when the adversary is adaptive, no randomized algorithm can improve the competitive ratio of the optimal deterministic algorithm. Thus, our work completely resolves the question of what competitive ratios can be achieved by randomized algorithms for the matroid buyback problem.

## 1 Introduction

Imagine a seller allocating a limited inventory (e.g. impressions of a banner ad on a specified website at a specified time in the future) to a sequence of potential buyers who arrive sequentially, submit bids at their arrival time, and expect allocation decisions to be made immediately after submitting their bid. An informed seller who knows the entire bid sequence can achieve much higher profits than an uninformed seller who discovers the bids online, because of the possibility that a very large bid is received after the uninformed seller has already allocated the inventory. A number of recent papers [1,2] have proposed a model that offsets this possibility by allowing the uninformed seller to cancel earlier allocation decisions, subject to a penalty which is a fixed fraction of the canceled bid value. This option of canceling an allocation and paying a penalty is referred to as *buyback*, and we refer to online allocation problems with a buyback option as *buyback problems*.

Buyback problems have both theoretical and practical appeal. In fact, Babaioff et al. [1] report that this model of selling was described to them by the ad marketing group at a major Internet software company. Constantin et al. [2] cite numerous other applications including allocation of TV, radio, and newspaper advertisements; they also observe that advance booking with cancellations is a common practice in the airline industry, where limited inventory is

---

\* This research was supported by NSF CAREER Award CCF-0643934.

oversold and then, if necessary, passengers are “bumped” from flights and compensated with a penalty payment, often in the form of credit for future flights.

Different buyback problems are distinguished from each other by the constraints that express which sets of bids can be simultaneously accepted. In the simplest case, the only constraint is a fixed upper bound on the total number of accepted bids. Alternatively, there may be a bipartite graph whose two vertex sets are called *bids* and *slots*, and a set of bids may be simultaneously accepted if and only if each bid in the set can be matched to a different slot using edges of the bipartite graph. Both of these examples are special cases of the *matroid buyback problem*, in which there is a matroid structure on the bids, and a set of bids may be simultaneously accepted if and only if they constitute an independent set in this matroid. Other types of constraints (e.g. knapsack constraints) have also been studied in the context of buyback problems [1], but the matroid buyback problem has received the most study. This is partly because of its desirable theoretical properties — the offline version of the problem is computationally tractable, and the online version admits an online algorithm whose payoff is identical to that of the omniscient seller when the buyback penalty is zero — and partly because of its well-motivated special cases, such as the problem of matching bids to slots described above.

As is customary in the analysis of online algorithms, we evaluate algorithms according to their competitive ratio: the worst-case upper bound on the ratio between the algorithm’s (expected) payoff and that of an informed seller who knows the entire bid sequence and always allocates to an optimal feasible subset without paying any penalties. The problem of deterministic matroid buyback algorithms has been completely solved: a simple algorithm was proposed and analyzed by Constantin et al. [23] and, independently, Babaioff et al. [4], and it was recently shown [1] that the competitive ratio of this algorithm is optimal for deterministic matroid buyback algorithms, even for the case of rank-one matroids (i.e., selling a single indivisible good). However, this competitive ratio can be strictly improved by using a randomized algorithm against an oblivious adversary. Babaioff et al. [1] showed that this result holds when the buyback penalty factor is sufficiently small, and they left open the question of determining the optimal competitive ratio of randomized algorithms — or even whether randomized algorithms can improve on the competitive ratio of the optimal deterministic algorithm when the buyback factor is large.

Our work resolves this open question by supplying a randomized algorithm whose competitive ratio (against an oblivious adversary) is optimal for all values of the buyback penalty factor. We present the algorithm and the upper bound on its competitive ratio in Section 3 and the matching lower bound in Section 4. Our algorithm is also much simpler than the randomized algorithm of [1], avoiding the use of stationary renewal processes. It may be viewed as an online randomized reduction that transforms an arbitrary instance of the matroid buyback problem into a specially structured instance on which deterministic algorithms are guaranteed to perform well. Our matching lower bound relies on defining

and analyzing a suitable continuous-time analogue of the single-item buyback problem.

*Adaptive adversaries.* In this paper we analyze randomized algorithms with an oblivious adversary. If the adversary is adaptive<sup>1</sup>, then no randomized algorithm can achieve a better competitive ratio than that achieved by the optimal deterministic algorithm. This fact is a direct consequence of a more general theorem asserting the same equivalence for the class of *request answer games* (Theorem 2.1 of [5] or Theorem 7.3 of [6]), a class of online problems that includes the buyback problem.<sup>2</sup>

*Strategic considerations.* In keeping with [14], we treat the buyback problem as a pure online optimization with non-strategic bidders. For an examination of strategic aspects of the buyback problem, we refer the reader to [2].

*Related work.* We have already discussed the work of Babaioff et al. [14] and of Constantin et al. [23] on buyback problems. Prior to this aforementioned work, several earlier papers considered models in which allocations, or other commitments, could be cancelled at a cost. Bialogorsky et al. [7] studied such “opportunistic cancellations” in the setting of a seller allocating  $N$  units of a good in a two-period model, demonstrating that opportunistic cancellations could improve allocative efficiency as well as the seller’s revenue. Sandholm and Lesser [8] analyzed a more general model of “leveled commitment contracts” and proved that leveled commitment never decreases the expected payoff to either contract party. However, to the best of our knowledge, the buyback problem studied in this paper and its direct precursors [12,3,4] is the first to analyze commitments with cancellation costs in the framework of worst-case competitive analysis rather than average-case Bayesian analysis.

## 2 Preliminaries

Consider a matroid<sup>3</sup>  $(\mathcal{U}, \mathcal{I})$  where  $\mathcal{U}$  is the ground set and  $\mathcal{I}$  is the set of independent subsets of  $\mathcal{U}$ . We will assume that the ground set  $\mathcal{U}$  is identified with the set  $\{1, \dots, n\}$ . There is a bid value  $v_i \geq 0$  associated to each element  $i \in \mathcal{U}$ . The information available to the algorithm at time  $k$  ( $1 \leq k \leq n$ ) consists of the first  $k$  elements of the bid sequence — i.e. the subsequence  $v_1, v_2, \dots, v_k$  — and the restriction of the matroid structure to the first  $k$  elements. (In other words, for every subset  $S \subseteq \{1, 2, \dots, k\}$ , the algorithm knows at time  $k$  whether  $S \in \mathcal{I}$ .)

<sup>1</sup> A distinction between *adaptive offline* and *adaptive online* adversaries is made in [5,6]. When we refer to an adaptive adversary in this paper, we mean an adaptive offline adversary.

<sup>2</sup> The definition of request answer games in [6] requires that the game must have a minimization objective, whereas ours has a maximization objective. However, the proof of Theorem 7.3 in [6] goes through, with only trivial modifications, for request answer games with a maximization objective.

<sup>3</sup> See [9] for the definition of a matroid.

At any step the algorithm can choose a subset  $S^k \subseteq S^{k-1} \cup \{k\}$ . This set  $S^k$  must be an independent set, i.e  $S^k \in \mathcal{I}$ . Hence the final set held by the algorithm is  $R = S^n$ . The algorithm must perform a buyback for every element of  $B = (\cup_{i=1}^n S^i) \setminus S^n$ . For any set  $S \subseteq \mathcal{U}$  let  $\text{val}(S) = \sum_{i \in S} v_i$ . Finally we define the payoff of the algorithm as  $\text{val}(R) - f \cdot \text{val}(B)$ .

### 3 Randomized Algorithm against Oblivious Adversary

This section gives a randomized algorithm with competitive ratio  $-W\left(\frac{-1}{e(1+f)}\right)$  against an oblivious adversary. Here  $W$  is Lambert’s  $W$  function<sup>4</sup>, defined as the inverse of the function  $z \mapsto ze^z$ . The design of our randomized algorithm is based on two insights:

1. Although the standard greedy online algorithm for picking a maximum-weight basis of a matroid can perform arbitrarily poorly on a worst-case instance of the buyback problem, it performs well when the ratios between values of different matroid elements are powers of some scalar  $r > 1 + f$ . (We call such instances “ $r$ -structured.”)
2. There is a randomized reduction from arbitrary instances of the buyback problem to instances that are  $r$ -structured.

#### 3.1 The Greedy Algorithm and $r$ -Structured Instances

**Definition 1.** Let  $r > 1$  be a constant. An instance of the matroid buyback problem is  $r$ -structured if for every pair of elements  $i, j$ , the ratio  $v_i/v_j$  is equal to  $r^l$  for some  $l \in \mathbb{Z}$ .

**Lemma 1.** Consider the greedy matroid algorithm GMA that always sells elements and buys them back as necessary to maintain the invariant that the set  $S^k$  is a maximum-weight basis of  $\{1, 2, \dots, k\}$ . For  $r > 1 + f$ , when the greedy algorithm is executed on an  $r$ -structured instance of the matroid buyback problem, its competitive ratio is at most  $\frac{r-1}{r-1-f}$ .

*Proof.* As is well known, at termination the set  $S$  selected by GMA is a maximum-weight basis of the matroid. To give an upper bound on the total buyback payment, we define a set  $B(i)$  for each  $i \in \mathcal{U}$  recursively as follows: if GMA never sold to  $i$ , or sold to  $i$  without simultaneously buyback back any element, then  $B(i) = \emptyset$ . If GMA sold to  $i$  while buying back  $j$ , then  $B(i) = \{j\} \cup B(j)$ . By induction on the cardinality of  $B(i)$ , we find that the set  $\{v_x/v_i \mid x \in B(i)\}$  consists of distinct negative powers of  $r$ , so

$$\sum_{x \in B(i)} v_x \leq v_i \cdot \sum_{i=1}^{\infty} r^{-i} = \frac{v_i}{r-1}.$$

---

<sup>4</sup> Lambert’s  $W$  function is multivalued for our domain. We restrict to the case where  $W\left(\frac{-1}{e(1+f)}\right) \leq -1$ .

**Algorithm Filter(ALG):**

- 1: Initialize  $S = \emptyset$ .
- 2: **for**  $i = 1, 2, \dots, n$  **do**
- 3:   Observe  $v_i, w_i$ .
- 4:   Let  $x_i = 1$  with probability  $w_i/v_i$ ,  
       else  $x_i = 0$ .
- 5:   Present  $i$  with value  $w_i$  to ALG.
- 6:   **if** ALG sells to  $i$  **and**  $x_i = 1$  **then**
- 7:     Sell to  $i$ .
- 8:   **end if**
- 9:   **if** ALG buys back  $j$  **and**  $x_j = 1$  **then**
- 10:     Buy back  $j$ .
- 11:   **end if**
- 12: **end for**

**Algorithm RandAlg( $r$ ):**

- 1: **Given:** a parameter  $r > 1 + f$ .
- 2: Sample uniformly random  $u \in [0, 1]$ .
- 3: **for all** elements  $i$  **do**
- 4:   Let  $z_i = u + \lfloor \ln_r(v_i) - u \rfloor$ .
- 5:   Let  $w_i = r^{z_i}$ .
- 6: **end for**
- 7: Run Filter(GMA) on instances  $\mathbf{v}, \mathbf{w}$ .

**Fig. 1.** Randomized algorithms Filter(ALG) and RandAlg( $r$ )

By induction on the number of iterations of the main loop, the set  $\bigcup_{i \in S} B(i)$  consists of all the elements ever bought back by GMA; consequently, the total buyback payment is bounded by  $f \cdot \sum_{i \in S} \sum_{x \in B(i)} v_x \leq \frac{f}{r-1} \sum_{i \in S} v_i$ . Thus, the algorithm’s net payoff is at least  $1 - \frac{f}{r-1}$  times the value of the maximum weight basis.  $\square$

**3.2 The Random Filtering Reduction**

Consider two instances  $\mathbf{v}, \mathbf{w}$  of the matroid buyback problem, consisting of the same matroid  $(\mathcal{U}, \mathcal{I})$ , with its elements presented in the same order, but with different values: element  $i$  has values  $v_i, w_i$  in instances  $\mathbf{v}, \mathbf{w}$ , respectively. Assume furthermore that  $v_i \geq w_i$  for all  $i$ , and that both values  $v_i, w_i$  are revealed to the algorithm at the time element  $i$  arrives. Given a (deterministic or randomized) algorithm ALG which achieves expected payoff  $P$  on instance  $\mathbf{w}$ , we present in Figure 1 an algorithm Filter(ALG) achieving expected payoff  $P$  on instance  $\mathbf{v}$ .

**Lemma 2.** *The expected payoff of Filter(ALG) on instance  $\mathbf{v}$  equals the expected payoff of ALG on instance  $\mathbf{w}$ .*

*Proof.* For each element  $i \in \mathcal{U}$ , let  $\sigma_i = 1$  if ALG sells to  $i$ , and let  $\beta_i = 1$  if ALG buys back  $i$ . Similarly, let  $\sigma'_i = 1$  if Filter(ALG) sells to  $i$ , and let  $\beta'_i = 1$  if Filter(ALG) buys back  $i$ . Observe that  $\sigma'_i = \sigma_i x_i$  and  $\beta'_i = \beta_i x_i$  for all  $i \in \mathcal{U}$ , and that the random variable  $x_i$  is independent of  $(\sigma_i, \beta_i)$ . Thus,

$$\begin{aligned} \mathbf{E} \left[ \sum_{i \in \mathcal{U}} \sigma'_i v_i - (1 + f) \beta'_i v_i \right] &= \mathbf{E} \left[ \sum_{i \in \mathcal{U}} \sigma_i x_i v_i - (1 + f) \beta_i x_i v_i \right] \\ &= \sum_{i \in \mathcal{U}} \mathbf{E}[\sigma_i - (1 + f) \beta_i] \mathbf{E}[x_i v_i] = \sum_{i \in \mathcal{U}} \mathbf{E}[\sigma_i - (1 + f) \beta_i] w_i \\ &= \mathbf{E} \left[ \sum_{i \in \mathcal{U}} \sigma_i w_i - (1 + f) \beta_i w_i \right]. \end{aligned}$$

The left side is the expected payoff of Filter(ALG) on instance  $\mathbf{v}$  while the right side is the expected payoff of ALG on instance  $\mathbf{w}$ .  $\square$

### 3.3 A Randomized Algorithm with Optimal Competitive Ratio

In this section we put the pieces together, to obtain a randomized algorithm with competitive ratio  $-W\left(\frac{-1}{e(1+f)}\right)$  against oblivious adversary<sup>5</sup>. The algorithm RandAlg( $r$ ) is presented in Figure [11](#).

**Lemma 3.** *For all  $i \in \mathcal{U}$ , we have  $v_i \geq w_i$  and  $\mathbf{E}[w_i] = \frac{r-1}{r \ln(r)} v_i$ .*

*Proof.* The random variable  $\ln_r(v_i) - z_i$  is equal to the fractional part of the number  $\ln_r(v_i) - u$ , which is uniformly distributed in  $[0, 1]$  since  $u$  is uniformly distributed in  $[0, 1]$ . It follows that  $w_i/v_i$  has the same distribution as  $r^{-u}$ , which proves that  $v_i \geq w_i$  and also that

$$\mathbf{E}\left[\frac{w_i}{v_i}\right] = \int_0^1 r^{-u} du = -\frac{1}{\ln(r)} \cdot r^{-u} \Big|_0^1 = \frac{r-1}{r \ln(r)}. \quad \square$$

**Theorem 1.** *The competitive ratio of RandAlg( $r$ ) is  $\frac{r \ln(r)}{r-1-f}$ .*

*Proof.* Let  $S^* \subseteq \mathcal{U}$  denote the maximum-weight basis of  $(\mathcal{U}, \mathcal{I})$  with respect to the weights  $\mathbf{v}$ . Since the mapping from  $v_i$  to  $w_i$  is monotonic (i.e.,  $v_i \geq v_j$  implies  $w_i \geq w_j$ ), we know that  $S^*$  is also a maximum-weight basis of  $(\mathcal{U}, \mathcal{I})$  with respect to the weights  $\mathbf{w}$ <sup>6</sup>. Let  $v(S^*) = \sum_{i \in S^*} v_i$  and let  $w(S^*) = \sum_{i \in S^*} w_i$ .

The input instance  $\mathbf{w}$  is  $r$ -structured, so the payoff of GMA on instance  $\mathbf{w}$  is at least  $\frac{r-1-f}{r-1} w(S^*)$ . The modified weights  $w_i$  satisfy two properties that allow application of algorithm Filter(ALG): the value of  $w_i$  can be computed online when  $v_i$  is revealed at the arrival time of element  $i$ , and it satisfies  $w_i \leq v_i$ . By Lemma [2](#), the expected payoff of Filter(GMA) on instance  $\mathbf{v}$ , conditional on the values  $\{w_i : i \in \mathcal{U}\}$ , is at least  $\left(\frac{r-1-f}{r-1}\right) \cdot w(S^*)$ . Finally, by Lemma [3](#) and linearity of expectation,  $\mathbf{E}[w(S^*)] \geq \left(\frac{r-1}{r \ln(r)}\right) \cdot v(S^*)$ . The theorem follows by combining these bounds.  $\square$

The function  $f(r) = \frac{r \ln(r)}{r-1-f}$  on the interval  $r \in (1+f, \infty)$  is minimized when  $-\frac{r}{1+f} = W\left(\frac{-1}{e(1+f)}\right)$  and  $f(r) = -W\left(\frac{-1}{e(1+f)}\right)$ . This completes our analysis of the randomized algorithm RandAlg( $r$ ).

## 4 Lower Bound

We prove the lower bound against an oblivious adversary. The proof first reduces to a continuous version of the problem and then applies Yao’s Principle [10](#). A detailed version of the proof sketches can be found at [11](#).

<sup>5</sup> Note that the algorithm is written in an offline manner just for convenience and can be implemented as an online algorithm.

<sup>6</sup> There may be other maximum-weight basis of  $\mathbf{w}$  which were not maximum-weight basis of  $\mathbf{v}$ .

### 4.1 Reduction to Continuous Version

Consider a new problem. Time starts at  $t = 1$  and stops at time  $t = x$ , where  $x$  is not known to the algorithm. The algorithm at any instant in time can make a mark. The payoff of the algorithm is equal to the time at which it made its final mark minus  $f$  times the sum of times of marks before the final mark. We note that any algorithm for the single item buyback problem with competitive ratio  $c$  can be transformed into an algorithm for the continuous case with competitive ratio  $c \times (1 + \epsilon)$  for arbitrarily small  $\epsilon > 0$ , by discretizing time into small intervals. We prove lower bound for this new problem.

### 4.2 Lower Bound against Oblivious Adversaries

**Theorem 2.** *Any randomized algorithm for the continuous version of the single item buyback problem has competitive ratio at least  $-W\left(\frac{-1}{e(1+f)}\right)$ .*

The proof is an application of Yao’s Principle [10]. We give a one-parameter family of input distributions (parametrized by a number  $y > 1$ ) for the continuous version and prove that any deterministic algorithm for the continuous version of the problem must have a competitive ratio which tends to  $-W\left(\frac{-1}{e(1+f)}\right)$  as  $y \rightarrow \infty$ . Note that for the continuous version of the problem input is just stopping time  $x$ . For a given  $y > 1$ , let the probability density for the stopping times be defined as follows.

$$\begin{aligned} f(x) &= 1/x^2 \text{ if } x < y \\ f(x) &= 0 \text{ if } x > y \end{aligned} \tag{1}$$

Note that the above definition is not a valid probability density function, so we place a point mass at  $x = y$  of probability  $\frac{1}{y}$ . Hence our distribution is a mixture of discrete and continuous probability. For notational convenience let  $d(F(x)) = f(x)$  where  $F$  is the cumulative distribution function. Also let  $G(x) = 1 - F(x)$ . Any deterministic algorithm is defined by a set  $T = \{u_1, u_2, \dots, u_k\}$  of times at which it makes a mark (given that it does not stop before that time).

**Lemma 4.** *There exists an optimal deterministic algorithm described by the set  $T = \{1, w, w^2, \dots, w^{k-1}\}$  for some  $w, k$ .*

*Proof.* Let  $T = \{u_1, u_2, \dots, u_k\}$ . We prove that  $u_i = u_{i+1}^{(i-1)/i}$  for  $i \in [k - 1]$  by induction and it is easy to see that the claim follows from this. For lack of space we just prove the inductive case. Let  $u_0 = 0$  and  $u_{k+1} = \infty$ . Let  $P$  be the expected payoff of the algorithm.

Note that  $P = \sum_{i=1}^k \int_{u_i}^{u_{i+1}} (u_i - f \cdot \sum_{j=1}^{i-1} u_j) d(F(y))$ . We can rewrite the equation as  $P = \sum_{i=1}^k (u_i - (1+f) \cdot u_{i-1}) \cdot G(u_i)$ . If we differentiate  $P$  with respect to  $u_i$ , equate to 0, and solve, then we obtain the equation  $u_i^2 = u_{i-1} \cdot u_{i+1}$ . By induction we know that  $u_{i-1} = u_i^{(i-2)/(i-1)}$ . Substituting and solving we get the necessary equation. □

**Lemma 5.** *For any algorithm described by  $T = \{1, w, w^2, \dots, w^{k-1}\}$ , the competitive ratio is bounded below by a number which tends to  $-W\left(\frac{-1}{e(1+f)}\right)$  as  $y$  tends to  $\infty$ .*

*Proof.* For lack of space we just give a sketch here. Note that if  $V$  is the expected payoff of a prophet who knows the stopping time  $x$ , then  $V = 1 + \ln(y)$ . Also for any algorithm described by  $T = \{1, w, w^2, \dots, w^{k-1}\}$  we have that  $P = 1 + (k - 1) \cdot \frac{w-1-f}{w}$ . Hence if  $c$  is the competitive ratio then  $c = V/P$ . By simple manipulation we see that this is larger than a number which tends to  $-W\left(\frac{-1}{e(1+f)}\right)$  as  $y$  tends to  $\infty$ .  $\square$

## References

1. Babaioff, M., Hartline, J.D., Kleinberg, R.: Selling ad campaigns: online algorithms with buyback. In: Proc. 10th ACM Conference on Electronic Commerce (2009)
2. Constantin, F., Feldman, J., Muthukrishnan, S., Pál, M.: An online mechanism for ad slot reservations with cancellations. In: Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1265–1274 (2009)
3. Constantin, F., Feldman, J., Muthukrishnan, S., Pál, M.: An online mechanism for ad slot reservations with cancellations. In: Proc. 4th Workshop on Ad Auctions (2008)
4. Babaioff, M., Hartline, J.D., Kleinberg, R.: Online algorithms with buyback. In: Proc. 4th Workshop on Ad Auctions (2008)
5. Ben-David, S., Borodin, A., Karp, R., Tardos, G., Wigderson, A.: On the power of randomization in on-line algorithms. *Algorithmica* 11(1), 2–14 (1994)
6. Borodin, A., El-Yaniv, R.: *Online Computation and Competitive Analysis*. Cambridge University Press, Cambridge (1998)
7. Białogorsky, E., Carmon, Z., Fruchter, G.E., Gerstner, E.: Research note: Overselling with opportunistic cancellations. *Marketing Science* 18(4), 605–610 (1999)
8. Sandholm, T., Lesser, V.R.: Leveled commitment contracts and strategic breach. *Games and Economic Behavior* 35, 212–270 (2001)
9. Oxley, J.: *Matroid Theory*. Oxford University Press, Oxford (1992)
10. Yao, A.C.C.: Probabilistic computations: Toward a unified measure of complexity. In: Proceedings of the 18th Annual Symposium on Foundations of Computer Science, Washington, DC, USA, pp. 222–227. IEEE Computer Society, Los Alamitos (1977)
11. Ashwinkumar, B.V., Kleinberg, R.: Randomized online algorithms for the buyback problem. CoRR abs/0908.0043 (2009)



# Envy-Free Allocations for Budgeted Bidders

David Kempe<sup>1</sup>, Ahuva Mu'alem<sup>2</sup>, and Mahyar Salek<sup>1</sup>

<sup>1</sup> Department of Computer Science, University of Southern California,  
CA 90089-0781, USA

{dkempe, salek}@usc.edu

<sup>2</sup> Social and Information Sciences Laboratory, California Institute of Technology,  
Pasadena, CA 91106, USA

ahumu@yahoo.com

**Abstract.** We study the problem of identifying prices to support a given allocation of items to bidders in an envy-free way. A bidder will envy another bidder if she would prefer to obtain the other bidder's item at the price paid by that bidder. Envy-free prices for allocations have been studied extensively; here, we focus on the impact of *budgets*: beyond their *willingness* to pay for items, bidders are also constrained by their *ability* to pay, which may be lower than their willingness.

In a recent paper, Aggarwal et al. show that a variant of the Ascending Auction finds a feasible and bidder-optimal assignment and supporting envy-free prices in polynomial time so long as the input satisfies certain non-degeneracy conditions. While this settles the problem of finding a feasible allocation, an auctioneer might sometimes also be interested in a *specific* allocation of items to bidders. We present two polynomial-time algorithms for this problem, one which finds maximal prices supporting the given allocation (if such prices exist), and another which finds minimal prices. We also prove a structural result characterizing when different allocations are supported by the same minimal price vector.

## 1 Introduction

One of the most central and basic economic problems is the allocation of items to individuals. This is frequently accomplished via auctions, wherein the bidders communicate their values for the items to an auctioneer, who then decides on an allocation of items to bidders and prices to be paid.

An important property of an auction is that it be *envy-free*: no bidder wishes to receive one or more items assigned to other bidders at the price the other bidders are paying. If bidders were envious in this sense, the outcome of the auction might not be stable, or bidders might refuse to participate in the auction in the future. There has been a big surge in interest in envy-free allocations and pricing of items within the computer science community recently [9,12,6,2]. Much of the work focuses on the interplay between combinatorial structure among the item sets bidders are interested in and the revenue that can be extracted, usually with efficient computation.

In reality, bidders are not only constrained by their *willingness* to pay for items, but also by their *ability* to pay [5,8]. For instance, a bidder looking for a

house might have an extremely high valuation for a mansion, but nowhere near the resources to buy it at a price close to her valuation. Then, her envy will only be relevant if another bidder gets to purchase the mansion at a price *which this bidder could afford*.

Introducing budget limitations changes the problem significantly. For instance, there may now be feasible allocations which do not maximize social welfare, and efficient allocations may not be feasible any more. More generally, the structure of feasible allocations and matching prices becomes quite rich. In a recent paper, Aggarwal et al. [1] show that a variant of the Ascending Auction finds, in polynomial time, a feasible assignment and supporting *envy-free budget-friendly truthful* prices so long as the input satisfies certain non-degeneracy conditions. In fact, the allocation they find is *bidder-optimal*, in the sense that the price paid by every bidder is a lower bound on the price the bidder could pay for *any* feasible allocation and corresponding prices.

While this settles the problem of finding a feasible allocation, an auctioneer might sometimes also be interested in a *specific* allocation. For instance, there may be constraints not captured otherwise which prescribe that certain allocations are preferable from the auctioneer's point of view. Thus, an important and natural question is whether, given the bidders' valuations and budgets (as well as the auctioneer's reserve prices), a *given* allocation of items to bidders can be supported with envy-free prices.

In this paper, we give two polynomial-time algorithms for this problem, one which finds maximal envy-free prices supporting the given allocation (if such prices exist), and another which finds minimal prices. In particular, our algorithms show the existence of maximal and minimal price vectors. Both algorithms are based on label-relaxation schemes (of a dynamically constructed graph) in the style of the Bellman-Ford algorithm for shortest paths; in the case of the minimal prices, this algorithm has to be augmented by a further insight to prevent pseudo-polynomial running time. Furthermore, as a first step toward a more complete characterization of feasible allocations and the corresponding supporting envy-free budget-friendly prices, we give a combinatorial condition for minimal price vectors to be the same.

**Related Work.** Guruswami et al. [11] initiated the study of *envy-free revenue-maximization* for *non-budget-constrained unit-demand* bidders. If all items must be allocated, the maximum price vector can be found in polynomial time [13]. However, if some items can be omitted to increase competition, then this general problem becomes APX-hard; the current best approximation guarantee is  $O(\log n)$  [11]. Multi-unit *truthful* auctions for *budget-constrained bidders* with linear valuations were first studied by Borgs et al. [4]. They constructed a truthful randomized mechanism which asymptotically achieves revenue maximization. Dobzinski et al. [8] essentially show that a deterministic truthful Pareto-optimal auction exists if and only if budgets are public information. Additionally, for the case of an infinitely-divisible single good, no anonymous truthful mechanism can produce Pareto-optimal allocations if bidders are budget-constrained [8], whereas if randomization is allowed, such mechanisms do exist [3].

## 2 Model and Preliminaries

We consider a set  $M$  of  $n$  distinct indivisible items, and a set  $N$  of  $n$  bidders. Bidders are *unit-demand*, i.e., each bidder is interested in purchasing at most one item. Bidder  $i$ 's willingness to pay is captured by a valuation function  $v$ . Thus, bidder  $i$  has value  $v_i(j)$  for item  $j$ . Additionally, each bidder has an item-specific budget  $b_i(j)$ , indicating her *ability* to pay for item  $j$ : the maximum amount of money the bidder can afford for this item. A particularly natural special case is when  $b_i(j) = b_i$  for all  $j$ , i.e., bidder  $i$  is constrained by a fixed amount of money. However, our results hold in more generality. If  $b_i(j) \leq v_i(j)$  for at least one item  $j$ , we call bidder  $i$  *budget-constrained*, otherwise, bidder  $i$  is *non-budget-constrained*. For convenience, we denote  $v_i^{(0)}(j) = \min(v_i(j), b_i(j))$ .

Item  $j$  will be assigned price  $p_j$ ; we use  $\mathbf{p}$  to denote the vector of all prices. The prices may be constrained by the auctioneer: the auctioneer has *reserve prices*  $r_j \geq 0$  for items  $j$ , such that an item cannot be sold at a price less than  $r_j$ . In other words, a price vector  $\mathbf{p}$  is feasible only if  $\mathbf{p} \geq \mathbf{r}$ . Additionally, when  $p_j < b_i(j)$ , we say that bidder  $i$  *can afford* item  $j$  with prices  $\mathbf{p}$ . (We require strict inequality for technical convenience; among other things, it makes the notion of a minimal price vector well-defined.) When assigned item  $j$  at price  $p_j$ , bidder  $i$  derives a *utility* of  $u_i(j) = v_i(j) - p_j$  if  $p_j < b_i(j)$  and  $-\infty$  otherwise. Therefore, the utility is positive whenever  $p_j < v_i^{(0)}(j)$ .

In general, an *allocation*  $a$  is a partition  $A_1, \dots, A_n$  of the  $n$  items among the  $n$  bidders, where  $A_i$  is the set of items allocated to bidder  $i$ . Since we focus on unit-demand bidders, we are particularly interested in allocations that are *assignments*, in that  $|A_i| = 1$  for all  $i$ , i.e., each bidder gets exactly one item. In that case, we write  $a_i$  for the unique item assigned to bidder  $i$ .

**Definition 1 (Envy-Free Budget-Friendly Allocations, Supporting Prices).** *An allocation  $\mathbf{a}$  is envy-free budget-friendly if there exists a price vector  $\mathbf{p} \geq \mathbf{r}$  such that for every  $i = 1, \dots, n$ :*

1.  $p_{a_i} < b_i(a_i)$  (bidder  $i$  can afford the item allocated to her) and  $p_{a_i} \leq v_i(a_i)$  (bidder  $i$  derives non-negative utility from her item).
2.  $v_i(a_i) - p_{a_i} \geq v_i(j) - p_j$  for all items  $j$  with  $p_j < b_i(j)$ . That is, bidder  $i$  would not prefer another item she can afford over her own at the current prices.

A feasible price vector  $\mathbf{p}$  satisfying these conditions is said to support the allocation  $a$ .

The notion of envy-free budget-friendly allocations can be considered a generalization of a Walrasian Equilibrium [7,10] to budget-constrained bidders.<sup>1</sup> Unlike the case of non-budget-constrained bidders, there need not be any envy-free budget-constrained assignments (e.g., [14]). Furthermore, even when such assignments do exist, the efficient allocation might not be envy-free budget-friendly.

<sup>1</sup> We are mainly interested in assignments; therefore, we do not require that any unallocated items have zero price.

Formally, the input consists of the matrix of valuations  $V = (v_i(j))_{i,j}$ , the matrix of budget limits  $B = (b_i(j))_{i,j}$ , and an allocation  $\mathbf{a}$ . The goal is to identify a price vector  $\mathbf{p}$  supporting  $\mathbf{a}$ , or to conclude that no such price vector exists.

### 3 Polynomial-Time Algorithms

For simplicity, we assume that the desired allocation  $\mathbf{a}$  is  $a_i = i$  for all bidders. We then use  $p_i$  to denote the price of the item assigned to bidder  $i$ . We can also assume that  $v_i^{(0)}(i) \geq r_i$ ; otherwise, no supporting price vector exists.

Both of our algorithms for the assignment problem are based on the notion of an *envy graph*.

**Definition 2 (Envy Graph  $G_{\mathbf{p}}$ ).** *Given an arbitrary price vector  $\mathbf{p}$ , the envy graph  $G_{\mathbf{p}}$  has one node for each bidder, and a directed edge from bidder  $i$  to bidder  $j$  if and only if  $p_j < b_i(j)$ , i.e., if and only if bidder  $i$  could afford bidder  $j$ 's assigned item at the current prices. Whenever the edge  $(i, j)$  is present, it is labeled  $\lambda_{(i,j)} = v_i(i) - v_i(j)$ .*

Intuitively, the label captures how much bidder  $i$  “prefers” bidder  $j$ 's item over her own, if both were priced the same. (The more negative  $\lambda_{(i,j)}$  is, the more  $i$  prefers  $j$ 's item.) Notice that the edge labels are *independent* of the price vector  $\mathbf{p}$ , and only the existence or non-existence of edges depends on the prices. The following two simple insights lie at the heart of our algorithms:

**Proposition 1.** *Let  $P$  be any directed path from  $i$  to  $j$  in  $G_{\mathbf{p}}$ , and  $L = \sum_{e \in P} \lambda_e$  the sum of labels along the path.*

1. *Let  $\mathbf{p}$  be any price vector such that for every price vector  $\mathbf{p}'$  supporting the allocation  $\mathbf{a}$ , we have  $\mathbf{p} \leq \mathbf{p}'$  (component-wise). Then,  $p'_j \geq p_i - L$ .*
2. *Let  $\mathbf{p}$  be any price vector such that for every price vector  $\mathbf{p}'$  supporting the allocation  $\mathbf{a}$ , we have  $\mathbf{p} \geq \mathbf{p}'$  (component-wise). Then,  $p'_i \leq p_j + L$ .*

**Proof.** We prove the first statement — the second one is analogous. For any edge  $(u, v) \in P$ , envy-freeness of  $\mathbf{p}'$  implies that  $p'_v \geq p'_u - \lambda_{(u,v)}$ . Adding the inequalities for all edges  $e \in P$ , and using that  $p'_i \geq p_i$  now proves the claim. ■

By setting  $i = j$  in Proposition [1](#), we obtain the following simple corollary:

**Corollary 1.** *If  $\mathbf{p}$  is an envy-free price vector, then  $G_{\mathbf{p}}$  contains no negative cycles.*

#### 3.1 Finding Minimal Prices

The first part of Proposition [1](#) suggests a simple pseudo-polynomial algorithm for finding supporting minimal prices for an allocation (or concluding that no supporting prices exist). Algorithm [1](#) is a label relaxation algorithm in the style of the Bellman-Ford shortest paths algorithm.

---

**Algorithm 1.** Label Relaxation for Minimal Supporting Prices

---

- 1: Start with  $p_i = r_i$  for all  $i$ .
  - 2: **while** there is an edge  $(i, j) \in G_{\mathbf{p}}$  with  $p_i > p_j + \lambda_{(i,j)}$  **do**
  - 3:   Update  $p_j := \min(b_i(j), p_i - \lambda_{(i,j)})$ .
  - 4:   Remove any edge  $(u, j)$  with  $p_j \geq b_u(j)$  from  $G_{\mathbf{p}}$ .
  - 5: **if**  $p_i \geq b_i(i)$  for any  $i$  **then**
  - 6:   No supporting prices exist.
  - 7: **else**
  - 8:    $\mathbf{p}$  is a supporting price vector.
- 

The pseudo-polynomial running time results from negative cycles in  $G_{\mathbf{p}}$ . To speed up the algorithm, we will therefore choose the edge  $(i, j)$  in the **while** loop judiciously to break negative cycles fast. Let  $C$  be a negative cycle in  $G_{\mathbf{p}}$ , with nodes  $u_1, u_2, \dots, u_k$ . Let  $P_{ij}$  denote the unique path from  $u_i$  to  $u_j$  on  $C$ , and  $L_{ij} = \sum_{e \in P_{ij}} \lambda_{(u_i, u_j)}$  the total edge weight on  $P_{ij}$ . Intuitively, the update step from Algorithm 1 will have to continue until at least one of the edges  $(u_i, u_{i+1})$  is broken, because item  $i + 1$  is not affordable to bidder  $i$  any more. However, this may take pseudo-polynomial time. Our goal is to “fast-forward” the update steps along the cycle.

**Lemma 1.** *There exists a node  $u_i$  such that  $p_{u_i} > p_{u_j} + L_{ij}$  for all  $j$ .*

**Proof.** Suppose for contradiction that for each  $i$ , there exists a  $j(i)$  such that  $p_{u_i} \leq p_{u_{j(i)}} + L_{ij(i)}$ . Consider the graph on nodes  $u_i$  with an edge from  $u_i$  to  $u_{j(i)}$ . Because each node has an outgoing edge, this graph must contain some cycle  $C' = \{u_{i_1}, \dots, u_{i_\ell}, u_{i_{\ell+1}} = u_{i_1}\}$  such that  $p_{u_{i_r}} \leq p_{u_{i_{r+1}}} + L_{i_r, i_{r+1}}$  for all  $1 \leq r \leq \ell$ . Because each node appears once on the right and left side, after adding up these inequalities and canceling out, we obtain that  $\sum_{r=1}^{\ell} L_{i_r, i_{r+1}} \geq 0$ . But the sum is exactly the weight of going around  $C$  one or more times (following  $C'$ ), and thus negative, a contradiction. ■

If we update the node prices in the order  $u_2, u_3, \dots, u_k$ , it is easy to see by induction that (1) each node will need to be updated upon its turn, and (2)  $u_i$  will be updated to  $p_{u_1} - L_{1i}$ . Extending this observation to updates continuing around  $C$ , we can see the following:

**Proposition 2.** *If the algorithm has updated the prices going around  $C$ , and has updated node  $u_i$   $c$  times, then its new price is  $p'_{u_i} = p_{u_1} - cL - L_{1i} > p_{u_i}$ .*

Thus, we can determine the outcome of the update process as follows: For each  $i$ , let  $c_i = \lfloor \frac{b_{u_{i-1}}(u_i) - (p_{u_1} - L_{1i})}{L} \rfloor$  be the number of iterations around the cycle after which bidder  $u_{i-1}$  cannot afford item  $u_i$  any more (where  $u_0 = u_k$ ). Then, let  $j = \operatorname{argmin}_i c_i$ , with ties broken for the smallest  $i$ . According to Proposition 2 and the definition of  $j$ , if we update each  $u_i$  (for  $i \leq j$ )  $c_i$  times, and each  $u_i$

---

<sup>2</sup> An alternative proof reduces this statement to the well-known “Racetrack” puzzle. We thank Peter Winkler for this observation.

for  $i > j$   $c_i - 1$  times, then  $p'_{u_j} > b_{u_{j-1}}(u_j)$ , and  $p'_{u_i} \leq b_{u_{i-1}}(u_i)$  for all  $i \neq j$ . In particular, this means that the updates are consistent with an execution of the relaxation algorithm.

Thus, Algorithm 2 is a polynomial-time version of Algorithm 1.

---

**Algorithm 2.** Polynomial-Time Minimal Supporting Prices
 

---

- 1: Start with  $p_i = r_i$  for all  $i$ .
  - 2: **while**  $G_{\mathbf{p}}$  contains a negative cycle  $C$  **do**
  - 3:   Let  $u_1 \in C$  be a node satisfying Lemma 1, and  $C = \{u_1, \dots, u_k\}$ .
  - 4:   Compute  $L_{1i} = \sum_{j=1}^{i-1} \lambda_{(u_j, u_{j+1})}$  for all  $i$ .
  - 5:   Compute  $c_i = \lfloor \frac{b_{u_{i-1}}(u_i) - (p_{u_1} - L_{1i})}{L} \rfloor$  for all  $i$ .
  - 6:   Let  $j = \operatorname{argmin}_i c_i$ , ties broken for smallest  $i$ .
  - 7:   Update  $p'_{u_i} = p_{u_1} - L_{1i} - c_j L$  for  $i \leq j$ , and  $p'_{u_i} = p_{u_1} - L_{1i} - (c_j - 1)L$  for  $i > j$ .
  - 8:   Update  $\mathbf{p} = \mathbf{p}'$ , and update  $G_{\mathbf{p}}$ .
  - 9:   **if**  $p_i \geq b_i(i)$  for any  $i$  **then**
  - 10:     No supporting prices exist.
  - 11: **else**
  - 12:    $\mathbf{p}$  is a supporting price vector.
- 

The running time in each iteration is dominated by finding a negative cycle, which can be accomplished in time  $O(mn)$  by a simple extension of the Bellman-Ford algorithm. All other operations take time  $O(n)$ . Since each iteration of the **while** loop removes at least one edge, the total running time is at most  $O(m^2n)$ .

Proposition 1 implies by induction that in each iteration, the vector  $\mathbf{p}$  of the algorithm satisfies  $\mathbf{p} \leq \mathbf{p}'$  (component-wise) for any price vector  $\mathbf{p}'$  supporting  $\mathbf{a}$ . Thus, whenever Algorithm 1 outputs a price vector  $\mathbf{p}$ , we have that  $\mathbf{p} \leq \mathbf{p}'$  for any price vector  $\mathbf{p}'$  supporting  $\mathbf{a}$ . Because Algorithm 2 outputs the same final vector as Algorithm 1, we have proved:

**Corollary 2.** *If  $\mathbf{a}$  is an envy-free budget-friendly allocation for  $V, \mathbf{b}$ , then Algorithm 2 outputs the (unique) minimal price vector  $\mathbf{p}^-$  satisfying  $\mathbf{p}^- \leq \mathbf{p}'$  (component-wise) for all price vectors  $\mathbf{p}'$  supporting  $\mathbf{a}$ . In particular, there exists a unique minimal price vector supporting  $\mathbf{a}$ .*

**Maximal Prices.** It is possible to find maximal prices supporting  $\mathbf{a}$ . In this case, the procedure starts with prices  $p_i = v_i^{(0)}(i)$  and iteratively makes price-adjustment similar to Algorithm 1, except prices are *decreased* in response to envy. If there remains a negative cycle once the algorithm terminates, we deduce that no supporting prices exist. The algorithm can be shown to run in polynomial time even without fast-forwarding. Due to space constraints, the algorithm will be discussed in detail in the full version of this paper.

## 4 Affordability Graphs and Minimal Price Vectors

The structure of feasible allocations and corresponding supporting prices is much richer in the presence of budgets than for traditional envy-free auctions. If all

bidders are non-budget-constrained, an allocation is feasible if and only if it is efficient (i.e.,  $\sum_i v_i(a_i) \geq \sum_i v_i(a_{\pi(i)})$  for any permutation  $\pi$ ). A price vector supports either all allocations, or none of them [10]. However, once we introduce budgets, the situation changes significantly. The efficient allocation may not be feasible with budgets, while inefficient allocations are. Furthermore, there can be allocations  $\mathbf{a}, \mathbf{a}'$  with corresponding supporting prices  $\mathbf{p}, \mathbf{p}'$  such that  $\mathbf{p}$  does not support  $\mathbf{a}'$ , and vice versa. As a first step toward a complete characterization, we give a combinatorial condition for *minimal* price vectors to be the same. The condition is based on the concept of an affordability graph.

**Definition 3 (Affordability Graph  $H_{\mathbf{p}}$ ).** *The affordability graph  $H_{\mathbf{p}}$  is a bipartite graph on bidders and items, containing an edge  $(i, j)$  if and only if bidder  $i$  can afford item  $j$  at the prices  $\mathbf{p}$ , i.e.,  $p_j < b_i(j)$ .*

If  $\mathbf{p}$  is a *minimal* price vector,  $H_{\mathbf{p}}$  captures all of the essential information about  $\mathbf{p}$ , in the following sense (a generalization of Lemma 6 in [10]):

**Lemma 2.** *Let  $\mathbf{a}, \mathbf{a}'$  be two envy-free budget-friendly assignments, and  $\mathbf{p}, \mathbf{p}'$  the corresponding minimal supporting prices. Then  $\mathbf{p} = \mathbf{p}'$  if and only if  $H_{\mathbf{p}} = H_{\mathbf{p}'}$ . Furthermore, if  $H_{\mathbf{p}} = H_{\mathbf{p}'}$ , then the social welfare of all bidders is the same under  $(\mathbf{a}, \mathbf{p})$  and  $(\mathbf{a}', \mathbf{p}')$ , i.e.,  $\sum_i v_i(a_i) = \sum_i v_i(a'_i)$ .*

**Proof.** One direction is obvious: if  $\mathbf{p} = \mathbf{p}'$ , then the edge  $(i, j)$  is in  $H_{\mathbf{p}}$  if and only if it is in  $H_{\mathbf{p}'}$ . Hence,  $H_{\mathbf{p}} = H_{\mathbf{p}'}$ . For the converse direction, assume that  $H_{\mathbf{p}} = H_{\mathbf{p}'}$ . Because  $\mathbf{a}$  is envy-free and supported by  $\mathbf{p}$ , each bidder prefers her own assigned item to all items she can afford, i.e.,

$$v_i(a_i) - p_{a_i} \geq v_i(j) - p_j \tag{1}$$

for every item  $j$  with  $p_j < b_i(j)$ . Because  $\mathbf{a}'$  is an allocation, we can write  $j = a'_k$  for a (unique)  $k$  in the right-hand side above, obtaining:

$$v_i(a_i) - p_{a_i} \geq v_i(a'_k) - p_{a'_k} \tag{2}$$

for each  $k$  with  $p_{a'_k} < b_i(a'_k)$ . Because bidder  $i$  can afford item  $a'_i$  with the price vector  $\mathbf{p}'$ , and the affordability graphs are the same,  $i$  can also afford  $a'_i$  with prices  $\mathbf{p}$ . Thus, we can apply Inequality (2) with  $k = i$ , to obtain that  $v_i(a_i) - p_{a_i} \geq v_i(a'_i) - p_{a'_i}$ . Summing this inequality over all bidders  $i$ , and noticing that both  $\mathbf{a}$  and  $\mathbf{a}'$  are permutations, gives us that

$$\sum_i (v_i(a_i) - p_{a_i}) \geq \sum_i (v_i(a'_i) - p_{a'_i})$$

Adding  $\sum_i p_{a_i}$  on both sides shows that  $\sum_i v_i(a_i) \geq \sum_i v_i(a'_i)$ . A completely symmetric argument shows the opposite inequality, so we have proved that  $\sum_i v_i(a_i) = \sum_i v_i(a'_i)$ .

Subtracting  $\sum_i p_{a_i} = \sum_i p_{a'_i}$  on both sides implies that  $\sum_i (v_i(a_i) - p_{a_i}) = \sum_i (v_i(a'_i) - p_{a'_i})$ . If there were an  $i$  with  $v_i(a_i) - p_{a_i} > v_i(a'_i) - p_{a'_i}$ , then there would have to be some  $k$  with  $v_k(a_k) - p_{a_k} < v_k(a'_k) - p_{a'_k}$ , which would contradict

the fact that  $\mathbf{p}$  supports  $\mathbf{a}$ . Thus,  $v_i(a_i) - p_{a_i} = v_i(a'_i) - p_{a'_i}$  for all bidders  $i$ . Combining this with Inequality (II) we get that  $v_i(a'_i) - p_{a'_i} \geq v_i(j) - p_j$  for every item  $j$  with  $p_j < b_i(j)$ . Thus,  $\mathbf{p}$  supports the assignment  $\mathbf{a}'$ , and by the minimality of  $\mathbf{p}'$ , we get that  $\mathbf{p}' \leq \mathbf{p}$  component-wise. A symmetric argument shows that  $\mathbf{p} \leq \mathbf{p}'$ , and thus completes the proof. ■

**Acknowledgments.** We would like to thank Itai Ashlagi, Liad Blumrosen, Ron Lavi, John Ledyard, Debasis Mishra and Noam Nisan for helpful discussions and comments.

## References

1. Aggarwal, G., Muthukrishnan, S., Pál, D., Pál, M.: General auction mechanism for search advertising. In: WWW (2009)
2. Amar, L., Mu'alem, A., Stoesser, J.: On the importance of migration for fairness in online grid markets. In: GRID (2008)
3. Bhattacharya, S., Conitzer, V., Munagala, K., Xia, L.: Incentive compatible budget elicitation in multi-unit auctions, working paper (2009)
4. Chayes, J.T., Borgs, C., Immorlica, N., Mahdian, M., Saberi, A.: Multi-unit auctions with budget-constrained bidders. In: EC (2005)
5. Che, Y.-K., Gale, I.: Standard auctions with financially constrained bidders. *Review of Economic Studies* 65(1), 1–21 (1998)
6. Chen, N., Ghosh, A., Vassilvitskii, S.: Optimal envy-free pricing with metric substitutability. In: EC (2008)
7. Demange, G., Gale, D., Sotomayor, M.: Multi-item auctions. *Journal of Political Economy* 94(4), 863–872 (1986)
8. Dobzinski, S., Lavi, R., Nisan, N.: Multi-unit auctions with budget limits. In: FOCS (2008)
9. Goldberg, A.V., Hartline, J.D.: Envy-free auctions for digital goods. In: EC (2003)
10. Gul, F., Stacchetti, E.: Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87(1), 95–124 (1999)
11. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: SODA, pp. 1164–1173 (2005)
12. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: EC (2004)
13. Shapley, L.S., Shubik, M.: The assignment game I: The core. *Journal of Game Theory* 1(1), 111–130 (1972)
14. van der Laan, G., Yang, Z.: An ascending multi-item auction with financially constrained bidders. Tinbergen Institute Discussion Papers 08-017/1, Tinbergen Institute (February 2008)



# On the Inefficiency Ratio of Stable Equilibria in Congestion Games

Arash Asadpour and Amin Saberi

Department of Management Science and Engineering,  
Stanford University, Stanford, CA 94305  
{asadpour,saberi}@stanford.edu

**Abstract.** Price of anarchy and price of stability are the primary notions for measuring the efficiency (i.e. the social welfare) of the outcome of a game. Both of these notions focus on extreme cases: one is defined as the inefficiency ratio of the worst-case equilibrium and the other as the best one. Therefore, studying these notions often results in discovering equilibria that are not necessarily the most likely outcomes of the dynamics of selfish and non-coordinating agents.

The current paper studies the inefficiency of the equilibria that are most stable in the presence of noise. In particular, we study two variations of non-cooperative games: atomic congestion games and selfish load balancing. The noisy best-response dynamics in these games keeps the joint action profile around a particular set of equilibria that minimize the potential function. The inefficiency ratio in the neighborhood of these “stable” equilibria is much better than the price of anarchy. Furthermore, the dynamics reaches these equilibria in polynomial time.

Our observations show that in the game environments where a small noise is present, the system as a whole works better than what a pessimist may predict. They also suggest that in congestion games, introducing a small noise in the payoff of the agents may improve the social welfare.

## 1 Introduction

The inefficiency of equilibria in noncooperative games has been studied vastly in the algorithmic game theory community during the last decade. It has been observed that the rational behavior of selfish players can be worse than a centrally designed outcome, where the decency of the outcome is measured by its corresponding (expected) social welfare (or cost) either in the long run or a certain period of time. There have been different attempts to understand how much these outcomes might be different from each other. The two most used notions are *price of anarchy* to understand how inferior the worst equilibrium is (comparing to the best possible outcome), and *price of stability* to compare the worst possible gap between the best equilibrium and the optimal outcome.

These notions are representing some extreme equilibria of the game which might be unstable in the presence of noise. In other words, even a small noise in either the behavior or payoffs of the players may move the game quickly out

of these equilibria. Motivated by this observation, and following the literature on evolutionary game theory, we study the noisy best response dynamics. In such dynamics, with overwhelming probability players take the best response action, but with a small probability they may take another action. This may be attributed to actual noise in the payoff, to the noise in the perceived information from the environment, or simply to the failure in realizing the “exact” best response.

Noisy best response dynamics, introduced by Fudenberg and Kreps [12] have been used extensively in evolutionary game theory to model the behavior of agents. They also have been used in the context of experimental economics to rationalize the perceived data. The corresponding perturbed equilibrium in this context is called “Quantal Response Equilibrium (QRE)”. Logit dynamics, first appeared in Blume [5], is one of the most widely used model of noisy best response dynamics in evolutionary game theory and experimental economics. This dynamics defines a Markov chain with stationary measure  $\mu_\beta$  where the parameter  $\beta$  is the inverse of the noise in the system. The dynamics of this chain is similar to that of simulated annealing. It has a drift towards equilibria with lower potential function.

## 1.1 Our Results

We study unweighted finite (i.e. atomic) congestion games (Section 3) and selfish load balancing games with identical machines. (We omit the proofs for the load balancing games here due to the lack of space, and present them in the longer version.) Our main contribution is to show that for these games and under certain mild assumptions, the dynamics reaches the neighborhood of equilibria with the globally minimum potential function after a reasonable time and then spends most of the time there. We also characterize the inefficiency ratio in this neighborhood.

Let us explain the statements in the previous paragraph more precisely. Define  $\mathcal{A}_\epsilon$  to be the subset of the states of the game for which the potential function is within a factor  $1 + \epsilon$  of the minimum potential function. We make the following observations.

1. For every constant  $\epsilon > 0$ , the game reaches an state in  $\mathcal{A}_\epsilon$  in polynomial time in the number of players. We show that for any constant  $\epsilon > 0$ ,  $\lim_{\beta \rightarrow \infty} \mu_\beta(\mathcal{A}_\epsilon) = 1$ . Consequently, when the noise is sufficiently small, the dynamics spends almost all the time in  $\mathcal{A}_\epsilon$ . We call  $\mathcal{A}_\epsilon$  a **stable region** of the game.
2. Define the **Inefficiency Ratio of Stable Equilibria (IRSE)** of a non-cooperative potential game to be the “worst case inefficiency ratio of the equilibria that are potential function minimizers”. We will show that the inefficiency ratio of all the states in  $\mathcal{A}_\epsilon$  is at most  $\text{IRSE} + c\epsilon$  for a constant  $c$  independent of the number of players.

These two observations together suggest that IRSE is a proper notion of the inefficiency ratio of the most plausible and stable outcome of the game: by the

**Table 1.** IRSE vs. Price of Stability and Anarchy, (\*) means folk theorem

Game	Cost Function	PofS	IRSE	PofA
Load Balancing on identical machines	Makespan	1 (*)	$\geq \frac{19}{18}, \leq \frac{3}{2}$	$2(1 - \frac{1}{m+1})$ [11]
Unweighted Atomic Congestion games	Linear	$1 + \frac{1}{\sqrt{3}} \approx 1.578$ [6]	$1 + \frac{1}{\sqrt{3}}$	2.5 [8] [2]
	Degree $p$ Polynomial	$\Theta(p)$ (*)	$p + 1$	$p^{\Theta(p)}$ [8]
	Linear (Symmetric/Network)	$\frac{4}{3}$ [1]	$\frac{4}{3}$	$\geq 2$

first observation, we know that the game will quickly reaches some state in the stable region  $\mathcal{A}_\epsilon$ , and stays there almost all the time afterwards. On the other hand, the second observation implies that for small enough values of  $\epsilon$  the worst inefficiency ratio of any state in the stable region gets arbitrarily close to IRSE.

Our results on IRSE are summarized in Table 1. We use the well-known “variational inequalities” for deriving these results. They show that the inefficiency ratio of stable equilibria in such games are usually much better than the price of anarchy and often close to the price of stability. Note that IRSE interestingly coincides with the price of stability in atomic congestion games and atomic selfish routings with linear cost functions. Although this may raise the speculation that the IRSE is in fact equal to the price of stability, we will see that it is not true in general. The IRSE for load balancing games can be at least  $\frac{19}{18}$  while the price of stability is 1. Possibly, the most challenging part of our proofs is to bound the hitting time of  $\mathcal{A}_\epsilon$ . It has been argued that stable equilibria is a good measure for predicting the outcome of a play only if it takes reasonable time for the dynamics to reach it [10,15]. We will show that when the number of resources is constant, the dynamics reaches  $\mathcal{A}_\epsilon$  in polynomial time in the number of players. Our convergence times are all polynomial in the number of players, but might be exponential in the number of resources (e.g. size of the network in the network congestion games, or the number of machines in the selfish load balancing games). In the longer version of this paper we show that it is impossible to get a convergence time which is polynomial in the number of resources. In fact, we show that such a result cannot exist even in games with three players. Furthermore, note that due to results of [7] and [3], proving a fast convergence to the exact potential minimizer solution also seems hopeless.

One interesting implication of our results in the context of protocol design is that in such games introducing a small noise in the payoff of the agents may increase the social welfare. In other words, apart from relying on the actual noises in the environment, the designer can introduce some noise in the system (for example by adding random Gaussian delays to the routing systems). Our results imply that this noise, while most likely not hurting individuals, may have a considerable positive effect on the social welfare.

### 1.2 Related Works

There have been several aims to analyze “no-regret learning” algorithms, where each player is essentially running a regret minimization algorithm on her observed

payoffs. For a detailed description and more references we encourage the reader to study [14] and [4]. Our work however, is conceptually different from them. We do not assume that the players are running a learning algorithm. Instead, we consider myopic agents in the presence of the noise. In this sense, our work is in the line of the literature on the “evolutionary stable states (ESS)”. For a comprehensive survey on the topic one can refer to the book by Fudenberg and Levine [13]. One of the most related works to ours in this subject is the paper by Chung, et al [9] which considers price of stochastic anarchy for the load balancing games and shows its boundedness.

## 2 Noisy Best-Response Dynamics

Suppose that a set  $V = \{1, 2, \dots, n\}$  of players are playing a game in the course of time. Each player  $i$  has a set of pure strategies  $S_i$ . The set  $S$ , Cartesian product of  $S_i$ 's, denote the set of all pure strategy profiles. We show by  $S_{-i}$  the set of all strategy profiles of players other than  $i$ . For a strategy profile  $s$ , the utility of player  $i$  is shown by  $u_i^s$ . In a best-response dynamics (without noise) each player  $i$  would change her own pure strategy if it leaves her better off.

In this section we address *noisy best-response dynamics* in potential games. In particular, we study “logit dynamics”. The observations made in this section are heavily based on the classic works of Blume [5] and Fudenberg et. al. [13].

Noisy best-response dynamics are defined by a parameterized family of Markov chains  $\text{Pr}_\beta$  indexed by  $\beta$ . The parameter  $\beta \in \mathbb{R}^+$  indicates how noisy the system is, and can be interpreted as the inverse noise of the system. Each player updates its current strategy  $s_i$  to  $s'_i \in S_i$  with probability  $p_\beta(s'_i, s_{-i})$  where:

$$p_\beta(s'_i, s_{-i}) \propto e^{-\beta u_i^{s'_i}}. \tag{1}$$

In above,  $s' = (s'_i, s_{-i})$  is the same as  $s$  with only  $s_i$  replaced by  $s'_i$  (and hence,  $s_{-i} = s'_{-i}$ ). Note that when  $\beta = +\infty$ , the system behaves with a noise-free best-response dynamics. In the potential games, the update rule in Eq. (1) can be rephrased as below.

$$p_\beta(s'_i, s_{-i}) \propto e^{-\beta \Phi(s')}. \tag{2}$$

Here,  $\Phi$  is the potential function of the game and thus,  $\Phi(s') - \Phi(s) = u_i^{s'_i} - u_i^s$ . In weighted potential games there exists values of  $w_i$  such that  $\Phi(s') - \Phi(s) = w_i(u_i^{s'_i} - u_i^s)$ . In such games we consider the noisy best-response dynamics explained by Eq. (2).

This game defines a reversible Markov chain on  $S$  with respect to the stationary distribution  $\mu(s) \propto e^{-\beta \Phi(s)}$  (for more details, see [5] and [13]). It is immediate that as  $\beta \rightarrow \infty$  the measure of  $\mathcal{A}_\epsilon$  (for any constant  $\epsilon$ ) goes to 1.

An interesting implication of the fact that the inverse noise,  $\beta$ , is unbounded from above is that as the system approaches to best response (i.e.  $\beta \rightarrow \infty$ ), the inefficiency ratio of the system in the long run is determined by IRSE. Whereas in contrast, in the best response itself ( $\beta = \infty$ ), the inefficiency ratio might be as bad as the price of anarchy.

### 3 Unweighted Congestion Games

Congestion games are one of the important classes of noncooperative games. In this section we study selfish routing in networks. We note that all proofs can be extended to congestion games. In a selfish routing game, selfish players control their own traffic and want to route it through a congested network. We assume that we have a directed graph  $G = (V, E)$ ,  $k$  source and sink pairs (or *commodities*)  $(s_1, t_1), \dots, (s_k, t_k)$  each of them corresponding to one selfish agent. We will be considering “unweighted” instances where each agent  $i$  is aiming to send one unit of traffic from  $s_i$  to  $t_i$  through a single path. There also is a nonnegative continuous nondecreasing cost function  $c_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$  for each edge  $e$  in the graph. The cost that agent  $i$  bears is the total sum of the costs of the edges through the path she has chosen. Therefore, the social cost will be  $C(f) = \sum_i \sum_{e \in P_i} c_e(f_e) = \sum_{e \in E} c_e(f_e) f_e$ . It is well-known that this game is a potential game with the potential function  $\Phi$  defined as  $\Phi(f) = \sum_{e \in E} \sum_{i=1}^{f_e} c_e(i)$ .

For brevity, in this section we only work with the cost functions of the form  $c_e(x) = x^p$ . For general polynomials of degree  $p$ , the same price of anarchy and congestion results hold. Only the lowerbound for  $\beta$  will be multiplied by a factor of  $\aleph$ , where  $\aleph$  is the biggest absolute value of the coefficients of cost functions.

First, we study the rate of convergence of this dynamics to the stable region.

**Theorem 1.** *Consider an unweighted atomic selfish routing game being played on a graph  $G = (V, E)$  with  $K$  vertices and  $M$  edges where  $n$  units of traffic are being routed between their corresponding sources and sinks. Also, suppose that the cost functions are polynomials of degree at most  $p$ . Then for any constant  $\epsilon > 0$  there exists a value  $\beta_0(M, K, p, \epsilon)$ , so that for any  $\beta \geq \beta_0(M, K, p, \epsilon)$ , the hitting time of  $\mathcal{A}_\epsilon$  is at most  $Kn^3$ . Moreover, the Markov chain will almost always be in  $\mathcal{A}_\epsilon$  after hitting it.*

*Proof.* we show that there exists a constant  $c = c(M, K, p, \epsilon)$  such that for

- [Case I]: For any  $n \geq c$ , the hitting time of  $\mathcal{A}_\epsilon$  is at most  $Kn^3$ .
- [Case II]: For  $n < c$ , the hitting time is bounded by  $O(1)$ .

The idea of the proof for the **CASE I** is to show that the value of potential function is defining a random walk with a negative drift and hence, it will eventually end up near its minimum value. The similar argument does not hold in **CASE II**. Instead, we will prove that in this case the mixing time of the Markov chain is small. We also show that the measure of  $\mathcal{A}_\epsilon$  is almost 1. This clearly bounds the hitting time.

**CASE I.** We first need to prove that if no player can make a huge gain by changing her path, then the value of potential function is close to its optimal value.

**Lemma 1.** *Let  $f$  be an arbitrary flow and  $g$  be any equilibrium flow. Suppose that for some given value  $X$ , no player can change her path in  $f$  and decrease her cost by more than  $X$ . Then,*

$$\Phi(f) \leq \Phi(g) \left( 1 + \frac{pM^{p+1}(X + K2^{p+1})^{\frac{1}{p+1}}}{n^{\frac{p}{p+1}}} \right).$$

*Proof.* [Lemma 1.] Let  $P_i$  and  $P'_i$  be the path selected by agent  $i$  in  $f$  and  $g$ , respectively. Since  $g$  is an equilibrium, we have

$$\sum_{e \in P'_i} g_e \leq \sum_{e \in P_i} (g_e + 1)^p \leq \sum_{e \in P_i} (g_e^p + 2^p n^{p-1}).$$

Similarly, by the assumption about  $f$  we get

$$\sum_{e \in P_i} f_e \leq Xn^{p-1} + \sum_{e \in P'_i} (f_e + 1)^p \leq Xn^{p-1} + \sum_{e \in P'_i} (f_e^p + 2^p n^{p-1}).$$

Adding the two inequalities above and taking the sum for all values of  $i$ , we conclude  $\sum_e (g_e^p - f_e^p)(g_e - f_e) \leq n^p(X + K2^{p+1})$ . It means that the difference of potential functions can be bounded from above by  $n^p(n^p(X + K2^{p+1}))^{1/(p+1)}$ . Note that the potential function is at least  $\frac{n^{p+1}}{pM^{p+1}}$ . Substituting this value will complete the proof.  $\square$

Define  $c = c(M, K, p, \epsilon) := \frac{2^{p+1}M^{2p+2}K}{\epsilon^{p+1}}$ . For  $n \geq c$ , we define a region  $\mathcal{A}_\delta$  where

$$\delta = \frac{pM^{p+1}(X + K2^{p+1})^{\frac{1}{p+1}}}{n^{\frac{p}{p+1}}},$$

and  $X = Kn^{p-1}$ . We prove that if we are outside of the set  $\mathcal{A}_\delta$  we have a tendency to go back there. More formally, we show that there will be a negative drift for the value of  $\Phi$  of the current state of the Markov chain. By Lemma 1 there always exists a move with payoff at least  $X$ . Hence, the expected change in the value of potential function will be

$$\frac{-Xe^{\beta X} + \sum_p \Delta\Phi(p)e^{-\beta\Delta\Phi(p)}}{e^{\beta X} + \sum_p e^{-\beta\Delta\Phi(p)}},$$

where the sums are over all possible paths  $p$  and  $\Delta\Phi(p)$  is the payoff of switching to path  $p$ . Note that the total number of possible paths is at most  $K!$ . One can check that for  $\beta = \Omega(K \log K)$ , the above fraction is under  $-X/2$ . It shows that the random walk on the value of potential function has a negative drift of size at least  $X/2n$ .

Note that the initial value of potential function can be at most  $Mn^{p+1} \leq K^2n^{p+1}$ . It means that for  $\beta = \Omega(K \log K)$ , after at most  $O(K^2n^{p+2}/X) = O(Kn^3)$  steps we get to  $\mathcal{A}_\delta$ . On the other hand, during these  $O(Kn^3)$  steps, the probability that the value of the potential function reaches a distance  $2X$  from all of the states in  $\mathcal{A}_\delta$  is exponentially small. But note that the the potential function cannot be less than  $\frac{n^{p+1}}{pM^p}$ . Hence in these  $O(Kn^3)$  steps, w.h.p we will not

leave  $\mathcal{A}_{2\delta}$ . And therefore,  $\mathcal{A}_{2\delta}$  is a stable region. The statement of the theorem in this case follows by realizing that  $\delta = O(\epsilon)$ .

**CASE II.** Note that the the payoff of each move for any player is bounded between  $-Kn^p$  and  $Kn^p$ . The total number of moves for each person is bounded by  $K!$ . Therefore, the non-negative entries of the transition matrix are bounded from below by  $e^{-\beta Kn^p} / (K! e^{\beta Kn^p}) = e^{-2\beta Kn^p} / K!$ .

Each state of the Markov chain represents a joint pure strategy profile of the players. The total number of such strategy profiles is at most  $K!^n$ . So, in the stationary distribution, the probability of each state is bounded from below by  $e^{-\beta \max \Phi} / (K!^n e^{-\beta \min \Phi}) \leq e^{-\beta Mn^{p+1}} / K!^n$ .

To bound the mixing time of the Markov chain, we define paths between its states as follows: we fix an ordering of players, and to go from each state to another players change their paths one by one according to the ordering. It is easy to see that the maximum congestion is  $K!^{n-1}$  and also the maximum length of these paths is  $n$ . Thus, the mixing time is bounded by  $ne^{\beta K^2 n^{p+1}} K!^{2n} \leq e^{2\beta K^2 n^{p+1}}$ , which is  $O(1)$  since  $n < c$ .

As the final step, we prove that the measure of  $\mathcal{A}_\epsilon$  in the stationary distribution is large. Let  $Y$  be the optimum value of the potential function. The measure of  $\mathcal{A}_\epsilon$  will be at least  $e^{-\beta Y} / ((K!)^n e^{-(1+\epsilon)\beta Y}) = e^{\beta \epsilon Y} K!^{-n}$ . Note that the value of  $Y$  cannot be less than  $\frac{n^{p+1}}{pM^p}$ . Therefore, values of  $\beta \geq \frac{KM^p \log K}{\epsilon n^p}$  ensures a measure of  $1 - o(1)$  for  $\mathcal{A}_\epsilon$ . □

Lemma 2 establishes a relation between the cost of the flows that approximately minimize the potential function  $\Phi$  and a flow that minimizes the cost  $C$ . The proof for the degree  $p > 2$  polynomial case is straightforward and is done by bounding the inequality  $\Phi(f) \leq (p + 1)C(f)$  for every routing  $f$ . The proof for the linear case is quite similar to the one in [6] and we omit it here.

**Lemma 2.** *In the unweighted atomic selfish routing game with linear costs let  $f$  be the equilibrium flow that minimizes the potential function  $\Phi$  and  $g$  be the optimal flow. We have  $C(f) \leq (1 + \frac{1}{\sqrt{3}})C(g)$ . Also when the costs are polynomials of degree  $p$ , then  $C(f) \leq (p + 1)C(g)$ . Also, the ratio between the social cost of the flows in the stable region  $\mathcal{A}_\epsilon$  and the optimum social cost will be respectively  $(1 + \epsilon)(1 + \frac{1}{\sqrt{3}})$  and  $(1 + \epsilon)(p + 1)$ .*

Note that the price of anarchy in these cases can be as large as 2.5 and  $p^{\Theta(p)}$ , respectively. Also, when the game is symmetric (single-source/single-sink), [1] implies that the inefficiency of the best equilibrium is at most  $\frac{4}{3}$  (which interestingly coincides with the price of anarchy in non-atomic selfish routing). On the other hand, the exact upperbound for the price of anarchy is not known in this case (it is of course at most 2.5). Through a series of unweighted symmetric routing instances whose price of anarchy gets arbitrary close to 2 from below, we can show that the price of anarchy in this case is at least 2. The example can be found in the longer version of the paper.

Note that Theorem 1 implies that the game will eventually settles down in the stable region of these potential games. Lemma 2 bounds the ratio of the

worst social cost in this region to the optimum social cost of the game. Hence, combining Theorem 1 and Lemma 2 enables us to bound the inefficiency ratio of the stable equilibria in such games, which is the ultimate result of this work.

**Theorem 2.** *For symmetric unweighted routing games with linear cost functions, IRSE is exactly  $\frac{4}{3}$ . Also, for congestion games, IRSE is exactly  $1 + \frac{1}{\sqrt{3}}$  (resp. at most  $p + 1$ ) when the cost functions are linear (resp. polynomials of degree at most  $p$ ).*

**Acknowledgements.** We would like to thank Éva Tardos, Tim Roughgarden, and Ramesh Johari for fruitful discussions and useful comments.

## References

1. 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, pp. 295–304 (2004)
2. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: STOC 2005, pp. 57–66 (2005)
3. Awerbuch, B., Azar, Y., Epstein, A., Mirrokni, V.S., Skopalik, A.: Fast convergence to nearly optimal solutions in potential games. In: EC 2008, pp. 264–273 (2008)
4. Blum, A., Hajiaghayi, M., Ligett, K., Roth, A.: Regret minimization and the price of total anarchy. In: STOC 2008, pp. 373–338 (2008)
5. Blume, L.: The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 387–426 (1993)
6. 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)
7. Chien, S., Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. In: SODA 2007, pp. 169–178 (2007)
8. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: STOC 2005, pp. 67–73 (2005)
9. 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)
10. Ellison, G.: Learning, Local Interaction, and Coordination. *Econometrica* 61(5), 1047–1071 (1993)
11. Finn, G., Horowitz, E.: A linear time approximation algorithm for multiprocessor scheduling. *BIT* 19(3), 312–320 (1979)
12. Fudenberg, D., Kreps, D.: Learning Mixed Equilibria. *Games and Economic Behavior* 5, 320–367 (1993)
13. Fudenberg, D., Levine, D.K.: *Theory of Learning in Games*. MIT Press, Cambridge (1998)
14. Kleinberg, R., Piliouras, G., Tardos, É.: Multiplicative updates outperform generic no-regret learning in congestion games. In: STOC 2009, pp. 533–542 (2009)
15. Montanari, A., Saberi, A.: Convergence to Equilibrium in Local Interaction Games and Ising Models. In: FOCS 2009 (to appear, 2009)



# Betting on the Real Line

Xi Gao<sup>1</sup>, Yiling Chen<sup>1,\*</sup>, and David M. Pennock<sup>2</sup>

<sup>1</sup> Harvard University

{xgao,yiling}@eecs.harvard.edu

<sup>2</sup> Yahoo! Research

pennockd@yahoo-inc.com

**Abstract.** We study the problem of designing prediction markets for random variables with continuous or countably infinite outcomes on the real line. Our interval betting languages allow traders to bet on any interval of their choice. Both the call market mechanism and two automated market maker mechanisms, *logarithmic market scoring rule* (LMSR) and *dynamic parimutuel markets* (DPM), are generalized to handle interval bets on continuous or countably infinite outcomes. We examine problems associated with operating these markets. We show that the auctioneer's order matching problem for interval bets can be solved in polynomial time for call markets. DPM can be generalized to deal with interval bets on both countably infinite and continuous outcomes and remains to have bounded loss. However, in a continuous-outcome DPM, a trader may incur loss even if the true outcome is within her betting interval. The LMSR market maker suffers from unbounded loss for both countably infinite and continuous outcomes.

**Keywords:** Prediction Markets, Combinatorial Prediction Markets, Expressive Betting.

## 1 Introduction

Prediction markets are speculative markets created for forecasting random variables. In practice, they have been shown to provide remarkably accurate probabilistic forecasts [1,2]. Existing prediction markets mainly focus on providing an aggregated probability mass function for a random variable with finite outcomes or discretized to have finite outcomes. For example, to predict the future printer sales level, the value of which lies on the positive real line, Hewlett-Packard's sales prediction markets partition the range of the sales level into about 10 exclusive intervals, each having an assigned Arrow-Debreu security that pays off \$1 if and only if the future sales level falls into the corresponding interval [3]. The price of each security represents the market probability that the sales level is within the corresponding interval. The set of prices provides a probability mass function for the discretized random variable.

However, many random variables of interest have continuous or countably infinite outcome spaces. For example, the carbon dioxide emission level in a certain

---

\* Part of this work was done while Yiling Chen was at Yahoo! Research.

period of time can be thought of as a continuous random variable on the positive real line; the printer sales level can be treated as a random variable with countably infinite outcomes, taking positive integer values. Discretizing such random variables into finite outcomes can potentially hurt information aggregation, as market participants may have information that can not be easily expressed with the ex-ante specified discretization. It is desirable to provide more expressive betting languages so that market participants can express their information more accurately and preferably in the same way they possess it.

In this paper, we design and study prediction market mechanisms for predicting random variables with continuous or countably infinite outcomes on the real line. We provide betting languages that allow market participants to bet on any interval of their choice and create the security on the fly. We generalize both the call market mechanism and two automated market maker mechanisms, *logarithmic market scoring rules* (LMSR) [4] and *dynamic parimutuel markets* (DPM) [5,6], to handle interval bets on continuous or countably infinite outcomes, and examine problems associated with operating these markets. We show that the auctioneer's order matching problem can be solved in polynomial time for call markets. DPM can be generalized to deal with interval bets on both countably infinite and continuous outcomes and remains to have bounded loss. However, in a continuous-outcome DPM, a trader may incur loss even if the interval she bets on includes the true outcome. The LMSR market maker suffers from unbounded loss for both countably infinite and continuous outcomes. Due to space constraints, the Appendix is omitted and available upon request.

**Related Work.** Our work is situated in the broad framework of designing combinatorial prediction market mechanisms that provide more expressiveness to market participants. Various betting languages for permutation combinatorics have been studied for call markets, including subset betting and pair betting [7], singleton betting [8], and proportional betting [9]. Fortnow et al. [10] analyzed betting on Boolean combinatorics in call markets. For LMSR market makers, Chen et al. [11] showed that computing the contract price is #P-hard for subset betting, pair betting, and Boolean formulas of two events. In a tournament setting, pricing in LMSR becomes tractable for some restricted Boolean betting languages [12]. Yoopick is a combinatorial prediction market implementation of LMSR that allows traders to bet on point spreads of their choice for sporting events [13]. It is implemented as a LMSR with a large number of finite outcomes. Agrawal et al. [14] proposed Quad-SCPM, which is a market maker mechanism that has the same worst-case loss as a quadratic scoring rule market maker. Quad-SCPM may be used for interval bets on countably infinite outcomes since its worst-case loss does not increase with the size of the outcome space.

## 2 Background

In this section, we briefly introduce three market mechanisms that have been used by prediction markets to predict random variables with finite outcomes.

### 2.1 Call Markets

A call market is an auctioneer mechanism, where the auctioneer (market institution) risklessly matches received orders. In a call market, participants submit buy or sell orders for individual contracts. All orders are assembled at one time in order to determine a market clearing price at which demand equals supply. Buy orders whose bid prices are higher than the clearing price and sell orders whose ask prices are lower than the clearing price are accepted. All transactions occur at the market clearing price. Most call markets are bilateral — matching buy and sell orders of the same contract. For multi-outcome events, call markets can be multilateral — allowing participants to submit orders on different contracts and performing global order matching [15,16,17,18].

### 2.2 Logarithmic Market Scoring Rules

Let  $v$  be a random variable with  $N$  mutually exclusive and exhaustive outcomes. A logarithmic market scoring rule (LMSR) [4,19] is an automated market maker that subsidizes trading to predict the likelihood of each outcome. An LMSR market maker offers  $n$  contracts, each corresponding to one outcome and paying \$1 if the outcome is realized [4,20]. Let  $q_i$  be the total quantity of contract  $i$  held by all traders combined, and let  $\mathbf{q}$  be the vector of all quantities held. The market maker utilizes a *cost function*  $C(\mathbf{q}) = b \log \sum_{j=1}^N e^{q_j/b}$  that records the total amount of money traders have spent. A trader that wants to buy any bundle of contracts such that the total number of outstanding shares changes from  $\mathbf{q}$  to  $\tilde{\mathbf{q}}$  must pay  $C(\tilde{\mathbf{q}}) - C(\mathbf{q})$  dollars. Negative quantities encode sell orders and negative “payments” encode sale proceeds earned by the trader. At any time, the instantaneous price of contract  $i$  is  $p_i(\mathbf{q}) = \frac{e^{q_i/b}}{\sum_{j=1}^N e^{q_j/b}}$ , representing the cost per share of purchasing an infinitesimal quantity. An LMSR is built upon the *logarithmic scoring rule*,  $s_i(\mathbf{r}) = b \log(r_i)$ . It is known that if the market maker starts the market with a uniform distribution its worst-case loss is bounded by  $b \log N$ .

### 2.3 Dynamic Parimutuel Markets

A dynamic parimutuel market (DPM) [5,6] is a dynamic-cost variant of a parimutuel market. From a trader’s perspective, DPM acts as a market maker in a similar way to LMSR. There are  $N$  securities offered in the market, each corresponding to an outcome of the random variable. The cost function of the market maker, which captures the total money wagered in the market, is  $C(\mathbf{q}) = \kappa \sqrt{\sum_{j=1}^N q_j^2}$ , while the instantaneous price for contract  $i$  is  $p_i(\mathbf{q}) = \frac{\kappa q_i}{\sqrt{\sum_{j=1}^N q_j^2}}$ , where  $\kappa$  is a free parameter. Unlike in LMSR, the contract payoff in DPM is not a fixed \$1. If outcome  $i$  happens and the quantity vector at the end of the market is  $\mathbf{q}^f$ , the payoff per share of the winning security is  $o_i = \frac{C(\mathbf{q}^f)}{q_i^f} = \frac{\kappa \sqrt{\sum_{j=1}^N (q_j^f)^2}}{q_i^f}$ .

A nice property of DPM is that if a trader wagers on the correct outcome, she is guaranteed to have non-negative profit, because  $p_i$  is always less than or equal to  $\kappa$  and  $o_i$  is always greater than or equal to  $\kappa$ . Because the price functions

are not well-defined when  $\mathbf{q} = \mathbf{0}$ , the market maker must begin with a non-zero quantity vector  $\mathbf{q}^0$ . Hence, the market maker's loss is bounded by  $C(\mathbf{q}^0)$ .

### 3 Call Markets for Interval Betting

For a random variable  $X$  that has continuous or countably infinite outcomes on the real line, we consider the betting language that allows traders to bet on any interval  $(l, u)$  of their choice on the real line and create a security for the interval on the fly. The security pays off \$1 per share when the betting interval contains the realized value of  $X$ . For countably infinite outcomes, the interval is interpreted as a set of outcomes that lie within the interval.

Suppose that the range of  $X$  is  $(L, U)$  where  $L \in \mathbb{R} \cup \{-\infty\}$  and  $U \in \mathbb{R} \cup \{+\infty\}$ . Traders submit buy orders. Each order  $i \in O$  is defined by  $(b_i, q_i, l_i, u_i)$ , where  $b_i$  denotes the bid price for a unit share of the security on interval  $(l_i, u_i)$ , and  $q_i$  denotes the number of shares of the security to purchase at price  $b_i$ . We note  $l_i \geq L$  and  $u_i \leq U$ . Given a set of orders  $O$  submitted to the auctioneer, the auctioneer needs to decide which orders can be risklessly accepted. We consider the auctioneer's problem of finding an optimal match to maximize its worst-case profit given a set of orders  $O$ .

We first define the state space  $S$  to be the partition of the range of  $X$  formed by orders  $O$ . For any order  $i \in O$ ,  $(l_i, u_i)$  defines 2 boundary points of the partition. Let  $A = (\cup_{i \in O} l_i) \cup \{L\}$  be the set of left ends of all intervals in  $O$  and the left end of the range of  $X$ , and  $B = (\cup_{i \in O} u_i) \cup \{U\}$  be the set of right ends of all intervals in  $O$  and the right end of the range of  $X$ . We rank all elements of  $A$  and  $B$  in order of increasing values, and denote the  $i$ -th element as  $e_i$ . Clearly,  $e_1 = L$  and  $e_{|A|+|B|} = U$ . We formally define the state space  $S$  as follows.

**Definition 1.** Let  $s_i \in S$  be the  $i$ -th element of the state space  $S$  for all  $1 \leq i \leq (|A| + |B| - 1)$ . If  $e_i = e_{i+1}$ , then  $s_i = \{e_i\}$ . Otherwise,  $s_i = (e_i, e_{i+1})$  if both  $e_i \in A$  and  $e_{i+1} \in A$ ;  $s_i = (e_i, e_{i+1})$  if  $e_i \in A$  and  $e_{i+1} \in B$ ;  $s_i = [e_i, e_{i+1}]$  if  $e_i \in B$  and  $e_{i+1} \in A$ ; and  $s_i = [e_i, e_{i+1})$  if  $e_i \in B$  and  $e_{i+1} \in B$ .

Because  $|S| = |A| + |B| - 1$ ,  $|A| \leq |O| + 1$ , and  $|B| \leq |O| + 1$ , we have  $|S| \leq 2|O| + 1$ .

With the definition of states given orders  $O$ , we formulate the auctioneer's optimal match problem as a linear program, analogous to the one used for permutation betting [7].

**Definition 2 (Optimal Match).** Given a set of order  $O$ , choose  $x_i \in [0, 1]$  such that the following linear program is optimized.

$$\begin{aligned} \max_{x_i, c} \quad & c & (1) \\ \text{s.t.} \quad & \sum_i (b_i - I_i(s))q_i x_i \geq c, \forall s \in S \\ & 0 \leq x_i \leq 1, \quad \forall i \in O \end{aligned}$$

$I_i(s)$  is the indicator variable for whether order  $i$  is winning in state  $s$ .  $I_i(s) = 1$  if the order gets a payoff of \$1 in  $s$  and  $I_i(s) = 0$  otherwise. The variable  $c$  represents the worst-case profit for the auctioneer, and  $x_i \in [0, 1]$  represents the fraction of order  $i \in O$  that is accepted. As the number of structural constraints

is at most  $2|O|+1$  and the number of variables is  $|O|$ , (III) can be solved efficiently. We state it in the following theorem.

**Theorem 3.** *For call markets, the auctioneer’s optimal order matching problem for interval betting on countably infinite and continuous outcomes can be solved in polynomial time.*

Thus, if the optimal solution to (II) generates positive worst-case profit  $c$ , the auctioneer accepts orders according to the solution. Otherwise, when  $c \leq 0$ , the auctioneer rejects all orders.

When there are few traders in the market, finding a counterpart to trade in a call market may be hard and the market may suffer from the *thin market* problem. Allowing traders to bet on different intervals further exacerbates the problem by dividing traders’ attention among a large number of subsets of securities, making the likelihood of finding a multi-lateral match even more remote. In addition, call markets are zero-sum games and hence are challenged by the *no-trade theorem* [21]. In the next two sections, we examine market maker mechanisms, which not only provide infinite liquidity but also subsidize trading, for interval betting.

## 4 Dynamic Parimutuel Markets for Interval Betting

For interval betting in DPMs, traders also create a security on the fly by choosing an interval  $(l, u)$ . However, the payoff of the security is not fixed to be \$1. Instead, each share of the security whose interval contains the realized value of the random variable entitles its holder to an equal share of the total money in the market. We generalize DPM to allow for (but not limited to) interval betting on countably infinite and continuous outcomes. The problem that we consider is whether these mechanisms still ensure the bounded loss of the market maker.

### 4.1 Infinite-Outcome DPM

We generalize DPM to allow for countably infinite outcomes, and call the resulting mechanism *infinite-outcome DPM*. In an infinite-outcome DPM, the underlying forecast variable can have countably infinite mutually exclusive and exhaustive outcomes. Each state security corresponds to one potential outcome. An interval bet often includes a set of state securities. The market maker keeps track of the quantity vector of outstanding state securities, still denoted as  $\mathbf{q}$ , which is a vector of dimension  $\infty$ . The cost and price functions for the infinite-outcome DPM are  $C^I(\mathbf{q}) = \kappa\sqrt{\sum_{j=1}^{\infty} q_j^2}$ , and  $p_i^I(\mathbf{q}) = \frac{\kappa q_i}{\sqrt{\sum_{j=1}^{\infty} q_j^2}}$ . The payoff per winning security if outcome  $i$  happens is  $o_i^I = \frac{\kappa\sqrt{\sum_{j=1}^{\infty} (q_j^I)^2}}{q_i^I}$ .

The loss of the market maker in an infinite-outcome DPM is still her cost to initiate the market. The market maker needs to choose an initial quantity vector  $\mathbf{q}^0$  such that her loss  $C^I(\mathbf{q}^0)$  is finite. In practice, an infinite-outcome DPM market maker can start with a quantity vector that has only finite positive elements and all others are zeros, or use an infinite converging series. Whenever

a trader purchases a state security whose current price is zero or that has not been purchased before, the market maker begins to track quantity and calculate price for that security. Hence, infinite-outcome DPM can be operated as a finite-outcome DPM that can add new state securities as needed. The market maker does not need to record quantities and calculate prices for all infinite outcomes, but only for those having outstanding shares. Infinite-outcome DPM maintains the desirable price-payoff relationship of DPM — the payoff of a security is always greater than or equal to  $\kappa$  and its price is always less than or equal to  $\kappa$ .

### 4.2 Continuous-Outcome DPM

We then generalize DPM to allow for continuous outcomes, and call the resulting mechanism *continuous-outcome DPM*. The cost and price functions of a continuous-outcome DPM are  $C = \kappa \sqrt{\int_{-\infty}^{+\infty} q(y)^2 dy}$  and  $p(x) = \frac{\kappa q(x)}{\sqrt{\int_{-\infty}^{+\infty} q(y)^2 dy}}$ .

A trader can buy  $\delta$  shares of an interval  $(l, u)$ . The market maker then increases  $q(x)$  by  $\delta$  for all  $x \in (l, u)$ . The trader’s payment equals the change in value of the cost function. However, strictly speaking, function  $p(x)$  does not represent price, but is better interpreted as a density function. The instantaneous price for buying infinitely small amounts of the security for interval  $(l, u)$  is  $p_{(l,u)} = \int_l^u p(x)dx = \frac{\kappa \int_l^u q(x)dx}{\sqrt{\int_{-\infty}^{+\infty} q(y)^2 dy}}$ . If the realized value of the random variable is  $x^*$ , each share of a security on any interval that contains  $x^*$  has payoff  $o(x^*) = \frac{C}{q^f(x^*)} = \frac{\kappa \sqrt{\int_{-\infty}^{+\infty} q^f(y)^2 dy}}{q^f(x^*)}$ , where  $q^f(y)$  is the number of outstanding shares for securities whose interval contains  $y$  at the close of the market.

A continuous-outcome DPM market maker can choose an initial quantity distribution  $q^0(x)$  such that her loss is finite. However, the desirable price-payoff relationship that holds for the original DPM no longer holds for continuous-outcome DPM. A trader who bets on the correct outcome may still lose money. Theorem 4 states the price-payoff relationship for continuous-outcome DPM. Proof of the theorem is provided in the Appendix of the full paper.

**Theorem 4.** *The price per share for buying a security on interval  $(l, u)$  is always less than or equal to  $\kappa\sqrt{u-l}$ . If traders can bet on any non-empty open interval, the payoff per share is bounded below by 0. If traders could bet only on open intervals of size at least  $z$ , the payoff per share is bounded below by  $\frac{\kappa\sqrt{2z}}{2}$ .*

## 5 Logarithmic Market Scoring Rule for Interval Betting

For LMSR, we define the same interval betting language as in call markets. A trader can create a security by specifying an interval  $(l, u)$  to bet on. If the realized value of  $X$  falls into the interval, the security pays off \$1 per share. We generalize LMSR to allow countably infinite and continuous outcomes and study whether the market maker still has bounded loss.

LMSR for finite outcomes can be extended to accommodate interval betting on countably infinite outcomes simply by changing the summations in the price and cost functions to include all countably infinite outcomes. However, as the

LMSR market maker’s worst-case loss is  $b \log N$ , the market maker’s worst-case loss is unbounded as  $N$  approaches  $\infty$ .

We generalize LMSR to accommodate continuous outcome spaces. A logarithmic scoring rule for a continuous random variable is  $s(r(x)) = b \log(r(x))$  where  $x$  is the realized value for the random variable and  $r(x)$  is the reported probability density function for the random variable evaluated at  $x$ . Using an equation system similar to the one proposed by Chen and Pennock [20], we derive the corresponding price and cost functions for the continuous logarithmic scoring rule:  $C = b \log(\int_{-\infty}^{+\infty} e^{q(y)/b} dy)$ , and  $p(x) = \frac{e^{q(x)/b}}{\int_{-\infty}^{+\infty} e^{q(y)/b} dy}$ . Here,  $p(x)$  does not represent price, but is best interpreted as a density function. The instantaneous price for buying infinitely small amounts of the security for interval  $(l, u)$  is  $\int_l^u p(x) dx$ . If the interval  $(l, u)$  contains the realized value, one share of the security entitles its holder \$1 payoff.

However, by Theorem 5 the worst-case loss is still unbounded for a continuous LMSR market maker even with the restriction on the size of allowable intervals. Proof of Theorem 5 is presented in the Appendix of the full paper.

**Theorem 5.** *A continuous logarithmic market scoring rule market maker has unbounded worst-case loss, with or without the restriction that traders can bet only on intervals of size at least  $z$ .*

## 6 Conclusion and Future Directions

We study interval betting on random variables with continuous or countably infinite outcomes for call markets, DPM, and LMSR. We show that the auctioneer’s order matching problem in call markets can be solved in polynomial time for interval bets. DPM can be generalized to handle both countably infinite and continuous outcomes. Unfortunately, in a continuous-outcome DPM, a trader may incur loss even if her betting interval contains the true outcome. LMSR market makers, however, suffer from unbounded loss for both countably infinite and continuous outcomes.

One important future direction is to design automated market maker mechanisms with desirable properties, especially bounded loss, when handling continuous outcome spaces. In particular, it may be fruitful to explore interval bets with variable payoffs for outcomes within the interval. The interval contracts for call markets and LMSR give the same payoff as long as the outcome falls within the specified interval. Implicitly, this assumes that a trader’s prediction of the random variable is a uniform distribution over the given interval. Alternatively, it would be interesting to allow for the trader’s probability distribution of the random variable to take other shapes over the given interval, and hence to allow payoffs to vary correspondingly for outcomes within the interval.

## References

1. Berg, J.E., Forsythe, R., Nelson, F.D., Rietz, T.A.: Results from a dozen years of election futures markets research. In: Plott, C.A., Smith, V. (eds.) Handbook of Experimental Economic Results (2001)



2. Wolfers, J., Zitzewitz, E.: Prediction markets. *Journal of Economic Perspective* 18(2), 107–126 (2004)
3. Chen, K.Y., Plott, C.R.: Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working paper No. 1131, California Institute of Technology (2002)
4. Hanson, R.D.: Combinatorial information market design. *Information Systems Frontiers* 5(1), 107–119 (2003)
5. Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: *ACM Conference on Electronic Commerce, EC* (2004)
6. Mangold, B., Dooley, M., Dornfest, R., Flake, G.W., Hoffman, H., Kasturi, T., Pennock, D.M.: The tech buzz game. *IEEE Computer* 38(7), 94–97 (2005)
7. Chen, Y., Fortnow, L., Nikolova, E., Pennock, D.M.: Betting on permutations. In: *ACM Conference on Electronic Commerce (EC)*, pp. 326–335 (2007)
8. Ghodsi, M., Mahini, H., Mirrokni, V.S., ZadiMoghaddam, M.: Permutation betting markets: singleton betting with extra information. In: *ACM Conference on Electronic Commerce (EC)*, pp. 180–189 (2008)
9. Agrawal, S., Wang, Z., Ye, Y.: Parimutuel betting on permutations. In: Papadimitriou, C., Zhang, S. (eds.) *WINE 2008*. LNCS, vol. 5385, pp. 126–137. Springer, Heidelberg (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. *Decision Support Systems* 39(1), 87–104 (2004)
11. Chen, Y., Fortnow, L., Lambert, N., Pennock, D.M., Wortman, J.: Complexity of combinatorial market makers. In: *ACM Conference on Electronic Commerce (EC)*, pp. 190–199 (2008)
12. Chen, Y., Goel, S., Pennock, D.M.: Pricing combinatorial markets for tournaments. In: *ACM Symposium on Theory of Computing (STOC)*, pp. 305–314 (2008)
13. Goel, S., Pennock, D., Reeves, D.M., Yu, C.: Yoopick: A combinatorial sports prediction market. In: *AAAI*, pp. 1880–1881 (2008)
14. Agrawal, S., Delage, E., Peters, M., Wang, Z., Ye, Y.: A unified framework for dynamic pari-mutuel information market design. In: *ACM Conference on Electronic Commerce (EC)*, pp. 255–264 (2009)
15. Bossaerts, P., Fine, L., Ledyard, J.: Inducing liquidity in thin financial markets through combined-value trading mechanisms. *European Economic Review* (46), 1671–1695 (2002)
16. Lange, J., Economides, N.: A parimutuel market microstructure for contingent claims trading. NYU School of Business Discussion Paper No. EC-01-13 (2007)
17. Baron, K., Lange, J.: *Parimutuel Applications in Finance: New Markets for New Risks*. Palgrave Macmillan, Basingstoke (2005)
18. 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)
19. Hanson, R.D.: Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* 1(1), 1–15 (2007)
20. Chen, Y., Pennock, D.M.: A utility framework for bounded-loss market makers. In: *Conference on Uncertainty in Artificial Intelligence (UAI 2007)*, pp. 49–56 (2007)
21. Milgrom, P., Stokey, N.L.: Information, trade and common knowledge. *Journal of Economic Theory* 26(1), 17–27 (1982)



# Characterization of Strategy-Proof, Revenue Monotone Combinatorial Auction Mechanisms and Connection with False-Name-Proofness<sup>\*</sup>

Taiki Todo, Atsushi Iwasaki, and Makoto Yokoo

Graduate School of ISEE, Kyushu University, Fukuoka 819-0395, Japan  
todo@agent.is.kyushu-u.ac.jp, {iwasaki,yokoo}@is.kyushu-u.ac.jp

**Abstract.** A combinatorial auction mechanism consists of an allocation rule and a payment rule. There have been several studies on characterizing *strategy-proof* allocation rules. In particular, conditions called *weak-monotonicity* has been identified as a full characterization of them. On the other hand, *revenue monotonicity* is recognized as one of the desirable properties. A combinatorial auction mechanism is revenue monotone if a seller's revenue is guaranteed to weakly increase as the number of bidders grows. Though the property is quite reasonable, there exists virtually no work on the characterization. In this paper, we identified a simple condition called *summation-monotonicity*. We then proved that we can construct a strategy-proof, revenue monotone mechanism if and only if the allocation rule satisfies weak-monotonicity and summation-monotonicity. To the best of our knowledge, this is the first attempt to characterize revenue monotone allocation rules. In addition, we shed light on a connection between revenue monotonicity and false-name-proofness. In fact, we proved that, assuming a natural condition, revenue monotonicity is equivalent to false-name-proofness for single-item auctions.

## 1 Introduction

Mechanism design of combinatorial auctions has become an integral part of Electronic Commerce and a promising field for applying AI and agent technologies. Among various studies related to Internet auctions, those on combinatorial auctions have lately attracted considerable attention. Mechanism design is the study of designing a rule/protocol that achieves several desirable properties assuming that each agent/bidder hopes to maximize his own utility. One desirable property of a combinatorial auction mechanism is *strategy-proofness*. A mechanism is strategy-proof if, for each bidder, reporting his true valuation is a *dominant strategy*, i.e., an optimal strategy regardless of the actions of other bidders.

A combinatorial auction mechanism consists of an allocation rule that defines the allocation of goods for each bidder, and a payment rule that defines the payment of each winner. There have been many studies on characterizing strategy-proof social choice function (an allocation rule in combinatorial auctions) in

---

<sup>\*</sup> The full version of this paper [\[1\]](#) is provided in the author's web-site.

the literature of social choice theory. In particular, a family of *monotonicity* concepts was identified to characterize implementable social choice functions. For example, Lavi *et al.* proposed *weak-monotonicity* and showed that it is a necessary and sufficient condition for strategy-proof mechanisms when several assumptions hold on the domain of types [2]. Such a characterization of allocation rules is quite useful for developing/verifying strategy-proof mechanisms. These conditions are defined only on an allocation rule; i.e., if it satisfies such a condition, it is guaranteed that there exists an appropriate payment rule that achieves strategy-proofness. Thus, a mechanism designer can concentrate on the allocation rule when developing a new mechanism or verifying an existing one.

As a line of these studies, *false-name-proofness* is recognized as one of the desirable properties of combinatorial auctions. It is a kind of generalization of strategy-proofness in an environment where a bidder can use multiple identifiers, e.g., multiple e-mail addresses [3]. Todo *et al.* characterized false-name-proof allocation rules by a condition called *sub-additivity* [4]. Rastegari *et al.* mentioned a connection between false-name-proofness and revenue monotonicity [5], which is known as another desirable property of combinatorial auctions.

A mechanism is revenue monotone if a seller's revenue from an auction is guaranteed to weakly increase as the number of bidders grows. The property is quite reasonable, since growing number of bidders increases competition. However, though it is shown that even the theoretically well-founded Vickrey-Clarke-Groves (VCG) mechanism does not always achieve this property. Furthermore, there exists virtually no work on characterizing revenue monotone combinatorial auction mechanisms. One notable exception is the work by Rastegari *et al.* [5], who introduced a notion called *weak-maximality*, which is a weaker notion of Pareto efficiency, and proved that there exists no mechanism that achieves revenue monotonicity, strategy-proofness, and weak-maximality.

To the best of our knowledge, our paper is the first attempt to characterize revenue monotone allocation rules in combinatorial auctions. We first identify a condition called *summation-monotonicity* and prove that we can find an appropriate payment rule if and only if the allocation rule satisfies weak-monotonicity and summation-monotonicity. We have actually verified existing combinatorial auctions and found that several non-trivial mechanisms are not revenue monotone. In addition, we shed light on a connection between revenue monotonicity and false-name-proofness. We prove that false-name-proofness is equivalent to revenue monotonicity for single-item auctions, assuming a natural condition.

## 2 Preliminaries

Assume there exists a set of potential bidders  $\mathbb{N} = \{1, 2, \dots, n\}$  and a set of goods  $G = \{g_1, g_2, \dots, g_m\}$ . Let us define  $N \subseteq \mathbb{N}$  as the set of bidders participating in an auction. Each bidder  $i \in N$  has his preferences over each bundle or goods  $B \subseteq G$ . Formally, we model this by supposing that bidder  $i$  privately observes a parameter, or signal,  $\theta_i$  that determines his preferences. We refer to  $\theta_i$  as the *type* of bidder  $i$  and assume it is drawn from a set  $\Theta_i$ .

Let us denote the set of all possible type profiles as  $\Theta_N = \Theta_1 \times \dots \times \Theta_n$  and a type profile as  $\theta = (\theta_1, \dots, \theta_n) \in \Theta_N$ . Observe that type profiles always have one entry for every potential bidder, regardless of the set of participating bidders  $N$ . We use the symbol  $\emptyset$  in the vector  $\theta$  as a placeholder for each non-participating bidder  $i \notin N$ . Then, we represent  $\theta = (\theta_1, \dots, \theta_{i-1}, \emptyset, \theta_{i+1}, \dots, \theta_n)$  as  $\theta_{-i}$  if  $i \notin N$ . When a set of bidders  $N$  is participating in the auction, let us denote the set of possible type profiles reported by  $N$  as  $\Theta_N(\subseteq \Theta_N)$ . That is,  $\Theta_N$  is the set of all type profiles  $\theta$  for which  $\theta_i = \emptyset$  if and only if  $i \notin N$ .

We assume a valuation  $v$  is normalized by  $v(\theta_i, \emptyset) = 0$ , satisfies *free disposal*, i.e.,  $v(\theta_i, B') \geq v(\theta_i, B)$  for all  $B' \supseteq B$ , and satisfies *no externalities*, i.e., a valuation  $v$  is determined only by his obtained bundle. We call each  $\Theta_i$  that satisfies these conditions a *order based domain* in combinatorial auctions [2]. In other words, the domain of types  $\Theta_i$  is rich enough to contain all possible valuations. We require this assumption so that weak-monotonicity characterizes strategy-proofness.

A combinatorial auction mechanism  $\mathcal{M}$  consists of an allocation rule  $X$  and a payment rule  $p$ . When a set of bidders  $N$  participates, an allocation rule is defined as  $X : \Theta_N \rightarrow A_N$ , where  $A_N$  is a set of possible outcomes. Similarly, a payment rule is defined as  $p : \Theta_N \rightarrow \mathbb{R}_+^N$ . Let  $X_i$  and  $p_i$  respectively denote the bundle allocated to bidder  $i$  and the amount bidder  $i$  must pay.

For simplicity, we restrict our attention to a *deterministic* mechanism and assume that a mechanism is *almost anonymous*, i.e., obtained results are invariant under permutation of identifiers except for the cases of ties. We also assume that a mechanism satisfies consumer sovereignty [5], i.e., there always exists a type  $\theta_i$  for bidder  $i$ , where bidder  $i$  can obtain bundle  $B$  regardless of reported types  $\theta_{-i}$  by other agents except  $i$ . In other words, if bidder  $i$ 's valuation for  $B$  is high enough, then  $i$  can obtain  $B$ . Furthermore, we restrict our attention to *individually rational* mechanisms. A mechanism is individually rational if  $\forall N \subseteq \mathbb{N}, \forall i \in N, \forall \theta_i, \forall \theta_{-i}, v(\theta_i, X_i(\theta)) - p_i(\theta) \geq 0$  holds. This means that no participant suffers any loss in a dominant strategy equilibrium, i.e., the payment never exceeds the valuation of the allocated goods.

Let us introduce two desirable properties for combinatorial auction mechanisms: *strategy-proofness* and *revenue monotonicity*. First, we introduce the notion of *strategy-proofness*.

**Definition 1 (Strategy-proofness).** A combinatorial auction mechanism  $\mathcal{M}(X, p)$  is strategy-proof, if  $\forall N \subseteq \mathbb{N}, \forall i \in N, \forall \theta_{-i}, \forall \theta_i, \theta'_i, v(\theta_i, X_i(\theta)) - p_i(\theta) \geq v(\theta_i, X_i(\theta'_i, \theta_{-i})) - p_i(\theta'_i, \theta_{-i})$  holds.

In other words, a mechanism is strategy-proof if reporting true type maximizes his utility regardless of the other bidders' reports. Then, we introduce a condition called *weak-monotonicity*, which fully characterizes allocation rules in strategy-proof combinatorial auction mechanisms.

**Definition 2 (Weak-monotonicity [2]).** An allocation rule  $X$  satisfies weak-monotonicity if  $\forall N \subseteq \mathbb{N}, \forall i \in N, \forall \theta_{-i}, \forall \theta_i, \theta'_i, v(\theta_i, X_i(\theta_i, \theta_{-i})) - v(\theta'_i, X_i(\theta_i, \theta_{-i})) \geq v(\theta_i, X_i(\theta'_i, \theta_{-i})) - v(\theta'_i, X_i(\theta'_i, \theta_{-i}))$  holds.

**Theorem 1 ([2]).** *There exists an appropriate payment rule  $p$  such that a combinatorial auction mechanism  $\mathcal{M}(X, p)$  is strategy-proof, if and only if  $X$  satisfies weak-monotonicity.*

This theorem indicates that if an allocation rule is weakly monotone, we can always find an appropriate payment rule to truthfully implement the allocation rule. Thus, when designing an auction mechanism, we can concentrate on designing an allocation rule and forget about a payment rule, at least for a while.

Next, we introduce *revenue monotonicity*, a well-known desirable property in combinatorial auctions.

**Definition 3 (Revenue monotonicity [6]).** *A combinatorial auction mechanism  $\mathcal{M}(X, p)$  is revenue monotone if  $\forall N \subseteq \mathbb{N}, \forall \theta, \forall j \in \mathbb{N}$ , the following inequality holds:*

$$\sum_{i \in \mathbb{N}} p_i(\theta) \geq \sum_{i \in \mathbb{N} \setminus \{j\}} p_i(\theta_{-j}) \tag{1}$$

The left side of Eq. (1) indicates the seller’s revenue from the auction when the set of bidders  $N$  participates in the auction. The right side indicates the seller’s revenue when bidder  $j$  drops out. In other words, a combinatorial auction is revenue monotone if the seller’s revenue does not increase by dropping bidders.

### 3 Characterizing Revenue Monotone Allocation Rules

This section introduces a simple condition called summation-monotonicity that fully characterizes revenue monotone allocation rules.

**Definition 4 (Summation-monotonicity).** *An allocation rule  $X$  satisfies summation-monotonicity if  $\forall N \subseteq \mathbb{N}, \forall \theta, \forall j \in \mathbb{N}$ , the following inequality holds:*

$$\begin{aligned} \forall i \in \mathbb{N}, \forall \theta'_i \quad & \text{s.t. } X_i(\theta'_i, \theta_{-i}) \supseteq X_i(\theta), v(\theta'_i, X_i(\theta'_i, \theta_{-i})) = v(\theta'_i, X_i(\theta)), \\ \forall i \in \mathbb{N} \setminus \{j\}, \forall \theta''_i \quad & \text{s.t. } v(\theta''_i, X_i(\theta''_i, \theta_{-\{i,j\}})) = 0, \end{aligned}$$

$$\sum_{i \in \mathbb{N}} v(\theta'_i, X_i(\theta)) \geq \sum_{i \in \mathbb{N} \setminus \{j\}} v(\theta''_i, X_i(\theta_{-j})) \tag{2}$$

Note here that  $\theta_{-\{i,j\}}$  denotes a type profile excluding bidder  $i$  and  $j$ .

An intuitive explanation why summation-monotonicity holds for a strategy-proof and revenue monotone mechanism is as follows. Let us consider a combinatorial auction mechanism with two goods  $g_1$  and  $g_2$ . Assume that it allocates  $g_1$  to bidder 1 and  $g_2$  to bidder 2 when the set of bidders  $N$  participates. On the other hand, also assume that it allocates  $g_1$  to bidder 3 and  $g_2$  to bidder 4 when bidder  $j$  drops out from the auction. The two squares on the top of Figure (1) represents the total payments for bidder 1 and 2 and those on the bottom for bidder 3 and 4. The figure means that if the mechanism is revenue monotone,  $p_1(\theta) + p_2(\theta) \geq p_3(\theta_{-j}) + p_4(\theta_{-j})$  holds for the payment rule  $p$ .

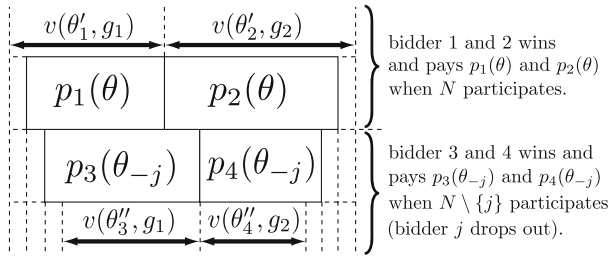


Fig. 1. Summation-monotonicity

The arrows on the top of Figure 1 indicates the left side of Eq. 2. Here,  $\theta'_1$  means the minimum type that bidder 1 obtains  $g_1$  or any superset, fixing other bidders' types than bidder 1. Under the mechanism,  $v(\theta'_1, g_1)$  must be greater than  $p_1(\theta)$ . Otherwise, 1 has an incentive not to participate in the auction and individual rationality is violated. Similarly,  $\theta'_2$  means the minimum type that bidder 2 obtains  $g_2$  or any superset, fixing other bidders' types than bidder 2 and  $v(\theta'_2, g_2)$  must be greater than  $p_2(\theta)$ .

The arrows on the bottom indicates the right side of Eq. 2. Here,  $\theta''_3$  means the maximum type that bidder 3 cannot obtain  $g_1$ , i.e., he obtains nothing, fixing other bidders' types than bidder 3. Under the mechanism,  $v(\theta''_3, g_1)$  must be smaller than  $p_3(\theta)$ . Otherwise, a bidder with  $\theta''_3$  as the true type has an incentive to pretend that his type is  $\theta_3$  and to obtain  $g_1$ . Similarly,  $\theta''_4$  means the minimum type that bidder 4 cannot obtain  $g_2$ , fixing other bidders' types than bidder 4 and  $v(\theta''_4, g_2)$  must be smaller than  $p_4(\theta_{-j})$ .

For these facts, summation-monotonicity must hold for strategy-proof, revenue monotone mechanisms. Furthermore, as long as summation-monotonicity and weak-monotonicity hold, we can find an appropriate payment rule  $p$  so that  $p_1(\theta) + p_2(\theta) \geq p_3(\theta_{-j}) + p_4(\theta_{-j})$  holds. Then, we derive the following theorem. For space reasons, we omit the proof of the theorem.

**Theorem 2.** *There exists an appropriate payment rule  $p$  such that a combinatorial auction mechanism  $\mathcal{M}(X, p)$  is strategy-proof and revenue monotone, if and only if  $X$  satisfies weak-monotonicity and summation-monotonicity.*

This theorem shows that, if an allocation rule satisfies weak-monotonicity and summation-monotonicity, we can always find a payment rule so that the obtained mechanism is strategy-proof and revenue monotone. If it does not satisfy weak-monotonicity or summation-monotonicity, it is impossible to find such a payment rule. Notice that the impossibility result for strategy-proof, revenue monotone mechanisms in [5] states that there exists no strategy-proof, revenue monotone mechanism that always achieves weak-maximality.

Thus, our proposed summation-monotonicity condition would enable us to design such a mechanism, that does not always achieve weak-maximality. The condition also enables us to verify whether a mechanism is revenue monotone. We demonstrate whether summation-monotonicity is satisfied in two allocation rules in the following claims.

**Table 1.** A Pareto efficient allocation rule is not summation-monotone

	$g_1$	$g_2$	$\{g_1, g_2\}$
bidder 1	7	0	7
bidder 2	0	0	8
bidder 3	0	7	7

*Claim.* A Pareto efficient allocation rule does not satisfy summation-monotonicity.

Assume there are three bidders 1, 2, and 3 and two goods  $g_1$  and  $g_2$  for sale. First, let us consider the situation where their reported types are given in Table 1. For an allocation rule that achieves Pareto efficiency,  $g_1$  is allocated to bidder 1 and  $g_2$  is allocated to bidder 3. Let bidder 1 have  $\theta'_1$  such that  $v(\theta'_1, g_1) = 1 + \epsilon$ . He is allocated  $g_1$ . Similarly, let bidder 3 have  $\theta'_3$  such that  $v(\theta'_3, g_2) = 1 + \epsilon$ . He is allocated  $g_2$ . Thus, by a Pareto efficient allocation rule, we obtain  $2(1 + \epsilon)$ , as the left side of Eq. 2 illustrates.

On the other hand, let us consider the situation where bidder 3 drops out from the auction. Since the allocation rule is Pareto efficient, bidder 2 obtains  $\{g_1, g_2\}$  if he has a greater value than 7 on  $\{g_1, g_2\}$ . Thus, if he has a type  $\theta''_2$  such that  $v(\theta''_2, \{g_1, g_2\}) = 7 - \epsilon$  and  $v(\theta''_2, g_1) = v(\theta''_2, g_2) = 0$ , he obtains no good, i.e.,  $X(\theta''_2) = \emptyset$ . On the other hand, for bidder 1, since he obtains no good under the Pareto efficient allocation rule,  $v(\theta''_1, \emptyset) = 0$ . Therefore, the right side of Eq. 2, that is, the maximum bid in which bidder 2 loses, is  $7 - \epsilon$ .

Thus, we obtain  $2(1 + \epsilon) < 7 - \epsilon$ , and the Pareto efficient allocation rule does not satisfy summation-monotonicity. In fact, by bidder 3’s dropping out, the seller’s revenue increases from 2 to 7, and revenue monotonicity fails.

*Claim.* The allocation rule in the Set mechanism satisfies summation-monotonicity.

The Set mechanism is one of the simplest mechanisms. It allocates all goods to a single bidder, namely, the bidder with the largest valuation for the grand bundle of all goods. Effectively, it sells the grand bundle as a single good, in a Vickrey/second-price auction. The allocation rule in the Set mechanism is described as follows:

$$X_i(\theta_i, \theta_{-i}) = \begin{cases} G & \text{if } v(\theta_i, G) \geq \max_{l \in N \setminus \{i\}} v(\theta_l, G) \\ \emptyset & \text{otherwise.} \end{cases}$$

Assume that bidder  $i$  wins when a set of bidders  $N$  participates. The left side  $v(\theta'_i, G)$  of Eq. 2, that is, the minimum bid in which bidder  $i$  still wins, satisfies  $v(\theta'_i, G) \geq \max_{l \in N \setminus \{i\}} v(\theta_l, G)$ .

On the other hand, let us consider the situation where bidder  $j$  drops out and assume that bidder  $k(\in N \setminus \{j\})$  wins. Then, the right side of Eq. 2 is the maximum bid in which bidder  $k$  loses. First, let us consider the case where  $j \neq i$  holds. In this case, winner doesn’t change by  $j$ ’s dropping out. That is,  $k = i$  holds. Then, the right side  $v(\theta''_i, G)$  of Eq. 2 satisfies  $v(\theta''_i, G) \leq \max_{l \in N \setminus \{i, j\}} v(\theta_l, G)$

and clearly  $\max_{l \in N \setminus \{i\}} v(\theta_l, G) \geq \max_{l \in N \setminus \{i, j\}} v(\theta_l, G)$  holds. Therefore, we obtain  $v(\theta'_i, G) \geq v(\theta''_i, G)$  and Eq. 2 holds.

Next, let us consider the case where  $j = i$  holds. Here, the winner changes by  $j$ 's dropping out. Then, the right side  $v(\theta''_k, G)$  of Eq. 2 satisfies  $v(\theta''_k, G) \leq \max_{l \in N \setminus \{i, k\}} v(\theta_l, G)$  and clearly  $\max_{l \in N \setminus \{i\}} v(\theta_l, G) \geq \max_{l \in N \setminus \{i, k\}} v(\theta_l, G)$  holds. Therefore, we obtain  $v(\theta'_i, G) \geq v(\theta''_k, G)$ , and Eq. 2 holds. Thus, this allocation rule satisfies summation-monotonicity.

## 4 Revenue Monotonicity and False-Name-Proofness

This section considers a connection between revenue monotonicity and false-name-proofness mentioned in 5. False-name-proofness is a kind of generalization of strategy-proofness in an environment where a bidder can use multiple identifiers, e.g., multiple e-mail addresses 3. A mechanism is *false-name-proof* if for each bidder, reporting his true valuations using a single identifier (although the bidder can use multiple identifiers) is a dominant strategy. It has been shown that the VCG is not false-name-proof and there exists no false-name-proof, Pareto efficient mechanism 3. *Todo et al.* characterized false-name-proof allocation rules by a condition called *sub-additivity*. They proved that we can find an appropriate payment rule if and only if the allocation rule simultaneously satisfies weak-monotonicity and sub-additivity 4.

**Theorem 3 (4).** *There exists an appropriate payment rule  $p$  such that a combinatorial auction mechanism  $\mathcal{M}(X, p)$  is false-name-proof, if and only if  $X$  simultaneously satisfies weak-monotonicity and sub-additivity.*

The sub-additivity condition is quite similar to our proposed summation-monotonicity condition. Indeed, it has been considered that there is some connection between revenue monotonicity and false-name-proofness 6,5. For example, Table 1, which provides an example where VCG does not achieve revenue monotonicity, also provides an example that it does not achieve false-name-proofness. Let us consider a situation where bidder  $1'$ , who has a type  $\theta_{1'} = (0, 0, 14)$ , uses two identifiers 1 and 3. Since VCG allocates  $g_1$  and  $g_2$  to bidders 1 and 3, respectively, bidder  $1'$  obtains  $\{g_1, g_2\}$  and pays 2. On the other hand, when only two bidders  $1'$  and 2 participate in the auction, i.e., when bidder  $1'$  does not use false identifiers, bidder  $1'$  obtains  $\{g_1, g_2\}$  and pays 8.

As the above example shows, increasing the number of participating bidders by, or not by false identifiers reduces the seller's revenue. Therefore, it seems that a sub-additive allocation rule coincides with a summation-monotone one, and vice versa. However, in general, it is not always true. It is straightforward to make a counter example, though this paper does not give a rigorous proof. Nevertheless, for single-item auctions, we prove that false-name-proofness is equivalent to revenue monotonicity, assuming the following natural condition.

**Assumption 1.** *For any set of participating bidder  $N$  and for any bidder  $j(\in N)$ , if a mechanism allocates a good to a bidder when  $N \setminus \{j\}$  participates, it always allocates the good to someone including the bidder when  $N$  participates.*

We believe that introducing Assumption [11](#) is quite natural. For a seller, it is undesirable that a good is no longer allocated when more bidders join the auction. Now, let us introduce the next theorem. For space reasons, we omit the proof.

**Theorem 4.** *Under Assumption [7](#), a single-item auction mechanism is false-name-proof if and only if it is strategy-proof and revenue monotone.*

## 5 Conclusions and Future Works

We identified a simple condition called *summation-monotonicity*, which characterizes strategy-proof, revenue monotone allocation rules in combinatorial auctions. We proved that we can construct a strategy-proof, revenue monotone mechanism if and only if the allocation rule satisfies weak-monotonicity and summation-monotonicity.

To the best of our knowledge, this is the first attempt to characterize revenue monotone allocation rules. To demonstrate the power of our characterization, we verified existing combinatorial auction mechanisms and found that several non-trivial mechanisms are not revenue monotone. In addition, we shed light on a connection between revenue monotonicity and false-name-proofness and proved that assuming a natural condition, revenue monotonicity is equivalent to false-name-proofness for single-item auctions.

In future works, we hope to design a novel deterministic strategy-proof, revenue monotone combinatorial auction mechanism, that does not always achieve weak-maximality, since such a deterministic mechanism has not been proposed yet, although a randomized mechanism has been [7](#).

## References

1. Todo, T., Iwasaki, A., Yokoo, M.: Characterization of strategy-proof, revenue monotone combinatorial auction mechanisms and connection with false-name-proofness (2009), [http://agent.is.kyushu-u.ac.jp/~todo/wine2009\\_full.pdf](http://agent.is.kyushu-u.ac.jp/~todo/wine2009_full.pdf)
2. Lavi, R., Mu'alem, A., Nisan, N.: Towards a characterization of truthful combinatorial auctions. In: FOCS, pp. 574–583 (2003)
3. 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)
4. Todo, T., Iwasaki, A., Yokoo, M., Sakurai, Y.: Characterizing false-name-proof allocation rules in combinatorial auctions. In: AAMAS, pp. 265–272 (2009)
5. Rastegari, B., Condon, A., Leyton-Brown, K.: Revenue monotonicity in deterministic, dominant-strategy combinatorial auctions. Technical report, UBC CS Technical Report TR-2009-09 (2009)
6. Rastegari, B., Condon, A., Leyton-Brown, K.: Revenue monotonicity in combinatorial auctions. In: AAI, pp. 122–127 (2007)
7. Rastegari, B., Condon, A., Leyton-Brown, K.: Stepwise randomized combinatorial auctions achieve revenue monotonicity. In: SODA, pp. 738–747 (2009)



# 2D-TUCKER Is PPAD-Complete

Dömötör Pálvölgyi

Ecole Polytechnique Fédérale de Lausanne, Switzerland

dom@cs.elte.hu

<http://www.cs.elte.hu/~dom>

**Abstract.** Tucker's lemma states that if we triangulate the unit disc centered at the origin and color the vertices with  $\{1, -1, 2, -2\}$  in an antipodal way (if  $|z| = 1$ , then the sum of the colors of  $z$  and  $-z$  is zero), then there must be an edge for which the sum of the colors of its endpoints is zero. But how hard is it to find such an edge? We show that if the triangulation is exponentially large and the coloring is determined by a deterministic Turing-machine, then this problem is **PPAD**-complete which implies that there is not too much hope for a polynomial algorithm.

## 1 Introduction

Papadimitriou defined in [3] the complexity class **PPAD** which is a special class of search problems. It contains the problems which are reducible to LEAFD, a total search problem defined as follows.

**Definition 1** (LEAFD). *We are given the description of a deterministic Turing machine  $M$  which for every input  $v \in \{0, 1\}^n$  returns an ordered pair from the set  $\{0, 1\}^n \cup \{\text{no}\}$  in time  $\text{poly}(n)$ . These will be denoted by  $M_{in}(v)$  and  $M_{out}(v)$ . This defines a directed graph on  $V = \{0, 1\}^n$  such that  $uv \in E$  if  $M_{out}(u) = v$  and  $M_{in}(u) = v$ . This graph is a collection of directed cycles and paths. We also require that  $M_{in}(0^n) = \text{no}$ , meaning that  $0^n$  is a leaf (or an isolated node). The parity argument implies that in the former case there must be another leaf. The input is  $0^n$  (apart from the description of  $M$ ), the output of the search problem is another leaf (or  $0^n$  if it is an isolated node).*

**Remark 2.** *We can suppose that  $M$  is equipped with a standard built-in mechanism that checks its running time and if  $M$  would run too long, it halts and outputs no. It can also guarantee  $M_{in}(0^n) = \text{no}$ . It can be easily checked whether the description of  $M$  has this property and if not, then the output of the search problem can also be violation. In the problems that we will define later, we similarly allow the output to be violation if  $M$  violates one of the properties that we require.*

**Remark 3.** *This problem is almost in **TFNP**, the class of total search problems verifiable in polynomial time. We do not want to define this class here (see [2]). The reason why it is not in the class with this definition is that the verification time will depend on the running time of  $M$ , so it is not bounded by some fixed*

polynomial like it is in the case of SAT. An alternative definition would be to define a class called LEAFD-C where the running time of  $M$  would be bounded by  $n^c$  or to define  $M$  as a boolean circuit. It is not the goal of this paper to go deeper in this problem.

For the definition of reduction among search problems, we direct the reader to the original paper of Papadimitriou [3]. It was also shown there that the search versions of many well-known theorems that use some kind of parity argument belong to **PPAD**, moreover, many are also complete for this class. The following analogue of Sperner's lemma was shown to be **PPAD**-complete by Chen and Deng [1]. (When we write  $x \in \{0, 1\}^n$ , we also mean the number in base two that it represents.)

**Definition 4** (2D-SPERNER). *We are given the description of a deterministic Turing machine  $M$  which for every input  $(u, v) \in \{0, 1\}^{2n}$  such that  $u + v \leq 2^n$  returns either 1, 2 or 3 in time  $\text{poly}(n)$ . Furthermore,  $M(0, 0) = 1$ ,  $M(2^n, 0) = 2$ ,  $M(0, 2^n) = 3$ , for all  $i < 2^n$   $M(0, i) \neq 3$ ,  $M(i, 0) \neq 2$  and for all  $i + j = 2^n$   $M(i, j) \neq 1$ . The output (whose existence is guaranteed by Sperner's lemma) is  $(u, v) \in \{0, 1\}^{2n}$  for which  $M(u, v)$ ,  $M(u+1, v)$  and  $M(u, v+1)$  are all different.*

One can similarly define 3D-SPERNER and other higher dimensional analogues. It is also possible to define a continuous version, which can be denoted by 2D-BROUWER, the interested reader is again directed to [3] where it is also shown that all these variants are equivalent to LEAFD and thus are **PPAD**-complete.

**Definition 5** (2D-TUCKER). *We are given the description of a deterministic Turing machine  $M$  which for every input  $(u, v) \in \{0, 1\}^{2n}$  returns either 1,  $-1$ , 2 or  $-2$  in time  $\text{poly}(n)$ . Furthermore, for all  $i$   $M(0, i) = -M(2^n, 2^n - i)$  and  $M(i, 0) = -M(2^n - i, 2^n)$ . The output (whose existence is guaranteed by Tucker's lemma) is  $(u, v) \in \{0, 1\}^{2n}$  and  $(u', v') \in \{0, 1\}^{2n}$  for which  $|u - u'| \leq 1$ ,  $|v - v'| \leq 1$  and  $M(u, v) = -M(u', v')$ .*

**Remark 6.** *Tucker's lemma is often stated in a slightly different way, more similar to Sperner's, and it requires the square to be triangulated. The above search problem is clearly easier than the triangulated one, so when we prove a hardness result about 2D-TUCKER, that also implies the hardness of the triangulated version, so our results hold for both cases.*

The respective higher dimensional and continuous versions are denoted by 3D-TUCKER and 2D-BORSUK-ULAM, the interested reader is again directed to [3] where it is shown that the higher dimensional version is **PPAD**-complete and 2D-BORSUK-ULAM is equivalent to 2D-TUCKER. The **PPAD**-completeness of 2D-TUCKER was posed as an open problem both in [3] and in [1]. In this note we prove this result.

**Theorem 7.** *2D-TUCKER is **PPAD**-complete.*

## 2 Reduction of LEAFD to 2D-TUCKER

It was shown in [3] that  $2D-TUCKER \in \mathbf{PPAD}$ , to prove hardness, we will reduce LEAFD to it. The reduction is surprisingly easy and only uses technics similar to the ones appearing already in [1].

We will call the vertices of the grid *points* and the vertices of the graph generated by  $M$  simply *vertices*. We say that two points are *neighbors* if their distance is  $\leq \sqrt{2}$  (meaning there is a little square that has both as its vertex). We call two points *negated* if the sum of their colors is zero.

The goal is, that given any  $M$  that generates an input for LEAFD, we want to produce a coloring  $c$  of the points of the  $20 \cdot 2^{2n} \times 20 \cdot 2^{2n}$  grid with colors  $\pm\{1, 2\}$  such that if one finds two negated neighbors, then we can find a leaf in the graph generated by  $M$ .

The idea is that for every vertex we reserve a part of the grid and if there are two negated neighbors in a reserved part, that will imply that the vertex to which this part belongs to is a leaf (there cannot be negated neighbors outside the parts reserved for vertices). Most part of the square is filled with 1's, the edges are represented by tubes of  $-2, -1, 2$  going from one reserved part to the other, and these tubes are disjoint (if two tubes would cross, we slightly modify them in the vicinity of the crossing so as they evade each other). Unfortunately it is quite ugly to give a precise description of this construction by words, we advice the reader to consult the Figures which might be sufficient even without reading the text to understand the whole reduction.

For a vertex  $v_i$  (where the indices are an arbitrary enumeration of the  $2^n$  vertices with  $v_0$  being  $0^n$ ) we reserve a part close to the left side of the square,  $V_i = [8, \dots, 10] \times [20i2^n + 10, \dots, 20(i + 1)2^n - 10]$ . For different  $i$ 's, these parts are disjoint and are above each other. We also reserve a part for every possible edge of the graph. For the possible  $v_i v_j$  edge we reserve a part that connects the lower half of  $V_i$  and the upper half of  $V_j$  via a  $\sqcap$  shape<sup>1</sup>,  $E_i = [11, \dots, 20i2^n + 10 + 10j] \times [20i2^n + 10 + 10j, \dots, 20i2^n + 10 + 10j + 2] \cup [20i2^n + 10 + 10j, \dots, 20i2^n + 10 + 10j + 2] \times [20i2^n + 10 + 10j, \dots, 20j2^n + 10 \cdot 2^n + 10 + 10i + 2] \cup [11, \dots, 20i2^n + 10 + 10j + 2] \times [20j2^n + 10 \cdot 2^n + 10 + 10i, \dots, 20j2^n + 10 \cdot 2^n + 10 + 10i + 2]$ . These regions are mainly disjoint, every intersection  $E_i \cap E_j$  is a little square, far from the other edges. If  $v_i v_j$  is an edge, then we fill out this tube of thickness 3 with  $-2, -1, 2$ , with the  $-1$ 's being in the middle, the  $-2$ 's being in the bottom when leaving  $v_i$  and in the top when entering  $v_j$ , the  $2$ 's being in the top when leaving  $v_i$  and in the bottom when entering  $v_j$ . (We deal with the intersections of filled out tubes later). Remember that most of the square is filled out with 1's, so if a point does not belong to a part reserved to an edge or vertex, then its color is 1. This way we do not create any negated neighbors outside of the parts reserved for vertices, since the boundaries of the tubes are always  $\pm 2$ 's. Inside  $V_i$ , if  $v_h v_i$  and  $v_i v_j$  are both edges, we fill out the vertical tube of thickness 3 leading from where the tube of the edge from  $v_h$  enters down to where the tube of the edge to

---

<sup>1</sup> We suggest to skip the following ugly description and just read the properties in the next sentence.

-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	2
-1	2	2	2	2	2	2	2	2	2	2	2
-1	2	1	1	1	1	1	1	1	1	1	1
-1	2	1									1
-1	2	1									1
-1	2	1		2	2	2	2				1
-1	2	1		2	-1	-1	-1				1
-1	2	1	1	2	-1	-2	-2				1
-1	2	2	2	2	-1	-2					1
-1	-1	-1	-1	-1	-1	-2					1
-2	-2	-2	-2	-2	-2	-2					1
-2	1	1	1	1	1	1	1	1	1	1	1

	-2	-2	-2	-2	-2	-2	-2	-2			
$V_j$	-1	-1	-1	-1	-1	-1	-1	-1	-2		
	2	2	2	2	2	2	2	-1	-2		
								2	-1	-2	
								2	-1	-2	
								2	-1	-2	
								2	-1	-2	
								2	-1	-2	
	2	2	2	2	2	2	2	-1	-2		
$V_i$	-1	-1	-1	-1	-1	-1	-1	-1	-2		
	-2	-2	-2	-2	-2	-2	-2	-2	-2		

Fig. 1. The boundary “going” into the leaf  $v_0$  and An edge  $v_i v_j$

-2	-2	-2	-2	-2	-2	-2	-2	-2			
-2	-1	-1	-1	-1	-1	-1	-1	-1	from $V_h$		
-2	-1	2	2	2	2	2	2	2			
-2	-1	2									
-2	-1	2									
-2	-1	2									
-2	-1	2									
-2	-1	2	2	2	2	2	2	2			
-2	-1	-1	-1	-1	-1	-1	-1	-1	to $V_j$		
-2	-2	-2	-2	-2	-2	-2	-2	-2			

			-2	-1	2	2	2				
			-2	-1	-1	-1	2				
			-2	-2	-2	-1	2				
						-2	-1	2			
2	2	2	2	2	2	-2	-1	2	2	2	2
-1	-1	-1	-1	2	2	-2	-1	-1	-1	-1	-1
-2	-2	-2	-1	2	2	-2	-2	-2	-2	-2	-2
			-2	-1	2						
			-2	-1	2	2	2				
			-2	-1	-1	-1	2				
			-2	-2	-2	-1	2				
			-2	-1	2						

Fig. 2. The part  $V_i$  with edge from  $v_h$  and to  $v_j$  and Handling a crossing





Fig. 3. The graph of the path  $v_0 v_3 v_1 v_4 v_2$  before and after handling the crossings

$v_j$  starts  $([8, \dots, 10] \times [20i2^n + 10 + 10j, \dots, 20i2^n + 10 \cdot 2^n + 10 + 10h + 2])$  with  $-2, -1, 2$  such that the  $-1$ 's are in the middle, the  $-2$ 's are to the left and the  $2$ 's are to the right. This way again do not create any negated neighbors. If  $v_i$  is a leaf, then leave it filled out with  $1$ 's (which gives negated neighbors) except for  $v_0$ . To  $v_0$  we "drive in" the boundary of the square, we set  $c(m, m) = 2$ ,  $c(0, 0) = c(0, 1) = -2$ , for all  $1 < i$   $c(0, i) = c(i, m) = -1$ , for  $0 < i < m$   $c(i, m - 1) = 2$ , for  $2 < i < m$   $c(1, i) = 2$  and continue this tube to inside  $v_0$  and from there to the start of the tube of the edge to its only neighbor.

We have almost solved the problem, the only thing left that we must handle is if two filled out tubes cross. In this case we can simply modify the tubes in the vicinity of their crossing such that we do not create negated neighbors. If for example the edge  $ab$  would cross  $cd$ , then we modify the tubes such that we obtain an  $ad$  and a  $bc$  edge (see Figures). Of course these will not really be tubes leading from  $V_a$  to  $V_d$  and from  $V_b$  to  $V_c$  because we are handling several crossings, but that does not matter for us. We only want to preserve the conditions that the colors are easy to determine and that there are no negated neighbors outside the parts reserved for vertices.

Now the color of any point can be determined by a finite number of computations of  $M$  (we can easily decide from the coordinates of any point whether it belongs to a part reserved for a vertex, an edge, to a crossing or to the remaining part of the grid). If we find two negated neighbors, they must be in a part reserved for a vertex that is a leaf in the original graph. This finishes the reduction.

### 3 Remarks and Acknowledgment

The same argument works to solve 2D-SPERNER which slightly simplifies the proof of [1].

An interesting question would be to determine the complexity of the so-called octahedral Tucker's lemma (here the dimension would be a part of the input in unary), which might tell something about the complexity of necklace splitting among two thieves with a lot of different kinds of beads. Since this theorem is not so widely known and can be stated in a purely combinatorial way, we state it here.

**Lemma 8.** *(Octahedral Tucker's lemma) If for any set-pair  $A, B \subset [n]$ ,  $A \cap B = \emptyset$ ,  $A \cup B \neq \emptyset$  we have a  $\lambda(A, B) \in \pm[n - 1]$  color, such that  $\lambda(A, B) = -\lambda(B, A)$ , then there are two set-pairs,  $(A_1, B_1)$  and  $(A_2, B_2)$  such that  $A_1 \subset A_2$ ,  $B_1 \subset B_2$  and  $\lambda(A_1, B_1) = -\lambda(A_2, B_2)$ .*

### Acknowledgement

I would like to thank Jarek Byrka for discussions and reading the first version of this paper.

## References

1. Chen, X., Deng, X.: On the complexity of 2D discrete fixed point problem. In: 33rd International Colloquium on Automata, Languages and Programming, pp. 489–500 (2006)
2. Megiddo, N., Papadimitriou, C.H.: On total functions, existence theorems, and computational complexity. *Theoretical Computer Science* 81(2), 317–324 (1991)
3. Papadimitriou, C.: On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. System Sci.* 48, 498–532 (1994)

# Bidder Optimal Assignments for General Utilities<sup>\*</sup>

Paul Dütting, Monika Henzinger, and Ingmar Weber

Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland  
{paul.duetting,monika.henzinger,ingmar.weber}@epfl.ch

**Abstract.** We study the problem of matching bidders to items where each bidder  $i$  has general, strictly monotonic utility functions  $u_{i,j}(p_j)$  expressing her utility of being matched to item  $j$  at price  $p_j$ . For this setting we prove that a bidder optimal outcome always exists, even when the utility functions are non-linear and non-continuous. Furthermore, we give an algorithm to find such a solution. Although the running time of this algorithm is exponential in the number of items, it is polynomial in the number of bidders.

## 1 Introduction

In two-sided matching markets buyers are to be matched to items and the seller receives a monetary compensation from the buyers. Such markets have been studied for several decades [1,2]. They have seen a surge of interest with the spread of sponsored search auctions [3,4], where advertisers (bidders) are competing for the available advertising slots (items). Our research is also motivated by sponsored search but our results are not specific to this setting in any way.

A solution is *bidder optimal* if it gives each bidder the highest possible utility. For the case where each bidder  $i$  has a utility function  $u_{i,j}(p_j)$  for item  $j$  that drops linearly in the price  $p_j$  it has long been known how to find such solutions [5,6]. Recently, an algorithm was presented that also copes with per-bidder-item reserve prices (where a certain minimum price  $p_j \geq r_{i,j}$  is required if bidder  $i$  is matched to item  $j$ ) and per-bidder-item maximum prices (where bidder  $i$  can pay no more than  $m_{i,j}$  for item  $j$ ) [7]. This algorithm requires the input to be in “general position”, which e.g. requires that all reserve and maximum prices are different. Our results do not require this assumption.

In [8,9,10] the *existence* of bidder optimal solutions was shown for general, strictly monotonic utility functions, as long as the utility functions are *continuous*. However, no algorithm was given to find such a solution. For the special case of piece-wise linear functions an algorithm was presented in [11,12], where the arguments used to prove termination leads to a time bound exponential in the number of items. The authors then argue that arbitrary continuous functions can be uniformly approximated by such piece-wise linear functions. However,

---

<sup>\*</sup> This work was conducted as part of a EURYI scheme award (see <http://www.esf.org/euryi/>).

neither the approximation accuracy nor the running time are analyzed and the relationship between the two is unclear. Also, such a uniform approximation does not exist for *discontinuous* utility functions.

Our main contributions are two-fold. First, we prove that even for general, strictly monotonic discontinuous utility functions a bidder optimal outcome always exists. Second, we give an algorithm to find such an outcome. The running time of this algorithm is polynomial in the number of bidders  $n$  but exponential in the number of items  $k$ . However, for sponsored search  $k$  can be viewed as constant or, at least,  $k \ll n$ .

Our model includes per-bidder-item reserve prices  $r_{i,j}$  and also per-bidder reserve utilities  $o_i$ , usually referred to as *outside options*. We can also model per-bidder-item max prices  $m_{i,j}$  with a discontinuous drop of  $u_{i,j}(p_j)$  at the price  $p_j = m_{i,j}$ . Other settings that can be modeled are interest rates where, e.g., up to a certain point a bidder can still pay from her own pocket but for higher prices she has to borrow money from a bank, leading to a faster drop in utility. Similarly, settings where the bidder is “risk averse” and loses utility faster for higher prices due to an associated higher variance can be modeled.

Note that both for our and related previous results the utility function is a function of the price only. Other inter-item dependencies cannot be modeled. Such dependencies include a drop in utility if some other bidder gets a particular item or a positive utility only if a bidder gets a particular set of items. For a survey concerning such combinatorial auctions we refer the reader to [13].

## 2 The Assignment Problem

The problem is defined as follows: We are given a set  $I$  of  $n$  bidders and a set  $J$  of  $k$  items. We use letter  $i$  to denote a bidder and letter  $j$  to denote an item. Each bidder  $i$  has a *utility function*  $u_{i,j}(p_j)$  for each item  $j$  expressing her utility of being matched to item  $j$  at price  $p_j$ . We assume that (i) the utility functions  $u_{i,j}(\cdot)$  are strictly monotonically decreasing and (ii) for the outside options  $o_i$  (defined below) there exist *threshold values*  $\bar{p}_{i,j}$  s.t.  $u_{i,j}(\bar{p}_{i,j}) \leq o_i$ . We do *not* assume that  $u_{i,j}(\cdot)$  is (globally) continuous, but we do require that (iii) it is locally right-continuous, i.e. that  $\forall x : \lim_{\epsilon \rightarrow 0^+} u_{i,j}(x + \epsilon) = u_{i,j}(x)$ .<sup>1</sup>

We want to compute a *matching*  $\mu \subseteq I \times J$ , where we require *all* bidders to be matched. To ensure that this is possible, even if  $k < n$ , we introduce symbolic *dummy items*. If bidder  $i$  is matched to a dummy item  $j$  then this represents that  $i$  is actually not matched. In the *proofs of existence* of bidder optimal solutions we will assume that there are  $n$  dummy items, one for each bidder, in our *algorithm* we directly deal with the case of an unmatched bidder as the running time would suffer significantly from an increase of the (small) number of items  $k$  to  $n + k$ . To distinguish between a match to a dummy and a real item we call the latter case *properly matched*. We use  $\mu(i)$  to denote the

<sup>1</sup> At the end of Section 3 we show that all three of these requirements are necessary for a bidder optimal solution to exist.



item that is matched to bidder  $i$  in  $\mu$  and  $\mu(j)$  to denote the bidder, if any, that is matched to item  $j$ . Note that we do *not* require all items to be matched.

We say that a matching  $\mu$  with prices  $p = (p_1, \dots, p_k)$  is *feasible* if (i)  $p_j \geq r_{i,j}$  for all  $(i, j) \in \mu$ , where  $r_{i,j}$  is a *reserve price*, and (ii)  $u_{i,\mu(i)}(p_{\mu(i)}) \geq o_i$ , where  $o_i$  is an *outside option*.<sup>2</sup> As mentioned, we have one (conceptual) dummy item  $j$  for each bidder  $i$  s.t.  $u_{i,j}(x) = o_i - x$  and  $u_{i',j}(x) = o_{i'} - 1 - x$  for all  $i' \neq i$ , as well as  $r_{i,j} = 0$  and  $r_{i',j} = o_{i'}$  for all  $i' \neq i$ .<sup>3</sup> We say that the outcome  $(\mu, p)$  is *stable* if for all  $(i, j) \in I \times J$  we have  $u_{i,j}(p_j) \leq u_{i,\mu(i)}(p_{\mu(i)})$ , i.e. each bidder gets an item which, at the prices of the outcome, is one of her first choices. We will also refer to prices  $p$  as feasible and/or stable if a corresponding feasible and/or stable matching  $\mu$  exists. The triple  $(u_{i,j}(\cdot), r_{i,j}, o_i)$  along with the implicitly assumed sets  $I$  and  $J$  constitute the *input* of the assignment problem. Note that our definition of stability is slightly *stronger* than the definition in [7] in the sense that our stability implies their stability, which also involves the  $r_{i,j}$ .

Our goal is to find a *bidder optimal* solution. We say that  $(\mu, p)$  is bidder optimal if it is both feasible and stable and for every other feasible and stable  $(\mu', p')$  we have that  $u_{i,\mu(i)}(p_{\mu(i)}) \geq u'_{i,\mu'(i)}(p'_{\mu'(i)})$  for all  $i$ .

Given the input  $(u'_{i,j}(\cdot), r_{i,j}, o_i)$  let  $(\mu', p)$  be the outcome of a mechanism and analogously for  $u''_{i,j}(\cdot)$ . We say that the mechanism is *truthful* if for every bidder  $i$  with utility functions  $u_{i,1}(\cdot), \dots, u_{i,k}(\cdot)$  and any two input matrices of utility functions  $u'$  and  $u''$  with  $u'_{i,j}(\cdot) = u_{i,j}(\cdot)$  for all  $j$  and  $u'_{k,j}(\cdot) = u''_{k,j}(\cdot)$  for all  $k \neq i$  and all  $j$  we have that  $u_{i,\mu'(i)}(p'_{\mu'(i)}) \geq u_{i,\mu''(i)}(p''_{\mu''(i)})$  for all  $i$ . Note that this definition only involves the utility functions and not  $r_{i,j}$  or  $o_i$ . We assume that the  $r_{i,j}$ , which are a property of the sellers, cannot be falsified by the bidders. Furthermore, it is easy to see that misreporting the  $o_i$  is not beneficial to  $i$ . Overreporting can only lead to a missed chance of being assigned an item and underreporting can lead to a utility below the true outside option.

### 3 Existence of a Bidder Optimal Solution

**Theorem 1.** *For any input  $(u_{i,j}(\cdot), r_{i,j}, o_i)$  to the assignment problem there exists a bidder optimal outcome  $(\mu^*, p^*)$ .*

So far this result was only known for *continuous* utility functions [8,9,10]. Our proof consists of the following steps: Lemma 1, which can be proven by contradiction, shows that lowest feasible and stable prices are sufficient for bidder optimality. Lemma 2 shows that any two feasible and stable outcomes  $(\mu, p)$  and  $(\mu', p')$  can be combined to give a new feasible and stable solution with prices  $\min(p, p')$ . Lemma 3 then asserts that the infimum  $p^*$  over all feasible and stable prices, if it exists, corresponds to a feasible and stable outcome  $(\mu^*, p^*)$ . Finally,

<sup>2</sup> The second part of the feasibility definition is often referred to as *individual rationality* of the bidders. Similarly, the reserve price condition can be referred to as individual rationality of the sellers with bidder-dependent payoffs.

<sup>3</sup> The intuition is that in any feasible and stable outcome the price of item  $j$  can be as low as zero and so bidder  $i$  will have a utility of at least  $o_i$ .

Lemma 4 finishes the proof as it gives the existence of at least one feasible and stable outcome, establishing the existence of an infimum by Lemma 2.

**Lemma 1.** *If  $(\mu^*, p^*)$  is feasible and stable and  $p_j^* \leq p_j$  for all  $j$  and any  $(\mu, p)$  that is feasible and stable, then  $(\mu^*, p^*)$  is bidder optimal.*

**Lemma 2.** *Any two outcomes  $(\mu, p)$  and  $(\mu', p')$  which are feasible and stable for the input  $(u_{i,j}(\cdot), r_{i,j}, o_i)$  can be combined s.t. there exists a matching  $\hat{\mu}$  which together with the prices  $\hat{p} = \min(p, p')$  is feasible and stable for the same input.*

Even though the setting in [10] is for *continuous* utility functions and for a slightly weaker definition of stability involving both bidder *and* seller, their proof of this particular lemma (their Lemma 2) goes through unchanged. Concretely, it is shown that each bidder  $i$  gets item  $\hat{\mu}(i)$  at a price  $\hat{p}_{\hat{\mu}(i)}$  corresponding to  $u_{i,\hat{\mu}(i)}(\hat{p}_{\hat{\mu}(i)}) = \max(u_{i,\mu(i)}(p_{\mu(i)}), u_{i,\mu'(i)}(p'_{\mu'(i)}))$ . In other words, the two outcomes are stitched together in the best possible way for all bidders.

Lemma 2 implies that if there are any feasible and stable prices then there are unique infimum prices  $p^* = \inf\{p : p \text{ are feasible and stable prices}\}$ . It remains to show that (i)  $p^*$  corresponds to a feasible and stable outcome  $(\mu^*, p^*)$  (Lemma 3) and (ii) that there is at least one feasible and stable outcome (Lemma 4).

**Lemma 3.** *If there exists a feasible and stable outcome  $(\mu, p)$  matching all bidders, then there exists a feasible and stable outcome  $(\mu^*, p^*)$  matching all bidders with lowest prices. I.e. no other feasible and stable outcome  $(\mu', p')$  matching all bidders can have  $p'_j < p_j^*$  for any  $j$ .*

Our proof uses the following definitions: Let  $F_p \subseteq I \times J$  be the *first choice graph* at prices  $p$  which contains an edge from bidder  $i$  to item  $j$  if and only if  $j \in \operatorname{argmax}_{j'} u_{i,j'}(p_{j'})$ . Note that  $(\mu, p)$  is *stable* if and only if  $\mu \subseteq F_p$ . Let  $\tilde{F}_p \subseteq F_p$  denote the subset of *feasible* edges  $(i, j)$  where  $p_j \geq r_{i,j}$  and  $u_{i,j}(p_j) \geq o_i$ .<sup>4</sup> For  $i \in I$  and  $j \in J$  we define  $F_p(i) = \{j : \exists(i, j) \in F_p\}$  and  $F_p(j) = \{i : \exists(i, j) \in F_p\}$ . For  $T \subseteq I$  and  $S \subseteq J$  we define  $F_p(T) = \cup_{i \in T} F_p(i)$  and  $F_p(S) = \cup_{j \in S} F_p(j)$ . We define  $\tilde{F}_p(i)$ ,  $\tilde{F}_p(j)$ ,  $\tilde{F}_p(T)$ , and  $\tilde{F}_p(S)$  analogously. We call a (possibly empty) set  $S \subseteq J$  *strictly overdemanding* for prices  $p$  wrt  $T \subseteq I$  if (i)  $\tilde{F}_p(T) \subseteq S$  and (ii)  $\forall R \subseteq S, R \neq \emptyset : |\tilde{F}_p(R) \cap T| > |R|$ . Using Hall's Theorem [14] one can show that a feasible and stable matching exists for given prices  $p$  if and only if there is no strictly overdemanding set of items  $S$ .

*Proof.* If we assume that there exists at least one feasible and stable outcome, then Lemma 2 shows that there exist unique infimum prices  $p^*$ .

For a contradiction suppose that there is no matching  $\mu^*$  s.t.  $(\mu^*, p^*)$  is feasible and stable. Then, by Hall's Theorem, there must be a set  $T$  of bidders s.t.  $\tilde{F}_{p^*}(T)$  is strictly overdemanding for prices  $p^*$  wrt  $T$ .

In any feasible and stable outcome  $(\hat{\mu}, \hat{p})$  we have  $\hat{p}_j \geq p_j^*$  for all items  $j$  and, thus, the overdemand for the items in  $\tilde{F}_{p^*}(T)$  can only be resolved if (i) at least

<sup>4</sup> The second feasibility condition is redundant as if  $u_{i,j}(p_j) < o_i$  then bidder  $i$  strictly prefers her dummy item whose price can always be assumed to be 0 in any bidder optimal solution. See Lemma 4.

one of the bidders  $i \in T$  has a feasible first choice item  $j \in J \setminus F_{p^*}(T)$  under  $\hat{p}$  or (ii) for some item  $j \in F_{p^*}(T) \setminus \tilde{F}_{p^*}(T)$  we have that  $\hat{p}_j \geq r_{i,j}$ . Case (i) corresponds, for each pair  $(i, j) \in T \times J \setminus F_{p^*}(T)$ , to a price increase relative to  $p^*$  of  $s_j^i = \inf\{x \geq 0 : u_{i,j}(p_j^* + x) \leq \max_{j' \in J \setminus F_{p^*}(T)} u_{i,j'}(p_{j'}^*)\}$ , which is  $> 0$  and contained in the set itself as  $u_{i,j}(\cdot)$  is right-continuous.<sup>5</sup> Case (ii) corresponds, for each pair  $(i, j) \in I \times F_{p^*}(T) \setminus \tilde{F}_{p^*}(T)$ , to a price increase relative to  $p^*$  of  $f_j^i = r_{i,j} - p_j^*$ , which is also  $> 0$ . Let  $\delta_j^i = \min(s_j^i, f_j^i)$  if  $j \in F_{p^*}(i) \setminus \tilde{F}_{p^*}(i)$  and let  $\delta_j^i = s_j^i$  otherwise. Then  $\Sigma_\delta = \min_i \sum_j \delta_j^i$  is a lower bound on the sum of the price increases for any feasible and stable outcome  $(\hat{\mu}, \hat{p})$ .

Lemma 2, however, shows that for any  $\epsilon > 0$  there exist feasible and stable prices  $p'$  s.t.  $|p'_j - p_j^*| < \epsilon$  for all items  $j$ . For  $\epsilon = \Sigma_\delta/|J|$  this gives a contradiction to the fact that the price increases corresponding to  $\Sigma_\delta$  were required by any feasible and stable solution. We conclude that there exists at least one matching  $\mu^*$  s.t.  $(\mu^*, p^*)$  is feasible and stable.  $\square$

**Lemma 4.** *For the assignment problem there are unique lowest stable prices  $p^*$  s.t. any other feasible and stable outcome  $(\mu, p)$  has  $p_j \geq p_j^*$  for all  $j \in J$ .*

*Proof.* By Lemma 3 we know that we only have to show the existence of some stable outcome  $(\mu^*, p^*)$ . Set  $p_j^* = \max_i(\bar{p}_{i,j})$  for the non-dummy items  $j$  (where the  $\bar{p}_{i,j}$ 's are the threshold values defined above) to ensure that all bidders have (at most) a utility of  $o_i$  and then match all bidders to dummy items at a price of 0. As no real item is matched and all utilities are  $o_i$  this is feasible. It is also stable as all real items have prices so high that no bidder strictly prefers them over a dummy item.  $\square$

Finally, we show that all three conditions on the utility functions (see Section 2) are required to guarantee the existence of a bidder optimal solution:

*Strict monotonicity:* Consider a setting with three bidders and two items and the following utility functions:  $u_{1,1}(x) = u_{3,2}(x) = 1 - x$ ,  $u_{1,2}(x) = u_{3,1}(x) = -x$  and  $u_{2,1}(x) = u_{2,2}(x) = 2$  if  $x \leq 1$  and  $u_{2,1}(x) = u_{2,2}(x) = 3 - x$  otherwise. All  $r_{i,j} = o_i = 0$ . Then one feasible and stable outcome is  $\mu = \{(1, 1), (2, 2)\}$ ,  $p = (0, 1)$  whereas another is  $\mu = \{(2, 1), (3, 2)\}$ ,  $p = (1, 0)$ . In neither of the two settings can the price for the item with price 0 be lowered any further without upsetting stability. The first outcome is strictly preferred by bidder 1, whereas the second is strictly preferred by bidder 2.

*Eventual drop in utility to  $o_i$ :* Consider two bidders who both have the utility function  $u_{i,1}(x) = 1/(1 + x)$  for a single item. Again,  $r_{i,1} = o_i = 0$ . Then no matter how large  $p_1$  is, both bidders will still strictly prefer the item over being unmatched.

*Right continuity:* Consider two bidders who both have the following utility function for a single item:  $u_{i,1}(x) = 2 - x$  if  $x \leq 1$  and  $u_{i,1}(x) = -x$  otherwise.

<sup>5</sup> This no longer holds without the requirement of right-continuity as discussed at the end of this section.

Then a price of  $p_1 \leq 1$  will not be stable, as both bidders strictly prefer the item over being unmatched. So any stable price needs to satisfy  $p_1 > 1$  and this set no longer contains its infimum. If we change the first condition of the utility function to  $x < 1$ , ensuring right-continuity, then the price  $p_1 = 1$  is stable, even though the item cannot be assigned to either of the two bidders.

**Truthfulness.** If the reserve prices  $r_j$  depend only on the items, then Theorem 2 in [10] shows that any mechanism that computes a bidder optimal outcome is truthful. If the reserve prices  $r_{i,j}$  also depend on the bidders, then this is no longer true: Consider a setting with two bidders and two items and the following utility functions, reserve prices, and outside options:  $u_{1,1}(x) = 6 - x$ ,  $u_{1,2}(x) = 5 - x$ ,  $u_{2,1}(x) = 6 - x$ ,  $u_{2,2}(x) = 6 - x$ ,  $r_{1,1} = 2$ ,  $r_{1,2} = 0$ ,  $r_{2,1} = 1$ ,  $r_{2,2} = 2$ , and  $o_1 = o_2 = 0$ . The bidder optimal outcome is  $\mu = \{(1, 1), (2, 2)\}$ ,  $p = (2, 2)$ . If Bidder 2 reports  $u_{2,2}(x) = 0$ , then the bidder optimal outcome is  $\mu = \{(1, 2), (2, 1)\}$ ,  $p = (1, 0)$ . This gives Bidder 2 a strictly higher utility.

### 4 Algorithm for General Utilities

Here we present an algorithm which *directly* computes a bidder optimal outcome for general utility functions, as opposed to settling for a piecewise-linear approximation [11][12]. Our algorithm assumes that computations of the “inverse utility function”  $u_{i,j}^{-1}(x) = \min\{p : u_{i,j}(p) \leq x\}$  take constant time. If  $u_{i,j}(\cdot)$  is continuous then  $u_{i,j}^{-1}(\cdot)$  is indeed the inverse function. More generally, it is merely a one-sided inverse function satisfying  $u_{i,j}^{-1}(u_{i,j}(p)) = p$ .

**Description of the Algorithm.** Try out all possible matchings in which all bidders, but not necessarily all items, are matched. For a particular matching  $\mu$  try all possible ways of ordering the (up to)  $k$  properly matched bidders. For each ordering initialize *lower bounds* on the prices  $\forall j : b_j = \min_i r_{i,j}$ . Execute the following steps for every properly matched bidder  $i$  according to the current ordering: Fix the price of item  $\mu(i)$  to be  $p_{\mu(i)} = \max(b_{\mu(i)}, r_{i,\mu(i)})$  and update  $b_{\mu(i)} = p_{\mu(i)}$ . If  $u_{i,\mu(i)}(p_{\mu(i)}) < o_i$  then abort the check of the particular ordering.<sup>6</sup> If there exists a previously considered bidder  $i'$  where  $u_{i,\mu(i)}(p_{\mu(i)}) < u_{i,\mu(i')}(b_{\mu(i')})$  then also abort the check of the particular ordering.<sup>7</sup> If neither of these two cases happens, update the vector of price bounds by setting  $\forall j : b_j = \max(b_j, u_{i,\mu(i)}^{-1}(b_{\mu(i)}))$ . Once all properly matched bidders have been considered, go through all bidders matched to dummy items. If for such a bidder  $i$  there exists a matched item  $j = \mu(i')$  where  $u_{i,j}(p_j) > o_i$  then abort. Otherwise, set  $p_j = \max(p_j, u_{i,j}^{-1}(o_i))$  for all items  $j$ .<sup>8</sup> After considering all matchings and orderings, output any pair  $(\mu, p)$  corresponding to the lowest found prices  $p$ .

<sup>6</sup> The price  $p_{\mu(i)}$  is now already too high to allow a feasible matching of  $i$  to  $\mu(i)$ .  
<sup>7</sup> We assumed that  $i'$  got her item at the current price bound  $p_{\mu(i')} = b_{\mu(i')}$  but this price would no longer be stable wrt to  $i$ .  
<sup>8</sup> Note that these price updates can only affect unmatched items. For all matched items  $\mu(i)$  the price was finalized when their bidder  $i$  was considered.

**Lemma 5.** *If the above algorithm terminates without aborting then the returned output  $(\mu, p)$  is feasible and stable.*

*Proof.* When a particular matched bidder  $i$  is considered the price  $p_{\mu(i)}$  will not rise anymore until abortion or until a new matching/ordering is tried. During the consideration of  $i$  we update the prices of all other items to ensure that  $i$ 's utility is highest for item  $\mu(i)$ . As the prices for the *other* items might increase further this ensures that on termination we have  $u_{i,j}(p_j) \leq u_{i,\mu(i)}(p_{\mu(i)})$  for all  $j$ . Prices are feasible as whenever we match a bidder  $i$  to an item  $j$  we set  $p_j = \max(b_j, r_{i,j})$ . Utilities are feasible as when considering  $i$  we only continue if  $i$ 's utility is non-negative and this does not change until termination. As for unmatched bidders, we ensure that they have a utility of no more than  $o_i$  for any real item and so they are stable.  $\square$

If one requires  $p_j \geq b_j$  in addition to feasibility, then the  $b_j$  can simply be absorbed into the reserve price by setting  $r'_{i,j} = \max(r_{i,j}, b_j)$ . So whenever the algorithm updates the bounds  $b_j$  while considering bidder  $i$  this is conceptually a restart of the algorithm with a new input, where reserve prices are changed and bidder  $i$  and item  $\mu(i)$  have been removed.

**Lemma 6.** *If we know that for the input  $(u_{i,j}(\cdot), r_{i,j}, o_i)$  a feasible and stable outcome exists, then every bidder optimal outcome  $(\mu, p)$  matches at least one bidder  $i$  to an item  $j$  at price  $p_{\mu(i)} = r_{i,\mu(i)}$ .*

*Proof.* Theorem 1 in [10], whose proof does not require continuity, shows that given a bidder optimal outcome there is no other feasible (but not necessarily stable) solution where all bidders are strictly better off. But if *all* prices of matched items  $j$  satisfied  $p_j > r_{\mu(j),j}$  then by setting all prices for the matched items to  $p_j = r_{\mu(j),j}$  all bidders would benefit as we assume that utility functions are strictly decreasing in the price and prices  $p_j$  are still feasible.  $\square$

Lemma 6 is the main ingredient to reconstruct the lowest stable prices for a given matching. It lets us remove one bidder  $i$  at a time where at each step we can ensure that  $p_{\mu(i)}$  is not affected by removals of bidders in future iterations.

**Theorem 2.** *If for  $\mu$  currently being tried by the algorithm there exist feasible and stable prices then the algorithm will find the lowest stable prices  $p$  for  $\mu$ .*

*Proof.* By induction on the number of bidders  $n$ . If  $n = 1$  then in the bidder optimal outcome the (unique) bidder  $i$  gets item  $\mu(i) \in \operatorname{argmax}_{j'} u_{i,j'}(\max(b_{j'}, r_{i,j'}))$  at price  $\max(b_{\mu(i)}, r_{i,\mu(i)}) = r_{i,\mu(i)}$  and the price of all items  $j' \neq \mu(i)$  must be at least  $u_{i,j'}^{-1}(u_{i,\mu(i)}(p_{\mu(i)}))$  to guarantee stability for bidder  $i$ . These are also the prices the algorithm will return if it does not abort. Now suppose the result holds for all  $t \leq n$  and we want to prove the claim for  $n + 1$ . By Lemma 6 we know that at least one item  $j$  is sold for the price  $p_j = \max(b_j, r_{i,j})$  to some bidder  $i$ , where  $b_j$  is the current lower bound. As we try out all possible orderings, this bidder  $i$  will also be selected first in an ordering and hence obtain item  $j$  at price  $p_j$ . To ensure that bidder  $i$  does not prefer a different item we must have

increased the price bounds of all items  $j' \neq j$  to  $b_{j'} = \max(b_{j'}, u_{i,j'}^{-1}(u_{i,j}(p_j)))$  for any stable solution in which bidder  $i$  gets item  $j$  at price  $p_j$ . These are exactly the lower price bounds the algorithm ensures. Given these lower price bounds and the fact that prices only increase, bidder  $i$  is stable for the prices computed for this ordering. Hence we can remove bidder  $i$  and are left with a new instance with  $n$  bidders and one less item. By induction, we find the lowest feasible and stable prices for this problem. As (i) we try out *all* possible orderings of bidders and (ii) for the correct ordering we obtain the lowest feasible and stable prices  $p$  respecting the initial lower bound  $b_j = \min_i r_{i,j}$  for all  $j$ , we obtain the theorem.  $\square$

Theorem 2 together with Lemma 1 shows that the algorithm outputs a bidder optimal outcome.

**Running Time.** There are  $O(k!(n+k)^k)$  different matchings of  $n$  bidders to up to  $k$  items and there are  $O(k!) = O(k^k)$  permutations of the up to  $k$  properly matched bidders. Computing the price updates for a given matching-bidder pair takes time  $O(nk)$ . Hence the overall running time is  $O((n+k)^k \cdot k^{2k+1} \cdot n)$ .

## References

1. Shapley, L.S., Shubik, M.: The assignment game: The core I. *Journal of Game Theory* 29, 111–130 (1972)
2. Roth, A.E., Sotomayor, M.A.O.: Two-sided matching: A study in game-theoretic modeling and analysis. Cambridge University Press, Cambridge (1990)
3. 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)
4. Lahaie, S., Pennock, D.M., Saberi, A., Vohra, R.V.: Sponsored Search Auctions. In: *Algorithmic Game Theory*, pp. 699–716. Cambridge University Press, Cambridge (2007)
5. Leonard, H.B.: Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy* 91, 461–479 (1983)
6. Demange, G., Gale, D., Sotomayor, M.: Multi-item auctions. *Journal of Political Economy* 94(4), 863–872 (1986)
7. Aggarwal, G., Muthukrishnan, S., Pál, D., Pál, M.: General auction mechanism for search advertising. In: *World Wide Web Conference*, pp. 241–250 (2009)
8. Crawford, K., Knoer: Job matching with heterogenous firms and workers. *Econometrica* 49, 437–450 (1981)
9. Quinzii, M.: Core and competitive equilibria with indivisibilities. *International Journal of Game Theory* 13(1), 41–60 (1984)
10. Demange, G., Gale, D.: The strategy structure of two-sided matching markets. *Econometrica* 53(4), 873–888 (1985)
11. Alkan, A.: Existence and computation of matching equilibria. *European Journal of Political Economy* 5, 285–296 (1989)
12. Alkan, A., Gale, D.: The core of the matching game. *Games and economic behavior* 2(3), 203–212 (1990)
13. Cramton, P., Shoham, Y., Steinberg, R.: *Combinatorial Auctions*. MIT Press, Cambridge (2006)
14. Hall, P.: On representatives of subsets. *Journal of the London Mathematical Society* 10, 26–30 (1935)

# A New Ranking Scheme of the GSP Mechanism with Markovian Users<sup>\*</sup>

Xiaotie Deng<sup>1</sup> and Jiajin Yu<sup>2,\*\*</sup>

<sup>1</sup> City University of Hong Kong  
csdeng@cityu.edu.hk

<sup>2</sup> Fudan University  
jiajinyu@fudan.edu.cn

**Abstract.** Sponsored search auction is used by most search engines to select ads to display on the web page of a search result, according to advertisers' bidding prices. The income of this targeted advertising business is a big part of the revenue of most search engines. The most widely used approach to choose ads is the Generalized Second Price (GSP) auction, choosing the  $i$ -th highest bidder to display at the  $i$ -th most favorable position and charging the  $(i + 1)$ -st highest bidding price. Most previous works about GSP auction are based on the separation assumption: the probability a user will click on an ad is composed of two independent parts: a quality factor of the ad itself and a position factor of the slot of the ad. The previous model does not include the externality an ad may bring to the other ads. We study a GSP auction in a Markovian user model where the externality is considered by modeling a user's probability behavior when he reads ad list. In particular, we propose a new ranking scheme for the bidders. We prove Nash equilibrium always exists in the auction and study the efficiency of the auction by theoretical analysis and simulation. We compare our results with social optimum and previous approaches. Comparison shows that our scheme approximates the social optimum and improves previous approaches under various circumstances.

## 1 Introduction

Targeted advertising with search results is *the* major source of income for most search engines. Here is a common scenario. A user submits a query to the search engine. The search engine returns a web page of search results which is shown in the user's web browser. Besides the search results, several query-related ads known as the *sponsored links* are displayed from top to bottom on the right side of the web page. If a user clicks an ad, he will be navigated to the advertiser's site. The search engine will get paid by the advertiser for this click.

---

\* Work reported in this article is supported by a research grant (Project No: CityU 112909) of Research Grants Council of Hong Kong SAR.

\*\* Work of this author was also supported by a grant from the Shanghai Leading Academic Discipline Project (project number B114) and the Shanghai Committee of Science and Technology of China (no. 08DZ2271800).



For a particular query or keyword, many advertisers would want to have a slot in the sponsored links. Most search engines use an auction to select the advertisers that get slots and the positions of their ads. An approach known as *Generalized Second Price* (GSP) auction has been widely used by search engines like Google, Yahoo!, MSN and has been studied extensively in the academia [1,2,3]. In a GSP auction, each bidder submits a bid representing the maximum amount he is willing to pay for a click. The search engine ranks bidders in the decreasing order of their bids, and the ads of the first  $k$  bidders in this ranking will be displayed along the search results. The order of the ads from top to bottom is the same as the decreasing order of the bids. It is known that Nash equilibrium exists in the GSP auction [2,3].

Most previous works of sponsored search auction are based on the separation assumption: the probability that an ad will get a click is composed of two separate parts [1], a quality factor based on the advertiser itself, and a position factor based on the position of the ad. This assumption does not model the impact an ad in a higher slot might have on an ad in a lower slot. Recently a Markovian user model is proposed [4,5] to capture the externality that one ad can bring to other ads. Basically the idea is that when a user sees a particular ad, not only he will click the ad based on the ad's quality, but the ad will also impact the user's willingness to read the ads below. More formally, in this model, when a user sees an ad  $i$ , he will click it with a probability  $r_i$ . The user may also want to read other ads below ad  $i$ . The model uses another parameter  $q_i$  of ad  $i$  to model the probability that a user will scan down. Based on this model, an auction selects several winning advertisers, puts them on the ad slots, and charges them when clicks occur.

*Our Work.* In this paper, we study a GSP-like auction under a Markovian user model. In the spirit of GSP (ie. ranking bidders and selecting the first  $k$  ones), we propose a particular ranking of advertisers based on their bids and parameters ( $r_i$  and  $q_i$ ) to approximate optimal efficiency. We prove that the auction always has Nash equilibrium by an explicit characterization of one Nash equilibrium in the auction. To show that the auction approximates optimal efficiency, we show the efficiency of our equilibrium is no less than a factor of the optimal efficiency. We also compare our auction to a previous one [6] and show that the efficiency of our auction can be lower bounded by a factor of the efficiency of one equilibrium in the previous one. We supplement our theoretical analysis by simulation results, which show our auction approximates the optimum and improves the previous one under various conditions.

*Related Work.* Aggarwal et.al [4], and Kempe and Mahdian [5] independently proposed the sponsored search auction with Markovian users. Both of their works mainly focused on how to compute the optimal solution with respect to the bidders. Giotis and Karlin [6] showed that when advertisers are ordered decreasingly by  $r_i b_i$ , Nash equilibrium always exists in the GSP auction with Markovian users. They also gave price of anarchy and price of stable analysis of the efficiency of their GSP auction.



## 2 Model

Let  $B = \{1, 2, \dots, n\}$  be a set of bidders (or advertisers). For each bidder  $i$ , we use two parameters to model the behavior of a user when he read bidder  $i$ 's ad. One is the "quality" factor  $r_i$ , the probability that a user clicks the ad after he sees it. The other is the "externality" factor  $q_i$ , the probability that a user will read the ads below no matter whether he clicks bidder  $i$ 's ad or not. A user starts reading the ad at the top of the sponsored links and clicks bidder  $i$ 's ad with probability  $r_i$ . Then the user has probability  $q_i$  to look the ads below the current one, no matter he clicks the current ad or not. That also means the user will leave the sponsored links area with probability  $1 - q_i$  after finishing reading bidder  $i$ 's ad. Each bidder  $i$  also has a private value  $v_i$ : the expected value he can get when a user enters his website through a click on his ad.

When an auction starts, each bidder submits a bid  $b_i$  that is the maximum amount he is willing to pay for one click. Given the bids and the parameters of the bidders, the auction selects some of them to put on the ad slots of the web page. Assume we have  $k$  slots in a web page. The auction computes an assignment, an injective function  $\pi : \{1, \dots, k\} \rightarrow B$ , such that bidders  $\pi(1), \pi(2), \dots, \pi(k)$  are assigned ad slots from 1 to  $k$  (slots are numbered from top to bottom). According to the model described above, the expected value of the bidder at slot  $j$  is

$$\left( \prod_{i=1}^{j-1} q_{\pi(i)} \right) r_{\pi(j)} v_{\pi(j)}.$$

We define the *efficiency* of an assignment  $\pi$  as

$$V(\pi) = r_{\pi(1)} v_{\pi(1)} + q_{\pi(1)} r_{\pi(2)} v_{\pi(2)} + \dots + \left( \prod_{i=1}^{k-1} q_{\pi(i)} \right) r_{\pi(k)} v_{\pi(k)} \quad (1)$$

In the GSP auction, search engine ranks bidders by their biddings decreasingly and assigns ad slots to the first  $k$  bidders [3][2]. In the weighted GSP auction [1] bidders are ranked decreasingly by  $w_i b_i$ , in which  $w_i$  can be viewed as the "quality" of the bidder  $i$ . In this paper, we study a particular ranking scheme in the Markovian user model. The search engine ranks bidders by  $r_i b_i / (1 - q_i)$  decreasingly and selected the first  $k$  bidders to put on the ad slots. Each bidder  $\pi(i)$  displayed on the web page shall pay

$$p_{\pi(i)} = \frac{r_{\pi(i+1)} b_{\pi(i+1)} (1 - q_{\pi(i)})}{(1 - q_{\pi(i+1)}) r_{\pi(i)}} \quad (2)$$

for one click, which is the minimal value he must bid to keep his current position in the assignment. Given the assignment  $\pi$ , the expected utility of bidder at slot  $i$  is

$$u_{\pi(i)} = \left( \prod_{i=1}^{k-1} q_{\pi(i)} \right) r_{\pi(i)} (v_{\pi(i)} - p_{\pi(i)}) \quad (3)$$

### 3 Nash Equilibrium

In order to simplify the notation, we write  $f_i = r_i v_i / (1 - q_i)$  as the *adjusted value*, and  $g_i = r_i b_i / (1 - q_i)$  as the *adjusted bids*. In this section, we number bidders in the decreasing order of their adjusted values. We find a Nash equilibrium that always exists in the auction. The order of the bidders in the equilibrium is the same as the decreasing order of their adjusted value.

**Theorem 1.** *A Nash equilibrium exists in this auction where the adjusted bids are*

$$g_i = \begin{cases} f_i & \text{if } i = 1 \\ f_{i-1}(1 - q_i) + g_{i+1}q_i & \text{if } 1 < i \leq k \\ f_i & \text{if } i > k \end{cases} \tag{4}$$

*Proof.* The recursive definition of adjusted bids in Equation 4 is similar to the ones found in [13], a convex combination of value and bid. Here value and bid are adjusted by bidder’s parameter. In most cases, the bid of one agent depends on the bid of the agent immediately below him and the value of the agent immediately above him.

It can be verified that by the definition of Equation 4, adjusted bids are sorted in the decreasing order and each bidder has non-negative utilities. In a Nash equilibrium, any bidder who gets an ad slot does not want be placed in another position by changing his bidding. To write it explicitly, the adjusted values and adjusted bids should satisfy the following constraints where  $1 \leq s < t \leq k$ .

$$f_s - g_{s+1} \geq \left( \prod_{i=s+1}^t q_i \right) (f_s - g_{t+1}) \quad f_t - g_s \leq \left( \prod_{i=s}^{t-1} q_i \right) (f_t - g_{t+1})$$

The verification of these two inequalities are straightforward by following the recursive definition of adjusted bids.

For any bidder  $t$  who does not get a slot in the equilibrium, to win a slot with nonnegative utility, their adjusted values  $f_t$  must be no less than any adjusted winning bids  $g_i$ . Recall the least among them is  $g_k = (1 - q_k)f_{k-1} + q_k f_{k+1} \geq f_t$  for any  $t > k$ . Therefore bidders who are not assigned ad slots have no incentive to raise their bids. □

### 4 Efficiency Analysis

In this section, we will compare the efficiency of one equilibrium in our auction to the optimal efficiency and the efficiency of one equilibrium in the auction proposed previously in [6]. There might exist many equilibria in our auction (so as in the auction in [6]), but here we only study the one proved to exist in Theorem 1. Let the assignment of this equilibrium be  $\pi_a$ . We denote the optimal assignment as  $\pi^*$  and the assignment where bidders are ordered by  $r_i v_i$  as  $\pi_b$ . Giotis and Karlin proved that assignment  $\pi_b$  is an equilibrium [6] when bidders

are ordered by  $r_i b_i$  decreasingly and payment of a click on the ad at slot  $i$  is  $r_{i+1} b_{i+1} / r_i$ . In the following discussion, we sometimes treat an assignment  $\pi$  as a sequence of bidders to simplify the discussion. We will use  $O_a$  as the set of bidders that are assigned slots in assignment  $\pi_a$ . Set  $O^*, O_b$  are defined similarly. And we write  $\phi_i = \prod_{j=1}^{i-1} q_j$ . First we list two lemmas about the structure of the optimal solution proved before.

**Lemma 1.** [4] *In the optimal assignment  $\pi^*$ , the bidders are sorted by the decreasing order of  $r_i v_i / (1 - q_i)$ , but not necessarily the first  $k$  greatest ones.*

**Lemma 2.** [4] *For any bidder  $i$  that is assigned an ad slot in the optimal solution, if some bidder  $j$  is not in the assignment, while  $r_j v_j \geq r_i v_i$  and  $r_j v_j / (1 - q_j) \geq r_i v_i / (1 - q_i)$ , then replacing bidder  $i$  by bidder  $j$  will get an assignment whose efficiency is not worse.*

The following lemma is implied in the proof of Lemma 1 in [4]. We will use this property quite often in the following.

**Lemma 3.** [4] *Let bidder  $i$  and  $j$  be two adjacent bidders in an assignment  $\pi$ , where  $\pi(i) < \pi(j)$  and  $r_i v_i / (1 - q_i) \leq r_j v_j / (1 - q_j)$ . Swapping the positions of two bidders will not decrease the efficiency of the assignment.*

**Proposition 1**

$$V(\pi_a) \geq \frac{\prod_{i \in O_a \setminus O^*} (1 - q_i)}{\prod_{j \in O^* \setminus O_a} (1 - q_j)} V(\pi^*)$$

*Proof.* Let bidder  $i$  be the first one in  $\pi_a$  but not in  $\pi^*$ . By Lemma 3, by swapping the position of bidder  $i$  and the position of the bidder below him, we decrease the efficiency of the assignment. We swap bidder  $i$  like this until it reaches at the end of the sequence. The efficiency at this time is

$$V_1 = r_1 v_1 + \dots + \phi_{i-1} r_{i-1} v_{i-1} + \sum_{s=i+1}^k (\phi_s / q_i) r_s v_s + (\phi_{k+1} / q_i) r_i v_i,$$

and  $V(\pi_a) \geq V_1$ . Let bidder  $j$  be the first one in  $\pi^*$  but not in  $\pi_a$ , then  $r_i v_i / (1 - q_i) \geq r_j v_j / (1 - q_j)$ , and we have

$$V_2 = r_1 v_1 + \dots + \phi_{i-1} r_{i-1} v_{i-1} + \sum_{s=i+1}^k (\phi_s / q_i) r_s v_s + (\phi_{k+1} / q_i) r_j v_j \frac{1 - q_i}{1 - q_j} \leq V_1.$$

Since bidder  $i$  is not in the optimum, we have  $r_j v_j \geq r_i v_i$  by Lemma 2. Therefore

$$V_3 = \frac{1 - q_i}{1 - q_j} \left( r_1 v_1 + \dots + \phi_{i-1} r_{i-1} v_{i-1} + \sum_{s=i+1}^k (\phi_s / q_i) r_s v_s + (\phi_{k+1} / q_i) r_j v_j \right) \leq V_2.$$

We now get a new assignment where all bidders below bidder  $i$  in assignment  $\pi_a$  move up one slot and bidder  $j$  is at the bottom slot. Since bidder  $j$  is not

in  $O_a$ ,  $r_j v_j / (1 - q_j)$  is the smallest among  $k$  bidders in the new assignment. Therefore in the new assignment these  $k$  bidders are again sorted decreasingly by  $r_i v_i / (1 - q_i)$ . Starting from this new assignment, we again replace one bidder  $i'$  in  $O_a \setminus O^*$  by a bidder  $j'$  in  $O^* \setminus O_a$  in the same way, and get the assignment whose efficiency is no greater than a  $(1 - q_{i'}) / (1 - q_{j'})$ -factor of  $V_3$ . In this way, we will get the exactly optimal assignment at last.  $\square$

We give two examples to show our bound is almost tight, even in the restricted setting [6] where bidders never submit bids exceeding their true valuations.

*Example 1.* The auction has two bidders and one slot. Bidder 1 has parameters  $r_1 v_1 = 1, 1 - q_1 = \epsilon < 1$ . Bidder 2 has parameters  $r_2 v_2 = X > 1, 1 - q_2 = \epsilon X + \delta < 1$ . Note that  $r_1 v_1 / (1 - q_1) \geq r_2 v_2 / (1 - q_2)$ . Therefore the Nash equilibrium will have efficiency 1; while the efficiency of optimal solution is  $X$ .

*Example 2.* Bad efficiency can also happen in the restricted setting [6], where bidders never submit bids exceeding their true valuations. Consider the auction of two bidders and one slot in Example [1] in the restricted setting. If the adjusted bid of bidder 1 is lower than the adjusted bid of bidder 2, bidder 1 always has an incentive to bid higher to win the auction. Bidder 1 can ensure his winning by making his adjusted bidding  $r_1 b_1 / (1 - q_1)$  greater than  $r_2 v_2 / (1 - q_2)$ . This again gives an equilibrium whose efficiency is 1 while the optimal is  $X$ .

Now we prove that the efficiency of assignment  $\pi_a$  can also be lower bounded by a factor of efficiency of assignment  $\pi_b$ . The proof is similar to the proof of Proposition [1], though a little modification is needed to deal with the order of bidders in assignment  $\pi_b$ .

**Proposition 2**

$$V(\pi_a) \geq \frac{\prod_{j \in O_a \setminus O_b} (1 - q_j)}{\prod_{i \in O_b \setminus O_a} (1 - q_i)} V(\pi_b)$$

*Proof.* By Lemma [3], we know whenever there is a bidder  $i$  whose  $r_i v_i / (1 - q_i)$  is smaller than the bidder below him, we can swap positions of these two and increase the efficiency. Therefore, starting from assignment  $\pi_b$ , we sort the bidders in  $\pi_b$  in the decreasing order of  $r_i v_i / (1 - q_i)$ . By then we get an assignment  $\pi'_b$  with more efficiency. Then we use the approach in the proof of Proposition [1]. Delete bidder  $i$  who is the first one in the new assignment  $\pi'_b$  but not in  $\pi_a$  and add bidder  $j$  who is the first one in  $\pi_a$  but not in  $\pi_b$  to the end of the sequence. Here we need show that  $r_j v_j / (1 - q_j)$  is no less than  $r_i v_i / (1 - q_i)$  and  $(1 - q_i) / (1 - q_j)$  is no less than 1. First note that bidder  $i$  is in  $\pi_b$  but bidder  $j$  is not, therefore  $r_i v_i \geq r_j v_j$ . For assignment  $\pi_a$ , since bidder  $i$  is not in  $O_a$  but bidder  $j$  is,  $r_i v_i / (1 - q_i) \leq r_j v_j / (1 - q_j)$  and  $(1 - q_i) / (1 - q_j) \geq 1$ .

The efficiency of the new assignment increases at least a factor of  $(1 - q_i) / (1 - q_j)$  by a similar argument to the proof of Proposition [1]. To maintain the non-decreasing order of  $r_i v_i / (1 - q_i)$  of the assignment, we might need to move up bidder  $j$  some slots, but it will not decrease the efficiency. Then we can do the next round until no bidder can be swapped out. Each time, we increase the efficiency by at least a factor of  $(1 - q_i) / (1 - q_j)$ .  $\square$

**Proposition 3.**  $V(\pi_a) \leq kV(\pi_b)$ .

*Proof.* Let bidder  $y$  be the one with the greatest  $r_i v_i$  in the assignment  $\pi_a$ . We have  $V(\pi_a) = r_{\pi_a(1)} v_{\pi_a(1)} + \dots \leq k r_y v_y$ . Let bidder  $x$  be the one with the greatest  $r_i v_i$  among all bidders. By the definition of  $\pi_b$ , bidder  $x$  is at the first slot in its ranking. Therefore  $V(\pi_a) \leq k r_y v_y \leq k r_x v_x \leq k V(\pi_b)$ .  $\square$

## 5 Simulations

In the previous section, we lower bounded the efficiency of an equilibrium that exist in our auction by a factor of the optimal efficiency. We also lower bounded the efficiency of that equilibrium by a factor of the efficiency of a similar equilibrium in the auction proposed in [6]. In this section, we compare these efficiencies by simulation. By running simulated auctions, we first evaluate the ratio  $V(\pi_a)/V(\pi^*)$ , then we evaluate the ratio  $V(\pi_a)/V(\pi_b)$ .


### 5.1 Simulation Parameters

In every single simulation, we run an auction with 50 bidders and 8 slots. Each bidder has three parameters,  $v_i, r_i, q_i$ , which are drawn independently from a normal distribution with mean 0.5. The deviation of the distribution is what we want to study. As we saw in the examples in the previous section, if the  $q_i$  in the different assignments vary a lot, the lower bound can be very bad. Therefore we run simulations with different deviations (but each time, all three parameters have the same deviation). We want to see how the ratio changes under different deviations. For each deviation, we run the simulation 200 times and take the average as the ratio under this deviation. The range of the deviation is  $[0, 0.3]$  and we take 100 deviations uniformly on this range.

### 5.2 Results

**Efficiency of Assignment  $\pi_a$  and Optimum.** Figure (1a) plots  $V(\pi_a)/V(\pi^*)$  under different deviations. From the figure, we see that when the deviation is small, the efficiency of assignment  $\pi_a$  is nearly as good as the optimal assignment. When the deviation becomes larger, the efficiency of assignment  $\pi_a$  begins to drop, but not significantly. At last when the deviation goes to 0.3, the efficiency of assignment  $\pi_a$  is still more than 90% of the optimal assignment.

**Efficiency of Assignment  $\pi_a$  and  $\pi_b$ .** Figure (1b) plots  $V(\pi_a)/V(\pi_b)$  under different deviations. From the figure, we see that if the deviation is small, assignment  $\pi_a$  and assignment  $\pi_b$  have basically the same efficiency. When the deviation is large,  $V(\pi_a)$  becomes greater than  $V(\pi_b)$ . When the deviation comes to 0.3,  $V(\pi_a)$  is greater than  $V(\pi_b)$  for about 50% percent. The ratio does not seem to converge when deviation becomes large. From Figure 1b, we also see that  $V(\pi_a)$  is greater than  $V(\pi_b)$  on average under all deviations, although the lower bound in Proposition 2 can be arbitrary small. The result suggests the new auction might have better efficiency than the auction in [6] on average under some circumstances.



**Fig. 1.** Simulation Results

## 6 Conclusion

In this paper, we proposed a new ranking of bidders in the GSP-like auction with Markovian users. We proved GSP auction with this new ranking, where bidders are ordered by  $r_i v_i / (1 - q_i)$ , must have an equilibrium. We lower bounded the efficiency of one equilibrium in our auction by a factor of the optimal efficiency and a factor of the efficiency of an equilibrium of a GSP auction proposed in [6]. We did some simulations whose results suggested the new ranking approximates social optimum and improves the previous scheme in various circumstances.

## References

1. Aggarwal, G., Goel, A., Motwani, R.: Truthful Auctions for Pricing Search Keywords. In: Proceedings 7th ACM Conference on Electronic Commerce (EC 2006), Ann Arbor, Michigan, USA, June 11-15, 2006, pp. 1–7 (2006)
2. 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)
3. Varian, H.R.: Position auctions. *International Journal of Industrial Organization* 25(6), 1163–1178 (2007)
4. Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M.: Sponsored Search Auctions with Markovian Users. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 621–628. Springer, Heidelberg (2008)
5. Kempe, D., Mahdian, M.: A Cascade Model for Externalities in Sponsored Search. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 585–596. Springer, Heidelberg (2008)
6. Giotis, I., Karlin, A.R.: On the Equilibria and Efficiency of the GSP Mechanism in Keyword Auctions with Externalities. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 629–638. Springer, Heidelberg (2008)

# Mediated Equilibria in Load-Balancing Games<sup>★</sup>

Joshua R. Davis<sup>1</sup>, David Liben-Nowell<sup>1</sup>, Alexa Sharp<sup>2</sup>, and Tom Wexler<sup>2</sup>

<sup>1</sup> Carleton College, Northfield, MN  
joshuardavis@q.com, dlibenno@carleton.edu

<sup>2</sup> Oberlin College, Oberlin, OH  
{alexa.sharp,tom.wexler}@oberlin.edu

**Abstract.** *Mediators* are third parties to whom the players in a game can delegate the task of choosing a strategy; a mediator forms a *mediated equilibrium* if delegating is a best response for all players. Mediated equilibria have more power to achieve outcomes with high social welfare than Nash or correlated equilibria, but less power than a fully centralized authority. Here we begin the study of the power of mediation by using the mediation analogue of the price of stability—the ratio of the social cost of the best mediated equilibrium BME to that of the socially optimal outcome OPT. We focus on load-balancing games with social cost measured by weighted average latency. Even in this restricted class of games, BME can range from as good as OPT to no better than the best correlated equilibrium. In unweighted games BME achieves OPT; the weighted case is more subtle. Our main results are (1) that the worst-case ratio BME/OPT is at least  $(1 + \sqrt{2})/2 \approx 1.2071$  (and at most  $1 + \phi \approx 2.618$  [3]) for linear-latency weighted load-balancing games, and that the lower bound is tight when there are two players; and (2) tight bounds on the worst-case BME/OPT for general-latency weighted load-balancing games. We also give similarly detailed results for other natural social-cost functions.

## 1 Introduction

The recent interest in algorithmic game theory by computer scientists is in large part motivated by the recognition that the implicit assumptions of traditional algorithm design are ill-suited to many real-world settings. Algorithms are typically designed to generate solutions that can be implemented by some centralized authority. But often no such centralized authority exists; solutions arise through the interactions of self-interested, independent agents. Thus researchers have begun to use game theory to model these competitive, decentralized situations.

One classic example is the paper of Koutsoupias and Papadimitriou [16], who consider the effect of decentralizing a standard load-balancing problem. In the resulting game, each job is controlled by a distinct player who selects a machine to serve her job so as to minimize delay. The authors compare the social cost

---

\* All authors were supported in part by NSF grant CCF-0728779. This work was also supported by grants from Carleton College, Oberlin College, and Denison University. Thanks to Bobby Kleinberg for suggesting this line of inquiry to us.

(expected maximum delay) of the Nash equilibria of this game to that of a centrally designed optimal solution. The maximum of these ratios is the *price of anarchy* of the game, quantifying the worst-case cost of decentralized behavior.

One can imagine a continuum indexing the amount of power that a centralized authority has in implementing solutions to a given problem, from utter impotence (leading to a potentially inefficient Nash equilibrium) to dictatorial control (leading to the socially optimal outcome). For example, consider a weak authority who can *propose* a solution simultaneously to all players, but who has no power to enforce it. The players would agree to such a proposal only if it were a Nash equilibrium—but the authority could propose the *best* Nash equilibrium. The ratio of the cost of the best Nash equilibrium to the global optimum is the *price of stability*, which may be much better than the price of anarchy.

A *correlator* is a more powerful authority, in that it is not required to broadcast the entire proposed solution; it signals each player individually with a suggested action, chosen from some known joint probability distribution. The resulting stable outcomes are called *correlated equilibria* [2]. Any Nash equilibrium is a correlated equilibrium, but often a correlator can induce much better outcomes.

A *mediator* [1,19,21,22,23,24,26] is an authority who offers to act on behalf of the players; any player may *delegate* to the mediator the responsibility of choosing a strategy. In a *mediated equilibrium*, all players prefer to delegate than to play on their own behalf. The strategies that the mediator selects for the delegating players may be correlated; moreover, the distribution from which the mediator draws these strategies may depend on which players have opted to delegate. A mediator can enforce an equilibrium by threatening to have the delegating players “punish” any player departing from mediation. Any correlated equilibrium can be represented as a mediated equilibrium, but the converse is not true; mediators are more powerful than correlators.

*The present work: mediated load-balancing games.* In this paper, we begin to quantify the powers and limitations of mediators. We consider the mediation analogue of the price of stability: how much less efficient than the globally optimal outcome OPT is the best mediated equilibrium BME? (While one could ask questions analogous to the price of anarchy instead, the spirit here is that of a well-intentioned central authority who would aim for the best, not the worst, outcome within its power.) We initiate this study in the context of *load-balancing games*. Each player  $i$  controls a job that must be assigned to a machine. Each machine  $j$  has a nonnegative, nondecreasing latency function  $f_j(x)$ , and each player incurs a *cost* of  $f_j(\ell_j)$  for choosing machine  $j$ , where  $\ell_j$  is the total load of jobs on machine  $j$ . We split load-balancing games into classes along two dimensions:

- *unweighted* vs. *weighted*: in weighted games, job  $i$  has weight  $w_i$  and experiences cost  $f_j(\sum_{i' \text{ uses } j} w_{i'})$  on machine  $j$ ; in unweighted games all  $w_i = 1$ .
- *linear* vs. *general* latencies: in linear games,  $f_j(x) = a_j \cdot x$  for  $a_j \geq 0$ ; for general latencies  $f_j$  can be an arbitrary nonnegative, nondecreasing function.

The social cost is measured by the weighted average latency experienced by the jobs; see Section 6 for results using other social cost functions.



	unweighted jobs	weighted jobs
<b>linear latencies</b>	BME = OPT [19] BCE $\leq 4/3 \cdot$ BME [tight] (Lemma 1 [4])	BME $\leq 2.618 \cdot$ OPT, BCE $\leq 2.618 \cdot$ BME [3] $n = 2$ : BME $\leq 1.2071 \cdot$ OPT (Thm. 2) BCE $\leq 4/3 \cdot$ BME (Thm. 2) [both tight for $n = 2$ ]
<b>general latencies</b>	BME = OPT [19] BCE $\leq n \cdot$ BME [tight] (Lemma 1)	BME $\leq \Delta \cdot$ OPT [tight] (Thm. 3) BCE $\leq \Delta \cdot$ BME [tight] (Thm. 3)

**Fig. 1.** Summary of our results for weighted-average-latency social cost. Here OPT is the socially optimal outcome, BME (BCE) the best mediated (correlated) equilibrium,  $n$  the number of jobs, and  $\Delta$  the ratio of total job weight to smallest job weight.

Load-balancing games are appealing for this work for two reasons. First, they include cases in which mediators can achieve OPT and cases in which they cannot even better the best Nash equilibrium BNE. Second, the prices of anarchy and stability, and corresponding measures of correlated equilibria, are well understood for these games and many of their variants [3,4,5,6,15,16,17,25]. Most relevant for what follows are an upper bound of  $1 + \phi \approx 2.618$  on the price of anarchy in weighted linear games [3] and a tight upper bound of  $4/3$  on the price of stability in unweighted linear games [4].

We extend this line of work to mediated equilibria with the following results. (Figure 1 summarizes those for the weighted-average-latency social cost.)

- In the unweighted case, the BME is optimal, regardless of the latency functions' form. This result follows from a recent theorem of Monderer and Tennenholtz [19], which in fact holds for any symmetric game. See Section 3.
- In weighted linear-latency games with two players, we give tight bounds on the best solution a mediator can guarantee: a factor of  $(1 + \sqrt{2})/2 \approx 1.2071$  worse than OPT but  $4/3$  better than the best correlated equilibrium BCE. Thus mediators lie strictly between dictators and correlators. See Section 4.
- In weighted nonlinear-latency games, mediated equilibria provide no worst-case improvement over correlated or even Nash equilibria. See Section 5.
- We also analyze mediation under two other social cost functions that have been considered in the literature: (i) the maximum latency of the jobs; and (ii) the average latency, unweighted by the jobs' weights. See Section 6.

*Related work.* Koutsoupias and Papadimitriou initiated the study of the price of anarchy in load-balancing games, considering weighted players, linear latencies, and the maximum (rather than average) social cost function [16]. A substantial body of follow-up work has improved and generalized their initial results [6,7,9,18]. See [12] and [27] for surveys. A second line of work takes social cost to be the sum of players' costs. Lücking et al. [11,17] measure the price of anarchy of mixed equilibria in linear and convex routing games in this setting. Awerbuch et al. [3] consider both the unweighted and weighted cases on general networks. Suri et al. [15,25] examine the effects of asymmetry in these games.

Caragiannis et al. [4] give improved bounds on the price of anarchy and stability. Christodoulou and Koutsoupias [5,6] bound the best- and worst-case correlated equilibria in addition to improving existing price of anarchy and stability results.

Other aspects of correlated equilibria have been explored recently, including their existence [13] and computation [13,14,20]. Mediated equilibria have developed in the game theory literature over time; see Tennenholtz [26] for a summary. Mediated equilibria have been studied for position auctions [1], for network routing games [23,24], and in the context of social choice and voting [21,22]. Strong mediated equilibria have also been considered [19,24].

## 2 Notation and Background

An  $n$ -player,  $m$ -machine load-balancing game is defined by a nondecreasing latency function  $f_j : [0, \infty) \rightarrow [0, \infty]$  for each machine  $j \in \{1, \dots, m\}$ ; and a weight  $w_i > 0$  for each player  $i \in \{1, \dots, n\}$ . We consider games in which every job has access to every machine: a pure strategy profile  $\mathbf{s} = \langle s_1, \dots, s_n \rangle$  can be any element of  $\mathcal{S} := \{1, \dots, m\}^n$ . The load  $\ell_j$  on a machine  $j$  under  $\mathbf{s}$  is  $\sum_{i:s_i=j} w_i$ , and the latency of machine  $j$  is  $f_j(\ell_j)$ . The cost  $c_i(\mathbf{s})$  to player  $i$  under  $\mathbf{s}$  is  $f_{s_i}(\ell_{s_i})$ . Pure Nash equilibria exist in all load-balancing games [10,8]. A load-balancing game is linear if each  $f_j$  is of the form  $f_j(x) = a_j \cdot x$  for some  $a_j \geq 0$  and unweighted if each  $w_i = 1$ . Machine  $j$  is dominated by machine  $j'$  for player  $i$  if, no matter what machines the other  $n - 1$  players use, player  $i$ 's cost is lower using machine  $j'$  than using machine  $j$ .

A nonempty subset of the players is called a coalition. A mediator is a collection  $\Psi$  of probability distributions  $\psi_T$  for each coalition  $T$ , where the probability distribution  $\psi_T$  is over pure strategy profiles for the players in  $T$ . The mediated game  $M_\Gamma^\Psi$  is a new  $n$ -player game in which every player either participates in  $\Gamma$  directly by choosing a machine in  $S := \{1, \dots, m\}$  or participates by delegating. That is, the set of pure strategies in  $M_\Gamma^\Psi$  is  $Z = S \cup \{s_{\text{med}}\}$ . If the set of delegating players is  $T$ , then the mediator plays the correlated strategy  $\psi_T$  on behalf of the members of  $T$ . In other words, for a strategy profile  $\mathbf{z} = \langle z_1, z_2, \dots, z_n \rangle$  where  $T := \{i : z_i = s_{\text{med}}\}$  and  $\bar{T} := \{i : z_i \neq s_{\text{med}}\} = \{i : z_i \in S\} = \{1, \dots, n\} - T$ , the mediator chooses a strategy profile  $\mathbf{s}_T$  according to the distribution  $\psi_T$ , and plays  $s_i$  on behalf of every player  $i \in T$ ; meanwhile, each player  $i$  in  $\bar{T}$  simply plays  $z_i$ . The expected cost to player  $i$  under the strategy profile  $\mathbf{z}$  is then given by  $c_i(\mathbf{z}) := \sum_{\mathbf{s}_T} c_i(\mathbf{s}_T, \mathbf{z}_{\bar{T}}) \cdot \psi_T(\mathbf{s}_T)$ . (The mediators described here are called minimal mediators in [19], in contrast to a seemingly richer class that allow more communication from players to the mediator.)

A mediated equilibrium for  $\Gamma$  is a mediator  $\Psi$  such that the strategy profile  $\langle s_{\text{med}}, s_{\text{med}}, \dots, s_{\text{med}} \rangle$  is a pure Nash equilibrium in  $M_\Gamma^\Psi$ . Every probability distribution  $\psi'$  over the set of all pure strategy profiles for  $\Gamma$  naturally corresponds to a mediator  $\Psi$ , where the probability distribution  $\psi_T$  for a coalition  $T$  is the marginal distribution for  $T$  under  $\psi'$ —that is,  $\psi_T(\mathbf{s}_T) = \sum_{\mathbf{s}' : \mathbf{s}'_T = \mathbf{s}_T} \psi'(\mathbf{s}')$ . If  $\psi'$  is a correlated equilibrium then this  $\Psi$  is a mediated equilibrium.

The social cost of a strategy profile  $\mathbf{s}$  is the total (or, equivalently, average) cost of the jobs under  $\mathbf{s}$ , weighted by their sizes—that is,  $\sum_i w_i \cdot c_i(\mathbf{s})$ . (We discuss

other social cost functions in Section 6.) We denote by OPT the (cost of the) profile  $\mathbf{s}$  that minimizes the social cost. We denote the worst Nash equilibrium—the one that maximizes social cost—by WNE, and the best Nash (correlated, mediated) equilibrium by BNE (BCE, BME). Note that  $\text{OPT} \leq \text{BME} \leq \text{BCE} \leq \text{BNE} \leq \text{WNE}$  because every Nash equilibrium is a correlated equilibrium, etc. The *price of anarchy* is  $\text{WNE}/\text{OPT}$ , and the *price of stability* is  $\text{BNE}/\text{OPT}$ .

### 3 Unweighted Load-Balancing Games

Although the unweighted case turns out to have less interesting texture than the weighted version, we start with it because it is simpler and allows us to develop some intuition. We begin with an illustrative example:

*Example 1.* There are  $n$  unweighted jobs and two machines  $L$  and  $R$  with latency functions  $f_L(x) = 1 + \varepsilon$  for any load, and  $f_R(x) = 1$  for load  $x > n - 1$  and  $f_R(x) = 0$  otherwise.

For each player,  $R$  dominates  $L$ , so  $\langle R, R, \dots, R \rangle$  is the unique correlated and Nash equilibrium, with social cost  $n$ . Consider the following mediator  $\Psi$ . When all  $n$  players delegate, the mediator picks uniformly at random from the  $n$  strategy profiles in which exactly one player is assigned to  $L$ . When any other set of players delegates, those players are deterministically assigned to  $R$ . If all players delegate under  $\Psi$ , each player’s expected cost is  $(1 + \varepsilon)/n$ ; if any player deviates, then that player will incur cost at least 1. Thus  $\Psi$  is a mediated equilibrium. Its cost is only  $1 + \varepsilon$ , which is optimal, while  $\text{BNE} = \text{BCE} = n$ .

In fact, this “randomize among social optima” technique generalizes to all unweighted load-balancing games—in any such game,  $\text{BME} = \text{OPT}$ . This is a special case of a general theorem of Monderer and Tennenholtz [19] about mediated equilibria robust to deviations by coalitions. (See also [24].)

Example 1 shows that with nonlinear latency functions BCE may be much worse than OPT, even in the unweighted 2-machine case. But even linear unweighted load balancing has a gap between BCE and OPT, even in the 2-job, 2-machine case. The following example demonstrates the gap.

*Example 2 (Caragiannis et al. [4]).* There are two (unweighted) jobs and two machines  $L$  and  $R$  with latency functions  $f_L(x) = x$  and  $f_R(x) = (2 + \varepsilon) \cdot x$ .

Here  $\text{BCE} = 4$  and  $\text{OPT} = 3 + \varepsilon$ . (Machine  $L$  dominates  $R$ ; no player can be induced to use  $R$  in any correlated equilibrium.) We can show that this example is tight with respect to the gap between BME and BCE, using a result on linear unweighted load-balancing games of Caragiannis et al. [4] and the “randomize among social optima” mediation technique. We can also show a tight bound for unweighted nonlinear latency load-balancing games (details omitted for space).

**Lemma 1.** *In  $n$ -player unweighted load-balancing games:*

- for games with linear latency functions,  $\text{BCE} \leq 4/3 \cdot \text{BME}$ . This bound is tight.
- for not-necessarily-linear latency functions,  $\text{BCE} \leq n \cdot \text{BME}$ . This is tight.

We now have a complete picture for unweighted load balancing: a tight bound on the gap between BME and BCE and the theorem that  $\text{BME} = \text{OPT}$ .

## 4 Weighted Linear Load-Balancing Games

We now turn to weighted load-balancing games, where we find a richer landscape of results: among other things, cases in which BME falls strictly between BCE and OPT. We begin with the linear-latency case. (All proofs are omitted due to space.)

**Theorem 2.** *In any 2-machine, 2-job weighted game with linear latencies:*

1.  $\text{BCE}/\text{BME} \leq 4/3$ . This bound is tight for an instance with weights  $\{1, 1\}$  and with latency functions  $f_L(x) = x$  and  $f_R(x) = (2 + \varepsilon) \cdot x$ .
2.  $\text{BME}/\text{OPT} \leq \frac{1+\sqrt{2}}{2}$ . This bound is tight for an instance with weights  $\{1, 1 + \sqrt{2}\}$  and with latency functions  $f_L(x) = x$  and  $f_R(x) = (1 + 2\sqrt{2}) \cdot x$ .

The worst case for BCE/BME is actually unweighted—in fact, Example 2. This result fully resolves the 2-player, 2-machine case with linear latency functions. Included in this class of games are instances in which  $\text{BCE} > \text{BME} > \text{OPT}$ . One concrete example is with weights  $\{1, 1 + \sqrt{2}\}$ ,  $f_L(x) = x$ , and  $f_R(x) = \frac{3+3\sqrt{2}}{2} \cdot x$ , when  $\text{BCE}/\text{BME} = \frac{20+4\sqrt{2}}{23} \approx 1.1155$  and  $\text{BME}/\text{OPT} = \frac{14\sqrt{2}-1}{17} \approx 1.1058$ .

Adding additional machines to a 2-player instance does not substantively change the results (there is no point in either player using anything other than the two “best” machines), but the setting with  $n \geq 3$  players requires further analysis, and, it appears, new techniques. Recent results on the price of anarchy in linear load-balancing games [3,4,6] imply an upper bound of  $1 + \phi \approx 2.618$  on  $\text{BME}/\text{OPT}$  for any number of players  $n$ , where  $\phi$  is the golden ratio. We believe that the worst-case ratio of  $\text{BME}/\text{OPT}$  does not decrease as  $n$  increases. (Consider an  $n$ -player instance in which  $n - 2$  players have jobs of negligible weight and the remaining 2 players have jobs as in Theorem 2.) However, we do not have a proof that  $\text{BME}/\text{OPT}$  cannot worsen from  $\frac{1+\sqrt{2}}{2} \approx 1.2071$  as  $n$  grows; nor do we have a 3-job example for which  $\text{BME}/\text{OPT}$  is worse than  $\frac{1+\sqrt{2}}{2}$ . The major open challenge emanating from our work is to close the gap between the upper bound ( $\text{BME}/\text{OPT} \leq 2.618$ ) and our bad example ( $\text{BME}/\text{OPT} = 1.2071$ ) for general  $n$ .

## 5 Weighted Nonlinear Load-Balancing Games

We now consider weighted load-balancing games with latency functions that are not necessarily linear. We know from Lemma 1 that even in unweighted cases the power of Nash and correlated equilibria is limited. The weighted setting is even worse: the price of anarchy is unbounded, even if we restrict our attention to pure equilibria. Consider two identical machines, with latencies  $f(x) = 0$  for  $x \leq 5$  and  $f(x) = 1$  for  $x \geq 6$ . There are four jobs, two of size 3 and two of size 2. A solution with cost zero exists (each machine has one size-2 and one size-3 job), but putting the two size-3 jobs on one machine and the two size-2 jobs on the other is a pure Nash equilibrium too. We can show that the price of stability is better in this setting, but in general BME is no better than BNE:

**Theorem 3.** *In any  $n$ -player weighted load-balancing game with job weights  $\{w_1, \dots, w_n\}$  (and not necessarily linear latency functions),  $\text{BNE} \leq \Delta \cdot \text{OPT}$ , where  $\Delta := \sum_i w_i / \min_i w_i$  is the ratio of total job weight to smallest job weight. Thus  $\text{BME} \leq \Delta \cdot \text{OPT}$  and  $\text{BCE} \leq \Delta \cdot \text{BME}$ . Both bounds are tight.*

## 6 Other Social-Cost Functions

Thus far we have discussed the social cost function  $\text{sc}_{\text{avg}}(\mathbf{s}) := \sum_i w_i \cdot c_i(\mathbf{s})$  exclusively. Two other social cost functions have received attention in the literature: the maximum latency  $\text{sc}_{\text{max}}(\mathbf{s}) := \max_i c_i(\mathbf{s})$  and the unweighted average latency  $\text{sc}_{\text{unavg}}(\mathbf{s}) := \sum_i c_i(\mathbf{s})$ . Under  $\text{sc}_{\text{max}}$ , in any load-balancing game  $\text{BME} = \text{OPT} = \text{BNE}$ : starting from  $\text{OPT}$ , run best-response dynamics (BRD) until it converges; no BRD step increases the maximum load, so the resulting Nash equilibrium is still  $\text{OPT}$ . Thus mediation is uninteresting under  $\text{sc}_{\text{max}}$ .

The behavior of  $\text{sc}_{\text{unavg}}$  turns out to be similar to that of  $\text{sc}_{\text{avg}}$ . For nonlinear latencies, an analogue to Theorem 3 states that  $\text{BCE} \leq n \cdot \text{BME}$  and  $\text{BME} \leq n \cdot \text{OPT}$  ( $\text{OPT}$  is at least the maximum cost  $x$  experienced by a job in  $\text{OPT}$ ; running BRD from  $\text{OPT}$  yields a Nash equilibrium where each job experiences cost at most  $x$ ); the examples from Theorem 3 and Example 1 both remain tight. The 2-job linear case is also qualitatively similar; however, in contrast to the  $\text{sc}_{\text{avg}}$  setting (where there is a bound of  $\text{BME}/\text{OPT} \leq 2.618$  for  $n$ -player games), even mediators cannot enforce outcomes that are close to  $\text{OPT}$  under  $\text{sc}_{\text{unavg}}$  as the number of players grows, even in linear-latency games. (Our construction also demonstrates that none of  $\text{BCE}$ ,  $\text{BNE}$ , and  $\text{WNE}$  can provide constant approximations to  $\text{OPT}$ .)

**Theorem 4.** *Under  $\text{sc}_{\text{unavg}}$ , in linear-latency weighted load-balancing games:*

- for 2 jobs and 2 machines,  $\text{BCE}/\text{BME} \leq \frac{4}{3}$  (this is tight for Example 2) and  $\text{BME}/\text{OPT} \leq \frac{2+4\sqrt{2}}{7} \approx 1.0938$  (this is tight for Theorem 2's example).
- for  $n$  jobs and 2 machines,  $\text{BME}/\text{OPT}$  is not bounded by any constant.

## 7 Future Directions

In this paper we have begun to analyze the power of mediators in the spirit of price of stability, focusing on load-balancing games under the weighted average latency social cost function. We have a complete story for unweighted games and for weighted games with general latency functions. The biggest open question is the gap between  $\text{BME}$  and  $\text{OPT}$  in  $n$ -player weighted linear games. We know that for all such games  $\text{BME}/\text{OPT} \leq 2.618$  [3], and that there exist examples in which  $\text{BME}/\text{OPT} = 1.2071$ . What is the worst-case  $\text{BME}/\text{OPT}$  for  $n \geq 3$  players? In particular, it may be helpful to understand better the connection between  $\text{sc}_{\text{unavg}}$  and  $\text{sc}_{\text{avg}}$ : it was unexpected that the same instance is the worst case for both functions in the 2-player case (Theorem 2 and Theorem 4).

The broader direction for future research, of course, is to characterize the power of mediators in games beyond load balancing. For example, the upper bound of  $\text{BME}/\text{OPT} \leq 2.618$  in weighted linear load-balancing games comes from an upper bound on the price of anarchy in congestion games, a more general class of games. It is an interesting question as to how much better mediated equilibria are than correlated equilibria in, say, linear-latency weighted congestion games.

## References

1. Ashlagi, I., Monderer, D., Tennenholtz, M.: Mediators in position auctions. In: EC 2007, pp. 279–287 (2007)
2. Aumann, R.: Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55(1), 1–18 (1987)
3. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: STOC, pp. 57–66 (2005)
4. 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)
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. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: STOC 2005, pp. 67–73 (2005)
7. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *ACM Trans. on Algorithms* 3(1), 4 (2007)
8. 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)
9. Fotakis, D., Kontogiannis, S., Spirakis, P.: Selfish unsplittable flows. *Theoretical Computer Science* 348(2), 226–239 (2005)
10. 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)
11. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B., Rode, M.: Nash equilibria in discrete routing games with convex latency functions. *J. Comput. Syst. Sci.* 74(7), 1199–1225 (2008)
12. Gairing, M., Lücking, T., Monien, B., Tiemann, K.: Nash equilibria, the price of anarchy and the fully mixed Nash equilibrium conjecture. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 51–65. Springer, Heidelberg (2005)
13. Hart, S., Mas-Colell, A.: A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5), 1127–1150 (2000)
14. Kakade, S., Kearns, M., Langford, J., Ortiz, L.: Correlated equilibria in graphical games. In: EC 2003, pp. 42–47 (2003)
15. Kothari, A., Suri, S., Tóth, C., Zhou, Y.: Congestion games, load balancing, and price of anarchy. In: Workshop on Combinatorial and Algorithmic Aspects of Networking 2004 (2004)
16. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563. Springer, Heidelberg (1999)
17. Lücking, T., Mavronicolas, M., Monien, B., Rode, M.: A new model for selfish routing. *Theoretical Computer Science* 406(3), 187–206 (2008)
18. Mavronicolas, M., Spirakis, P.: The price of selfish routing. *Algorithmica* 48(1), 91–126 (2007)
19. Monderer, D., Tennenholtz, M.: Strong mediated equilibrium. *Artificial Intelligence* 173(1), 180–195 (2009)

20. Papadimitriou, C., Roughgarden, T.: Computing equilibria in multi-player games. *Journal of the ACM* 55(3), 1–29 (2008)
21. Peleg, B., Procaccia, A.: Implementation by mediated equilibrium. Discussion paper 463, Center for the Study of Rationality, Hebrew University (2007)
22. Peleg, B., Procaccia, A.: Mediators and truthful voting. Working paper (2008)
23. Rozenfeld, O., Tennenholtz, M.: Strong and correlated strong equilibria in monotone congestion games. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) *WINE 2006*. LNCS, vol. 4286, pp. 74–86. Springer, Heidelberg (2006)
24. Rozenfeld, O., Tennenholtz, M.: Routing mediators. In: *IJCAI 2007* (2007)
25. Suri, S., Tóth, C., Zhou, Y.: Selfish load balancing and atomic congestion games. *Algorithmica* 47(1), 79–96 (2007)
26. Tennenholtz, M.: Game-theoretic recommendations: Some progress in an uphill battle. In: *Proc. Intl. Conf. on Autonomous Agents and Multiagent Systems 2008* (2008)
27. Vöcking, B.: Selfish load balancing. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. (eds.) *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)



# On the Impact of Strategy and Utility Structures on Congestion-Averse Games

Thomas Voice<sup>1</sup>, Maria Polukarov<sup>1</sup>, Andrew Bye<sup>2</sup>, and Nicholas R. Jennings<sup>1</sup>

<sup>1</sup> School of Electronics and Computer Science, University of Southampton, UK  
{tdv, mp3, nrj}@ecs.soton.ac.uk

<sup>2</sup> Hewlett-Packard Laboratories, Bristol, UK  
andrew.bye@hp.com

**Abstract.** Recent results regarding *games with congestion-averse utilities* (or, *congestion-averse games*—CAGs) have shown they possess some very desirable properties. Specifically, they have pure strategy Nash equilibria, which may be found in polynomial time. However, these results were accompanied by a very limiting assumption that each player is capable of using *any* subset of its available set of resources. This is often unrealistic—for example, resources may have complementarities between them such that a minimal number of resources is required for any to be useful. To remove this restriction, in this paper we prove the existence and tractability of a pure strategy equilibrium for a much more general setting where each player is given a matroid over the set of resources, along with the bounds on the size of a subset of resources to be selected, and its strategy space consists of all elements of this matroid that fit in the given size range. Moreover, we show that if a player strategy space in a given CAG does not satisfy these matroid properties, then a pure strategy equilibrium need not exist, and in fact the determination of whether or not a game has such an equilibrium is NP-complete. We further prove analogous results for each of the congestion-averse conditions on utility functions, thus showing that current assumptions on strategy and utility structures in this model cannot be relaxed anymore.

## 1 Introduction

Congestion games—in which self-interested players strategically choose from a common set of resources and derive individual utilities that depend on the total congestion on each resource—are fundamental to a wide range of applications. Examples include resource and task allocation, firm competition for production processes, routing problems, network design, and other kinds of resource sharing scenarios in distributed systems [8][15][16]. Such games are important because Rosenthal [15] showed that they always possess Nash equilibria in pure strategies. This follows by a potential function argument [10], implying that such an equilibrium can be reached in a natural way when players iteratively (and unilaterally) improve their strategies in response to the others' choices. However, such a sequence of (even maximal, or, best) responses may take an exponential number of iterations, as is shown in [4]. In fact, it is PLS-complete to compute a pure strategy equilibrium for general congestion games. Motivated by this fact, much recent effort in algorithmic game theory has gone into study of interesting subclasses of congestion games that are computationally tractable. These, for example, include *singleton* (or, *resource selection*) *congestion games* [5], in which each player



is restricted to choosing a single resource, and a more general class of *matroid congestion games* [1], in which players choose among the bases of a matroid over the set of resources. Note, however, that in both cases only strategies (i.e., sets) of the same cardinality are allowed (respectively, 1 or the rank of the matroid).

Congestion games have been extensively studied in a variety of contexts in computer science and economics, giving rise to several extensions of the original model. In particular, the models for *local-effect games* [6], *ID-congestion games* [9], *player-specific* [7] and *weighted congestion games* [7]—in which a player’s payoff is effected by the number, identities or weights of players choosing its selected or neighboring resources—have been considered. However, such games have been constrained to use utility functions that are linear sums with respect to resources, and assumed full reliability and synchronicity of services. More recently, additional generalizations [11][12][13][14] dealt with the possibility that resources may fail to execute their assigned tasks, or with the actual order in which the tasks are executed, thus incorporating non-linear and non-additive utility functions in the context of congestion games. These models, however, assumed mutual independence among the resources and imposed particular structures on the players’ strategy spaces.

Generalising beyond problem classes with desirable properties raises the important question of developing meaningful criteria that have to be satisfied in order to guarantee that these properties are still present in the generalised model. Given the fact that congestion games have pure-strategy equilibria, we are interested in the question of how far such a sufficient criterion for the existence of pure strategy Nash equilibria can go, and which properties would ensure polynomial complexity of such equilibria.

To this end, Bye et al [2] provided a very general framework that, in particular, includes the abovementioned models of congestion games with faulty or asynchronous resources and player-specific congestion games in the superclass called *games with congestion-averse utilities* (or, *congestion-averse games*—CAGs). In a CAG, the payoff of a player is determined by the vector of resource congestion (thus capturing the possibility of mutual dependencies among the resources), via any real-valued function that satisfies certain “congestion-averse” conditions—i.e., *monotonicity*, *submodularity* and *independence of irrelevant alternatives*. The authors proved the existence of a pure strategy Nash equilibrium in these games and provided a polynomial time algorithm for its computation. This result was based on *the single profitable move property* (SPMP) of these games, implying that a strategy profile is a Nash equilibrium if and only if it is stable under adds, drops or switches with a single resource. The congestion-averse assumptions have been shown to be minimal to guarantee the existence of this property. However, the question of necessity of these assumptions for the existence of a pure strategy Nash equilibrium remained open. Also, though the model of CAGs captures a wide range of important scenarios, it assumes that a player is capable of using any subset of its accessible set of resources, which is unrealistic in many real-life situations, in which, for example, only certain sets of resources are useful in combination.

**Our contribution.** Given this motivation, we show that the analysis of Bye et al [2] can be generalised even further, towards what we call *Matroid Congestion-Averse Games*, or MCAGs, in which the set of strategies of each player consists of all the sets within a certain size range from a matroid. This, in particular, includes (but is not

restricted to) the possibilities of having a full power set over any subset of resources (as in CAGs [2]), or any other matroid, or having a set of elements of a fixed size—for instance, singletons (as in resource selection games [5]) or bases of a matroid (as in matroid congestion games [1]). For this setting, we prove that all such games have the SPMP and possess at least one equilibrium in pure strategies. We give an algorithm which converges on an equilibrium, with time limits polynomial in the number of players and resources in the game. Essentially, we extend all the previous results on CAGs, weakening the bounds on algorithm running time. We further complete these results by showing that under various relaxations of the matroid or congestion-averse assumptions, these properties are no longer present, and in fact the determination of whether or not a game has a pure strategy equilibrium is NP-complete. The proofs are omitted, due to space limitations.

## 2 Notation and Background

Consider a *congestion setting* (or, *domain*) with a set  $\mathbf{N} = \{1, \dots, N\}$  of players and a set  $\mathbf{R} = \{r_1, \dots, r_R\}$  of resources. A player  $i$ 's strategy is to choose a subset of resources from  $\mathbf{R}$ , and every  $N$ -tuple of strategies  $\sigma = (\sigma_i)_{i \in \mathbf{N}}$  corresponds to an  $R$ -dimensional congestion vector  $h(\sigma) = (h_r(\sigma))_{r \in \mathbf{R}}$  where  $h_r(\sigma)$  is the number of players who select resource  $r$  (we drop the profile to give  $h_r$  when it's clear which profile is under consideration). For any player  $i \in \mathbf{N}$ , its *personalised* vector of congestion,  $h^i(\sigma)$ , is defined to be a vector in  $\mathbb{N}^R$  that coincides with  $h(\sigma)$  for all the resources that have been selected by  $i$  and that has zero entries for all of its unselected resources:  $h_r^i(\sigma) = h_r(\sigma)$  if  $r \in \sigma_i$  and  $h_r^i(\sigma) = 0$  otherwise. For  $h \in \mathbb{N}^R$ , its “support”,  $S(h) \subseteq \{1, \dots, R\}$ , is defined as  $\{j : h_{r_j} > 0\}$ . The utility of player  $i$  in a congestion setting is given by a function  $U_i : \mathbb{N}^R \rightarrow \mathbb{R}$  that assigns a real value to a (personalised) vector of congestion.  $\square$

**Games with congestion-averse utilities.** A utility function is congestion-averse if it (i) monotonically decreases with respect to increasing congestion, (ii) is submodular in that the “better” collection of resources a player uses—the less incentive it has to add new resources, and (iii) is independent of irrelevant alternatives (i.e., if a player “prefers” one resource over another at their current congestion levels, then it still does so no matter what other changes are made to any other resources). Formally, given a profile  $\sigma$  and a set of *elementary changes* (or, *single moves*) defined on  $\sigma$  as follows:

- **add**  $A_i(r)$ —player  $i$  adds an unselected resource  $r$ :  $\sigma'_i = \sigma_i \cup \{r\}$ ,
- **drop**  $D_i(r)$ —player  $i$  drops a selected resource  $r$ :  $\sigma'_i = \sigma_i \setminus \{r\}$ ,
- **switch**  $S_i(r_+ \leftarrow r_-)$ —player  $i$  switches resources by adding resource  $r_+$  and dropping resource  $r_-$  (note that  $S_i(r_+ \leftarrow r_-) = A_i(r_+) + D_i(r_-)$   $\square$ ), a utility function  $U : \mathbb{N}^R \rightarrow \mathbb{R}$  is said to be *congestion-averse* if it satisfies:

<sup>1</sup> Note that the player's utility only depends on the numbers of players choosing each resource but not on their identities—that is, we consider *anonymous* settings (see [3] for results on approximating equilibria in anonymous games).

<sup>2</sup> Here and in what follows, “+” should be understood to mean sequential execution, read left-to-right. We also use this notation to indicate elementary changes applied to strategy profiles: e.g.,  $\sigma + D$  denotes a drop applied to profile  $\sigma$ .

- **monotonicity:** If  $S(h) = S(h')$  and  $\forall r, h_r \geq h'_r$ , then  $U(h) \leq U(h')$ ;
- **submodularity:** Improving a resource selection by either (i) profitable switches, (ii) extending the set of utilised resources or (iii) reducing congestion on them does not make new adds more profitable, or drops less profitable; likewise, unprofitable switches, deleting the resources or increasing the congestion does not make drops more profitable, or adds less profitable. Equivalently, for any  $h, h'$  and  $h''$  such that  $|S(h)| = 1$  and  $S(h) \not\subseteq S(h'), S(h'')$ ,

$$U(h + h') - U(h') \leq U(h + h'') - U(h'')$$

- if either (i)  $|S(h') \setminus S(h'')| = |S(h'') \setminus S(h')| = 1$  and  $U(h') \geq U(h'')$ , (ii)  $S(h'') \subseteq S(h')$  and  $h_{j''} = h_{j'} \forall j \in S(h'')$ , or (iii)  $S(h') = S(h'')$  and  $h' \leq h''$ ;
- **independence of irrelevant alternatives:** If  $S_i(r_+ \leftarrow r_-)$  is a profitable switch for player  $i$  given profile  $\sigma$ , then it is profitable for  $i$  from any other profile  $\sigma'$  satisfying  $r_- \in \sigma'_i, r_+ \notin \sigma'_i, h_{r_-}(\sigma) = h_{r_-}(\sigma')$  and  $h_{r_+}(\sigma) = h_{r_+}(\sigma')$ .

A *congestion-averse game* (CAG) is a game in the congestion domain with congestion-averse utility functions, where each player  $i \in \mathbf{N}$  has a subset  $\mathbf{R}_i \subseteq \mathbf{R}$  of  $R_i \in \mathbf{N}$  accessible resources, and its strategy space,  $\Sigma_i$ , is a power set of  $\mathbf{R}_i$ .

### 3 Matroid Congestion-Averse Games

In this section, we extend the model of congestion-averse games to encompass more general and complex player strategy spaces, which we loosely build on *matroids*. Before we give a formal definition of such games, we briefly introduce matroids.<sup>3</sup>

**Definition 1.** A *matroid*,  $M$ , is a collection of subsets of some set of elements  $X$ , with the property that if some  $Y \subseteq X$  is in  $M$  then all subsets of  $Y$  are in  $M$ . Further, if  $V \in M$  is such that  $|V| < |Y|$  then there exists some  $a \in Y \setminus V$  such that  $V \cup \{a\} \in M$ .

A *matroid congestion-averse game* (MCAG) is now defined as a game in the congestion domain with congestion-averse utility functions, over strategy spaces where each player  $i$  is given a matroid  $M_i$  and integers  $n_i \leq m_i$ , and its strategy space consists of all the subsets  $X \in M_i$  such that  $n_i \leq |X| \leq m_i$ . More precisely,

**Definition 2.** An *MCAG*  $\Gamma = (\mathbf{N}, \mathbf{R}, (U_i(\cdot))_{i \in \mathbf{N}})$  consists of a set  $\mathbf{N}$  of  $N \in \mathbf{N}$  players, a set  $\mathbf{R}$  of  $R \in \mathbf{N}$  resources, and for each player  $i$  a matroid  $M_i$  over  $\mathbf{R}$ , integers  $n_i \leq m_i$ , and a congestion-averse utility function  $U_i : \mathbb{N}^{\mathbf{R}} \rightarrow \mathbb{R}$ . The strategy space for each player  $i \in \mathbf{N}$  is the set of all the subsets  $X \in M_i$  satisfying  $n_i \leq |X| \leq m_i$ , and its payoff from a strategy profile  $\sigma$  is  $u_i(\sigma) = U_i(h^i(\sigma))$ , where  $h^i(\sigma)$  is  $i$ 's personalised vector of congestion as determined by  $\sigma$ .

*Remark 1.* Note that our strategy structures cover (but are not restricted to) the possibilities of having a power set over any subset of resources (as in CAGs), or any other matroid (full or incomplete), or having a set of elements of a fixed size—for example, a set of singletons (as in resource selection games) or a set of bases of a matroid (as in matroid congestion games).

<sup>3</sup> For a detailed discussion of matroids, we refer the reader to [17].

Interestingly, as we show below, the CAG technique based on particular “ladders” of elementary changes, appears to be universal enough to capture the matroid case. Recall, however, that this method builds heavily on several properties, including (i) the single profitable move property, (ii) the existence of a strategy profile which is stable to both switches and adds (and a method to find such), and (iii) the possibility to rank resources by their attraction to a player, that easily follows from the fact that any elementary change is available to any player at any profile. For MCAGs, however, the existence of these properties is not at all obvious. We start with their proofs in the following subsection.

### 3.1 Preliminary Results

**The single profitable move property.** We first show that the matroid congestion-averse games have the SPMP—the *single profitable move property*, implying that a profile is in equilibrium if and only if it does not admit profitable elementary changes. We begin with a lemma.

**Lemma 1.** *Given a strategy profile of an MCAG, a player’s strategy is a (strict) best response within the subspace of strategies of the same size, if and only if no (strictly) profitable switch from this strategy is available to the player. Indeed, given a profile  $\sigma$ , suppose some player  $i$  has any alternative strategy  $\sigma'_i$  such that  $|\sigma_i| = |\sigma'_i|$  and player  $i$  would (strictly) prefer  $\sigma'_i$  over  $\sigma_i$ . Then, there is some  $r_+ \in \sigma'_i \setminus \sigma_i$  and  $r_- \in \sigma_i \setminus \sigma'_i$  such that the switch  $S(r_+ \leftarrow r_-)$  is (strictly) profitable for player  $i$  at  $\sigma$ .*

We can now prove the single profitable move property for MCAGs.

**Theorem 1.** *Given an MCAG, a strategy profile  $\sigma$  is a Nash equilibrium if and only if there are no (maximal) strictly profitable switches, drops or adds.*

A strategy profile  $\sigma$  is termed as *A-stable* (*D-stable*, *S-stable*) if it admits no maximally profitable adds (drops, switches); likewise for *AS-stable*, *DS-stable* and so on. Thus, the SPMP states that a profile is in equilibrium if and only if it is ADS-stable. The SPMP has been used to develop techniques for finding pure strategy equilibria in CAGs and several of their subclasses. These methods used particular dynamics of elementary changes that initialised with a strategy profile which is either AS- or DS-stable. Indeed, since in CAGs a player is allowed to use any subset from its set of accessible resource, the existence of such a profile follows trivially—all the players just play the full or the empty set. However, this strategy may not be available for a general MCAG. Nevertheless, as we shall show, every such a game possesses a strategy profile (or, a “state”) which is stable under adds and switches.

**Finding an AS-stable state.** First, we prove the following theorem.

**Theorem 2.** *Given an MCAG, consider  $P$  the set of pairs consisting of a single resource and a congestion level on that resource. For each player  $i$ , there exists a ranking function  $V_i(\cdot)$  on  $P$  such that for any congestion vector  $h$  and strategy  $\sigma$ , if there is a switch  $S(r_j \leftarrow r_k)$  available to player  $i$  then  $V_i(r_j, h_j + 1) \leq V_i(r_k, h_k)$  if and only if the switch  $S(r_j \leftarrow r_k)$  is profitable.*

**Corollary 1.** *Any best response dynamics within the reduced space where each player is restricted to the maximal size strategies, will terminate in an AS-stable state within  $N^2 R^2$  moves.*

*Remark 2.* In games with fixed size strategies—like singleton or matroid congestion games—drops are never available. So, any AS-stable strategy profile is also a Nash equilibrium. For such games, the above corollary proves the existence of, and provides a method for finding, a pure strategy Nash equilibrium.

**Dynamics.** Given the existence of an AS-stable state, we now explore the convergence for matroids of the drop- and swap-dynamics, as defined in [2]. We start with a brief definition of drop- and swap-ladders, and then proceed and describe their properties in MCAGs.

**Definition 3.** A **drop ladder** is a sequence  $D_{i_0}(r_0) + S_{i_1}(r'_1 \leftarrow r_1) + \dots + S_{i_m}(r'_m \leftarrow r_m)$ , consisting of a maximally profitable drop followed by a sequence of  $m \geq 0$  maximally profitable switches, and a **swap ladder** is a drop-ladder followed by a maximally profitable add with its tail:  $D_{i_0}(r_0) + S_{i_1}(r_0 \leftarrow r_1) + \dots + S_{i_m}(r_{m-1} \leftarrow r_m) + A_{i_{m+1}}(r_m)$ . The swap ladder is described as **minimal** if all intermediate strategy profiles before the last add were A-stable (i.e., if the add is performed at the first opportunity).

Note that the above definition of a swap ladder implies that a profitable add is made to the tail of the corresponding drop ladder. This is well defined by the following lemma.

**Lemma 2.** *Given a CAG, let  $\sigma$  be an AS-stable profile that possesses a drop ladder of length  $m$ , and let  $\sigma^k$  denote the result of applying the drop and the first  $k$  switches to  $\sigma$ . Suppose further that for each  $\sigma^k$  for  $1 \leq k < m$  there are no profitable adds. Then, for all  $1 \leq k \leq m$ , the only switches which are profitable at  $\sigma^k$  are those which “chain” with the previous switch or the initial drop, i.e. those who switch in the resource which was most recently dropped or switched out. Furthermore, if there is a profitable add  $A_i(r_+)$  for profile  $\sigma^m$  then  $r_+ = r_m$ .*

Thus, the result of a swap ladder possesses the same congestion vector as the original profile; as we shall see, this will imply that minimal swap ladders preserve AS-stability. The result will follow from Lemma 3 below:

**Lemma 3.** *Consider the sequence of adds, drops and switches that a single player makes in a sequence of minimal swap ladders. For each player, we rank the resources according to the ranking function defined in Theorem 2 using the fixed congestion levels present between swap ladders. Then, (i) if the ranks of resources a player has selected are put in decreasing order, then this set of values increases lexicographically with every switch; (ii) every add must add a resource that is strictly higher ranked than the resource most recently dropped; (iii) the ranks of dropped resources are non-decreasing.*

**Corollary 2.** *There can be no more than  $NR(R+2)$  elementary changes in total in any sequence of minimal swap ladders.*

### 3.2 Main Results

We are now ready to conclude the existence and tractability of a pure strategy Nash equilibrium in MCAGs. The following proposition, coupled with Corollary 2 implies our main result in Theorem 3.

**Proposition 1.** *Applying a minimal swap ladder to an AS-stable state preserves AS-stability.*

**Theorem 3.** *Every MCAG has a pure strategy Nash equilibrium.*

## 4 Necessity of the MCAG Model Assumptions

In this section we complete our results on congestion-averse games, by demonstrating that the strategy space and utility function assumptions in our model cannot be further relaxed. Specifically, we will show that if any of the (ranged) matroid or the congestion-averse properties is removed, a pure strategy equilibrium is not guaranteed to exist, and in fact the determination of whether or not a game has such an equilibrium is NP-complete.

### 4.1 Non-matroid Congestion-Averse Games

In an MCAG, each player has a strategy space which consists of all the sets within a certain size range from a matroid. This can be expressed equivalently with the following. Suppose a player has a strategy space  $S$ . Then, for all  $X \neq Y$  in  $S$ , the following two assumptions hold: (i) if  $|X| = |Y|$  then for each element  $x \in X$  there is an element  $y \in Y$  such that  $(X \setminus \{x\}) \cup \{y\} \in S$ ; (ii) if  $|X| < |Y|$  then all  $|X|$  element subsets of  $Y$  are in  $S$ . This is termed as *the ranged matroid property*. We will show that if we relax any of these assumptions then a pure strategy equilibrium is no longer guaranteed and, in fact, the determination of whether or not such an equilibrium exists is not tractable. This follows by a reduction argument using the 3-SAT problem which is known to be NP-complete.

**Theorem 4.** *In an MCAG setting, violation of either of the ranged matroid assumptions on strategy spaces may result in a game with no pure strategy equilibria. Moreover, it is in general NP-complete to determine whether a game possesses such an equilibrium.*

### 4.2 Non-congestion-Averse Utilities

Here we present similar results on the congestion-averseness assumptions on utility functions. Specifically, we show that violation of any of these assumptions (or even partial relaxation of submodularity) may result in a game with no pure strategy equilibrium, and reduce from the 3-SAT to show NP-hardness of the equilibrium existence decision problem.

**Theorem 5.** *In an MCAG setting, if any one of the congestion-averse conditions on utility functions is violated then a pure strategy Nash equilibrium is not guaranteed to exist. Moreover, there are instances of games in which it is NP-complete to determine its existence. This also applies to the case in which the submodularity assumption is only partially violated, that is either parts (i) and (iii) or part (ii) of the assumption hold.*



## 5 Conclusions

We investigated the impact of the strategy space and payoff function structures on games with congestion-averse utilities. We extended the previous positive results on the existence and tractability of pure strategy equilibria to the case, in which the set of strategies of each player consists of all the sets within a certain size range from a matroid. This covers a wide range of settings, including those where the players' strategies are represented by singletons, bases of matroids, and power sets over a set of accessible resources. Our result is tight in that the relaxation of the (ranged) matroid property or each of the congestion-averseness conditions may lead to a game without a pure strategy equilibrium, and it is in general NP-complete to determine the existence of such an equilibrium. Thus, we conclude that the current assumptions on strategy and utility structures in this model cannot be further relaxed.

## References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. *Journal of the ACM* 55(6) (2008)
2. Bye, A., Polukarov, M., Jennings, N.R.: Games with congestion-averse utilities. In: *SAGT 2009*, pp. 220–232 (2009)
3. Daskalakis, C., Papadimitriou, C.: Computing equilibria in anonymous games. In: *FOCS 2007*, pp. 83–93 (2007)
4. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure Nash equilibria. In: *STOC 2004*, pp. 604–612 (2004)
5. Jeong, S., McGrew, R., Nudelman, E., Shoham, Y., Sun, Q.: Fast and compact: A simple class of congestion games. In: *AAAI 2005*, pp. 489–494 (2005)
6. Leyton-Brown, K., Tennenholtz, M.: Local-effect games. In: *IJCAI 2003*, pp. 772–780 (2003)
7. Milchtaich, I.: Congestion games with player-specific payoff functions. *Games and Economic Behavior* 13, 111–124 (1996)
8. Milchtaich, I.: Congestion models of competition. *American Naturalist* 147(5), 760–783 (1996)
9. Monderer, D.: Solution-based congestion games. *Advances in Math. Economics* 8, 397–407 (2006)
10. Monderer, D., Shapley, L.S.: Potential games. *Games and Economic Behavior* 14, 124–143 (1996)
11. Penn, M., Polukarov, M., Tennenholtz, M.: Congestion games with failures. In: *EC 2005*, pp. 259–268 (2005)
12. Penn, M., Polukarov, M., Tennenholtz, M.: Congestion games with load-dependent failures: Identical resources. *Games and Economic Behavior* 67(1), 156–173 (2009)
13. Penn, M., Polukarov, M., Tennenholtz, M.: Random order congestion games. *Mathematics of Operations Research* 34(3), 706–725 (2009)
14. Penn, M., Polukarov, M., Tennenholtz, M.: Taxed congestion games with failures. *Annals of Mathematics and Artificial Intelligence* (to appear, 2009)
15. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
16. Roughgarden, T., Tardos, E.: How bad is selfish routing. *Journal of the ACM* 49(2), 236–259 (2002)
17. Schrijver, A.: *Combinatorial optimization: Polyhedra and efficiency*, volume B, *Matroids, Trees, Stable Sets*, ch. 36–39. Springer, Heidelberg (2003)

# On Strong Equilibria in the Max Cut Game

Laurent Gourvès and Jérôme Monnot

LAMSADE, CNRS FRE 3234 & Université de Paris-Dauphine  
{laurent.gourves,monnot}@lamsade.dauphine.fr

**Abstract.** This paper deals with two games defined upon well known generalizations of MAX CUT. We study the existence of a *strong equilibrium* which is a refinement of the Nash equilibrium. Bounds on the price of anarchy for Nash equilibria and strong equilibria are also given. In particular, we show that the MAX CUT game always admits a strong equilibrium and the strong price of anarchy is  $2/3$ .

## 1 Introduction

Suppose that  $n$  agents communicate via radio signals but only two distinct frequencies are available. In this scenario we are given a symmetric  $n \times n$  matrix which indicates, for each pair of agents, the strength of the interference that they experiment if they select the same frequency. We suppose that each agent chooses her frequency in order to minimize the sum of interferences that she experiments<sup>1</sup>, no matter what is the situation of the others. We use *strategic game theory* as a formal framework to study the following question: What would be the worst configuration that the selfish agents can reach compared to a solution where a central entity assigns frequencies optimally? When *Nash equilibria* – a situation where no agent can unilaterally deviate and benefit – are considered, this ratio is better known as the *price of anarchy* (PoA) [10]. It captures the performance of systems where selfish players interact without central coordination. Intuitively, a PoA far from 1 indicates that the system requires regulation.

Nash equilibria are considered as stable configurations. However a Nash equilibrium is not sustainable if the agents can realize that they all benefit if they perform a simultaneous deviation whereas any unilateral move is inefficient. The *strong equilibrium* introduced by Aumann [2] is a refinement of the Nash equilibrium where for every deviation by a group of agents, at least one member of the group does not benefit. The *strong price of anarchy* (SPoA) [1] is the PoA reduced to strong equilibria.

So, what are the PoA and SPoA of the above mentioned interference game? The game was already studied in [6,4,7]. It is defined upon the well known MAX CUT problem: Given a simple weighted graph, find a bipartition of the vertex set such that the weight of the edges having an endpoint in both parts of the partition, i.e. the cut, is maximum. In the MAX CUT game, a player's utility is her contribution to cut, i.e. the weight of the edges of the cut which are incident

---

<sup>1</sup> Or equivalently, maximize the sum of interferences that she does not experiment.



to her. The game always possesses a pure Nash equilibrium since it admits a potential function [15] (the weight of the cut). It is a kind of folklore that the PoA is 1/2. Up to our knowledge, nothing is known about the existence of a strong equilibrium and the SPoA of the MAX CUT game.

In this paper, we study two generalizations of the MAX CUT game which are similarly defined upon two generalizations of MAX CUT: NAE SAT and MAX  $k$ -CUT. An instance of the NAE SAT problem is a set of clauses, each of them being satisfied if its literals are not all true (or not all false) and one asks a truth assignment maximizing the weight of satisfied clauses. MAX CUT is equivalent to NAE SAT if each clause is made of two unnegated variables. In the NAE SAT game, every player tries to maximize the weight of satisfied clauses where she appears. A motivation of the NAE SAT game is given in the sequel. In the MAX  $k$ -CUT problem, one asks a  $k$  partition of the vertex set inducing a maximum weight cut. Its associated game is the interference game with  $k$  frequencies instead of 2.

## 2 Definitions and Notations

A *strategic game* is a tuple  $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$  where  $N$  is the set of players (we suppose that  $|N| = n$ ),  $A_i$  is the set of strategies of player  $i$  and  $u_i : \times_i A_i \rightarrow \mathbb{R}$  is player  $i$ 's utility function. A *pure state* or *pure strategy profile* of the game is an element of  $\times_i A_i$ . Although players may choose a probability distribution over their strategy set, we only consider pure strategy profiles in this paper. Players are supposed to be rational, i.e. each of them plays in order to maximize her utility.

Given a state  $a$ ,  $(a_{-i}, b_i)$  denotes the state where  $a_i$  is replaced by  $b_i$  in  $a$  while the strategy of the other players remains unchanged. A state  $a$  is a *Nash equilibrium* (NE) if there no player  $i \in N$  and a strategy  $b_i \in A_i$  such that  $u_i((a_{-i}, b_i)) > u_i(a)$ . Given two states  $a, a'$  and a coalition  $C \subseteq N$ ,  $(a_{-C}, a')$  denotes the state where  $a_i$  is replaced by  $a'_i$  in  $a$  for all  $i \in C$ . A state  $a$  is a *strong equilibrium* (SE) if there is no non-empty coalition  $C \subseteq N$  and a profile  $a' \in A$  such that  $u_i((a_{-C}, a')) > u_i(a)$  for all  $i \in C$ . A state  $a$  is an  *$r$ -strong equilibrium* ( $r$ -SE) if there is no non-empty coalition  $C \subseteq N$  of size at most  $r$  and a profile  $a' \in A$  such that  $u_i((a_{-C}, a')) > u_i(a)$  for all  $i \in C$ . Therefore a SE is a NE, a NE is a 1-SE and a SE is  $n$ -SE ( $n$  is the number of players).

The *price of anarchy* (PoA) measures the performance of decentralized systems [10] via its Nash equilibria. More formally, let  $\Gamma$  be a family of strategic games, let  $\gamma$  be an instance of  $\Gamma$ , let  $A_\gamma$  be the strategy space of  $\gamma$ , let  $Q : A_\gamma \rightarrow \mathbb{R}_+$  be a social function, let  $\mathcal{E}(\gamma)$  be the set of all pure Nash equilibria of  $\gamma$  and let  $o_\gamma$  be a social optimum for  $\gamma$  (i.e.  $o_\gamma = \operatorname{argmax}_{a \in A_\gamma} Q(a)$ ). The *pure price of anarchy* of  $\Gamma$  is  $\min_{\gamma \in \Gamma} \min_{a \in \mathcal{E}(\gamma)} Q(a)/Q(o_\gamma)$ . If  $\mathcal{SE}(\gamma)$  denotes the set of all strong equilibria of  $\gamma$  then the *strong price of anarchy* (SPoA) [1] is  $\min_{\gamma \in \Gamma} \min_{a \in \mathcal{SE}(\gamma)} Q(a)/Q(o_\gamma)$ . The  $r$ -SPoA is similarly defined when restricting ourselves to  $r$ -strong equilibria.

### 3 Related Work and Contribution

The MAX CUT game is a *game of congestion* [11]. Congestion games is a particular subclass of *potential games* [15] which are known to always possess a pure strategy NE. Any NE of the MAX CUT game corresponds to a local optimum whose computation is sometimes polynomial (cubic graphs [13], unweighted case) but PLS-complete in general [16]. Recently Christodoulou et al. [4] studied the rate of convergence to an approximate NE in the MAX CUT game and the social welfare of states obtained after a polynomial number of best response steps.

The MAX CUT game is close to the *party affiliation game* [6] and the *consensus game* [3]. The MAX  $k$ -CUT game is related to the model of migration studied by Quint and Shubik [12]. A land where several animals live is partitioned into  $k$  areas and each animal has to choose one. Two animals seeking the same resources (e.g. food or living conditions) compete if they share the same area. We assume that every kind of resource exists in each area. Then each animal migrates to the area where competition is minimum.

In a broader study on *clustering games* [7], Hoefler proved that the PoA of the *unweighted* MAX  $k$ -CUT game is  $(k - 1)/k$ . However, nothing is known about existence of a SE for this game and its SPoA. Up to our knowledge, nothing is known about the PoA of the NAE SAT game, the existence of a SE and the SPoA. However every game studied in this paper is a particular case of congestion games. Congestion games possess a SE in many situations (some of them are identified in [8,9,5,14]) but its existence is not always guaranteed.

In this paper we study the existence of a SE and the (S)PoA of the MAX  $k$ -CUT and NAE SAT games. Section 4 is devoted to the NAE SAT game. If each clause has two literals then we prove that any optimal solution is a SE and the SPoA is  $2/3$ . With more literals per clause, we show that no 2-SE is guaranteed while a pure NE, a 1-SE in fact, must exist. The PoA of the NAE SAT game is in general  $1/2$  and  $q/(q + 1)$  if each clause is made of exactly  $q \geq 3$  literals. Section 5 is devoted to the MAX  $k$ -CUT game. Our positive result states that any optimal solution is a 3-SE (when  $k = 2$ , an optimal cut is a SE by the result given for the NAE SAT game). Our negative result states that for  $k \geq 3$ , there is an instance with two distinct optimal cuts: one is a SE while the other is not a 4-SE. Before giving a conclusion, we show that the  $r$ -SPoA of the MAX CUT game is equal to  $1/2$  if  $r$  is bounded above by the square root of the number of players.

Due to space limitations, proofs are sometimes sketched or skipped.

### 4 The NAE SAT Game

Given a set  $X$  of boolean variables and a set  $\mathcal{C}$  of clauses, each of them being composed of at least two literals defined over  $X$  and a weight function  $w : \mathcal{C} \rightarrow \mathbb{R}_+$ , NAE SAT is to find a truth assignment  $\tau : X \rightarrow \{\mathbf{true}, \mathbf{false}\}$  such that the weight of NAE-satisfied clauses is maximum. A clause is NAE-satisfied if its literals are not all true or not all false (NAE= not all equal). In the following  $q$ -NAE SAT refers to the case where each clause has exactly  $q$  literals.

In the NAE SAT game, each variable is controlled by a selfish player with strategy **true** or **false**. A player’s utility is the weight of NAE-satisfied clauses where she appears. The social function is the weight of NAE-satisfied clauses.

As an application, imagine a population of animals cut into two groups (or gangs) denoted by  $T$  and  $F$ . Anyone can choose to live in  $T$  or  $F$  but not in both. In addition every individual  $i$  carries a set  $\gamma_i$  of genes (his *genotype*) that he wants to be ideally present in both groups. If  $i$  chooses  $T$  (resp.  $F$ ) then all his genes are in  $T$  (resp.  $F$ ) and exactly  $|\gamma_i \cap \bigcup_{j \in F} \gamma_j|$  (resp.  $|\gamma_i \cap \bigcup_{j \in T} \gamma_j|$ ) of his genes are in  $F$ . Then, in order to maximize the presence of his genotype,  $i$  prefers  $T$  if  $|\gamma_i \cap \bigcup_{j \in F \setminus \{i\}} \gamma_j| \geq |\gamma_i \cap \bigcup_{j \in T \setminus \{i\}} \gamma_j|$ , otherwise  $i$  prefers  $F$ . One can model the situation as a NAE SAT game: each animal  $i$  is a variable  $x_i$ , each gene  $g$  carried by at least two animals is a clause including a positive literal  $x_i$  iff  $g \in \gamma_i$ . Thus,  $i \in T$  (resp.  $i \in F$ ) means  $x_i$  is **true** (resp.  $x_i$  is **false**).

The NAE SAT game always has a pure Nash equilibrium since it can be defined as a congestion game. Then it is consistent to study its *pure* PoA.

**Theorem 1.** *The PoA of the NAE SAT game is*

- (i)  $q/(q + 1)$  if each clause has size exactly  $q$  with  $q \geq 3$
- (ii)  $1/2$  otherwise

Now we turn our attention to strong equilibria. We first show that every instance of the 2-NAE SAT game possesses a SE.

**Theorem 2.** *Every optimum of the 2-NAE SAT game is a SE.*

It follows that every optimum of the MAX CUT game is a SE since MAX CUT is equivalent to 2-NAE SAT if all literals are positive. When  $q \geq 3$ , the following result states that some instances of the  $q$ -NAE SAT game do not have a  $(q - 1)$ -SE (the existence of a 1-SE, i.e. a NE, is guaranteed).

**Theorem 3.** *For any  $q \geq 3$ , the existence of a  $(q - 1)$ -SE is not guaranteed for the  $q$ -NAE SAT game.*

Then it is consistent to study the *pure* SPoA of the 2-NAE SAT game.

**Theorem 4.** *The SPoA of the 2-NAE SAT game is  $2/3$ .*

*Proof.* Let  $I = (X, \mathcal{C})$  be an instance of 2-NAE-SAT where  $X$  is the set of variables and  $\mathcal{C}$  is the set of clauses weighted by  $w$ . Let  $\sigma$  (resp.  $\sigma^*$ ) a strong equilibrium (resp. an optimal truth assignment) of  $I$ . Without loss of generality, we assume that  $\sigma(x) = \mathbf{true}$  for all  $x \in X$ . Indeed if  $\sigma(x) = \mathbf{false}$  then one can replace every  $\bar{x}$  (resp.  $x$ ) by  $x$  (resp.  $\bar{x}$ ) and set  $\sigma(x) = \mathbf{true}$ .

Let  $A = \{x \in X : \sigma(x) = \sigma^*(x)\}$  and  $B = X \setminus A$ . In particular, we have  $\sigma^*(x) = \mathbf{true}$  for every  $x \in A$  and  $\sigma^*(x) = \mathbf{false}$  for every  $x \in B$ . Note that the truth assignment where every variable of  $A$  is set to **false** and every variable of  $B$  is set to **true** is also optimal. Indeed switching all variables of a clause does not change its status, i.e. it remains NAE-satisfied or NAE-unsatisfied.

Let us suppose that  $|A| = r$  and  $|B| = s$ . We rename the variables of  $A$  and  $B$  as follows. From now on  $A = \{a_1, \dots, a_r\}$  and  $B = \{b_1, \dots, b_s\}$ . If  $A_j$  denotes  $\{a_j, a_{j+1}, \dots, a_r\}$  for  $j = 1, \dots, r$  (resp.,  $B_j$  denotes  $\{b_j, b_{j+1}, \dots, b_s\}$  for  $j = 1, \dots, s$ ) then we suppose that the player associated with  $a_j$  (resp.,  $b_j$ ) does not benefit when every  $a \in A_j$  (resp., every  $b \in B_j$ ) plays **false** while the others play **true**. Notice that this renaming is well defined because  $\sigma$  is a strong equilibrium. Actually, when players in  $A_j$  form a coalition, then at least one player does not benefit because  $\sigma$  is a strong equilibrium.

We define some subsets of  $\mathcal{C}$  as follows:

- A clause  $c \in \mathcal{C}$  belongs to  $\zeta_j^A$  (resp.  $\zeta_j^B$ ) iff  $\sigma$  NAE-satisfies  $c$ ,  $c$  contains a literal defined upon  $a_j$  (resp.  $b_j$ ) and  $c$  is not NAE-satisfied by the truth assignment where the variables of  $\{a_j, a_{j+1}, \dots, a_r\}$  (resp.  $\{b_j, b_{j+1}, \dots, b_s\}$ ) are **false** while any other variable is **true**. Let  $\zeta^A = \bigcup_{j=1}^r \zeta_j^A$  and let  $\zeta^B = \bigcup_{j=1}^s \zeta_j^B$ .
- A clause  $c \in \mathcal{C}$  belongs to  $\chi_j^A$  (resp.  $\chi_j^B$ ) iff  $\sigma$  NAE-satisfies  $c$ ,  $c$  contains a literal defined upon  $a_j$  (resp.  $b_j$ ) and  $c \notin \zeta_j^A$  (resp.  $c \notin \zeta_j^B$ ). Let  $\chi^A = \bigcup_{j=1}^r \chi_j^A$  and let  $\chi^B = \bigcup_{j=1}^s \chi_j^B$ .
- A clause  $c \in \mathcal{C}$  belongs to  $\alpha_j^A$  (resp.  $\alpha_j^B$ ) iff  $\sigma$  does not NAE-satisfy  $c$ ,  $c$  contains a literal defined upon  $a_j$  (resp.  $b_j$ ) and  $c$  is NAE-satisfied by the truth assignment where the variables of  $\{a_j, a_{j+1}, \dots, a_r\}$  (resp.  $\{b_j, b_{j+1}, \dots, b_s\}$ ) are **false** while any other variable is **true**. Let  $\alpha^A = \bigcup_{j=1}^r \alpha_j^A$  and let  $\alpha^B = \bigcup_{j=1}^s \alpha_j^B$ .
- A clause  $c \in \mathcal{C}$  belongs to  $\beta_j^A$  (resp.  $\beta_j^B$ ) iff  $\sigma$  does not NAE-satisfy  $c$ ,  $c$  contains a literal defined upon  $a_j$  (resp.  $b_j$ ) and  $c \notin \alpha_j^A$  (resp.  $c \notin \alpha_j^B$ ). Let  $\beta^A = \bigcup_{j=1}^r \beta_j^A$  and let  $\beta^B = \bigcup_{j=1}^s \beta_j^B$ .

In what follows,  $w(C)$  denotes the weight of a given set of clauses  $C$ . Let us give some intermediate properties.

*Property 1.*  $\zeta_j^A \cap \zeta_{j'}^A = \emptyset$  for all  $j, j'$  such that  $1 \leq j < j' \leq r$  and  $\zeta_j^B \cap \zeta_{j'}^B = \emptyset$  for all  $j, j'$  such that  $1 \leq j < j' \leq s$ .

*Property 2.*  $\sigma^*$  does not NAE-satisfy any clause  $c \in \alpha^A \Delta \alpha^B$ .

*Property 3.*  $\sigma^*$  does not NAE-satisfy any clause  $c \in \beta^A \cup \beta^B$ .

*Property 4.*  $\sigma^*$  does not NAE-satisfy any clause  $c \in \zeta^A \cap \zeta^B$ .

Using Properties (2), (3), (4) and  $\mathcal{C} = \alpha^A \cup \alpha^B \cup \beta^A \cup \beta^B \cup \zeta^A \cup \zeta^B \cup \chi^A \cup \chi^B$  we can give the following bound on  $\mathcal{Q}(\sigma^*)$ :

$$\begin{aligned} \mathcal{Q}(\sigma^*) &\leq w(\alpha^A \cup \alpha^B) + w(\beta^A \cup \beta^B) + w(\zeta^A \cup \zeta^B) + w(\chi^A \cup \chi^B) \\ &\quad - (w(\alpha^A \Delta \alpha^B) + w(\beta^A \cup \beta^B) + w(\zeta^A \cap \zeta^B)) \\ &= w(\alpha^A \cap \alpha^B) + w(\zeta^A \Delta \zeta^B) + w(\chi^A \cup \chi^B) \end{aligned} \tag{1}$$

The value of  $\mathcal{Q}(\sigma)$  is as follows:

$$\mathcal{Q}(\sigma) = w(\zeta^A \cup \zeta^B \cup \chi^A \cup \chi^B) = w(\zeta^A \cup \zeta^B) + w(\chi^A \cup \chi^B) \tag{2}$$

Take any variable  $a_j \in A$ . The utility of the associated player in the SE  $\sigma$  is  $w(\zeta_j^A) + w(\chi_j^A)$ . This utility becomes  $w(\alpha_j^A) + w(\chi_j^A)$  if each player in the coalition  $\{a_j, \dots, a_r\}$  sets his variable to **false**. By construction  $a_j$  does not benefit. Therefore  $w(\zeta_j^A) + w(\chi_j^A) \geq w(\alpha_j^A) + w(\chi_j^A)$  which is equivalent to  $w(\zeta_j^A) \geq w(\alpha_j^A)$ . Summing up this inequality for  $j = 1$  to  $r$  and using Property **□**, we obtain:

$$w(\zeta^A) = \sum_{j=1}^r w(\zeta_j^A) \geq \sum_{j=1}^r w(\alpha_j^A) \geq w(\alpha^A) \geq w(\alpha^A \cap \alpha^B) \tag{3}$$

One can conduct the same analysis and obtain:

$$w(\zeta^B) = \sum_{j=1}^s w(\zeta_j^B) \geq \sum_{j=1}^s w(\alpha_j^B) \geq w(\alpha^B) \geq w(\alpha^A \cap \alpha^B) \tag{4}$$

Using inequalities **(3)** and **(4)**, we get:

$$\begin{aligned} w(\zeta^A) + w(\zeta^B) &\geq 2w(\alpha^A \cap \alpha^B) \\ w(\zeta^A) + w(\zeta^B) + 2w(\zeta^A \Delta \zeta^B) &\geq 2w(\alpha^A \cap \alpha^B) + 2w(\zeta^A \Delta \zeta^B) \\ 2w(\zeta^A \cup \zeta^B) + w(\zeta^A \Delta \zeta^B) &\geq 2w(\alpha^A \cap \alpha^B) + 2w(\zeta^A \Delta \zeta^B) \\ 3w(\zeta^A \cup \zeta^B) &\geq 2w(\alpha^A \cap \alpha^B) + 2w(\zeta^A \Delta \zeta^B) \\ 3w(\zeta^A \cup \zeta^B) + 2w(\chi^A \cup \chi^B) &\geq 2w(\alpha^A \cap \alpha^B) + 2w(\zeta^A \Delta \zeta^B) + 2w(\chi^A \cup \chi^B) \\ 3\mathcal{Q}(\sigma) &\geq 2\mathcal{Q}(\sigma^*) \end{aligned}$$

A tight example is composed of three clauses of weight one:  $x_1 \vee x_2, x_3 \vee x_4$  and  $x_1 \vee x_3$ . If  $\sigma(x_1) = \sigma(x_2) = \mathbf{true}$  and  $\sigma(x_2) = \sigma(x_4) = \mathbf{false}$  then  $\sigma$  is a SE NAE-satisfying the first two clauses. Indeed the utility of  $x_2$  and  $x_4$  is maximum in this configuration (every clause where they appear is NAE-satisfied) so they have no incentive to deviate. So when  $\sigma(x_2) = \sigma(x_4) = \mathbf{false}$ , it is not difficult to see that both  $x_1$  and  $x_3$  have the same utility as in  $\sigma$ , whatever they play. If  $\sigma(x_1) = \sigma(x_4) = \mathbf{true}$  and  $\sigma(x_2) = \sigma(x_3) = \mathbf{false}$  then the three clauses are NAE-satisfied. □

It follows that the SPoA of the MAX CUT game is  $2/3$  (the tight example is made of positive literals so it can be represented as an instance of MAX CUT). When restricting ourselves to instances of the NAE SAT game which admit a SE, the proof of Theorem **□** can be extended to prove that the SPoA is  $q/(q + 1)$  if each clause has size exactly  $q$  and  $2/3$  otherwise (it suffices to give adjusted proofs of Properties **□** to **□**).

### 5 The MAX $k$ -CUT Game

Given a graph  $G = (V, E)$  and a weight function  $w : E \rightarrow \mathbb{R}_+$ , MAX  $k$ -CUT is to partition  $V$  into  $k$  sets  $V_1, V_2 \dots V_k$  such that the sum of the weight of the edges

having their endpoints not in the same part of the partition is maximum. The MAX  $k$ -CUT game is defined as follows. Each vertex is controlled by a player with strategy set  $\{1, 2, \dots, k\}$ . A player's utility is the total weight of the edges incident to her and such that her neighbor has a different strategy. The social function  $Q$  for a state  $a$  is  $\sum_{\{[i,j] \in E: a_i \neq a_j\}} w([i, j])$ .

The MAX  $k$ -CUT game always has a pure Nash equilibrium since it can be easily defined as a congestion game but an alternative proof is to observe that an optimal  $k$ -cut is a NE (it is known that an optimum is a NE for MAX CUT). In [7] it is shown that the PoA of the *unweighted* MAX  $k$ -CUT game is  $\frac{k-1}{k}$  and one can easily extend the result to the weighted case.

Now we investigate the existence of a SE for the MAX  $k$ -CUT game. The MAX CUT game ( $k = 2$ ) always admits a SE since an optimal cut must be a SE. It is a corollary of Theorem 2. When  $k \geq 3$ , our positive result is that an optimal cut of the MAX  $k$ -CUT game is a 3-SE (proof by contradiction).

**Theorem 5.** *Every optimum of the MAX  $k$ -CUT game is a 3-SE.*

The following result states that we can not go beyond  $r = 3$  to prove that *any* optimal cut is an  $r$ -SE.

**Proposition 1.** *An optimum of the MAX 3-CUT game is not necessarily a 4-SE.*

Hence an optimum of the MAX  $k$ -CUT game is not necessarily a SE but the existence of a SE is not compromised because the instance we found to state Proposition 1 admits two optima, one is not a 4-SE whereas the other is a SE.

To conclude this section, one can be interested in bounding the SPoA of the MAX CUT game if only coalitions of limited size are conceivable, i.e. the  $q$ -SPoA. The following result shows that, even if  $q$  is large, the  $q$ -SPoA drops to  $1/2$ .

**Theorem 6.** *For any  $\varepsilon > 0$  and  $q = O(|V|^{1/2-\varepsilon})$  where  $|V|$  is the number of nodes, the  $q$ -SPoA of the MAX CUT game is  $1/2$ .*

## 6 Concluding Remarks

We investigated two games which generalize MAX CUT and the focus was on strong equilibria, their existence and how far they are from socially optimal configurations. Some questions are left open.

For the  $q$ -NAE SAT game where  $q \geq 3$ , we presented an instance without any  $(q - 1)$ -SE but can we guarantee that there is an  $r$ -SE for some  $1 < r < q - 1$ ? For example, is there a 2-SE when  $q \geq 4$ ? Another interesting direction would be to characterize instances which possess a SE.

For the MAX  $k$ -CUT game, we showed that a 3-SE exists but can we go further? Though any optimum is not guaranteed to be a 4-SE, it is possible that only some optima are 4-SE. Actually we conjecture that the MAX  $k$ -CUT game always possesses a SE. If it is true then what would the SPoA?

## References

1. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: Proc. of SODA 2007, pp. 189–198 (2007)
2. Aumann, R.: Acceptable points in general cooperative n-person games. In: Contribution to the Theory of Games. Annals of Mathematics Studies, vol. IV, 40, pp. 287–324 (1959)
3. Balcan, M.-F., Blum, A., Mansour, Y.: Improved equilibria via public service advertising. In: Proc. of SODA 2009, pp. 728–737 (2009)
4. Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and Approximation in Potential Games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 349–360. Springer, Heidelberg (2006)
5. Epstein, A., Feldman, M., Mansour, Y.: Strong equilibrium in cost sharing connection games. *Games and Economic Behavior*, vol 67(1), 51–68 (2008)
6. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The complexity of pure Nash equilibria. In: Proc. of STOC 2004, pp. 604–612 (2004)
7. Hoefer, M.: Cost Sharing and Clustering under Distributed Competition, Ph.D Thesis, Universität Konstanz (2007)
8. Holzman, R., Law-Yone, N.: Strong Equilibrium in Congestion Games. *Games and Economic Behavior* 21(1-2), 85–101 (1997)
9. Holzman, R., Law-Yone, N.: Network structure and strong equilibrium in route selection games. *Mathematical Social Science* 46, 193–205 (2003)
10. Koutsoupias, E., Papadimitriou, C.: Worst Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
11. Monderer, D., Shapley, L.S.: Potential Games. *Games and Economic Behavior* 14, 124–143 (1996)
12. Quint, T., Shubik, M.: A model of migration, mimeo (1994)
13. Poljak, S.: Integer Linear Programs and Local Search for Max-Cut. *SIAM J. Comput.* 24(4), 822–839 (1995)
14. Rozenfeld, O., Tennenholtz, M.: Strong and Correlated Strong Equilibria in Monotone Congestion Games. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 74–86. Springer, Heidelberg (2006)
15. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2(1), 65–67 (1973)
16. Schäffer, A.A., Yannakakis, M.: Simple Local Search Problems That are Hard to Solve. *SIAM J. Comput.* 20(1), 56–87 (1991)

# Stability and Convergence in Selfish Scheduling with Altruistic Agents

Martin Hoefer\* and Alexander Skopalik\*\*

Dept. of Computer Science, RWTH Aachen University, Germany  
{mhoefer,skopalik}@cs.rwth-aachen.de

**Abstract.** In this paper we consider altruism, a phenomenon widely observed in nature and practical applications, in the prominent model of selfish load balancing with coordination mechanisms. Our model of altruistic behavior follows recent work by assuming that agent incentives are a trade-off between selfish and social objectives. In particular, we assume agents optimize a linear combination of personal delay of a strategy and the resulting social cost. Our results show that even in very simple cases a variety of standard coordination mechanisms are not robust against altruistic behavior, as pure Nash equilibria are absent or better response dynamics cycle. In contrast, we show that a recently introduced TIME-SHARING policy yields a potential game even for partially altruistic agents. In addition, for this policy a Nash equilibrium can be computed in polynomial time. In this way our work provides new insights on the robustness of coordination mechanisms. On a more fundamental level, our results highlight the limitations of stability and convergence when altruistic agents are introduced into games with weighted and lexicographical potential functions.

## 1 Introduction

One of the most fundamental scenarios in algorithmic game theory are selfish load balancing models [19]. Since the seminal paper by Koutsoupias and Papadimitriou [16] they have attracted a large amount of interest. The reasons are central applications in distributed processing, conceptual simplicity, and that they contain in a nutshell many prominent challenges in designing distributed systems for selfish participants. A fundamental assumption in the vast majority of previous work is that all agents are selfish. Their goals are restricted to optimizing their direct personal delay. However, this assumption has been repeatedly questioned by economists and psychologists. In experiments it has been observed that participants' behavior can be quite complex and contradictory to selfishness [17, 18]. Various explanations have been given for this phenomenon, e.g. senses of fairness [8], reciprocity among agents [13], or spite and altruism [6, 18].

---

\* Supported by a fellowship within the Postdoc-Program of the German Academic Exchange Service (DAAD) and by DFG through UMIC Research Centre at RWTH Aachen University.

\*\* Supported in part by the German Israeli Foundation (GIF) under contract 877/05.



In this paper, we consider altruism in non-cooperative load balancing games. It is natural to study the effects of an important phenomenon like altruism in a core scenario of algorithmic game theory. Our model of altruism is similar to the one used recently in [3, 14] and related to the study of coalitional stability concepts [10, 11], although we do not require agent cooperation in our model. Instead, each agent  $i$  is assumed to be partly selfish and partly altruistic. Her incentive is to optimize a linear combination of personal cost and social cost, given by the sum of cost values of all agents. The strength of altruism of each agent  $i$  is captured by her altruism level  $\beta_i \in [0, 1]$ , where  $\beta_i = 0$  results in a purely selfish and  $\beta_i = 1$  in a purely altruistic agent.

We consider altruistic agents in various types of scheduling games resulting from coordination mechanisms [4]. In these games agents are tasks, and each task chooses to allocate one out of several machines. For a machine the coordination mechanism is a local scheduling policy that determines the schedule of the tasks which choose to allocate the machine. Quite a number of policies have been proposed [1, 2, 4, 5, 15], mostly with the objective to minimize the price of anarchy [16] for makespan social cost. In addition to modelling a natural phenomenon, altruistic agents yield a measure of robustness for these mechanisms. Our results provide an interesting distinction between the studied policies in terms of stability and convergence properties. In addition, they also shed some light on an interesting and more fundamental aspect. Previously, we studied altruists in atomic congestion games [14], which have an exact potential function. For atomic games, there are a number of special cases, in which a potential function argument guarantees existence of pure Nash equilibria and convergence of better response dynamics even for games with altruists. These cases include games with linear delay functions, or  $\beta$ -uniform agents that all have the same altruism level  $\beta$ . In this paper we analyze altruism in arguably the most basic games with weighted and lexicographical potential functions, and we expect our results to hold similarly e.g. for other coordination mechanisms based on lexicographical improvement arguments [2]. After addition of altruists, potential functions are largely absent here, even for identical machines or  $\beta$ -uniform agents. In contrast, the very positive results for the TIME-SHARING policy rely on the existence of an exact potential for the original game and the construction is very similar to [14]. It is an interesting open problem to see if there is a connection between these cases, or if a general characterization of the existence of potential functions under altruistic behavior can be derived.

## 1.1 Our Results

We study altruistic agents with four different coordination mechanisms. At first in Section 3 we consider the classic MAKESPAN policy [16], which is probably the most widely studied policy and yields a weighted potential function. For altruistic agents we show that this favorable property breaks down. There are games without pure Nash equilibria, and deciding this property for a game is NP-hard, even on identical machines. In Section 4 we study simple ordering based policies like SHORTEST-FIRST and LONGEST-FIRST that yield a lexicographic potential

for non-altruistic users [15]. While for SHORTEST-FIRST on identical machines existence of a pure Nash equilibrium is guaranteed even for arbitrary altruism levels, the resulting games are no potential games as better response dynamics might cycle. For LONGEST-FIRST we additionally show that there are games without pure Nash equilibria. Finally, in Section 5 we consider the TIME-SHARING policy introduced in [5]. While the policy is somewhat similar to MAKESPAN, the results are completely different. For this policy we show the existence of a potential function, even for arbitrary altruism levels and unrelated machines. Thus, existence of pure Nash equilibria and convergence of better response dynamics is always guaranteed. In addition, we show how to compute a Nash equilibrium in polynomial time. Due to lack of space most proofs are omitted in this extended abstract.

## 2 Scheduling with Coordination Mechanisms

We consider scheduling games with coordination mechanisms [4]. A scheduling game  $G$  consists of a set  $N$  of  $n$  agents and a set  $M$  of  $m$  machines. Each agent  $i \in N$  is a *task* and picks as a strategy the machine it wants to be processed on. In the case of identical machines, task  $i$  has processing time  $p_i$  on every machine. In case of related machines there is a speed factor  $s_j$  for machine  $j$ , and the processing time of  $i$  on  $j$  becomes  $p_i/s_j$ . For unrelated machines there is a separate processing time  $p_{ij}$  for every task  $i$  and machine  $j$ .

The strategy choices of the tasks result in a schedule  $s : N \rightarrow M$ , an assignment of every task to exactly one machine. On each machine there is a *coordination mechanism*, i.e. a sequencing policy that sequences the tasks and assigns starting and finishing time for each task. We assume here that machines cannot preempt tasks, but depending on the mechanism a machine might process one or more tasks simultaneously. For a given sequencing policy SP on the machines, we define the social cost of a schedule as  $c^{SP}(s) = \sum_j f_j(s)$ , where  $f_j(s)$  is finishing time of task  $j$  in schedule  $s$ . To model altruism we use for each task  $i$  the *altruism level*  $\beta_i$  [3, 14]. If  $\beta_i > 0$ , we call task  $i$  an altruist. If  $\beta_i = 1$  we call task  $i$  a pure altruist, if  $\beta_i = 0$  we call him an egoist. The individual cost of a task  $i$  incorporates the effect on the social cost:  $c_i^{SP}(s) = \beta_i c^{SP}(s) + (1 - \beta_i) f_i(s) = f_i(s) + \beta_i \sum_{j \neq i} f_j(s)$ . A pure Nash equilibrium of the game is a schedule, in which no task can decrease his individual cost with a unilateral strategy change. Clearly, if all tasks are pure altruists, then every game on unrelated machines has a pure Nash equilibrium and every sequential better response dynamics converges.

## 3 Makespan and Random Policies

The first and most widely studied policy is the MAKESPAN policy [16], in which all tasks on one machine are processed simultaneously and finish at the same time. In the RANDOM policy [15] tasks are ordered in a random order and then processed consecutively in this order. Obviously, RANDOM and MAKESPAN

are equivalent in terms of (expected) finishing times on identical and related machines.

MAKESPAN induces a weighted potential game. Let  $\ell_j = \sum_{i : s_i=j} p_{ij}$  be the load of tasks choosing machine  $j$ . For identical machines the weighted potential is  $\Phi(s) = \sum_{j=1}^m \ell_j^2$ . For a task  $i$  we have  $c_i^{MS}(s) - c_i^{MS}(s'_i, s_{-i}) = \frac{1}{p_i}(\Phi(s) - \Phi(s'_i, s_{-i}))$ . This potential is easily extended to related machines [7]. For the MAKESPAN policy it is shown in [9] that for a population of only egoists best response dynamics can take  $O(2^{\sqrt{n}})$  steps to converge to a pure Nash equilibrium. For identical machines there is a scheduling of tasks to reach a Nash equilibrium with better response dynamics in polynomial time. In addition, there are polynomial time algorithms to compute Nash equilibria on related machines and instances with link restrictions [9,12].

Including altruists provides a quite different set of results. We observe that even if there is only one altruist, existence of a pure Nash equilibrium is not guaranteed.

**Proposition 1.** *There is a game on two identical machines with the MAKESPAN or RANDOM policy, one altruist, and appropriately many egoists that has no pure Nash equilibrium.*

*Proof.* Consider a game with two machines, one pure altruist with  $p_1 = 5$ , and four egoists with  $p_2 = 10, p_3 = p_4 = p_5 = 1$ . Assume there is a pure Nash equilibrium. In an equilibrium, task 2 chooses a different machine than task 1. The tasks 3, 4, and 5 choose a different machine than task 2. However, task  $p_1$  would choose the machine with only task 2, which leads to a contradiction. For an altruist with general  $\beta_1 > 0$ , we add approximately  $1/\beta_1$  many egoists with  $p_i = 1$  and scale  $p_1$  and  $p_2$  accordingly in order to preserve the argument.  $\square$

In addition, we can show that it is NP-hard to decide if a pure Nash equilibrium exists. The reduction is from PARTITION.

**Theorem 2.** *It is weakly NP-hard to decide if a game on three identical machines with MAKESPAN and one pure altruist has a pure Nash equilibrium.*

*Proof.* We reduce from PARTITION. An instance  $\mathcal{I}$  is given as  $(a_1, \dots, a_n) \in \mathbb{N}^n$  and  $\mathcal{I} \in \text{PARTITION}$  if and only if  $\exists I \subset \{1, \dots, n\}$  with  $\sum_{i \in I} a_i = \sum_{j \in \{1, \dots, n\} \setminus I} a_j$ . We first reduce a given instance  $\mathcal{I} = (a_1, \dots, a_n)$  to an instance  $\mathcal{I}' = (a_1, \dots, a_n, a_{n+1}, \dots, a_{n+8})$  with  $a_{n+1} = \dots = a_{n+8} = \sum_{i \in \{1, \dots, n\}} a_i$ . Clearly  $\mathcal{I} \in \text{PARTITION}$  if and only if  $\mathcal{I}' \in \text{PARTITION}$ .

In a second step we construct a scheduling game  $G_{\mathcal{I}'}$  that has a pure Nash equilibrium if and only if  $\mathcal{I}' \in \text{PARTITION}$ . The game consists of three machines and  $n + 8 + 2$  tasks. The processing time  $p_i$  of task  $1 \leq i \leq n + 8$  is  $a_i$ . Task  $n + 9$  has processing time  $p_{n+9} = \sum_{1 \leq j \leq n+8} a_j$  and task  $n + 10$  has processing time  $p_{n+10} = \frac{1}{2} \sum_{1 \leq j \leq n+8} a_j$ . All tasks are pure egoists except for task  $n + 10$  who is a pure altruist.

If  $\mathcal{I} \in \text{PARTITION}$ , there is an  $I \subset \{1, \dots, n + 8\}$  with  $\sum_{i \in I} a_i = \frac{1}{2} \sum_{1 \leq j \leq n+8} a_j$ . Scheduling all tasks  $i \in I$  on machine one, all tasks  $j \in$

$\{1, \dots, n + 8\} \setminus I$  on machine two, and the remaining tasks  $n + 9$  and  $n + 10$  on machine three is a pure Nash equilibrium. If  $\mathcal{I} \notin \text{PARTITION}$ , one can show that there is no pure Nash equilibrium.  $\square$

## 4 Policies with Global Ordering

Probably the simplest policy with global ordering is the SHORTEST-FIRST policy, in which each machine orders tasks shortest-first depending on their processing time and processes them consecutively in this order. There is a lexicographic potential [4, 15], and every better response dynamics in a game of only egoists converges to a pure Nash equilibrium. In addition, there is a scheduling of better response moves such that a Nash equilibrium is reached in polynomial time [15]. In addition, for identical machines this pure Nash equilibrium is essentially unique and coincides with the social optimum. This implies that for identical machines and SHORTEST-FIRST there always exists a pure Nash equilibrium for any altruistic population of tasks.

**Proposition 3.** *For a game on identical machines with SHORTEST-FIRST policy there is always a pure Nash equilibrium for any altruistic population of tasks.*

For a population of pure altruists suboptimal Nash equilibria can evolve. This means that the presence of altruists actually deteriorates the social cost of stable solutions.

**Proposition 4.** *The price of anarchy in scheduling games with SHORTEST-FIRST and only pure altruists is at least  $9/8$ .*

Let us further examine convergence properties of best-response dynamics. We use the above game to construct a cycling sequence even for uniform altruists, for any  $\beta \in (0, 1)$ .

**Theorem 5.** *Best-response dynamics do not converge to a pure Nash equilibrium, even for two identical machines with SHORTEST-FIRST, for (1) three egoists and one pure altruist; or (2) four  $\beta$ -uniform altruists, for every  $\beta \in (0, 1)$ .*

In the remainder of this section we briefly discuss another simple ordering policy, namely LONGEST-FIRST. For entirely egoistic populations this policy yields a potential game for identical and related machines. It has recently been shown that for three unrelated machines LONGEST-FIRST does not guarantee a pure Nash equilibrium [5]. When it comes to heterogeneous populations, it is possible to show that even on identical machines pure Nash equilibria can be absent.

**Theorem 6.** *There are games that have no pure Nash equilibrium on two identical machines with LONGEST-FIRST policy and (1) one altruist, and five egoists; or (2) six  $\beta$ -uniform altruists, for any  $\beta \in (0, 1/3)$ .*

## 5 Time-Sharing Policy

In contrast to the previous results, we show here that there is a policy closely related to MAKESPAN and SHORTEST-FIRST, for which stabilization is robust against arbitrary altruistic behavior. The TIME-SHARING policy is inspired by generalized processor sharing. It has recently been studied as a coordination mechanism in [5]. All tasks are started simultaneously, and all tasks are processed in equal shares by the machine. When the smallest task is finished, the machine is shared in equal parts by the remaining tasks, and so on. For a population of only egoists the policy yields an exact potential function, even on unrelated machines. The potential function can be rewritten as the sum of completion times  $c^{SF}(s)$  for the same assignment and the SHORTEST-FIRST policy. This turns out to be the sum of completion times  $c^{TS}(s)$  for TIME-SHARING with a correction term. Using straightforward calculation it is possible to show

$$\Phi(s) = c^{SF}(s) = \frac{1}{2} \left( c^{TS}(s) + \sum_i p_{i,s_i} \right) .$$

This allows us to derive the following result.

**Theorem 7.** *For any population of tasks on unrelated machines with the TIME-SHARING policy, a pure Nash equilibrium always exists and any better response dynamics converges.*

*Proof.* We can construct a weighted potential using  $\Phi$  and add a set of correction terms. This is essentially the same approach as for the case of linear delays in [14]. In particular, we get

$$\Phi_w(s) = \Phi(s) - \sum_i p_{i,s_i} \cdot \frac{\beta_i}{1 + \beta_i} = \frac{1}{2} \left( c^{TS}(s) + \sum_i p_{i,s_i} \cdot \frac{1 - \beta_i}{1 + \beta_i} \right) .$$

Suppose task  $i$  switches from  $s_i$  to  $s'_i$ . We denote the resulting states by  $s$  and  $s' = (s'_i, s_{-i})$ . Then,

$$\begin{aligned} c_i^{TS}(s) - c_i^{TS}(s') &= (1 - \beta_i)(f_i(s) - f_i(s')) + \beta_i(c^{TS}(s) - c^{TS}(s')) \\ &= (1 - \beta_i)(\Phi(s) - \Phi(s')) + \beta_i(c^{TS}(s) - c^{TS}(s')) \\ &= \frac{1 + \beta_i}{2} \cdot (c^{TS}(s) - c^{TS}(s')) + \frac{1 - \beta_i}{2} \cdot (p_{i,s_i} - p_{i,s'_i}) \\ &= (1 + \beta_i) \cdot (\Phi_w(s) - \Phi_w(s')) . \quad \square \end{aligned}$$

This implies existence of pure Nash equilibria and convergence of every better response dynamics. In addition, we show that computing a Nash equilibrium can be done in polynomial time.

**Theorem 8.** *For any population of tasks on unrelated machines with the TIME-SHARING policy, a pure Nash equilibrium can be computed in polynomial time.*

## References

1. Azar, Y., Jain, K., Mirrokni, V.: (Almost) optimal coordination mechanisms for unrelated machine scheduling. In: Proc. 19th Symp. Discrete Algorithms (SODA), pp. 323–332 (2008)
2. Caragiannis, I.: Efficient coordination mechanisms for unrelated machine scheduling. In: Proc. 20th Symp. Discrete Algorithms (SODA), pp. 815–824 (2009)
3. Chen, P., Kempe, D.: Altruism, selfishness, and spite in traffic routing. In: Proc. 9th Conf. Electronic Commerce (EC), pp. 140–149 (2008)
4. Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination mechanisms. *Theor. Comput. Sci.* 410(36), 3327–3336 (2009)
5. Dürr, C., Thang, N.: Non-clairvoyant scheduling games. In: Proc. 2nd Intl. Symp. Algorithmic Game Theory, SAGT (to appear, 2009)
6. Eshel, I., Samuelson, L., Shaked, A.: Altruists, egoists and hooligans in a local interaction model. *Amer. Econ. Rev.* 88(1), 157–179 (1998)
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. Fehr, E., Schmidt, K.: A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics* 114, 817–868 (1999)
9. 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)
10. 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)
11. Fotakis, D., Kontogiannis, S., Spirakis, P.: Atomic congestion games among coalitions. *ACM Trans. Algorithms* 4(4) (2008)
12. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B.: Computing Nash equilibria for scheduling on restricted parallel links. In: Proc. 36th Symp. Theory of Computing (STOC), pp. 613–622 (2004)
13. Gintis, H., Bowles, S., Boyd, R., Fehr, E.: *Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life*. MIT Press, Cambridge (2005)
14. Hoefer, M., Skopalik, A.: Altruism in atomic congestion games. In: Proc. 17th European Symposium on Algorithms (ESA), pp. 179–189 (2009)
15. Immorlica, N., Li, L., Mirrokni, V., Schulz, A.: Coordination mechanisms for selfish scheduling. *Theor. Comput. Sci.* 410(17), 1589–1598 (2009)
16. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Proc. 16th Symp. Theoretical Aspects of Computer Science (STACS), pp. 404–413 (1999)
17. Ledyard, J.: Public goods: A survey of experimental research. In: Kagel, J., Roth, A. (eds.) *Handbook of Experimental Economics*, pp. 111–194. Princeton University Press, Princeton (1997)
18. Levine, D.: Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics* 1, 593–622 (1998)
19. Vöcking, B.: Selfish load balancing. In: Nisan, N., Tardos, É., Roughgarden, T., Vazirani, V. (eds.) *Algorithmic Game Theory*, ch. 20. Cambridge University Press, Cambridge (2007)

# Gaming Dynamic Parimutuel Markets

Qianya Lin<sup>1,\*</sup> and Yiling Chen<sup>2</sup>

<sup>1</sup> City University of Hong Kong, Hong Kong SAR

<sup>2</sup> Harvard University, Cambridge, MA, USA

**Abstract.** We study the strategic behavior of risk-neutral non-myopic agents in *Dynamic Parimutuel Markets* (DPM). In a DPM, agents buy or sell shares of contracts, whose future payoff in a particular state depends on aggregated trades of all agents. A forward-looking agent hence takes into consideration of possible future trades of other agents when making its trading decision. In this paper, we analyze non-myopic strategies in a two-outcome DPM under a simple model of incomplete information and examine whether an agent will truthfully reveal its information in the market. Specifically, we first characterize a single agent's optimal trading strategy given the payoff uncertainty. Then, we use a two-player game to examine whether an agent will truthfully reveal its information when it only participates in the market once. We prove that truthful betting is a Nash equilibrium of the two-stage game in our simple setting for uniform initial market probabilities. However, we show that there exists some initial market probabilities at which the first player has incentives to mislead the other agent in the two-stage game. Finally, we briefly discuss when an agent can participate more than once in the market whether it will truthfully reveal its information at its first play in a three-stage game. We find that in some occasions truthful betting is not a Nash equilibrium of the three-stage game even for uniform initial market probabilities.

## 1 Introduction

Prediction markets are used to aggregate dispersed information about uncertain events of interest and have provided accurate forecasts of event outcomes, often outperforming other forecasting methods, in many real-world domains [1,2,3,4,5,6,7,8]. To achieve its information aggregation goal, a prediction market for an uncertain event offers contracts whose future payoff is tied to the event outcome. For example, a contract that pays off \$1 per share if there are more than 6,000 H1N1 flu cases confirmed in U.S. by August 30, 2009 and \$0 otherwise can be traded to predict the likelihood of the specified activity level of H1N1 flu.

Most market mechanisms used by prediction markets, including *continuous double auctions* (CDA) and *market scoring rules* (MSR) [9,10], trade contracts whose payoff in each state is fixed, as in the above example. Contracts in *dynamic parimutuel markets* (DPM) [11,12], however, have variable payoff that depends

---

\* Part of this work was done while Qianya Lin was visiting Harvard University.



on the aggregated trades of all market participants. The payoff uncertainty makes DPM a mechanism that admits more speculation and strategic play.

As the goal of prediction markets is to aggregate information, it is important to understand whether and how participants reveal their information in the market. In this paper, we study the strategic behavior of risk-neutral non-myopic agents in a two-outcome DPM under a simple setting of incomplete information, with the intent to understand how forward-looking agents reveal their information in DPM and whether they will reveal their information truthfully. We first characterize a single agent's optimal trading strategy given payoff uncertainty. Then, we consider a two-player two-stage dynamic game where each player only participates once in DPM, to examine whether the first player has incentives to misreport its information to mislead the second player and obtain higher profit even if it can only play once. We prove that truthful betting is a Nash equilibrium of the two-stage game for uniform initial market probabilities. We show that there exists some initial market probabilities at which the first player has incentives to mislead the other agent in the two-stage game. Finally, we discuss a three-stage game in which an agent can participate more than once. We find that the truthful betting is not a Nash equilibrium of the three-stage game in some occasions even for uniform initial market probabilities.

**Related Work.** Chen et al. [13] provide a specific example of a two-player two-stage game in DPM where the second player is perfectly informed and show that the first player may sometimes choose not to trade. Our work does not assume perfectly informed agents and we characterize non-myopic strategies in more general settings. The example of Chen et al. is a special case of our results for the two-player two-stage game. Nikolova and Sami [14] use a projection game to study DPM. They show that a rational agent will never hold shares of both outcomes in a two-outcome DPM when short sales are not allowed. This is consistent with our characterization of a single agent's optimal strategy given payoff uncertainty. Bu et al. [15] study the strategies of a myopic agent who believes that the contract payoff in the future is the same as the payoff if the market closes right after the agent's trade in a DPM. Our work focuses on forward-looking agents who take into consideration of the payoff uncertainty when making their trading decisions.

Some theoretical attempts have been made to characterize non-myopic strategies in other markets, including logarithmic market scoring rule (LMSR) [13, 16, 17], financial markets (i.e. CDA) [18, 19, 20], and parimutuel markets [21]. In all these markets, agents may have incentives to misreport their information. Ostrovsky [22] provides a separability condition that contracts need to satisfy to guarantee market convergence to full information aggregation at a perfect Bayesian equilibrium in LMSR and CDA.

## 2 Dynamic Parimutuel Markets

A dynamic parimutuel market (DPM) [11, 12] is a dynamic-cost variant of a parimutuel market. Suppose an uncertain event of interest has  $n$  mutually



exclusive outcomes. Let  $\Omega$  denote the outcome space. A DPM offers  $n$  contracts, each corresponding to an outcome. As in a parimutuel market, traders who wager on the true outcome split the total pool of money at the end of the market. However, the price of a single share varies dynamically according to a price function, hence incentivizing traders to reveal their information earlier.

DPM operates as a market maker. Let  $q_\omega$  be the total number of shares of contract  $\omega$  that have been purchased by all traders. We use  $\mathbf{q}$  to denote the vector of outstanding shares for all contracts. The DPM market maker keeps a cost function,  $C(\mathbf{q}) = \sqrt{\sum_{\omega \in \Omega} q_\omega^2}$ , that captures total money wagered in the market, and an instantaneous price function for contract  $\omega$ ,  $p_\omega = \frac{q_\omega}{\sqrt{\sum_{\psi \in \Omega} q_\psi^2}}$ .

A trader who buys contracts and changes the outstanding shares from  $\mathbf{q}$  to  $\tilde{\mathbf{q}}$  pays the market maker  $C(\tilde{\mathbf{q}}) - C(\mathbf{q})$ . The market probability on outcome  $\omega$  is  $\pi_\omega = \frac{q_\omega^2}{\sum_{\psi \in \Omega} q_\psi^2}$ . In DPM, market price of a contract does not represent the market probability of the corresponding state. Instead,  $\pi_\omega = p_\omega^2$ .

If outcome  $\omega$  is realized, each share of contract  $\omega$  gets an equal share of the total market capitalization. Its payoff is  $o_\omega = \frac{\sqrt{\sum_{\psi \in \Omega} (q_\psi^f)^2}}{q_\omega^f}$ , where  $q_\omega^f$  is the outstanding shares of contract  $\omega$  at the end of the market. All other contracts have zero payoff. As the value of  $\mathbf{q}^f$  is not known before the market closes,  $o_\omega$  is not fixed while the market is open. The relation of the final market price, final market probability, and the contract payoff when outcome  $\omega$  is realized, is  $o_\omega = \frac{1}{p_\omega^f} = \frac{1}{\pi_\omega^f}$ , where  $p_\omega^f$  and  $\pi_\omega^f$  denote the last market price and market probability before the market closes.

As a market maker mechanism, DPM offers infinite liquidity. Because the price function is not defined when  $\mathbf{q} = 0$ , the market maker subsidizes the market by starting the market with some positive shares. The subsidy turns DPM into a positive-sum game and can circumvent the *no-trade theorem* [23] for zero-sum games. Tech Buzz Game [12] used DPM as its market mechanism and market probabilities in the game have been shown to offer informative forecasts for the underlying events [24].

### 3 Our Setting

We consider a simple incomplete information setting for a DPM in this paper. There is a single event whose outcome space contains two discrete mutually exclusive states  $\Omega = \{Y, N\}$ . The eventual event outcome is picked by Nature with prior probability  $P(Y) = P(N) = \frac{1}{2}$ . The DPM offers two contracts, each corresponding to one outcome. There are two players in the market. Each player  $i$  receives a piece of private signal  $c_i \in \{y_i, n_i\}$ . The signal is independently drawn by Nature conditional on the true state. In other words, signals are conditionally independent,  $P(c_i, c_j | \omega) = P(c_i | \omega)P(c_j | \omega)$ . The prior probabilities and the signal distributions are common knowledge to all players.

We further assume that player's signals are symmetric such that  $P(y_i | Y) = P(n_i | N)$  for all  $i$ . With this, we define the *signal quality* of player  $i$  as  $\theta_i =$

$P(y_i|Y) = P(n_i|N)$ . The signal quality  $\theta_i$  captures the likelihood for agent  $i$  to receive a “correct” signal. Without loss of generality, we assume  $\theta_i \in (\frac{1}{2}, 1]$ . The above conditions, together with conditional independence of signals, imply that for two players  $i$  and  $j$  we have  $P(Y|c_i, y_j) > P(Y|c_i, n_j)$ ,  $P(N|c_i, n_j) > P(N|c_i, y_j)$ ,  $P(y_i|y_j) > P(n_i|y_j)$ , and  $P(n_i|n_j) > P(y_i|n_j)$  for  $c_i \in \{y_i, n_i\}$ .

Short sell is not allowed in the market. Agents are risk neutral and participate in the market sequentially. We also assume that they do not possess any shares at the beginning of the market and have unlimited wealth.

### 4 Optimal Trading Strategy of a Single Agent

We first consider a single agent’s optimal trading strategy given the payoff uncertainty in DPM. We assume that the agent only trades once in DPM. Let  $P(\omega, s)$  be the agent’s subjective probability that the event outcome will be  $\omega$  and the set of information available to the last trader is  $s$ . The market probability at the end of the market will reflect all available information. Hence the final price for contract  $\omega$  when  $s$  is available is  $\pi_\omega^f(s) = P(\omega|s)$ .

The agent compares the current price of a contract,  $p_\omega$ , with its expected future payoff. Note that the future payoff of contract  $\omega$  in state  $\omega$  only relates to the final market probability  $\pi_\omega^f$  and does not relate to the process of reaching it. The expected future payoff of contract  $\omega$  is  $\varphi_\omega = \sum_s \frac{P(\omega, s)}{\sqrt{\pi_\omega^f(s)}} = \sum_s \frac{P(\omega, s)}{\sqrt{P(\omega|s)}} = \sum_s P(s)\sqrt{P(\omega|s)}$ . We have the following lemma.

**Lemma 4.1.**  $\sum_\omega \varphi_\omega^2 \leq 1$ .

Suppose the agent purchases  $\Delta \mathbf{q}$  and changes the outstanding shares from  $\mathbf{q}$  to  $\tilde{\mathbf{q}} = \mathbf{q} + \Delta \mathbf{q}$ . The market prices before and after the trade are  $\mathbf{p}$  and  $\tilde{\mathbf{p}}$  respectively. Theorem 4.2 characterizes the agent’s optimal purchases when it attempts to maximize its expected profit.

**Theorem 4.2.** *In a two-outcome DPM, if a risk-neutral agent maximizes its expected profit by purchasing  $\Delta \mathbf{q} \geq 0$ , the following conditions must satisfy:*

1. For any contract  $\omega$ , if  $p_\omega < \varphi_\omega$ , then  $\Delta q_\omega > 0$  and  $\tilde{p}_\omega = \varphi_\omega$ .
2. For any contract  $\omega$ , if  $p_\omega > \varphi_\omega$ , then  $\tilde{p}_\omega \geq \varphi_\omega$  and when the inequality is strict,  $\Delta q_\omega = 0$ .
3. For any contract  $\omega$ , if  $p_\omega > \varphi_\omega$ ,  $\tilde{p}_\omega = \varphi_\omega$ , and  $\Delta q_\omega > 0$ , there exists an equivalent  $\Delta \mathbf{q}' \geq 0$  with  $\Delta q'_\omega = 0$  that satisfies conditions 1 and 2 and have the same expected profit as  $\Delta \mathbf{q}$ .
4. If  $p_Y > \varphi_Y$  and  $p_N > \varphi_N$ ,  $\Delta \mathbf{q} = \mathbf{0}$ .

Theorem 4.2 means that, in a two-outcome DPM, when  $\sum_\omega \varphi_\omega^2 < 1$ , the optimal strategy for an agent is to buy shares of the contract whose current price is lower than its expected payoff and drive its price up to its expected payoff. When  $\sum_\omega \varphi_\omega^2 = 1$ , it’s possible to achieve the desired market prices by purchasing both contracts, but this is equivalent, in terms of expected profit, to the strategy that

only purchases the contract whose current price is lower than its expected payoff. Thus, the optimal strategy of an agent is to buy shares for the contract whose current price is too low. We now give the optimal shares that an agent would purchase and its optimal expected profit in the following theorem.

**Theorem 4.3.** *In a two-outcome market, when  $\frac{q_\omega}{\sqrt{q_\omega^2 + q_{\bar{\omega}}^2}} < \varphi_\omega \leq 1$ , a trader with expected payoff  $\varphi_\omega$  for contract  $\omega$  will purchase  $\Delta q_\omega^* = \frac{\varphi_\omega}{\sqrt{1 - \varphi_\omega^2}} q_{\bar{\omega}} - q_\omega$  to maximize his expected profit, where  $q_\omega$  is the current outstanding shares for outcome  $\omega$  in the market, and  $q_{\bar{\omega}}$  is the outstanding shares for the other outcome. His optimal expected profit is  $U(\Delta q_\omega^*) = \sqrt{q_\omega^2 + q_{\bar{\omega}}^2} - q_\omega \varphi_\omega - q_{\bar{\omega}} \sqrt{1 - \varphi_\omega^2}$ . When  $\frac{q_\omega}{\sqrt{q_\omega^2 + q_{\bar{\omega}}^2}} > \varphi_\omega$  and  $\frac{q_{\bar{\omega}}}{\sqrt{q_\omega^2 + q_{\bar{\omega}}^2}} > \varphi_{\bar{\omega}}$ , the trader does not purchase any contract.*

Because of the payoff uncertainty, an agent who only trades once in DPM will not change the market price to its posterior probability as in CDA or MSR, but will change the market price to  $(\varphi_Y, \sqrt{1 - \varphi_Y^2})$  if purchasing contract  $Y$  and to  $(\sqrt{1 - \varphi_N^2}, \varphi_N)$  if purchasing contract  $N$ . The corresponding market probabilities are  $(\varphi_Y^2, 1 - \varphi_Y^2)$  and  $(1 - \varphi_N^2, \varphi_N^2)$  respectively. It is possible that the agent’s optimal strategy is to not trade in the market. This happens when the current price  $\frac{q_\omega}{\sqrt{q_\omega^2 + q_{\bar{\omega}}^2}}$  is greater than  $\varphi_\omega$  for all  $\omega$ .

## 5 Two-Player Games

When analyzing a single agent’s optimal strategy in the previous section, we do not consider the possibility that the agent’s behavior may affect the set of information available to later traders. In DPM, as agents can infer information of other agents from their trading decisions, an agent who plays earlier in the market may mislead those who play later and affects the information sets of the later traders. In this section, we use two-player games to study this issue.

### 5.1 Two-Player Two-Stage Game

We first consider a two-player two-stage game, the Alice-Bob game, to examine whether an agent will try to affect the expected contract payoff by misreporting its own information. In the Alice-Bob game, Alice and Bob are the only players in the market. Each of them can trade only once. Alice plays first, followed by Bob. We are interested in whether there exists an equilibrium at which Alice fully reveals her information in the first stage if she trades and Bob infers Alice’s information and acts based on both pieces of information in the second stage. In particular, at the equilibrium, when having signal  $c_A \in \{y_A, n_A\}$ , Alice believes that  $\varphi_\omega(c_A) = \sum_{c_B} P(c_B|c_A) \sqrt{P(\omega|c_A, c_B)}$  and plays her optimal strategy according to Theorem 4.3. Bob with signal  $c_B$  believes that  $\varphi_\omega(c_A, c_B) = \sqrt{P(\omega|c_A, c_B)}$  if Alice trades in the first stage and  $\varphi_\omega(c_B) = \sqrt{P(\omega|c_B, \text{Alice doesn't trade})}$  if Alice doesn’t trade, and plays his optimal strategy according to Theorem 4.3. We call such an equilibrium a *truthful betting* equilibrium. Note that the truthful

betting equilibrium does not guarantee full information aggregation at the end of the game, because if Alice does not trade her information is not fully revealed. Since we are interested in the strategic behavior of agents, we also assume that  $\varphi_\omega(y_A) \neq \varphi_\omega(n_A)$  to rule out degenerated cases.


In the rest of this section, we show that truthful betting equilibrium exists when the initial market probability is uniform, but does not exist with some other initial market probabilities.

**Truthful Betting Equilibrium.** We assume that the market starts with uniform initial market probability, i.e.  $\pi_Y = \pi_N = \frac{1}{2}$ . This means that the initial market prices are  $p_Y = p_N = \frac{1}{\sqrt{2}}$ . As signals are symmetric, we have the following lemma.

**Lemma 5.1.** *Suppose all information is revealed after Bob’s play. Alice’s expected payoff for contract Y when she has signal  $y_A$  equals her expected payoff for contract N when she has signal  $n_A$ . That is,  $\varphi_Y(y_A) = \varphi_N(n_A)$ , where  $\varphi_Y(y_A) = \sum_{c_B} P(c_B|y_A)\sqrt{P(Y|c_B, y_A)}$  and  $\varphi_N(n_A) = \sum_{c_B} P(c_B|n_A)\sqrt{P(N|c_B, n_A)}$ .*

We use  $\varphi$  to denote both  $\varphi_Y(y_A)$  and  $\varphi_N(n_A)$ . Theorem 5.2 characterizes the truthful betting equilibrium for the game with uniform initial market probability.

**Theorem 5.2.** *In a two-outcome DPM with uniform initial market probability, truthful betting is a Bayesian Nash equilibrium for the Alice-Bob game. At the equilibrium, Alice does not trade if  $\varphi \leq \frac{1}{\sqrt{2}}$ . If  $\varphi > \frac{1}{\sqrt{2}}$ , Alice purchases contract Y and changes the price for Y to  $\varphi$  if she has  $y_A$ , and purchases contract N and changes the price for N to  $\varphi$  if she has  $n_A$ . If Alice trades, Bob infers her signal and changes the market probability to the posterior probability conditional on both signals. If Alice does not trade, Bob changes the market probability to the posterior probability conditional on his own signal.*



**Fig. 1.** Signal Qualities and Alice’s Expected Payoff

Figure 1 plots the iso-value lines of  $\varphi$  as a function of  $\theta_A$  and  $\theta_B$ . The leftmost curve is  $\varphi(\theta_A, \theta_B) = \frac{1}{\sqrt{2}}$ . The value of  $\varphi$  increases as the curve moves toward the right. As the initial market price is  $\frac{1}{\sqrt{2}}$  for both outcomes, the curve  $\varphi(\theta_A, \theta_B) = \frac{1}{\sqrt{2}}$  gives the boundary that at the equilibrium Alice trades in the first stage. The shaded area gives the range of signal qualities that Alice is better off not trading at the equilibrium. When  $\theta_B = 1$ , that is Bob is perfectly informed, Alice won’t trade if her signal quality  $\theta_A$  is less than  $\frac{1}{\sqrt{2}}$ . This is consistent with the example given by Chen et al. [13].

**Non-existence of the Truthful Betting Equilibrium.** In Alice-Bob game, truthful betting is not a Nash equilibrium for arbitrary initial market probability.

**Theorem 5.3.** *In a two-outcome DPM, there exists some initial market probabilities where truthful betting is not a Nash equilibrium for the Alice-Bob game.*

The intuition is that if the initial market price for one contract is very low that Alice will purchase the the contract no matter which signal she gets, Alice may pretend to have a different signal by purchasing less when she should buy more if being truthful. If Bob is misled, this increases the expected payoff per share of the contract and hence can increase Alice's expected total profit even if she purchases less. This is very different from other market mechanisms. In both CDA and MSR, if a player only plays once in the market, disregard of whether there are other players behind it, the player will always play truthfully.

## 5.2 Two-Player Three-Stage Game

In the Alice-Bob game, Alice may not play truthfully in order to mislead Bob and achieve a higher expected payoff per share, but she does not directly make profits from Bob's uninformed trades. Now we consider a three-stage game, the Alice-Bob-Alice game, where Alice can play a second time after Bob's play. Truthful betting equilibrium in this game means that both players play their truthful betting equilibrium strategies of the Alice-Bob game in the first two stages and Alice does nothing in the third stage. Clearly, if Alice has incentives to deviate from truthful betting in the Alice-Bob game, she will also deviate in the Alice-Bob-Alice game, because playing a second time allows Alice to gain more profit by capitalizing on Bob's uninformed trades. Even for settings where truthful betting is a Bayesian Nash equilibrium for the Alice-Bob game, a truthful betting equilibrium may not exist for the Alice-Bob-Alice game. For example, with uniform initial market probability, if  $\theta_A = 0.6$ ,  $\theta_B = 0.8$  and Alice has  $y_A$ , Alice is better off pretending to have  $n_A$  given Bob believes that she plays truthfully. In contrast, in LMSR when agents have conditionally independent signals, truthful betting is the unique perfect Bayesian equilibrium [13, 17].

## 6 Conclusion

Using a simple setting of incomplete information, we show that DPM admits more gaming than several other prediction market mechanisms due to its payoff uncertainty. We show that even when a player only participates once in the market, e.g. in an Alice-Bob game, it still has incentives to bluff and pretend to have a different signal. The bluffing behavior exists more generally when traders participate the market more than once.

## Acknowledgement

We would like to express our sincere thanks to Professor Xiaotie Deng with his kind support for Qianya Lin in visiting Harvard University.

## References

1. Forsythe, R., Nelson, F., Neumann, G.R., Wright, J.: Forecasting elections: A market alternative to polls. In: Palfrey, T.R. (ed.) *Contemporary Laboratory Experiments in Political Economy*, pp. 69–111. University of Michigan Press, Ann Arbor (1991)
2. Berg, J.E., Forsythe, R., Nelson, F.D., Rietz, T.A.: Results from a dozen years of election futures markets research. In: Plott, C.A., Smith, V. (eds.) *Handbook of Experimental Economic Results* (2001)
3. Chen, K.Y., Plott, C.R.: Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working paper (2002)
4. Cowgill, B., Wolfers, J.: EricZitzewitz: Using prediction markets to track information flows: Evidence from Google. Working paper (2008)
5. Thaler, R.H., Ziemba, W.T.: Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives* 2(2), 161–174 (1988)
6. Pennock, D.M., Lawrence, S., Giles, C.L., Nielsen, F.A.: The real power of artificial markets. *Science* 291, 987–988 (2002)
7. Debnath, S., Pennock, D.M., Giles, C.L., Lawrence, S.: Information incorporation in online in-game sports betting markets. In: *The 4th ACM Conference on Electronic Commerce, EC* (2003)
8. Chen, Y., Chu, C.H., Mullen, T., Pennock, D.M.: Information markets vs. opinion pools: An empirical comparison. In: *The 6th ACM Conference on Electronic Commerce, EC* (2005)
9. Hanson, R.D.: Combinatorial information market design. *Information Systems Frontiers* 5(1), 107–119 (2003)
10. Hanson, R.D.: Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* 1(1), 1–15 (2007)
11. Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: *The 5th ACM Conference on Electronic Commerce, EC* (2004)
12. Mangold, B., Dooley, M., Dornfest, R., Flake, G.W., Hoffman, H., Kasturi, T., Pennock, D.M.: The tech buzz game. *IEEE Computer* 38(7), 94–97 (2005)
13. 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)
14. Nikolova, E., Sami, R.: A strategic model for information markets. In: *The Seventh Annual ACM Conference on Electronic Commerce, EC* (2007)
15. Bu, T.M., Deng, X., Lin, Q., Qi, Q.: Strategies in dynamic pari-mutuel markets. In: Papadimitriou, C., Zhang, S. (eds.) *WINE 2008*. LNCS, vol. 5385, pp. 138–153. Springer, Heidelberg (2008)
16. Dimitrov, S., Sami, R.: Non-myopic strategies in prediction markets. In: *The Nineth ACM Conference on Electronic Commerce, EC* (2008)
17. Chen, Y., Dimitrov, S., Sami, R., Reeves, D.M., Pennock, D.M., Hanson, R.D., Fortnow, L., Gonen, R.: Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica* (2009)
18. Allen, F., Gale, D.: Stock-price manipulation. *The Review of Financial Studies* (5), 503–529 (1992)
19. Chakraborty, A., Yilmaz, B.: Manipulation in market order models. *Journal of Financial Markets* 7(2), 187–206 (2004)

20. Kumar, P., Seppi, D.J.: Futures manipulation with ‘cash settlement’. *Journal of Finance* (47), 1485–1502 (1992)
21. Koessler, F., Noussair, C., Ziegelmeyer, A.: Parimutuel betting under asymmetric information. *Journal of Mathematical Economics*, 733–744 (2008)
22. Ostrovsky, M.: Information aggregation in dynamic markets with strategic traders. Working paper (2009)
23. Milgrom, P., Stokey, N.L.: Information, trade and common knowledge. *Journal of Economic Theory* 26(1), 17–27 (1982)
24. Chen, Y., Pennock, D.M., Kasturi, T.: An empirical study of dynamic pari-mutuel markets: Evidence from the Tech Buzz Game. In: *Workshop on Web Mining and Web Usage Analysis (WebKDD)*. LNCS. Springer, Heidelberg (2008)

# News Posting by Strategic Users in a Social Network

Mangesh Gupte<sup>1</sup>, MohammadTaghi Hajiaghayi<sup>2</sup>, Lu Han<sup>1</sup>, Liviu Iftode<sup>1</sup>,  
Pravin Shankar<sup>1</sup>, and Raluca M. Ursu<sup>3</sup>

<sup>1</sup> Department of Computer Science, Rutgers University, Piscataway, NJ 08854  
{mangesh,luhan,iftode,spravin}@cs.rutgers.edu

<sup>2</sup> AT&T Labs – Research, Florham Park, NJ 07932  
hajiagha@research.att.com

<sup>3</sup> Mount Holyoke College, South Hadley, MA 01075  
ursu22r@mtholyoke.edu

**Abstract.** We argue that users in social networks are *strategic* in how they post and propagate information. We propose two models — greedy and courteous — and study information propagation both analytically and through simulations. For a suitable random graph model of a social network, we prove that news propagation follows a threshold phenomenon, hence, “high-quality” information provably spreads throughout the network assuming users are “greedy”. Starting from a sample of the Twitter graph, we show through simulations that the threshold phenomenon is exhibited by both the greedy and courteous user models.

## 1 Introduction

Online social networks have become an increasingly popular medium for sharing information such as links, news and multimedia among users. The average Facebook user has 120 friends, and more than 30 million users update their status at least once each day [2]. More than 5 billion minutes are spent on Facebook each day (worldwide). As a direct consequence of these trends, social networks are fast overtaking traditional web as the preferred source of information [1].

In the early days of social networks, users tended to post predominantly personal information. Such information typically did not spread more than one hop, since only immediate friends were interested in it. Over time, online social networks have metamorphosed into a forum where people post information such as news that they deem to be of common interest. For example, during the recent Iran elections, traditional news media acknowledged the power and influence of social networks such as Twitter [3, 4].

Prior work has studied various aspects of information sharing on social networks. Domingos and Richardson studied the question of how to determine the set of nodes in a network that will most efficiently spread a piece of information for marketing purposes [5, 6]. Kempe, Kleinberg and Tardos proposed a discrete optimization formulation for this. Several recent studies focused on gathering intuition about influence spread from real-world data [9]. In [13], Leskovec, Singh and Kleinberg studied the patterns of cascading recommendations in social



networks by looking at how individuals recommend the products they buy in an online-retailer recommendation system. In [10], Leskovec, Backstrom and Kleinberg developed a framework for the dynamics of news propagation across the web. Morris studied games where each player interacts with a small set of neighbors [16]. He proved conditions under which the behavior adopted by a small set of users will spread to a large fraction of the network.

An aspect that has been overlooked so far is to understand *why* users post information such as news or links on social networks. In this paper, we posit that users in a social network have transitioned from being passive entities to *strategic users* who weigh in various factors (such as how interested their friends will be in the news) to decide whether to post. This trend leads to several interesting questions, such as: What factors do users consider when deciding whether to post an item? How does information diffuse over the social network based on user strategies.

Our main result states that, assuming strategic users, the spread of news over an online social network exhibits a *threshold behavior*: The news spreads to a significant fraction of the network if its “quality” is larger than a certain threshold that depends on how aggressive users are about posting news. If the quality is smaller than this threshold, only a sub-linear number of nodes in the network post the news.

The key contributions of this paper are:

1. We initiate the study of information propagation in social networks assuming strategic users.
2. We propose two models for strategic user behavior, greedy and courteous.
3. Assuming social networks can be modeled as certain random graphs, we prove that there is a threshold behavior when greedy users fully disseminate information.
4. We present a simulation study based on a real graph crawled from the Twitter social network, and show that the threshold phenomenon holds in both strategic models of user behavior.

In what follows, we provide a detailed description of our results. We start by defining the user model.

## 2 Strategic User Model

We propose a simple game to model the behavior of users posting news on social networks like Twitter and Facebook. For a particular user  $u$  in the network, whenever  $u$  first sees a previously unseen news item, she has the option of either posting it or not posting it. Her utility is 0 if she does not post it. If she does, then her utility depends on (i) The set  $I_u = \{\text{Neighbors who are interested in the news}\}$  and (ii) The set  $S_u = \{\text{Neighbors who, } u \text{ knows, have already seen the news before}\}$ .<sup>1</sup> Let  $N_u$  denote the set of  $u$ 's neighbors. We propose two particular forms for her utility:

<sup>1</sup>  $u$  might not know the true set of neighbors who have seen the news. She knows that a friend has seen the news only if a mutual friend posted it. This also means that we assume that every user knows which of her friends are themselves friends.

**Greedy Strategy:** The utility is additive and for every neighbor who likes the news (irrespective of whether the neighbor has seen it before or not), she gets utility  $+a$  and for every neighbor who does not, she gets utility  $-b$ . In this case, her decision to post only depends on  $a, b$  and  $f_u = \frac{|I_u|}{|N_u|}$ . User  $u$  posts only if her utility is positive, that is, the fraction  $f$  of users who like the news satisfies  $a \cdot f_u - b(1 - f_u) > 0 \iff f_u > \frac{b}{a+b}$ . Let us define  $t = \frac{b}{a+b}$ . In Section 4, we analyze this behavior and show that it depends critically on  $t$ .

**Courteous Strategy:** The main difference from the greedy strategy is that the user does not want to spam her friends. We model this by saying that if more than a  $c$  fraction of her friends have already seen the news before, she gets a large negative utility, when she posts the item. If a user does not post an item, then her utility is 0. In case the fraction  $\frac{|S_u|}{|N_u|} \leq c$ , then her utility is the same as in the greedy case. In particular, she gets utility  $+a$  for every neighbor who likes the news and has not seen it before (the set  $I_u \setminus S_u$ ), and she gets utility  $-b$  for every neighbor who does not like it (the set  $I_u^c \setminus S_u$ ). Hence, her strategy in this case is to post if the fraction of neighbors who have seen the news  $\frac{|S_u|}{|N_u|} \leq c$  and if the fraction  $\left(f_u = \frac{|I_u \setminus S_u|}{|N_u \setminus S_u|}\right)$  of neighbors in  $S_u^c$  who are interested in the news is  $\geq t$ . Note that, in this utility function, if a larger number of the user's neighbors have posted the news, she is *less* likely to post it. In section 5, we show simulation results for this behavior on a small sample of the Twitter Graph.

### 3 Notation and Preliminaries

Given a real symmetric matrix  $P$  with  $0 \leq p_{i,j} \leq 1$ , denote by  $G(n, P)$  a random graph where edge  $(i, j)$  exists with probability  $p_{i,j}$ . For notational convenience we shall denote by  $W_P = (V, E_P)$  the deterministic weighted graph with  $V$  as the vertex set and  $P$  as its adjacency matrix. Note that  $p_{i,j}$  gives the weight of edge  $(i, j)$  in  $W_P$ .

Leskovec et al show that social graphs, such as Autonomous Systems on the Internet, the citation graph for high energy physics from arXiv, and U.S. Patent citation database, can be modeled well using Stochastic Kronecker Graphs [11].  $G(n, P)$  is a generalization of Stochastic Kronecker Graphs [11, 14] as well as the Erdős and Rényi model of Random Graphs  $G(n, p)$  [7], and the model of random graphs with a given degree sequence [15]. In the following subsection, we prove properties of the graph  $G(n, P)$  that will be used in the next section to analyze when a certain news spreads across the network.

#### 3.1 Properties of $G(n, P)$

**Definition 1 (Density of a cut).** Given a graph  $G = (V, E)$ , define the density of the cut  $(S, V - S)$  as  $\frac{|E(S, V - S)|}{|S||V - S|}$  where  $E(S, V - S)$  denotes the set of edges between  $S$  and  $V - S$ . The partition  $(S, V - S)$  that minimizes this density is called the Sparsest Cut in the graph.

**Definition 2 ( $\alpha$ -balanced cut).** Given a graph  $G = (V, E)$  and a cut  $(S, V - S)$ , the cut is  $\alpha$ -balanced if and only if  $\min\{|S|, |V - S|\} \geq \alpha|V|$ .

Hence, by a sparsest  $\alpha$ -balanced cut we mean the cut with the minimum density over all cuts that are  $\alpha$ -balanced. We start with a lemma from Mahdian and Xu [14] that proves when  $G$  is connected.

**Lemma 3.** If the size of the min-cut in  $W_P$  is  $> c \log n$ , then with high probability, the sampled graph  $G \sim G(n, P)$  is connected [2].

Given  $U \subseteq V$ , we shall denote by  $G[U]$  the subgraph induced by  $U$ . The induced subgraph of  $W_P$  is denoted by  $W_P[U]$ . The following lemma gives a sufficient condition on the existence of a giant component in the graph  $G(n, P)$ .

**Lemma 4.** If there exists  $U \subseteq V$ , of size  $\Theta(n)$ , such that the sparsest  $\alpha$ -balanced cut of  $W_P[U]$  has density  $> \frac{c}{|U|}$ , then there is a giant connected component in  $G[U]$  of size  $\Theta(n)$  with high probability.

## 4 Analysis of a Model for Strategic User Behavior

We analyze the *greedy* strategy defined in Section 2 when it is played over a random graph  $G(n, P) = (V, E)$ . Our results also apply to Stochastic Kronecker Graphs [3]. According to this model, a user posts only if the fraction of interested neighbors is  $> t$ . We assume that for a given news, each user in the network likes it with probability  $q$ , which is independent of everything else. Probability  $q$  could model the quality of the news item or the inherent interest the subject generates [4]. Throughout, we assume that  $q$  and  $t$  are constants that do not depend on the number of nodes  $n$ .

We color nodes in  $G$  yellow if they are interested in the news and blue if they are not. A yellow node is *responsive* if more than a  $t$  fraction of its neighbors are interested in the news. Color responsive nodes red. We denote these sets by  $Y, B$  and  $R$ , respectively. Note that  $R \subseteq Y$ .  $G[R]$  is the graph induced by the red vertices. We are interested in the structure of the graph  $G[R]$ . We prove the following results in the next two subsections :

**Proposition 5.** Suppose  $G(n, P)$  is such that the min-cut in  $W_P$  is of weight  $\geq c \log n$  and that  $\log n$  random nodes in the network initially see the news. If  $q > t$ , with high probability, almost all nodes interested in the news will post it. On the other hand, if  $q < t$  then only a sub-linear number of the vertices will post the news.

Next, we give a condition on the sparsity of  $G$ , which gives a weaker result on the spread of the news.

<sup>2</sup> Throughout this paper, with high probability means with probability  $1 - o(1)$ .

<sup>3</sup> Due to space constraints, we leave the formal statement of the result and the proof of this claim to the Technical Report [8].

<sup>4</sup> It is an interesting question to relax this assumption, since in a typical social network we might expect nodes with similar interests to be clustered together.

**Proposition 6.** *Suppose  $G(n, P)$  is such that the sparsest  $\alpha$ -balanced cut in  $W_P$  has density  $\geq c/n$  and that  $\log n$  random nodes in the network initially see the news. If  $q > t$  then, with high probability, a constant fraction of the nodes interested in it will post it.*

### 4.1 The Connectedness of $G[R]$

We prove Proposition 5 in this subsection. Throughout this section, we shall assume that  $G \sim G(n, P)$  where  $P$  is such that the min-cut of  $W_P$  has weight  $\geq c \log n$ , for a large enough constant  $c$ . We start by looking at what happens to the min-cut in the sampled graph.

**Lemma 7.** *With high probability, the min-cut of the subgraph  $G[Y]$  of  $G$  induced by the yellow vertices has size  $> c' \log n$  for some constant  $c'$ .*

Now, we shall prove the main theorem of this section. We prove that  $G[R]$  is connected by using the fact that its min-cut is large.

**Theorem 8.** *If  $q > t$  and  $G \sim G(n, P)$  where  $P$  is such that the min-cut of  $W_P$  has weight  $\geq c \log n$ , then, with high probability, every vertex in  $Y$  also belongs to  $R$  and so  $G[R]$  is connected. When  $q < t$ , then, with high probability  $G[R]$  only contains  $o(n)$  vertices.*

*Proof.* Let  $Y_q$  be a random variable that takes value 1 with probability  $q$  and 0 otherwise. For a node  $v$ , let  $d(v)$  denote its degree in  $G$ .

Case 1:  $q > t$ :  $Pr[v \notin R] = Pr\left[\sum_{i=1}^{d(v)} Y_q < qd(v) - (q - t)d(v)\right]$   
 $\leq \exp\left(-\frac{((q-t)d(v))^2}{2 \cdot q \cdot d(v)}\right) \leq n^{-\frac{(q-t)^2 c}{2 \cdot q}}$ . This follows from Chernoff Bounds and the facts that  $Y_q$  are independent,  $\mathbb{E}[Y_q] = q$  and  $d(v) \geq c \log n$ . We apply the union bound to get that with probability  $\geq 1 - n^{1-\frac{(q-t)^2 c}{2 \cdot q}} \geq 1 - 1/n$  (when  $c > \frac{2q}{(q-t)^2}$ ) all nodes in  $Y$  also belong to  $R$  and, from Lemmas 3 and 7, we get that  $G[R]$  is connected.

Case 2:  $q < t$ :  $Pr[v \in R] = Pr\left[\sum_{i=1}^{d(v)} Y_q \geq qd(v) + (t - q)d(v)\right]$   
 $\leq \exp\left(-\frac{((t-q)d(v))^2}{2(qd(v) + (t-q)/3)}\right) \leq n^{-\frac{(t-q)^2 c}{t}}$ . Hence,  $\mathbb{E}[|R|] \leq |Y| \cdot n^{-\frac{(t-q)^2 c}{t}} \leq o(n)$ . So,  $R$  only contains a sub-linear number of nodes of  $G$ .

*Proof (of Proposition 5).* If  $q < t$ , then the proposition follows directly from Theorem 8. When  $q > t$ , Theorem 8 tells us that  $G[R]$  is connected. If any node in  $R$  receives the news, it will be propagated to all the nodes. However, the probability that none of the nodes in  $R$  get the news is  $\leq (1 - \frac{|R|}{n})^{\log n} = O(n^{-q}) = o(1)$ . Hence, with high probability, almost all the nodes interested in the news post it.

In the next subsection, we shall identify conditions on the distribution from which  $G$  is sampled, which are enough to show when a constant fraction of the nodes who are interested in the news actually receive it.

### 4.2 Existence of a Giant Component in $G[R]$

We now prove Proposition 6. In this section, we shall assume that  $G \sim G(n, P)$ , where  $P$  is such that the sparsest  $\alpha$ -balanced cut of  $W_P$  has density  $\geq c/n$ . We prove that the size of the sparsest-cut in  $G[Y]$  is not small.

**Lemma 9.** Consider the subgraph  $G[Y]$  induced by the yellow vertices. With high probability,  $G[Y]$  has a subgraph of size  $\Theta(n)$  whose sparsest  $\alpha$ -balanced cut has density  $\geq \frac{c'}{|Y|} \geq \frac{c'}{n}$ , for some constant  $c'$ .

**Theorem 10.** Let  $G$  be a random graph sampled from the distribution  $G(n, P)$  where the density of the sparsest  $\alpha$ -balanced cut in the graph  $W_P$  is greater than  $c/n$ . If  $q > t$ , then every yellow node is red with probability  $> 1 - \epsilon_c$ . Further, the induced graph  $G[R]$  has a giant connected component of size  $\Theta(n)$  with high probability.

*Proof.* Let  $Y_q$  be a random variable that takes value 1 with probability  $q$  and 0 otherwise. Let us denote the degree of node  $v$  by  $d(v)$ . Since  $(v, V \setminus v)$  is a cut,  $\frac{d(v)}{1 \cdot (n-1)} \geq \frac{c}{n} \implies d(v) \geq c$ .

For  $v \in Y$ ,  $Pr[v \notin R] = Pr\left[\sum_{i=1}^{d(v)} Y_q < qd(v) - (q-t)d(v)\right] \leq \exp\left(-\frac{((q-t)d(v))^2}{2 \cdot q \cdot d(v)}\right) \leq e^{-\frac{(q-t)^2 c}{2 \cdot q}} = \epsilon_c$ . Hence, a constant fraction  $f \geq 1 - \epsilon_c$  of the vertices in  $Y$  belong to  $R$ . From this and from Lemmas 4 and 9, we can prove that  $G[R]$  also has a giant connected component with high probability.

*Proof (of Proposition 6).* By Theorem 10,  $G[R]$  contains a giant component  $C$  of size  $(1 - \epsilon_c)n$ . If any node in  $C$  receives the news, it will be propagated to all the nodes in  $C$ . However, the probability that none of the nodes in  $C$  get the news is  $\leq (1 - \epsilon_c)^{\log n} = o(1)$ . Hence, a constant fraction of the nodes interested in the news actually receive it with high probability.

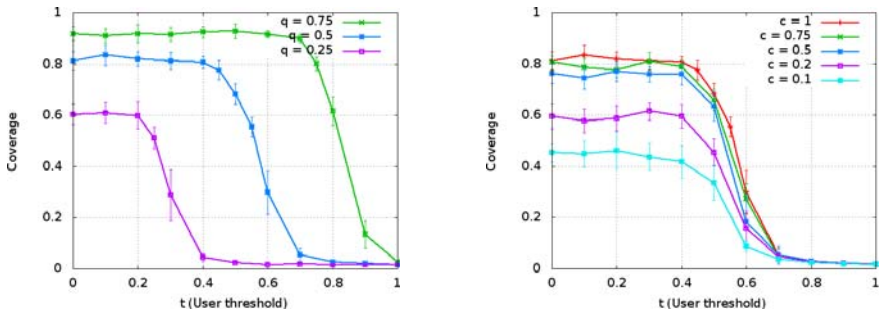
## 5 Simulation Results

In this section, we present results from the simulation of our two strategic user models over a partial crawl of the Twitter social network 5.

The dataset is obtained by means of a partial crawl of the Twitter social network. Since we are interested in link postings by typical users, we remove hubs (users with more than 500 friends) from our graph. The resulting graph we use in our experiments consists of 5978 users with 110152 directed friendship edges. Each simulation starts with a set of seed nodes (of size  $\log n$ ), and in each round, nodes decide whether to post the link using one of the two models described earlier in the paper.

We define coverage as the fraction of interested users who get to see the news item. Figure 1 plots how the coverage varies for the *greedy* and *courteous*

<sup>5</sup> All the data was obtained in accordance with Twitter’s terms of use.



(a) *Greedy* strategy: The coverage exhibits a step behavior when  $q = t$ . (b) *Courteous* strategy. ( $q = 0.5$ ). The coverage decreases logarithmically as the user is more courteous.

**Fig. 1.** Coverage of *greedy* and *courteous* strategies.  $q$  is the probability with which each user likes the link.  $t$  and  $c$  are the thresholds for the greedy and the courteous strategies, respectively.

strategies. For the courteous strategy, we fixed  $q = 0.5$ . Each simulation was repeated 10 times, and the standard deviations are indicated in the form of error bars.

**Greedy Strategy:** Figure 1(a) shows that, for all values of  $q$ , the coverage exhibits a step behavior and the step happens around  $q = t$ , which is consistent with Theorem 10. For different values of  $q$ , the percentage of coverage decreases with  $q$ . This is true even when  $t = 0$ , which means that the size of the connected components drops when we sample the graph. The Densification Power Law ( $E \propto N^k$ ) [12] that has been observed in other social networks would predict this behavior.

**Courteous Strategy:** Figure 1(b) shows the effect of the courteous strategy on the coverage. The parameter  $c$  indicates the threshold of neighbors that have already seen the link. A courteous user posts the link only if less than  $c$  fraction of neighbors have already seen the link. The figure shows that the coverage decreases logarithmically as the user is more courteous. This means that even when the users are courteous, if  $q > t$  then, news can still reach a reasonable fraction of the graph.

## 6 Conclusion and Future Work

We proposed the model of strategic user behavior in online social networks for news posting, defined two user models (greedy and courteous), presented formal analysis of the greedy model and simulated both models on the data set we collected from Twitter. We propose the following directions:

*Mine the Twitter data set:* Search for patterns in the way users post news links in order to validate the model and provide further insights about user strategies over social networks.

*Analyse the Courteous Strategy and Multiple Strategies:* We leave open a formal proof that the courteous strategy also exhibits threshold behavior. Further, we want to test whether similar results hold in a social network in the presence of multiple user strategies.

*Design a framework for advertisement on online social networks of strategic users:* We believe that the strategic user model has applications in advertising and marketing. We plan to investigate incentive schemes for marketers to encourage strategic users to advertise products over social networks.

**Acknowledgments.** We would like to thank S. Muthu Muthukrishnan for the inspiring discussions. We would also like to thank the anonymous reviewers. This work has been supported in part by the NSF grant CNS-0831268 and by the Summer 2009 DIMACS/DyDAn REU Program.

## References

- [1] Compete – Site Comparison, <http://siteanalytics.compete.com/facebook.com+google.com/?metric=uv>
- [2] Facebook Statistics, <http://www.facebook.com/press/info.php?statistics>
- [3] Journalism Rules Are Bent in News Coverage From Iran, <http://www.nytimes.com/2009/06/29/business/media/29coverage.html>
- [4] Twitter Revolution in Iran, [http://money.cnn.com/video/news/2009/06/22/\\_twitter\\_revolution\\_iran\\_pentagon.cnnmoney](http://money.cnn.com/video/news/2009/06/22/_twitter_revolution_iran_pentagon.cnnmoney)
- [5] Domingos, P., Richardson, M.: Mining the network value of customers. In: Proc. 7th ACM SIGKDD (2001)
- [6] Domingos, P., Richardson, M.: Mining knowledge-sharing sites for viral marketing. In: Proc. 8th ACM SIGKDD (2002)
- [7] Erdos, P., Renyi, A.: On the evolution of random graphs. Publ. Math. Inst. Hung. Acad. Sci. 5 (1960)
- [8] Gupte, M., Hajiaghayi, M., Han, L., Iftode, L., Shankar, P., Ursu, R.M.: News posting by strategic users in a social network. Technical Report DCS-tr-657, Rutgers University, Department of Computer Science (October 2009)
- [9] Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: Proc. 9th ACM SIGKDD (2003)
- [10] Leskovec, J., Backstrom, L., Kleinberg, J.: Meme-tracking and the dynamics of the news cycle. In: Proc. 15th ACM SIGKDD (2009)
- [11] Leskovec, J., Faloutsos, C.: Scalable modeling of real graphs using kronecker multiplication. In: International Conference on Machine Learning, ICML (2007)
- [12] Leskovec, J., Kleinberg, J., Faloutsos, C.: Graphs over time: Densification laws, shrinking diameters and possible explanations. In: Proc. 11th ACM SIGKDD (2005)
- [13] Leskovec, J., Singh, A., Kleinberg, J.: Patterns of influence in a recommendation network. In: Ng, W.-K., Kitsuregawa, M., Li, J., Chang, K. (eds.) PAKDD 2006. LNCS (LNAI), vol. 3918, pp. 380–389. Springer, Heidelberg (2006)
- [14] Mahdian, M., Xu, Y.: Stochastic kronecker graphs. In: Bonato, A., Chung, F.R.K. (eds.) WAW 2007. LNCS, vol. 4863, pp. 179–186. Springer, Heidelberg (2007)
- [15] Molloy, M., Reed, B.: A critical point for random graphs with a given degree sequence. Random Structures and Algorithms 6 (1995)
- [16] Morris, S.: Contagion. Review of Economic Studies 67 (2000)

# Erratum to: Prediction Mechanisms That Do Not Incentivize Undesirable Actions

Peng Shi, Vincent Conitzer, and Mingyu Guo

Department of Computer Science, Duke University,  
Durham, NC 27708, USA  
ps46@duke.edu, {conitzer,mingyu}@cs.duke.edu

S. Leonardi (Ed.): WINE 2009, LNCS 5929, pp. 89–100, 2009.  
© Springer-Verlag Berlin Heidelberg 2009

---

**DOI 10.1007/978-3-642-10841-9\_66**

Definition 10 of the paper entitled “Prediction Mechanisms That Do Not Incentivize Undesirable Actions” contains an error. Here is the correction:

**Definition 10.** *We say a proper scoring rule provides incentive  $c$  if the agent cannot guarantee within  $c$  of the optimal expected payment by giving some constant dummy report  $\mathbf{r}$ . (More precisely,  $\forall$  dummy report  $\mathbf{r}$ ,  $\exists$  distribution  $\mathbf{p}$  under which reporting truthfully instead of reporting the dummy value  $\mathbf{r}$  pays off by at least  $c$ :  $\tilde{S}(\mathbf{p}, \mathbf{p}) - \tilde{S}(\mathbf{r}, \mathbf{p}) \geq c$ .) A one-round prediction mechanism guarantees incentive  $c$  if for each agent  $j$  and each combination of others’ reports  $\mathbf{r}_{-j}$ , the corresponding proper scoring rule provides incentive  $c$ .*

---

The original online version for this chapter can be found at  
[http://dx.doi.org/10.1007/978-3-642-10841-9\\_10](http://dx.doi.org/10.1007/978-3-642-10841-9_10)

---



# Author Index

- Abdallah, Chaouki T. 430  
Adler, Ilan 471  
Alaei, Saeed 398  
Apt, Krzysztof R. 483  
Arcaute, Esteban 220  
Arthur, David 101  
Asadpour, Arash 545  
Ashwinkumar, B.V. 529  
Aziz, Haris 438
- Bhalgat, Anand 362  
Briest, Patrick 42  
Byde, Andrew 600
- Caragiannis, Ioannis 475  
Cardinal, Jean 125  
Chakraborty, Sourav 79  
Chakraborty, Tanmoy 362  
Chen, Ning 447  
Chen, Yiling 553, 623  
Conitzer, Vincent 89, 244, E1  
Corbo, Jacomo 505  
Crichigno, Jorge 430
- Daskalakis, Constantinos 471  
Davis, Joshua R. 591  
Demaine, Erik D. 125  
Deng, Xiaotie 410, 521, 583  
Devanur, Nikhil 79  
Díaz, Josep 455  
Dikshit, Devansh 386  
Dimitrov, Stanko 491  
Dütting, Paul 575
- Elkind, Edith 447  
Emek, Yuval 268  
Epelman, Marina 491  
Epstein, Leah 67, 232
- Fanelli, Angelo 55  
Feldman, Jon 374  
Feldman, Michal 268, 351  
Feldmann, Rainer 280  
Fiat, Amos 498  
Fiorini, Samuel 125  
Fotakis, Dimitris 316
- Gairing, Martin 184  
Gao, Xi 553  
Ghosh, Arpita 208  
Gkatzelis, Vasilis 316  
Goel, Sharad 196  
Gomes, Renato 172  
Gourvès, Laurent 292, 608  
Gravin, Nick 447  
Gualà, Luciano 42  
Guo, Mingyu 89, 244, E1  
Gupte, Mangesh 632
- Hajiaghayi, MohammadTaghi 632  
Han, Lu 632  
Harks, Tobias 463  
Henzinger, Monika 575  
Hoefer, Martin 18, 42, 616
- Iftode, Liviu 632  
Immorlica, Nicole 172  
Iwasaki, Atsushi 561
- Jennings, Nicholas R. 600  
Joret, Gwenaël 125
- Kaklamanis, Christos 475  
Kanellopoulos, Panagiotis 475  
Kaporis, Alexis C. 316  
Kempe, David 537  
Khanna, Sanjeev 362  
Khoury, Joud 430  
Kleiman, Elena 67, 232  
Kleinberg, Robert 529  
Klimm, Max 463  
Korula, Nitish 374  
Krause, Kate 430  
Kumar, Ravi 513  
Kyropoulou, Maria 475
- Lachish, Oded 438  
Lahaie, Sébastien 196  
Liben-Nowell, David 591  
Lin, Qianya 623  
Lindner, Claudia 149  
Lu, Pinyan 30, 137

- Mahdian, Mohammad 513  
 Malekian, Azarakhsh 398  
 Markakis, Evangelos 172, 483  
 Mavronicolas, Marios 280  
 McAfee, Preston 208  
 Mestre, Julián 67  
 Mirrokni, Vahab S. 18, 374  
 Mitsche, Dieter 455  
 Möhring, Rolf H. 463  
 Monien, Burkhard 280  
 Monnot, Jérôme 292, 608  
 Moscardelli, Luca 55  
 Motwani, Rajeev 101  
 Mu'alem, Ahuva 537  
 Muthukrishnan, S. 1, 374  
  
 Newman, Ilan 125  
 Nguyen, Thành 160  
  
 Pál, Martin 374  
 Pálvölgyi, Dömötör 569  
 Papadimitriou, Christos H. 471  
 Papineni, Kishore 208  
 Paterson, Mike 438  
 Penna, Paolo 256  
 Pennock, David M. 553  
 Polukarov, Maria 600  
  
 Qi, Qi 410, 521  
  
 Reeves, Daniel M. 244  
 Röglin, Heiko 18  
 Rothe, Jörg 149  
 Rozenfeld, Ola 113  
 Rustagi, Navin 455  
  
 Saberi, Amin 545  
 Saia, Jared 455  
 Salek, Mahyar 537  
 Sami, Rahul 491  
 Savani, Rahul 438  
 Sayedi, Amin 513  
 Schoppmann, Florian 256  
 Shankar, Pravin 632  
 Sharma, Aneesh 101  
  
 Sharp, Alexa 591  
 Shi, Peng 89, E1  
 Silvestri, Riccardo 256  
 Skopalik, Alexander 616  
 Spirakis, Paul G. 316  
  
 Tardos, Éva 17, 160  
 Telelis, Orestis A. 292  
 Teng, Shang-Hua 18  
 Tennenholtz, Moshe 113  
 Tessler, Ran 351  
 Thain, Nithum 304  
 Todo, Taiki 561  
  
 Ursu, Raluca M. 632  
  
 van Stee, Rob 232  
 Vassilvitskii, Sergei 196, 208, 220  
 Vazirani, Vijay V. 422  
 Ventre, Carmine 42  
 Vetta, Adrian 304  
 Vidali, Angelina 340  
 Voice, Thomas 600  
 Vorobeychik, Yevgeniy 505  
  
 Wang, Lei 422  
 Wang, Yajun 137  
 Weber, Ingmar 575  
 Weimann, Oren 125  
 Wexler, Tom 591  
 Widmayer, Peter 256  
 Wilf, Yoav 351  
 Wilkens, Christopher A. 328  
 Wingarten, Amiram 498  
  
 Xu, Ying 101  
  
 Yadati, Narahari 386  
 Yokoo, Makoto 561  
 Young, H. Peyton 13  
 Yu, Jiajin 583  
  
 Zhang, Jie 410  
 Zhou, Yuan 137