

Spyros Kontogiannis
Elias Koutsoupias
Paul G. Spirakis (Eds.)

LNCS 6386

Algorithmic Game Theory

Third International Symposium, SAGT 2010
Athens, Greece, October 2010
Proceedings

 Springer

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Germany

Madhu Sudan

Microsoft Research, Cambridge, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbruecken, Germany

Spyros Kontogiannis Elias Koutsoupias
Paul G. Spirakis (Eds.)

Algorithmic Game Theory

Third International Symposium, SAGT 2010
Athens, Greece, October 18-20, 2010
Proceedings

Volume Editors

Spyros Kontogiannis
Lecturer
University of Ioannina
Department of Computer Science
45110 Ioannina, Greece
E-mail: kontog@cs.uoi.gr

Elias Koutsoupas
Professor
National and Kapodistrian University of Athens
Department of Informatics and Telecommunications
Panepistimiopolis, Ilissia, 15784 Athens, Greece
E-mail: elias@di.uoa.gr

Paul G. Spirakis
Professor
Patras University
Research Academic Computer Technology Institute
26504 Rion-Patra, Greece
E-mail: spirakis@cti.gr

Library of Congress Control Number: 2010935286

CR Subject Classification (1998): I.6, H.5.3, J.1, K.6.0, H.3.5, G.1.2, F.2.2

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

ISSN 0302-9743
ISBN-10 3-642-16169-3 Springer Berlin Heidelberg New York
ISBN-13 978-3-642-16169-8 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 2010
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper 06/3180

Preface

The present volume was devoted to the third edition of the International Symposium on Algorithmic Game Theory (SAGT), an interdisciplinary scientific event intended to provide a forum for researchers as well as practitioners to exchange innovative ideas and to be aware of each other's efforts and results. SAGT 2010 took place in Athens, on October 18–20, 2010. The present volume contains all contributed papers presented at SAGT 2010 together with the distinguished invited lectures of Amos Fiat (Tel-Aviv University, Israel), and Paul Goldberg (University of Liverpool, UK). The two invited papers are presented at the beginning of the proceedings, while the regular papers follow in alphabetical order (by the authors' names).

In response to the call for papers, the Program Committee (PC) received 61 submissions. Among the submissions were four papers with at least one coauthor that was also a PC member of SAGT 2010. For these PC-coauthored papers, an independent subcommittee (Elias Koutsoupias, Paul G. Spirakis, and Xiaotie Deng) made the judgment, and eventually two of these papers were proposed for inclusion in the Scientific Program. For the remaining 57 (non-PC-coauthored) papers, the PC of SAGT 2010 conducted a thorough evaluation (at least 3, and on average 3.9 reviews per paper) and electronic discussion, and eventually selected 26 papers for inclusion in the Scientific Program.

An additional tutorial, “Games Played in Physics”, was also provided in SAGT 2010, courtesy of the academic research network *AlgoGames* (*Αλγοπαιγνιο*) of the University of Patras.

We wish to thank the creators of the EasyChair System, a free conference management system provided and supported by the group of Andrei Voronkov, which significantly assisted the work of the PC.

August 2010

Spyros Kontogiannis
Elias Koutsoupias
Paul G. Spirakis

Conference Organization

SAGT 2010 was organized jointly by the National & Kapodistrian University of Athens and the Research Academic Computer Technology Institute (CTI), in cooperation with the ACM Special Interest Group on Electronic Commerce (ACM-SIGECOM), the European Association for Theoretical Computer Science (EATCS), and the academic research network Algogames (*Αλγοπαιγνιο*) of the University of Patras.

Program Committee

Itai Ashlagi	Harvard Business School, USA
Xi Chen	University of Southern California, USA
Xiaotie Deng	City University of Hong Kong, Hong Kong
Kousha Etessami	University of Edinburgh, UK
Michal Feldman	Hebrew University of Jerusalem, Israel
Dimitris Fotakis	National Technical University of Athens, Greece
Paul W. Goldberg	University of Liverpool, UK
Rica Gonen	Yahoo! Research, USA
Haim Kaplan	Tel Aviv University, Israel
David Kempe	University of Southern California, USA
Spyros Kontogiannis	University of Ioannina, Greece
Annamaria Kovacs	Goethe University Frankfurt/Main, Germany
Elias Koutsoupas	National & Kapodistrian University of Athens, Greece (Co-chair)
Vangelis Markakis	Athens University of Economics and Business, Greece
Marios Mavronicolas	University of Cyprus, Cyprus
Aranyak Mehta	Google Research, USA
Seffi Naor	Technion - Israel Institute of Technology, Israel
Maher Said	Microsoft Research New England, USA
Michael Schapira	Yale University and UC Berkeley, USA
Paul G. Spirakis	Computer Technology Institute, Greece (Co-chair)
Bernhard von Stengel	London School of Economics and Political Science, UK
Berthold Vöcking	RWTH Aachen University, Germany

Steering Committee

Elias Koutsoupas	National & Kapodistrian University of Athens, Greece
Marios Mavronicolas	University of Cyprus, Cyprus
Dov Monderer	Technion, Israel
Burkhard Monien	University of Paderborn, Germany
Christos Papadimitriou	University of California at Berkeley, USA
Giuseppe Persiano	Università degli Studi di Salerno, Italy
Paul G. Spirakis	Computer Technology Institute, Greece (Chair)
Berthold Vöcking	RWTH Aachen, Germany

Organizing Committee

Rozina Efstathiadou	Computer Technology Institute, Greece
Lena Gourdoupi	Computer Technology Institute, Greece
Paul G. Spirakis	Computer Technology Institute, Greece (Co-chair)
Spyros Kontogiannis	University of Ioannina, Greece (Co-chair)
Elias Koutsoupas	National & Kapodistrian University of Athens, Greece

External Reviewers

Alaei, Saeed	Feldman, Moran	Meir, Reshef
Andersson, Daniel	Fu, Hu	Nazer Zadeh, Hamid
Anshelevich, Elliot	Gairing, Martin	Ochel, Marcel
Azar, Yossi	Hajiaghayi, Mohammad	Olonetsky, Svetlana
Babaioff, Moshe	Hansen, Thomas	Paes Leme, Renato
Bachrach, Yoram	Dueholm	Pierrakos, George
Ballester, Coralio	Harks, Tobias	Pountourakis, Emmanouil
Bateni, Mohammad Hossein	Hofer, Martin	Reijnierse, Hans
Brandt, Felix	Jain, Rahul	Roeglin, Heiko
Briest, Patrick	Jain, Shaili	Savani, Rahul
Bu, Tian-Ming	Kanellopoulos, Panagiotis	Sayedi, Amin
Chen, Ning	Krysta, Piotr	Schulz, Andreas
Chen, Po-An	Lai, John	Schwartz, Roy
Chien, Steve	Lavi, Ron	Segal, Ilya
Christodoulou, Giorgos	Legriel, Julien	Shalom, Israel
Daskalakis, Constantinos	Lewin-Eytan, Liane	Skopalik, Alexander
Devanur, Nikhil	Li, Xiaodong	Sun, Wei
Dobzinski, Shahar	Lin, Henry	Sundarajan, Mukund
Epstein, Amir	Lucier, Brendan	Syrgkanis, Vasileios
Fabrikant, Alex	Mansour, Yishay	Sørensen, Troels Bjerre

Xia, Lirong
Yaniv, Jonathan
Yannakakis, Mihalis

Yu, Jiajin
Zhang, Jie

Zhang, Jinshan
Zohar, Aviv

Sponsoring Institutions

- National & Kapodistrian University of Athens, Greece
- Research Academic Computer Technology Institute (CTI), Greece
- University of Patras, Greece
- Computer Engineering & Informatics Department, University of Patras, Greece
- Greek Ministry of Education
- ICT Programme of the EU under contract number FP7-215270 (FRONTS)
- ERC/StG Programme of the EU under the contract number 210743 (RIMACO)

Table of Contents

When the Players Are Not Expectation Maximizers (Invited Talk)	1
<i>Amos Fiat and Christos Papadimitriou</i>	
How Do You Like Your Equilibrium Selection Problems? Hard, or Very Hard? (Invited Talk)	15
<i>Paul W. Goldberg</i>	
A Simplex-Like Algorithm for Fisher Markets	18
<i>Bharat Adsul, Ch. Sobhan Babu, Jugal Garg, Ruta Mehta, and Milind Sohoni</i>	
Nash Equilibria in Fisher Market	30
<i>Bharat Adsul, Ch. Sobhan Babu, Jugal Garg, Ruta Mehta, and Milind Sohoni</i>	
Partition Equilibrium Always Exists in Resource Selection Games	42
<i>Elliot Anshelevich, Bugra Caskurlu, and Ameya Hate</i>	
Mixing Time and Stationary Expected Social Welfare of Logit Dynamics	54
<i>Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale, and Giuseppe Persiano</i>	
Pareto Efficiency and Approximate Pareto Efficiency in Routing and Load Balancing Games	66
<i>Yonatan Aumann and Yair Dombb</i>	
On Nash-Equilibria of Approximation-Stable Games	78
<i>Pranjal Awasthi, Maria-Florina Balcan, Avrim Blum, Or Sheffet, and Santosh Vempala</i>	
Improved Lower Bounds on the Price of Stability of Undirected Network Design Games	90
<i>Vittorio Bilò, Ioannis Caragiannis, Angelo Fanelli, and Gianpiero Monaco</i>	
On the Rate of Convergence of Fictitious Play	102
<i>Felix Brandt, Felix Fischer, and Paul Harrenstein</i>	
On Learning Algorithms for Nash Equilibria	114
<i>Constantinos Daskalakis, Rafael Frongillo, Christos H. Papadimitriou, George Pierrakos, and Gregory Valiant</i>	

On the Structure of Weakly Acyclic Games	126
<i>Alex Fabrikant, Aaron D. Jagard, and Michael Schapira</i>	
A Direct Reduction from k -Player to 2-Player Approximate Nash Equilibrium	138
<i>Uriel Feige and Inbal Talgam-Cohen</i>	
Responsive Lotteries	150
<i>Uriel Feige and Moshe Tennenholtz</i>	
On the Existence of Optimal Taxes for Network Congestion Games with Heterogeneous Users	162
<i>Dimitris Fotakis, George Karakostas, and Stavros G. Kolliopoulos</i>	
Computing Stable Outcomes in Hedonic Games	174
<i>Martin Gairing and Rahul Savani</i>	
A Perfect Price Discrimination Market Model with Production, and a (Rational) Convex Program for It	186
<i>Gagan Goel and Vijay Vazirani</i>	
The Computational Complexity of Trembling Hand Perfection and Other Equilibrium Refinements	198
<i>Kristoffer Arnsfelt Hansen, Peter Bro Miltersen, and Troels Bjerre Sørensen</i>	
Complexity of Safe Strategic Voting	210
<i>Noam Hazon and Edith Elkind</i>	
Bottleneck Congestion Games with Logarithmic Price of Anarchy	222
<i>Rajgopal Kannan and Costas Busch</i>	
Single-Parameter Combinatorial Auctions with Partially Public Valuations	234
<i>Gagan Goel, Chinmay Karande, and Lei Wang</i>	
On the Efficiency of Markets with Two-Sided Proportional Allocation Mechanisms	246
<i>Volodymyr Kuleshov and Adrian Vetta</i>	
Braess's Paradox for Flows over Time	262
<i>Martin Macko, Kate Larson, and Euboš Šteskal</i>	
The Price of Anarchy in Network Creation Games Is (Mostly) Constant	276
<i>Matúš Mihalák and Jan Christoph Schlegel</i>	
Truthful Fair Division	288
<i>Elchanan Mossel and Omer Tamuz</i>	

No Regret Learning in Oligopolies: Cournot vs. Bertrand	300
<i>Uri Nadav and Georgios Piliouras</i>	
On the Complexity of Pareto-optimal Nash and Strong Equilibria	312
<i>Martin Hoefer and Alexander Skopalik</i>	
2-Player Nash and Nonsymmetric Bargaining Games: Algorithms and Structural Properties	323
<i>Vijay V. Vazirani</i>	
On the Inefficiency of Equilibria in Linear Bottleneck Congestion Games	335
<i>Bart de Keijzer, Guido Schäfer, and Orestis A. Telelis</i>	
Minimal Subsidies in Expense Sharing Games	347
<i>Reshef Meir, Yoram Bachrach, and Jeffrey S. Rosenschein</i>	
Author Index	359

When the Players Are Not Expectation Maximizers

Amos Fiat¹ and Christos Papadimitriou²

¹ School of Computer Science
Tel Aviv University
fiat@tau.ac.il

² Computer Science Division
University of California at Berkeley
christos@cs.berkeley.edu

Abstract. Much of Game Theory, including the Nash equilibrium concept, is based on the assumption that players are expectation maximizers. It is known that if players are risk averse, games may no longer have Nash equilibria ([116]). We show that

1. Under risk aversion (convex risk valuations), and for almost all games, there are no mixed Nash equilibria, and thus either there is a pure equilibrium or there are no equilibria at all, and,
2. For a variety of important valuations other than expectation, it is NP-complete to determine if games between such players have a Nash equilibrium.

1 Introduction

In 1950 John Nash proved that every game has a mixed equilibrium. Myerson [17] gives a plethora of reasons as to why Nash's theorem (and his proposed framework of rationality in normal form games) underlies the foundations of modern economic thought. In recent years a computationally inspired challenge to the concept of mixed Nash equilibrium has arisen, see, e.g., [74], and the universality of the concept has become questionable in face of intractability results. In this paper we pursue another line of critique of the Nash equilibrium. In particular, we show that Nash's Theorem *does not hold* if the players are not expectation maximizers, in that almost all games fail to have a mixed Nash equilibrium (Theorem 3), and that it is NP-hard¹ to tell those that do from those that do not (Theorem 5). To understand our results in context, we begin by reviewing the rich literature on risk in Economics.

1.1 A Brief History of Risk

In many ways, risk is a defining characteristic of the modern world, and the analytical problems associated with it were pointed out early. Nicolas Bernoulli

¹ NP-hardness is stronger evidence of intractability than the PPAD-completeness of finding a Nash Equilibrium ([74]).

(1687 – 1759) posed the famous “St. Petersburg Paradox” [22], exposing the inadequacy of expectation in decision-making, and some decades later his nephew Daniel Bernoulli (1700 – 1782) [2] provided a solution by proposing to distinguish between money and the *utility of money*, and to model risk aversion by a utility function that is concave.

Two centuries later, Emil Borel [3] and John von Neumann [19] initiated the study of strategic behavior, and, two decades later, von Neumann and Oskar Morgenstern published their “Theory of Games and Economic Behavior” [20], where they expounded their Expected Utility Theory (EUT). They postulated that the risk behavior of an agent can be modeled as a (risk) *valuation* V mapping lotteries (distributions over the reals) to the reals². If this valuation satisfies some plausible axioms equivalent to *linearity*, then the agent’s behavior can be captured by a utility function, and the agent behaves as a maximizer of the expectation of his utility. A few years later, John Nash extended the work of von Neumann and Morgenstern to non-zero sum non-cooperative games [18], and showed that a mixed strategy equilibrium always exists; note that Nash’s Theorem is stated in the context of EUT. EUT can capture both risk-seeking and risk-averting agent behavior by having a valuation function that is convex or concave, respectively³.

In 1948, Friedman and Savage [8] attempted to deal with criticism of expected utility and considered models where utility is either a concave or convex function of money. Portfolio theory, developed in much more empirical and less principled/axiomatic manner from 1950 onwards by Marschak [16], Markowitz [15],[14] and many others [24],[27], considered valuations (functions from distributions to the reals) of the form “expectation minus variance” or “expectation minus standard deviation” as a model of agent behavior in the face of financial risk. E.g., the optimal portfolio for a given expected value is the one with minimum variance. There is no way to cast such behavior within the framework of Expected Utility Theory.

Independently, in 1951 Maurice Allais [1] suggested that there are problems with “the American School”, *i.e.*, he raised issues with the von Neumann-Morgenstern EUT. One of his examples was indeed the “expectation minus variance” valuation, but he also gave other empirical arguments (the *Allais paradox*) strongly suggesting that real human behavior cannot be modeled within EUT (and, consequently, certainly not by assuming that agents are expectation maximizers).

After Allais, many non-EUT valuations were proposed to address problems such as the Allais paradox, see the expository articles [26],[12],[13]. One such model

² Actually, instead of risk valuations, von Neumann and Morgenstern postulated, and axiomatized, *preferences* between any two distributions; but the valuation formalism used in this paper is essentially equivalent.

³ There is a point of possible confusion here. By “concave” valuation we mean a function from lotteries to the reals that is concave *in the probabilities*, and *not* in the values. For example, the variance is concave in the probabilities but convex in the values.

in wide use today is *prospect theory* due to Kahneman and Tversky [10]. Prospect theory predicts that a loss of x is much more painful than a gain of x is pleasant, and, importantly, that probabilities undergo subjective modifications in the agents' calculation of expectations, not unlike our example valuations V^5 and V^6 below.

More directly related to our work, Ritzberger [23] showed that for expected utilities with rank dependent probabilities reflected risk aversion, mixed Nash Equilibria will disappear. Chen and Neilson [9] considered the flip side of this phenomena and gave conditions under which a pure strategy must exist (but this requires a compact set of pure strategies).

1.2 Nash Equilibria and Risk

Whereas EUT was proposed by von Neumann and Morgenstern as a preamble to their theory of games and strategic behavior, non-EUT approaches to risk were primarily considered in non-strategic settings such as finance. The question of how non-EUT valuations impact non-cooperative game theory and Nash equilibria in particular was raised only in the 1990s by Crawford [6], who noted that, while Nash's theory holds when the agents' valuations are concave (see our Theorem 1, stated and proved here for completeness and computational emphasis), there are simple games, such as the 2×2 zero-sum game shown in (1) below, that have no Nash equilibria if the agents have convex risk valuations.

To understand the broad range of possible attitudes of strategic agents towards risk, consider the following six valuations (functions mapping distributions to the reals) modeling plausible attitudes of agents towards risk:

- V^1 If an agent is a *pure expectation maximizer*, then his valuation V^1 maps any distribution to its expected value. This is the framework used in virtually all of Game Theory.
- V^2 Most people are *risk averse*. One way to capture this would be valuation V^2 , which assigns to each distribution over the reals the expectation minus the variance; this was proposed by Marschak, Markowitz, Allais, and others. We use V_2 as an exemplar of risk averse valuations; there are many variants of V^2 in which one subtracts from the expectation the standard deviation or some other increasing function of the variance, or a small multiple thereof.
- V^3 Some agents may be *risk-seeking*; for example, valuation V^3 evaluates a distribution by its expectation *plus* (an increasing function of) the variance.
- $V^4[\theta]$ An agent may be facing a costly life-saving medical procedure and his only interest in the game is to maximize the probability that the payoff is above the cost θ . This defines valuation $V^4[\theta]$.
- $V^5[p]$ Another agent may be interested in maximizing her "almost certain bottom line:" the amount of money she gets with probability at least $p = .95$, say. Let us call this valuation $V^5[p]$.
- V^6 Finally, somebody else evaluates any discrete distribution over the reals by the average between the maximum and the minimum value which occur with a nonzero probability. We call this V^6 .

Note: We are not proposing these six valuations as the only possible attitudes toward risk, or even as plausible or reasonable ones; they are here only to demonstrate the range of possibilities and fix ideas. Our three results hold for very broad classes of such valuations, delimited in their statement or the discussion following the proof. Also, the last two valuations fall into a very important class proposed in Kahneman and Tversky’s Prospect Theory [10], in which expected utility is maximized, albeit with the probabilities modified. Briefly, in Prospect Theory valuations are of the form $V = \sum_i u(x_i)\pi_i$, where u is an ordinary utility function, but the π_i ’s are *modified probabilities*. The modification is done through an increasing function $G : [0, 1] \mapsto [0, 1]$ with $G(0) = 0$ and $G(1) = 1$ that *modifies the cumulative probabilities*. That is, if we assume that $x_1 \leq x_2, \leq \dots \leq x_n$, π_i is defined as $G(\sum_{j=1}^i p_j) - G(\sum_{j=1}^{i-1} p_j)$. It is easy to see that $V^5[p]$ corresponds to the modifier function $G(x) = 0$ if $x \leq 1 - p$, and $G(x) = 1$ otherwise. And V^6 corresponds to $G(x) = \frac{1}{2}$ for $0 < x < 1$. Kahneman and Tversky speculate that “real” modifier functions, consistent with experiments, are steeply increasing at 0 and at 1, go through the $(\frac{1}{2}, \frac{1}{2})$ point, and are flat around it. The effect is that small probabilities of extreme payoffs are exaggerated. Notice that our function G defining V^6 is a stylized and exaggerated function of this form ($\pi_1 = \pi_n = \frac{1}{2}$).

Of these six risk valuations, V^1 is the one considered throughout Game Theory and, naturally, Nash’s Theorem holds in it. Of the others, $V^4(\theta)$ falls squarely within the purvey of EUT: Just map the agent’s payoffs to zero if they are less than θ and to one otherwise, and solve the resulting game. As it turns out, Nash’s Theorem is valid under the risk-seeking valuation V^3 as well (see Theorem 1 and Proposition 2); the reason is, V^3 is *concave* in the probabilities.

The other three valuations, however, break Nash’s Theorem. For example, consider the following game proposed by Crawford, which we call F :

$$\begin{array}{|c|c|} \hline 1, -1 & 0, 0 \\ \hline 0, 0 & r, -r \\ \hline \end{array} \quad (1)$$

If the agents evaluate any distribution of payoffs by a convex valuation such as V^2 , then Crawford observes that there are no Nash equilibria in this game, pure or mixed. This holds for $r \neq 1$; interestingly, if $r = 1$ then there is a mixed Nash equilibrium with both players V^2 (or even with one player with valuation V^2 and the other player is an expectation maximizer).

1.3 Our Results

1. In terms of ubiquity of the Nash Equilibria, we show that
 - (a) *Almost all games have no mixed Nash equilibria if the players are risk averse* (Theorem 3). By “almost all” we mean that games that do have mixed Nash equilibria form a set of measure zero in the space of all games, with utilities drawn at random; for example, any such games must have equality between certain payoff values. Moreover, even ϵ -Nash equilibria will not exist. Pure Nash equilibria may still exist; but only a $1 - \frac{1}{e}$ fraction of games have them [25]. This was known for rank dependent expected utility functions (Ritzberger, 1996, [23]).

- (b) Even if the underlying game has a mixed Nash equilibrium, arbitrarily small random errors by the players in interpreting the payoff matrix will lead to instability with high probability (Observation [4](#)).
2. Any given game may not have a Nash equilibrium. We show that
- (a) *It is NP-complete to determine if a two-person game with non-EUT player valuations has a Nash Equilibrium*[4](#). We show this for functions such as V^2 , $V^5[p]$, and V^6 , and we lay out broad conditions on the risk valuations under which our proof works (Corollary [6](#)).
- (b) In contrast, for concave valuations such as V^3 — Nash equilibria are guaranteed to exist and are PPAD-complete, *i.e.*, the same complexity as Nash Equilibria under the expected utility theory (Theorem [11](#)).

1.4 The Model

To avoid confusion, we use the terminology of payoffs and valuations rather than utility. Under expected utility theory, our “payoffs” are considered utilities and some of our valuation functions are also utilities, whereas others cannot be so expressed.

A k -player game G , where $k > 1$, consists of k finite sets of *strategies* S_1, \dots, S_k and k *payoff functions* p_1, \dots, p_k mapping $S = \prod_i S_i$ to \mathbb{R} . We denote by $\Delta[S_i]$ the set of *mixed strategies* for player i .

Given a k -tuple of mixed strategies $x = (x_1, \dots, x_k) \in \prod_i \Delta[S_i]$, for any combination of pure strategies $s = (s_1, s_2, \dots, s_k) \in \prod S_i$ define $q_x(s) = \prod_{i=1}^k x_i(s_i)$, where $x_i(s_i)$ is the probability player i plays pure strategy s_i . Every k -tuple of mixed strategies, x , defines a *strategy distribution*, $S(x)$, over $\prod S_i$, where the probability of $s \in \prod S_i$ being played is $q_x(s)$.

Let G be a k -player game, and V_1, \dots, V_k be valuations. Given a payoff function p_i , let $R_i = p_i(s)$, $s \in \prod_i S_i$. R_i is the range of possible payoffs for agent i over all combinations of pure strategies.

Given a k -tuple of mixed strategies $x = (x_1, \dots, x_k)$, for every player i the strategy-distribution $S(x)$ implies a *payoff-distribution*, $P_i(x)$, the support of $P_i(x)$ is a subset of R_i , and the probability of $a \in R_i$ is

$$q_{i,x}(a) = \sum_{s \in S | p_i(s) = a} q_x(s).$$

Again, for any $1 \leq i \leq k$ we have $\sum_{a \in R_i} q_{i,x}(a) = 1$

A (*risk*) *valuation* is any function from payoff distributions to the reals. The functions $V^1, V^2, V^3, V^4[\theta]$, $V^5[p]$, and V^6 mentioned in the introduction are indicative examples of important valuations in the literature. Ergo, if agent i has risk valuation V_i then the value of a k tuple of mixed strategies $x = (x_1, x_2, \dots, x_k)$ to player i is $V_i(P_i(x))$.

⁴ Our NP-completeness proof uses a new gadget, based on a generalized rock-paper-scissors game, which is arguably simpler than the construction in [5](#), we suspect that this gadget may prove useful in other contexts as well.

Fix i and the mixed strategies of all players but that of player i , $x_{-i} \in \prod_{j \neq i} \Delta(S_j)$. For fixed $x_{-i} \in \prod_{j \neq i} \Delta(S_j)$, $x_i \in \Delta(S_i)$, we use the shorthand notation

$$V^{x_{-i}}(x_i) = V(P_i(x_i, x_{-i})).$$

Finally, a k -tuple of mixed strategies $x = (x_1, \dots, x_k) \in \prod_i \Delta[S_i]$ is called a (V_1, \dots, V_k) -Nash equilibrium if for every player i and every mixed strategy $x'_i \in \Delta[S_i]$,

$$V_i^{x_{-i}}(x'_i) \leq V_i^{x_{-i}}(x_i).$$

That is, changing i 's mixed strategy in any way, and keeping all other mixed strategies the same, results in a distribution over payoffs whose valuation, with respect to player i , is no better than the valuation of the current distribution.

Examples: Consider the valuations V^1, \dots, V^6 introduced above, and let us concentrate on the risk averse valuation V_2 , expectation minus variance. Analyzing mixed Nash equilibria when players behave this way (or in any of the other five V^i 's save V^1 , expectation) is very tricky. Pure strategies in a mixed Nash equilibrium are not necessarily individually best responses to the other players' strategies. Also, because of the nonlinearity of the valuations considered, games are not invariant under translation or scaling by positive constants.

Suppose that Crawford's game (see Equation (II)) is played by a row player with valuation V^2 , and a column player who is an expectation maximizer (i.e., the column player is risk neutral). If $r = 1$, then a mixed Nash equilibrium exists in which both players randomize uniformly. As we shall see, this situation is a singular exception (see the proof of Theorem 3). But suppose that $r > 1$. From the point of view of the expectation maximizer (column player), as the probability that the row player plays down is increased from zero, playing left is the best response, up to a point in which left and right — *and any mixture in between* — are at a tie. From then on, right is the best response. A similar behavior is observed of the V^2 (row) player — with an important difference. As the column player increases from zero the probability of playing right, the row player prefers up, and at some point there is a tie between up and down.

The catch is that, because of the convexity of V^2 , the points in between are *not* best responses, and there is a discontinuous jump from up to down. As a result, the trajectories of the dynamics of the two players (the two best response maps) do not intersect, and the game has no equilibrium. The same behavior is observed if the row player's valuation is expectation minus any increasing function of the variance (say, a small multiple, or square root), and if both players are risk averse.

Intuitively, the reason that $r = 1$ is a singularity (and a mixed Nash Equilibrium does in fact exist) is because the payoff-distributions don't depend on the mixed strategy chosen by the row player. So, although the strategy-distributions are in fact different, the payoff-distributions are invariant to changes in the row strategy, and the valuation to the row player is therefore also invariant to changes in the row strategy. , i.e., $V^{x_{-row}}(x_{row}) = V^{x_{-row}}(x'_{row})$, for any $x_{row} = (q, 1-q)$, $x'_{row} = (q', 1-q')$, $0 \leq q, q' \leq 1$.

2 Valuation Convexity and Concavity

Valuation functions are defined on payoff distributions. We require (and use) a limited form of convexity/concavity, that of valuation V_i convex (resp. concave) with respect to mixed strategies of player i , $x_i \in \Delta(S_i)$.

A valuation function V_i is said to be convex (resp. concave) with respect to x_i if for every $k-1$ tuple of mixed strategies $x_{-i} \in \prod_{j \neq i} \Delta(S_j)$, for any two mixed strategies $x_i, x'_i \in \Delta(S_i)$, and for any $0 \leq \alpha \leq 1$.

$$\begin{aligned} V^{x_{-i}}(\alpha x_i + (1-\alpha)x'_i) &\leq \alpha V^{x_{-i}}(x_i) + (1-\alpha)V^{x_{-i}}(x'_i); \\ V^{x_{-i}}(\alpha x_i + (1-\alpha)x'_i) &\geq \alpha V^{x_{-i}}(x_i) + (1-\alpha)V^{x_{-i}}(x'_i); \end{aligned}$$

A valuation function V_i is said to be strictly convex, or strictly concave, respectively, if for any two payoff distributions $P_i(x_i, x_{-i}) \neq P_i(x'_i, x_{-i})$, and for any $0 < \alpha < 1$,

$$\begin{aligned} V^{x_{-i}}(\alpha x_i + (1-\alpha)x'_i) &< \alpha V^{x_{-i}}(x_i) + (1-\alpha)V^{x_{-i}}(x'_i); \\ V^{x_{-i}}(\alpha x_i + (1-\alpha)x'_i) &> \alpha V^{x_{-i}}(x_i) + (1-\alpha)V^{x_{-i}}(x'_i); \end{aligned}$$

We say that V_i is *efficiently concave* in $x_i \in \Delta(S_i)$, if, for any x_{-i} , and for any strictly concave polynomial $t(x_i)$, the (unique) point $\operatorname{argmax}_{x_i} V^{x_{-i}}(x_i) + t(x_i)$ can be computed in time polynomial in $|S_i|$, the representation of t , the total number of bits in the coefficients of $V(x_i)$, and the number of bits of precision required.

3 Computing Nash Equilibria

We now give a sufficient condition on the V_i 's for (V_1, \dots, V_k) -Nash equilibria to exist and be as easy to compute as ordinary Nash equilibria: it suffices for each V_i to be efficiently concave in x_i . This result is well known to economists; for example, von Neumann stated his minmax Theorem in terms of concave functions. Here we restate and prove it for computational emphasis and contrast with our main result that follows.

Theorem 1. *Let G be a k player game and let V_1, \dots, V_k be valuations, where $k > 1$. If each V_i is concave in x_i , then any k -player game has an (V_1, \dots, V_k) -Nash equilibrium. If in addition the V_i 's are efficiently concave, then the problem of finding an (V_1, \dots, V_k) -Nash equilibrium is in PPA.*

Proof. Define the following function ϕ from $\prod_i \Delta[S_i]$ to itself: $\phi(x_1, \dots, x_k) = (y_1, \dots, y_k)$, where for each i y_i is defined as follows:

$$y_i = \operatorname{argmax}_{z_i \in \Delta(S_i)} V_i^{x_{-i}}(z_i) - \|z_i - x_i\|^2.$$

Since V_i is concave in x_i , and $-\|z_i - x_i\|^2$ is a strictly concave polynomial in z_i , the argmax exists and is unique, and continuous as a function of x . Therefore ϕ is a continuous function from a compact convex set to itself, and so, by

Brouwer's fixpoint theorem, ϕ has a fixpoint. It is now easy to see, arguing by contradiction, that this fixpoint is an (V_1, \dots, V_k) -Nash equilibrium of G . The computational elaboration of the theorem follows easily from the fact that approximating Brouwer fixpoints is in PPA [21]. \square

We next point out that three of our original six valuation examples fall into this benign category.

Proposition 2. *The valuations V^1, V^3 , and $V^4[\theta]$ are efficiently concave.*

Proof. It is easy to see that V^1 and $V^4[\theta]$ are actually linear, and thus trivially efficiently concave. For V^3 we need to establish that the variance of the distribution $P_i(x)$ is concave in x_i . Recall that for any random variable X the variance is $E(X^2) - (E(X))^2$. Since the first term is linear in x_i , we concentrate on the second term. It is easy to see that, if X is the random variable for the payoff to player i , the Hessian of $-(E(X))^2$ with respect to x_i (variables $x_i[s]$, $s \in S_i$) is

$$\frac{\partial^2[-(E(X))^2]}{\partial x_i[s_i] \partial x_i[s'_i]} = -h_{s_i} \cdot h_{s'_i},$$

where $h_{s_i} = E_x[X|s_i]$, and similarly for $h_{s'_i}$. Now it is clear that the Hessian is the tensor product of vector $(h_{s_i} : s_i \in S_i)$ with itself, negated, and thus it is trivially a negative semi-definite matrix. Hence the variance is indeed a concave function of the probabilities, and so is risk seeking valuation V^3 . That it is efficiently concave is straightforward. \square

Note: There is a point of confusion here. It is well known, and often useful, that the variance is a *convex* function of the *values*. However, it turns out that it is also a concave function of the *probabilities*.

Notice that adding to the expectation, instead of the variance, a positive multiple, or any concave function, of the variance (such as the standard deviation), preserves the valuation's concavity. Hence the positive result stated in the theorem applies to a broad variety of risk-seeking valuations.

4 Games with No Equilibria

But what if the valuations are not concave — for example, convex like V^2 ? Risk averse agents typically have strictly convex valuations. We know from Crawford that Nash equilibria may not exist, but two questions come up immediately: How prevalent are such pathologies? And even if they are prevalent, can they at least be characterized and excluded? In this and the next section we answer both questions in the negative.

For a pure strategy $s \in S_i$ let x_i^s be the strategy for agent i that plays s deterministically. Call a game *in general position* if for any player i , for any mixed strategy of the other players x_{-i} , and for any two pure strategies $s, s' \in S_i$, the payoff distributions $P_i(x_i^s, x_{-i}) \neq P_i(x_i^{s'}, x_{-i})$. Note that $P_i(x_i^s, x_{-i}) \neq P_i(x_i^{s'}, x_{-i})$ does not imply that $V_i^{x_{-i}}(x_i^s) \neq V_i^{x_{-i}}(x_i^{s'})$.

Our first negative result states that, if the players are risk-averse, then a game in general position cannot have a mixed Nash equilibrium. This result has an evocative probabilistic interpretation: In the space of all games considered as tuples of tensors, games with mixed Nash equilibria form a set of measure zero, a lower-dimensional manifold. Since it is known that the probability that a game has a pure Nash equilibrium in this space is asymptotically $1 - \frac{1}{e}$, this means that about 37% of all games have no Nash equilibrium when the players are risk averse.

Theorem 3. *Let Γ be a game in general position, and suppose that player i 's valuation is strictly convex (as a function of $x_i \in \Delta[S_i]$). Then in any Nash equilibrium this player plays a pure strategy. In particular, if all agent valuations are strictly convex, there is no mixed Nash equilibrium.*

Proof. Fix the (possibly mixed) strategies of all agents but i , x^{-i} . Assume that agent i has a strictly mixed strategy, which, along with all other agent strategies, is in Nash Equilibrium.

As x_i is a strictly mixed strategy, then for all $s \in S_i$, $x_i(s) < 1$, and for at least two distinct $s, s' \in S_i$, $x_i(s) > 0$ and $x_i(s') > 0$. It follows that

$$V_i^{x^{-i}}(x_i) = V_i^{x^{-i}} \left(\sum_{s \in S_i} x_i^s \cdot x_i(s) \right) < \sum_{s \in S_i} x_i(s) V_i^{x^{-i}}(x_i^s) \leq \max_{s \in S_i} V_i^{x^{-i}}(x_i^s).$$

This is in contradiction to the assumption that x_i was in Nash Equilibrium (alternately that $x_i = \operatorname{argmax}_{z_i \in \Delta(S_i)} V_i^{x^{-i}}(z_i)$). \square

Theorem 3 is in no way a converse of Theorem 1: there are many valuations that are neither concave nor strictly convex — valuations V^5 and V^6 , for example. However, it is not hard to see that for these two there are games in which there are no Nash equilibria. Generalizing Theorem 2 so that it comes close to being a converse of Theorem 1 is an interesting open question.

It now remains to ask: when is a game in general position? A game is not in general position if for some player i , there exists a $k - 1$ tuple of mixed strategies for the other players, x_{-i} , such that for two different pure strategies for player i , $s_i, s'_i \in S_i$, we have that $P_i(x^{s_i}, x_{-i}) = P_i(x^{s'_i}, x_{-i})$. Ergo, the two distributions don't depend on the choice between s_i and s'_i . The $k - 1$ tuple, x_{-i} , determines at most $N = \prod_{j \neq i} |S_j|$ probabilities for the various payoffs available if player i plays s and at most N probabilities for the various payoffs if player i plays s' .

Observation 4. *A sufficient condition for the game to be in general position is if all payoffs for player i are distinct, i.e., $p_i(s_i, s_{-i}) \neq p_i(s'_i, s_{-i})$ for all $s_i, s'_i \in S_i$ and $s_{-i} \in \prod_{j \neq i} S_j$. The two distributions $P_i(x_i^s, x_{-i})$ and $P_i(x_i^{s'}, x_{-i})$ cannot be the same because the support sets are different. This is not a necessary condition, e.g., Crawford's game with $r \neq 1$ is in general position.*

This also implies that if a small random error is added to every payoff value then the game will be in general position with high probability. \square

5 The Complexity of Risk Aversion

Theorem 5. *Given a 2-player game and non-concave valuations V, V' , it is NP-complete to tell if the game has an (V, V') -Nash equilibrium.*

Proof. We give below the proof for the case in which both players' valuations are V^2 ; Subsequently, we give extensions and generalizations.

NP-completeness is via reduction from 3SAT. Given a 3SAT instance ϕ with n variables x_1, \dots, x_n and a set of clauses $C = \{c_1, \dots, c_m\}$, where each c_j is a subset of size three of the set of literals $L = \{+x_1, -x_1, +x_2, \dots, -x_n\}$, we construct a 2-player game G_ϕ as follows. Both players have the same strategy set $S_1 = S_2 = L \cup C \cup \{f_1, f_2\}$, and their utilities are as follows ($M = 4n^2$):

- For every variable $x_i \in V$,

$$\begin{aligned} p_1(+x_i, +x_i) &= p_2(+x_i, +x_i) = p_1(-x_i, -x_i) \\ &= p_2(-x_i, -x_i) = M. \end{aligned}$$

However,

$$\begin{aligned} p_1(+x_i, -x_i) &= p_2(+x_i, -x_i) = p_1(-x_i, +x_i) \\ &= p_2(-x_i, +x_i) = M - 2n. \end{aligned}$$

Also, for every two variables $x_i, x_j \in V$ with $i \neq j$, $p_1(\pm x_i, \pm x_j) = M + g(i, j)$, and $p_2(\pm x_i, \pm x_j) = M - g(i, j)$, where $g(i, j) = 1$ if $j = i + 1 \pmod n$, -1 if $j = i - 1 \pmod n$, and 0 otherwise.

The game restricted to L is a generalized rock-paper-scissors zero-sum game (with payoffs translated by M) in which the signs of the literals do no matter, except that both players are incentivized not to play opposite literals of the same variable. It is easy to see that there are 2^n (V^2, V^2)-Nash equilibria of this game. In each, the two players choose the same truth assignment (n literals, one for each variable) and play every literal with probability $\frac{1}{n}$. The V_2 valuation for this payoff distribution is $M - \frac{2}{n}$.

- The purpose of the strategies in C is to force the truth assignment chosen by the equilibrium in L to satisfy ϕ ; if it does not, there is a strategy in C (the violated clause) that breaks the equilibrium by presenting a better alternative. The utilities are as follows: For any $c, c' \in C$, $p_1(c, c') = p_2(c', c) = M - 2n$ (it is disadvantageous for both players to both play in C). Also, for any literal λ and any clause c , $p_1(\lambda, c) = p_2(c, \lambda) = M - 2n$ (it is also disadvantageous to play a literal if the opponent is playing a clause.) Now, to encode the 3SAT instance, for any literal λ and clause c such that $\lambda \in c$, $p_1(c, \lambda) = p_2(\lambda, c) = M - \frac{2}{n}$, whereas if $\lambda \notin c$ we have $p_1(c, \lambda) = p_2(\lambda, c) = M + n - \frac{2}{n}$.
- Finally, the last two strategies f_1, f_2 provide an alternative game, to be played if an equilibrium does not exist in $L \cup C$ — but this game is essentially Crawford's game Γ (see Equation (III)), known to have no Nash equilibria for risk averse players. In particular,

- $p_1(x, f_j) = p_2(f_j, x) = 0$ for all $x \in L \cup C, j = 1, 2$;
 - $p_1(f_j, x) = p_2(x, f_j) = M - \frac{2}{n}$ for all $x \in L \cup C, j = 1, 2$;
 - $p_1(f_1, f_1) = \frac{M}{2} + 1$;
 - $p_1(f_1, f_2) = \frac{M}{2}$;
 - $p_1(f_2, f_1) = \frac{M}{2}$;
 - $p_1(f_2, f_2) = \frac{M}{2} + 2$.
 - $p_2(f_1, f_1) = \frac{M}{2} - 1$;
 - $p_2(f_1, f_2) = \frac{M}{2}$;
 - $p_2(f_2, f_1) = \frac{M}{2}$;
 - $p_2(f_2, f_2) = \frac{M}{2} - 2$.
- Each player can choose either to play the strategies in $L \cup C$ (we call this “playing ϕ ”), or choose one of the strategies f_1, f_2 (called “playing f ”). If both players play f , notice that they end up playing a version of the game Γ described above, translated upwards by $\frac{M}{2}$.
- Suppose that the two players play f with probability x and y , respectively. Then it is easy to see that the V^2 value of the first player is approximately

$$\begin{aligned} \frac{\tilde{M}}{2}xy + \tilde{M}(1-y) - \frac{\tilde{M}^2}{4}xy - (1-y)\tilde{M}^2 \\ + \frac{\tilde{M}^2}{4}x^2y^2 + \tilde{M}^2(1-y)^2 - \tilde{M}^2xy(1-y), \end{aligned}$$

where by \tilde{M} we denote $M(1 \pm O(\frac{1}{n}))$. By looking at the polynomial in x that results if we ignore $O(\frac{1}{n})$ terms and normalize by M^2

$$\left(\frac{y^2}{4}\right) \cdot x^2 + \left(y^2 - \frac{3y}{4}\right) \cdot x - y(1-y),$$

we notice (just by looking at the quadratic term) that, for all values of $y \in (0, 1]$, it attains its maximum (in x) at either $x = 0$ or $x = 1$. Hence, for large enough n , and if $y > 0$, player 1 either chooses purely ϕ , or purely f . If $y = 0$, then it is easy to see that $x = 0$ as well, because otherwise player 2 would be better off playing f ($y = 1$). Hence, player 1 plays either purely ϕ or purely f . By symmetry, the same holds for player 2. We conclude that, for large enough n , the only possible (V^2, V^2) -Nash equilibria have either both players playing ϕ , or both playing f .

- But of course, since both players choosing f entails playing a translated version of game Γ , there can be no (V^2, V^2) -Nash equilibria in which both players play f . Also, one player playing ϕ and the other f cannot be an equilibrium (it is a disaster for the player playing ϕ). We conclude that in any (V^2, V^2) -Nash equilibrium both players must play ϕ .
- So, let us assume that both players play ϕ . If one of them plays a strategy in C , then the other player cannot be playing strategies in L (because the f strategies would fare better for the other player than those in L , which now have V^2 payoff strictly less than $M - \frac{2}{n}$). Hence, the other player must be playing only C , and so the first player’s payoff with the strategy in c is again lower than that of the f strategies, and this cannot be an equilibrium.

- We conclude that if an equilibrium exists its support is a subset of L , and hence it is a truth assignment. We claim that it is satisfying. Because, if not, there is a clause c that is not satisfied — that is, no literal in it is played by both players. In this case, c has V^2 payoff equal to $M + n - \frac{2}{n}$, and so the players should rather play c . We conclude that if the game has an (V^2, V^2) -equilibrium, then ϕ is satisfiable.
- Conversely, if ϕ is satisfiable, we claim that both players playing the literals in the satisfying truth assignment with probability $\frac{1}{n}$ is an (V^2, V^2) -Nash equilibrium. The (V^2, V^2) payoff of each player is $M - \frac{2}{n}$. Playing the f strategies instead would have the same payoff. And playing any strategy in C would have (V^2, V^2) payoff at most $M - \frac{2}{n}$, because at least one literal in the clause is played with probability $\frac{1}{n}$, and this brings the payoff down to $M - \frac{2}{n}$. \square

Other Valuations. Looking at the proof of Theorem 5 we note that the following holds:

Corollary 6. *Suppose that V and V' are risk valuations with the following properties:*

- *There are games with arbitrarily large payoffs that have no (V, V') -Nash equilibria;*
- *The only (V, V') -Nash equilibrium of the generalized rock-paper-scissors game, even if shifted by M , is the uniform play;*
- *For large enough M and for any game whose payoff matrices are within L_2 distance one from*

M, M	$0, M$
$M, 0$	M, M

the only (V, V') -Nash equilibria are pure.

Then it is NP-complete to tell if a game has a (V, V') -Nash equilibrium. \square

6 Discussion and Open Problems

No equilibria means that agents may be in a state of constant flux, and NP-completeness implies that they can't even realize their predicament. For a game matrix Γ and valuation functions $\{V_i\}$, if Nash Equilibria do not exist then there is some $f(\Gamma, \{V_i\})$ such that for all $\varepsilon < f(\Gamma, \{V_i\})$ no ε -Nash equilibrium exists. Thus, $f(\Gamma, \{V_i\})$ is a measure of the instability of the player dynamics and large values may help explain dramatic instabilities in certain games. It is easy to choose r for Crawford's game or jittered versions of matching pennies, with risk valuations V^2 , for which $f(\Gamma, \{V^2\})$ equal to some 10% of the total value — very strong motivation for chaotic behavior. Measuring the instability of a game played by risk-averse players, and comparing this to observed behavior in strategic games, even markets, seems interesting.

Also, as pointed out in Observation 4, incomplete or partial information about the payoffs, if interpreted as making some random error about the value of the payoffs, results in games in general position and no mixed equilibria. It seems promising to characterize the instability of the resulting game as a function of the magnitude of error. Intuitively, instability should grow with error size.

The possibility that mixed Nash equilibria may not exist makes weaker solution concepts, such as the *correlated equilibrium*, more attractive. Unfortunately, it is easy to see that even those may not exist. For example, Crawford's uneven matching pennies game played by players whose risk valuation is V_6 has no correlated equilibria. We conjecture that the circumstances under which correlated equilibria fail to exist are much less common, but that it is still NP-complete to tell if one exists or not.

One technical open problem suggested by this work is to determine the extent of the valuations for which NP-completeness holds. We have not strived to state the most general theorem possible here, but we believe that our proof can be generalized in many directions. Note that, by the corollary, NP-completeness holds even when all players except for one are expectation maximizers.

References

1. Allais, M.: Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica* 21(4), 503–546 (1953), <http://www.jstor.org/stable/1907921>
2. Bernoulli, D.: Exposition of a new theory on the measurement of risk. Commentaries of the Imperial Academy of Science of Saint Petersburg (1738); Republished *Econometrica* 22(1), 23–36 (1954), <http://www.jstor.org/stable/1909829>
3. Borel, E.: The theory of play and integral equations with skew symmetric kernels. In: Translated by Leonard J. Savage, University of Chicago, from *Comptes Rendus de l'Académie des Sciences*, December 19, vol. 173, pp. 1304–1308 (1921); Reprinted in *Econometrica* 21(1), 97–100 (1953), <http://www.jstor.org/stable/1906946>
4. Chen, X., Deng, X.: Settling the complexity of two-player nash equilibrium. In: FOCS 2006: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, pp. 261–272. IEEE Computer Society, Washington (2006)
5. Conitzer, V., Sandholm, T.: New complexity results about nash equilibria. *Games and Economic Behavior* 63(2), 621–641 (2008), <http://econpapers.repec.org/RePEc:eee:gamebe:v:63:y:2008:i:2:p:621-641>
6. Crawford, V.P.: Equilibrium without independence. *Journal of Economic Theory* 50(1), 127–154 (1990), <http://www.sciencedirect.com/science/article/B6WJ3-4CYH5N4-CF/2/0071841ad1283474807beca912feb550>
7. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a nash equilibrium. *ACM Commun.* 52(2), 89–97 (2009)
8. Friedman, M., Savage, L.J.: The utility analysis of choices involving risk. *Journal of Political Economy* 56 (1948), <http://econpapers.repec.org/RePEc:ucp:jpolec:v:56:y:1948:p:279>
9. Chen, H.-C.: Pure-strategy equilibria with non-expected utility players. *Theory and Decision* 46, 201–212 (1999), <http://www.ingentaconnect.com/content/klu/theo/1999/00000046/00000002/00152924>

10. Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291 (1979), <http://www.jstor.org/stable/1914185>
11. Karni, E., Safra, Z.: Dynamic consistency, revelations in auctions and the structure of preferences. *Review of Economic Studies* 56(3), 421–433 (1989), <http://econpapers.repec.org/RePEc:bla:restud:v:56:y:1989:i:3:p:421-33>
12. Machina, M.J.: Choice under uncertainty: Problems solved and unsolved. *The Journal of Economic Perspectives* 1(1), 121–154 (1987), <http://www.jstor.org/stable/1942952>
13. Machina, M.J.: Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature* 27(4), 1622–1668 (1989), <http://ideas.repec.org/a/aea/jecclit/v27y1989i4p1622-68.html>
14. Markowitz, H.: *Portfolio Selection: Efficient Diversification in Investments*. John Wiley, Chichester (1959)
15. Markowitz, H.: Portfolio selection. *The Journal of Finance* 7(1), 77–91 (1952), <http://www.jstor.org/stable/2975974>
16. Marschak, J.: Rational behavior, uncertain prospects, and measurable utility. *Econometrica* 18(2), 111–141 (1950), <http://www.jstor.org/stable/1907264>
17. Myerson, R.B., Myerson, R.B.: Nash equilibrium and the history of economic theory. *Journal of Economic Literature* 37, 1067–1082 (1999)
18. Nash, J.: Non-cooperative games. *The Annals of Mathematics* 54(2), 286–295 (1951), <http://www.jstor.org/stable/1969529>
19. Neumann, J.v.: Zur Theorie der Gesellschaftsspiele. *Math. Ann.* 100(1), 295–320 (1928), <http://dx.doi.org/10.1007/BF01448847>
20. Neumann, J.V., Morgenstern, O.: *Theory of Games and Economic Behavior*. Princeton University Press, Princeton (1944), <http://jmvidal.cse.sc.edu/library/neumann44a.pdf>
21. Papadimitriou, C.H.: On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.* 48(3), 498–532 (1994)
22. Translated by Richard J. Pulskamp, N.B.: Correspondence of nicolas bernoulli concerning the st. petersburg game (1713-1732) (1999), <http://cerebro.xu.edu/math/sources/montmort/stpetersburg.pdf>
23. Ritzberger, K.: On games under expected utility with rank dependent probabilities. *Theory and Decision* 40(1), 1–27 (1996)
24. Sharpe, W.F.: A simplified model for portfolio analysis. *Management Science* 9(2), 277–293 (1963), <http://www.jstor.org/stable/2627407>
25. Stanford, W.: The limit distribution of pure strategy nash equilibria in symmetric bimatrix games. *Mathematics of Operations Research* 21(3), 726–733 (1996)
26. Starmer, C.: Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* 38(2), 332–382 (2000), <http://ideas.repec.org/a/aea/jecclit/v38y2000i2p332-382.html>
27. Varian, H.: A portfolio of nobel laureates: Markowitz, miller and sharpe. *The Journal of Economic Perspectives* 7(1), 159–169 (1993), <http://www.jstor.org/stable/2138327>

How Do You Like Your Equilibrium Selection Problems? Hard, or Very Hard?

Paul W. Goldberg*

Dept. of Computer Science
University of Liverpool,
Ashton Street, Liverpool L69 3BX, UK
P.W.Goldberg@liverpool.ac.uk

Abstract. The PPAD-completeness of Nash equilibrium computation is taken as evidence that the problem is computationally hard in the worst case. This evidence is necessarily rather weak, in the sense that PPAD is only known to lie “between P and NP”, and there is not a strong prospect of showing it to be as hard as NP. Of course, the problem of finding an equilibrium that has certain sought-after properties should be at least as hard as finding an unrestricted one, thus we have for example the NP-hardness of finding equilibria that are socially optimal (or indeed that have various efficiently checkable properties), the results of Gilboa and Zemel [6], and Conitzer and Sandholm [3]. In the talk I will give an overview of this topic, and a summary of recent progress showing that the equilibria that are found by the Lemke-Howson algorithm, as well as related homotopy methods, are PSPACE-complete to compute. Thus we show that there are no short cuts to the Lemke-Howson solutions, subject only to the hardness of PSPACE. I mention some open problems.

1 Overview

There are two ways to view any algorithm for computing Nash equilibria. First, simply as a way to find a Nash equilibrium, one that is hopefully fast in practice, even though it may take exponential time in the worst case. Second, as a criterion for *equilibrium selection*, i.e. choosing some equilibrium that is considered to be preferable to others, in some sense a more plausible outcome. The latter viewpoint is especially relevant if the algorithm in question is somehow simple or decentralized. We can consider the problem of computing an equilibrium that is found by some specified algorithm, noting that we are not restricted to using that particular algorithm in order to find the equilibrium.

By way of example, in [7] we analyzed the classical Lemke-Howson algorithm for bimatrix games, in this context. The computational challenge is: Given a bimatrix game, find one of the solutions that could be computed using Lemke-Howson. Note that we are not asking about the complexity of the algorithm itself, which is already well-known to take exponential time in the worst case [9].

* Currently supported by EPSRC Grant EP/G069239/1.

Just consider the complexity of finding the Lemke-Howson solutions. Of course, the problem is PPAD-hard, simply due to the PPAD-completeness of finding any equilibrium, but in fact we show that restricting to the Lemke-Howson solutions makes the problem PSPACE-complete, and thus in a sense even harder to compute than the restricted equilibria of [6,3].

Homotopy methods. The survey paper [8] discusses homotopy methods in detail, including the Lemke-Howson algorithm, which falls within this framework. The general idea is that in trying to solve a game \mathcal{G} , start by constructing a “starting game” \mathcal{G}_0 which is a version of \mathcal{G} where the numerical payoffs have been changed so that there is some “obvious” Nash equilibrium. Then consider a continuum of games that lie between \mathcal{G}_0 and \mathcal{G} . (The usual choice of a continuum of intermediate games has games whose payoffs are weighted averages of those in \mathcal{G}_0 and \mathcal{G} .) Within these games, there exists a continuous path of Nash equilibria that starts at the one for \mathcal{G}_0 and ends at an equilibrium of \mathcal{G} . Thus, we have specified a unique equilibrium of \mathcal{G} , and implicitly a natural path-following algorithm for finding it.

As an equilibrium selection theory, homotopy methods are attractive since the starting game \mathcal{G}_0 can be considered as representing some kind of “prior belief” about the behaviour of the other player(s). However, we showed in [7] that the equilibrium identified by this procedure is PSPACE-complete to compute, and moreover, the result extends to Lemke-Howson, in which the choice of initially dropped label corresponds to a particular choice of \mathcal{G}_0 .

Path-following algorithms. The PSPACE-completeness result for Lemke-Howson solutions uses, quite intensively, the ideas developed in [4]. The problem END OF THE LINE that is used to characterise the complexity class PPAD, has a PSPACE-complete version in which you are required to compute the end-of-line obtained by following the path that begins at the known starting vertex of the graph. (Notice that this version is no longer (apparently) in NP since there is no obvious efficient test that a solution is correct.) This problem, called in [7] OTHER END OF THIS LINE, or OEOTL for short, is the one we reduce from. The proof proceeds by showing that a homotopy method —itself a path-following algorithm— captures the path-following approach to solving OEOTL.

The extension to Lemke-Howson requires us to design the game in such a way that all alternative solutions that might be produced, share features that efficiently encode a solution to generic instances of OEOTL.

Fictitious play. We conclude with an interesting open problem that is inspired by the PSPACE-completeness results described above. Fictitious Play is a good algorithm to consider in the capacity of equilibrium selection. This simple and intuitive procedure (see e.g. [5] or indeed Wikipedia) is known to converge to Nash equilibrium under certain sufficient conditions, although for some games it fails to converge. See [2] for a detailed discussion of why it is a natural and appealing algorithm to consider. The following problem looks natural and simple to state:

Input. A bimatrix game \mathcal{G}

Question. Does Fictitious Play converge, when applied to \mathcal{G} ? Assume both players start at their first strategies.

Further question. Compute the equilibrium, if indeed FP converges.

The results of [1] indicate that a naive simulation of FP requires exponential time, but it does not rule out the possibility of “short cuts” alluded to in the abstract to this talk, in the context of Lemke-Howson. On the other hand, I know of no upper bound for the above questions, not even an exponential one.

References

1. Brandt, F., Fischer, F., Harrenstein, P.: On the Rate of Convergence of Fictitious Play. In: Kontogiannis, S., Koutsoupias, E., Spirakis, P. (eds.) SAGT 2010. LNCS, vol. 6386, pp. 103–114. Springer, Heidelberg (2010)
2. Conitzer, V.: Approximation Guarantees for Fictitious Play. In: Proceedings of the 47th Annual Allerton Conference on Communication, Control and Computing (Allerton 2009), pp. 636–643 (2009)
3. Conitzer, V., Sandholm, T.: Complexity results about Nash equilibria. In: 18th International Joint Conference on Artificial Intelligence, IJCAI (2003)
4. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The Complexity of Computing a Nash Equilibrium. *SIAM Journal on Computing* 39(1), 195–259 (2009)
5. Fudenberg, D., Levine, D.K.: *The Theory of Learning in Games*. MIT Press, Cambridge (1998)
6. Gilboa, I., Zemel, E.: Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior* 1, 80–93 (1989)
7. Goldberg, P.W., Papadimitriou, C.H., Savani, R.: The Complexity of the Homotopy Method, Equilibrium Selection, and Lemke-Howson Solutions, Arxiv technical report 1006.5352 (2010)
8. Herings, P.J.-J., Peeters, R.: Homotopy methods to compute equilibria in game theory. *Economic Theory* 42(1), 119–156 (2010)
9. Savani, R., von Stengel, B.: Hard-to-Solve Bimatrix Games. *Econometrica* 74(2), 397–429 (2006)

A Simplex-Like Algorithm for Fisher Markets

Bharat Adsul¹, Ch. Sobhan Babu², Jugal Garg¹,
Ruta Mehta¹, and Milind Sohoni¹

¹ Indian Institute of Technology, Bombay
{adsul, jugal, ruta, sohoni}@cse.iitb.ac.in
² Indian Institute of Technology, Hyderabad
sobhan@iith.ac.in

Abstract. We propose a new convex optimization formulation for the Fisher market problem with linear utilities. Like the Eisenberg-Gale formulation, the set of feasible points is a polyhedral convex set while the cost function is non-linear; however, unlike that, the optimum is always attained at a vertex of this polytope. The convex cost function depends only on the initial endowments of the buyers. This formulation yields an easy simplex-like pivoting algorithm which is provably strongly polynomial for many special cases.

1 Introduction

Fisher and Arrow-Debreu market models are the two fundamental market models in mathematical economics. In this paper, we focus on the Fisher market model with linear utilities. An instance of this model consists of a set of buyers, a set of divisible goods, initial endowments, also referred to as the money owned by the buyers, quantities of the goods and (linear) utility functions of the buyers. The problem is to determine market equilibrium prices and allocation of the goods to buyers such that the market clears and the utility function for each buyer is maximized. Towards this, Eisenberg and Gale [6,10] formulated a remarkable convex optimization program whose optimal solution, more precisely, values of the primal and dual variables at an optimal solution, captures equilibrium allocation and prices.

Recently, many algorithmic results [4,5,9,11] pertaining to the computation of market equilibrium prices and allocation for the linear case of Fisher and Arrow-Debreu market models have been obtained. In [4], Deng et al. gave a strongly polynomial time algorithm for the Fisher market with either constant number of goods or constant number of buyers. Building on the Eisenberg-Gale program, Devanur et al. [5] developed a primal-dual type first polynomial time algorithm to solve the Fisher market model. A polynomial time algorithm for the more general Arrow-Debreu market is also presented in [9]. More recently, a strongly polynomial time algorithm for the Fisher market was given by Orlin [11]. A tantalizing open question is to formulate a linear program that captures the Fisher solution. A positive resolution of this question would, of course, imply a simplex-like algorithm for computing the same. This paper is an attempt towards this objective.

In this paper, we propose a novel convex optimization formulation for the Fisher market problem. In the Eisenberg-Gale formulation [6,10], the set of feasible points is a convex polytope which merely models the packing constraints and is oblivious to the parameters of the problem. Like the Eisenberg-Gale formulation, the set of feasible points in our formulation is also a convex polytope. However, unlike that, our convex polytope is defined in terms of the input parameters, specifically utilities and money, and is rich enough so as to ensure that the optimum is always attained at a vertex of this polytope. Furthermore, the convex cost function in our formulation depends only on the initial endowments of the buyers. There is another formulation, which maximizes a convex function under flow constraints, obtained by Shmyrev [12] and Birnbaum et al. [2], however this formulation also does not guarantee the optimum to be at a vertex.

We define *special vertices* in our polytope and every such vertex corresponds to the Fisher solution with a different endowment vector. We give a combinatorial characterization of special vertices and show that starting from any special vertex, there is a simplex-like path of special vertices where the cost function monotonically increases and it ends at a vertex corresponding to the Fisher solution. There may be many such paths of special vertices in the polytope. Using a simple pivoting rule, we give an algorithm, which traces one such path and show that this algorithm is strongly polynomial for many special cases. Two interesting cases are:

- Either the number of buyers or the goods is fixed.
- All the non-zero utilities are of the type α^k , where $\alpha > 0$ and $0 \leq k \leq M$ (M is bounded by a polynomial in the number of buyers and goods).

This algorithm is conceptually simple, much easier to implement and runs very fast in practice. In fact, these special cases seem sufficient to handle most practical situations. This is because, firstly, in practice, utilities are hardly exactly known, and secondly, as shown in [1] buyers have every reason to strategize and report fictitious utilities. The events that may occur in the algorithm, while finding the adjacent special vertex, are similar as in the DPSV algorithm [5], however one crucial difference is that the prices, DPSV algorithm computes at intermediate stages, may not occur at a vertex in the polytope. The DPSV algorithm may be interpreted as an interior point method in our formulation. Further, the utility of our formulation is also illustrated by its easy extension to incorporate transportation costs as well [8]. There seems no way to modify Eisenberg-Gale or Shmyrev formulations to capture the equilibrium solution for this extended model. Independently, Chakrabarty et al. [3] also give a similar formulation for this extended model along with an algorithm to compute ϵ -approximate equilibrium prices and allocations. However, the Fisher market with transportation cost may have irrational solutions, so the optimum solution may not be at a vertex.

Organization. The rest of the paper is organized as follows. In Section 2, we give a precise formulation of the Fisher market problem and introduce the new convex optimization program and analyze it. In Section 3, we discuss the simplex-like algorithm. In Section 4, we show that the algorithm is provably strongly

polynomial for many special cases. In Section 5, we summarize the number of pivoting steps taken by the algorithm on random instances of the Fisher market. Finally we conclude in Section 6.

2 New Convex Optimization Formulation

We begin with a precise description of the Fisher market model.

2.1 Problem Formulation

The input to the Fisher market problem is a set of buyers \mathcal{B} , a set of goods \mathcal{G} , a utility matrix $U = [u_{ij}]_{i \in \mathcal{B}, j \in \mathcal{G}}$, a quantity vector $\mathbf{q} = (q_j)_{j \in \mathcal{G}}$ and a money vector $\mathbf{m} = (m_i)_{i \in \mathcal{B}}$, where u_{ij} is the utility derived by buyer i from a unit amount of good j , q_j is the quantity of good j , and m_i is the money possessed by buyer i . Let $|\mathcal{B}| = m$ and $|\mathcal{G}| = n$. We assume that for every good j , there is a buyer i such that $u_{ij} > 0$ and for every buyer i , there is a good j such that $u_{ij} > 0$, otherwise we may discard those goods and buyers from the market.

The problem is to compute equilibrium prices $\mathbf{p} = [p_j]_{j \in \mathcal{G}}$ and allocations $X = [x_{ij}]_{i \in \mathcal{B}, j \in \mathcal{G}}$ such that they satisfy the following two constraints:

- **Market Clearing:** The demand equals the supply of each good, *i.e.*, $\forall j \in \mathcal{G}$, $\sum_{i \in \mathcal{B}} x_{ij} = q_j$ and $\forall i \in \mathcal{B}$, $\sum_{j \in \mathcal{G}} p_j x_{ij} = m_i$.
- **Optimal Goods:** Every buyer buys only those goods, which give her the maximum utility per unit of money (*bang per buck*), *i.e.*, if $x_{ij} > 0$ then $\frac{u_{ij}}{p_j} = \max_{k \in \mathcal{G}} \frac{u_{ik}}{p_k}$.

Note that, by scaling u_{ij} 's appropriately, we may assume that q_j 's are unit.

2.2 Convex Program

In this section, we introduce the new convex optimization program whose optimal solution captures the Fisher market equilibrium. Our convex program is described in Table 1, where p_j corresponds to the price of good j and z_{ij} corresponds to the money spent by buyer i on good j . At optimum, $\frac{1}{y_i}$ is the *bang per buck* of buyer i . We refer to the ambient space as the y - p - z -space.

Note that the feasible set O is a convex polytope in y - p - z -space and the cost function is independent of the variables z_{ij} . Let O_{aux} be the auxiliary polytope in the y - p -space defined by the constraints 1 to 4 and the related convex program (with the same cost function) be the auxiliary convex program.

Claim. $Pr(O) = O_{aux}$, where $Pr(O)$ is the projection of O onto the y - p -space.

Proof. Clearly, $Pr(O) \subseteq O_{aux}$, and for $O_{aux} \subseteq Pr(O)$, $Z = [z_{ij}]$ should be constructed for a given $(\mathbf{y}, \mathbf{p}) \in O_{aux}$. One way to do this is by constructing a max-flow network, where there is an edge from the source to every good $j \in \mathcal{G}$ with capacity p_j and from every buyer $i \in \mathcal{B}$ to the sink with capacity m_i . Further, there is an edge from every good $j \in \mathcal{G}$ to every buyer $i \in \mathcal{B}$ with ∞ capacity. Clearly, the max-flow gives the required z_{ij} 's. \square

Table 1. New Convex Program

$$\begin{aligned}
& \text{maximize} && \sum_{i \in \mathcal{B}} m_i \log y_i \\
& \text{subject to} && \\
\forall i \in \mathcal{B}, \forall j \in \mathcal{G} & : && u_{ij} y_i \leq p_j && (1) \\
& && \sum_{j \in \mathcal{G}} p_j \leq \sum_{i \in \mathcal{B}} m_i && (2) \\
\forall i \in \mathcal{B} & : && y_i \geq 0 && (3) \\
\forall j \in \mathcal{G} & : && p_j \geq 0 && (4) \\
\forall i \in \mathcal{B} & : && \sum_{j \in \mathcal{G}} z_{ij} \leq m_i && (5) \\
\forall j \in \mathcal{G} & : && \sum_{i \in \mathcal{B}} z_{ij} = p_j && (6) \\
\forall i \in \mathcal{B}, \forall j \in \mathcal{G} & : && z_{ij} \geq 0 && (7)
\end{aligned}$$

Therefore, in order to understand the optimality conditions, we may as well work with the KKT conditions for the auxiliary convex program. Let $x_{ij}, q, \mu_i, \lambda_j$ be the Lagrangian (dual) variables corresponding to the equations (1-4). An optimal solution must satisfy the KKT conditions in Table 2. \square

Table 2. KKT conditions

$$\begin{aligned}
\forall i \in \mathcal{B} & : && \frac{m_i}{y_i} = \sum_{j \in \mathcal{G}} u_{ij} x_{ij} - \mu_i && (8) \\
\forall i \in \mathcal{B}, \forall j \in \mathcal{G} & : && (u_{ij} y_i - p_j) x_{ij} = 0 && (9) \\
\forall j \in \mathcal{G} & : && -\sum_{i \in \mathcal{B}} x_{ij} - \lambda_j + q = 0 && (10) \\
& && (\sum_{j \in \mathcal{G}} p_j - \sum_{i \in \mathcal{B}} m_i) q = 0 && (11) \\
\forall i \in \mathcal{B}, \forall j \in \mathcal{G} & : && x_{ij}, \lambda_j, \mu_i, q \geq 0 && (12) \\
\forall j \in \mathcal{G} & : && -p_j \lambda_j = 0 && (13) \\
\forall i \in \mathcal{B} & : && -y_i \mu_i = 0 && (14)
\end{aligned}$$

Claim. At any optimum, $\mu_i = 0, \forall i \in \mathcal{B}$ and $\lambda_j = 0, \forall j \in \mathcal{G}$.

Proof. $\mu_i \neq 0 \Rightarrow y_i = 0 \Rightarrow$ the optimal solution has cost $-\infty$. However, we may easily construct a feasible point in the polytope, where the cost is some real value, therefore all μ_i 's are zero. Similarly, $\lambda_j \neq 0 \Rightarrow p_j = 0 \Rightarrow y_i = 0$, for some $i \in \mathcal{B}$. Hence, all λ_j 's are zero. \square

Putting $\mu_i = 0$ and $\lambda_j = 0$ in the KKT conditions (8-12), we get,

$$\forall i \in \mathcal{B} \quad : \quad m_i = \sum_{j \in \mathcal{G}} u_{ij} x_{ij} y_i \quad (15)$$

$$\forall i \in \mathcal{B}, \forall j \in \mathcal{G} \quad : \quad (u_{ij} y_i - p_j) x_{ij} = 0 \quad (16)$$

$$\forall j \in \mathcal{G} \quad : \quad \sum_{i \in \mathcal{B}} x_{ij} = q \quad (17)$$

$$\quad : \quad \left(\sum_{j \in \mathcal{G}} p_j - \sum_{i \in \mathcal{B}} m_i \right) q = 0 \quad (18)$$

$$\forall i \in \mathcal{B}, \forall j \in \mathcal{G} \quad : \quad x_{ij}, q \geq 0 \quad (19)$$

From (15-18), $\sum_{i \in \mathcal{B}} m_i = \sum_{i \in \mathcal{B}} \sum_{j \in \mathcal{G}} p_j x_{ij} = \sum_{j \in \mathcal{G}} \sum_{i \in \mathcal{B}} p_j x_{ij} = \sum_{j \in \mathcal{G}} p_j q \Rightarrow q = 1$

Proposition 1. *Let $(\mathbf{y}, \mathbf{p}) \in O_{aux}$ be an optimal solution to the auxiliary convex program. Then \mathbf{p} is a market equilibrium price.*

Proof. As $q = 1$, interpreting $X = [x_{ij}]$ as an allocation, we see that conditions (15-17) imply that the market clearing constraint holds at the price vector \mathbf{p} . Further, using condition 2, we have $x_{ij} > 0 \Rightarrow y_i u_{ij} = p_j$. As $(\mathbf{y}, \mathbf{p}) \in O_{aux}$, we also have, $\forall i \in \mathcal{B}, \forall j \in \mathcal{G} : u_{ij} y_i \leq p_j$. Putting these two together, it is easily verified that the optimal goods constraint is also satisfied. \square

Proposition 2.

- (i) *The auxiliary convex program admits a unique optimal solution.*
- (ii) *Equilibrium prices are unique and allocations form a polyhedral set.*

Proof. Part (i) follows from the fact that the cost function is strictly concave, and part (ii) follows from the KKT conditions. \square

Let $(\mathbf{y}, \mathbf{p}) \in O_{aux}$ be the unique optimum solution to the auxiliary convex program. Let $\mathbb{X} = \{X = [x_{ij}]_{i \in \mathcal{B}, j \in \mathcal{G}} \mid (\mathbf{y}, \mathbf{p}, X) \text{ satisfies (8-14)}\}$. Note that \mathbb{X} is a convex set. As argued in the proof of Proposition 1, we may think of $X \in \mathbb{X}$ as an equilibrium allocation and \mathbf{p} as the equilibrium price. Now, we define $Z = [z_{ij}]$ w.r.t. $X \in \mathbb{X}$ as $z_{ij} = x_{ij} p_j$, $\forall i \in \mathcal{B}, \forall j \in \mathcal{G}$. In other words, z_{ij} is the money spent by buyer i on good j at the equilibrium allocation X . We refer to Z as an equilibrium money allocation. It easily follows that $(\mathbf{y}, \mathbf{p}, Z)$ is an optimum solution to the main convex program. Note that there is an $X^a \in \mathbb{X}$ such that the bipartite graph $G = (\mathcal{B}, \mathcal{G}, E)$, where $E = \{(i, j) \in \mathcal{B} \times \mathcal{G} \mid x_{ij}^a > 0\}$, is acyclic. Let Z^a be the equilibrium money allocation w.r.t. X^a . The next proposition asserts that $(\mathbf{y}, \mathbf{p}, Z^a)$ is in fact a vertex of O .

Proposition 3. *The point $(\mathbf{y}, \mathbf{p}, Z^a)$ is a vertex of O .*

Proof. There are $mn + m + n$ variables in the convex program, and we show that there are $mn + m + n$ linearly independent tight constraints at $(\mathbf{y}, \mathbf{p}, Z^a)$ (refer to Theorem 3.3.3 in [7] for details).

Remark 4. The auxiliary program itself captures the equilibrium prices at the optimal solution, though not necessarily at one of its vertices. [7] has the detailed analysis of both the polytopes.

3 A Simplex-Like Algorithm

We begin with some notation. Henceforth, we denote the input to the Fisher market problem by (U, \mathbf{m}) . The set of buyers and the set of goods are implicit. We use g_j and b_i to denote the good j and buyer i respectively. For convenience, we assume that $u_{ij} > 0, \forall i \in \mathcal{B}, \forall j \in \mathcal{G}$.

Now, we turn our attention to the polytope O defined in the previous section. We have shown that there exists a vertex $v = (\mathbf{y}, \mathbf{p}, Z)$ of the polytope O which captures the equilibrium prices and an equilibrium money allocation. An important property of v is that $\forall i \in \mathcal{B}, \forall j \in \mathcal{G}, z_{ij}(u_{ij}y_i - p_j) = 0$. In other words, every buyer spends money only on her optimal goods.

Definition 5. A vertex $v = (\mathbf{y}, \mathbf{p}, Z)$ of O is called **special** if $z_{ij}(u_{ij}y_i - p_j) = 0, \forall i \in \mathcal{B}, \forall j \in \mathcal{G}$.

It is easy to see that if $v = (\mathbf{y}, \mathbf{p}, Z)$ is a special vertex, then it corresponds to a solution for an instance of the Fisher market problem. Namely, let $\mathcal{B}' = \{i \in \mathcal{B} \mid y_i \neq 0\}, \mathcal{G}' = \mathcal{G}$ and U' be U restricted to $\mathcal{B}' \times \mathcal{G}'$. Further, for $i \in \mathcal{B}'$, let $m'_i = \sum_{j \in \mathcal{G}} z_{ij}$. Clearly, v corresponds to a solution of (U', \mathbf{m}') .

3.1 Characterization of Special Vertices

Let $v = (\mathbf{y}, \mathbf{p}, Z)$ be a special vertex of O . W.l.o.g., we may assume that all y_i 's and all p_j 's are non-zero at v , because if $p_j = 0$ for some $j \in \mathcal{G}$ at v , then v is a trivial point, *i.e.*, all coordinates are zero, and if $y_k = 0$ for some $k \in \mathcal{B}$ at v , then there is an adjacent vertex $v' = (\mathbf{y}', \mathbf{p}', Z')$ to v , where $\mathbf{p}' = \mathbf{p}, Z' = Z, y'_i = y_i, \forall i \neq k$, and $y'_k = \min_{j \in \mathcal{G}} \frac{p_j}{u_{kj}}$.

Now we describe a combinatorial characterization of v . Towards this, we define $E(v)$ and $F(v)$ as follows:

$$E(v) = \{(i, j) \in \mathcal{B} \times \mathcal{G} \mid u_{ij}y_i = p_j\} \quad \text{and} \quad F(v) = \{(i, j) \in \mathcal{B} \times \mathcal{G} \mid z_{ij} > 0\}$$

The elements in $E(v)$ and $F(v)$ are called *tight* and *non-zero* edges respectively. By definition, $F(v) \subseteq E(v)$. Let $G(E(v), F(v))$ be the graph, whose vertices are the connected components C_1, C_2, \dots of the bipartite graph $(\mathcal{B}, \mathcal{G}, F(v))$, and there is an edge between C_i and C_j in $G(E(v), F(v))$, if there is at least one edge in $E(v) - F(v)$ between the corresponding components of $(\mathcal{B}, \mathcal{G}, F(v))$.

We say that buyer i belongs to a vertex C of $G(E(v), F(v))$, if buyer i lies in the corresponding component of $(\mathcal{B}, \mathcal{G}, F(v))$. We call a connected component of G as simply a component of G .

Definition 6. W.r.t. $v = (\mathbf{y}, \mathbf{p}, Z)$,

- **surplus** of buyer i is defined to be the non-negative value $m_i - \sum_{j \in \mathcal{G}} z_{ij}$.
- a buyer is called a **zero surplus** buyer if its surplus is zero, otherwise it is called a **positive surplus** buyer.
- a component of $(\mathcal{B}, \mathcal{G}, F(v))$ is called **saturated** if all buyers in that component are zero surplus buyers, otherwise it is called **unsaturated**.

- a vertex of $G(E(v), F(v))$ is called **saturated** if the corresponding component of $(\mathcal{B}, \mathcal{G}, F(v))$ is saturated, otherwise it is called **unsaturated**.

Theorem 7. v has following properties:

- Every component of $(\mathcal{B}, \mathcal{G}, F(v))$ contains at most one positive surplus buyer.
- Every component of $G(E(v), F(v))$ has at least one saturated vertex.

Proof. If a component of $(\mathcal{B}, \mathcal{G}, F(v))$ contains more than one positive surplus buyers, then the z_{ij} 's in that component may be modified such that the same set of inequalities are tight before and after the modification, *i.e.*, v is not a vertex.

Similarly, if a component of $G(E(v), F(v))$ does not have a saturated vertex, then the p_j 's in that component may be scaled uniformly such that the same set of inequalities are tight before and after the scaling, hence a contradiction. \square

Corollary 8. If (U, \mathbf{m}) are algebraically independent, then

- the bipartite graph $(\mathcal{B}, \mathcal{G}, E(v))$ is a forest. Hence there is at most one edge in $E(v) - F(v)$ between any two components of $(\mathcal{B}, \mathcal{G}, F(v))$.
- every component of $G(E(v), F(v))$ has exactly one saturated vertex.

Lemma 9. Let v be a special vertex of O . Then

- (i) $(\mathcal{B}, \mathcal{G}, F(v))$ is acyclic.
- (ii) If (U, \mathbf{m}) are algebraically independent, then $(\mathcal{B}, \mathcal{G}, E(v))$ is acyclic and the number of positive surplus buyers is $|E(v) - F(v)|$.

Proof. Since v is a vertex of O , therefore $(\mathcal{B}, \mathcal{G}, F(v))$ is acyclic. Part (ii) follows from Theorem 7 and Corollary 8. \square

3.2 Algorithm

In general, a simplex-like pivoting algorithm moves from a vertex to an adjacent vertex such that the cost function increases. Therefore, we first describe the AdjacentVertex procedure for the main convex program.

We assume that (U, \mathbf{m}) are algebraically independent¹. The AdjacentVertex procedure, given in Table 3, takes a special vertex v and outputs another special vertex v' adjacent to v , such that the cost function increases. If v is optimum, then it outputs $v' = v$. Otherwise, there is a component C of $G(E(v), F(v))$ containing an unsaturated vertex. Clearly C is a tree and there is exactly one saturated vertex, say C_s , in C (Corollary 8). We consider C as the rooted tree with root C_s . We pick an edge e between C_s and an unsaturated vertex, say C_u , in C . Let (b_i, g_j) be the edge in $E(v) - F(v)$ corresponding to e . There are two cases depending on where b_i belongs: C_s (Case 1) or C_u (Case 2).

Case 1: We get a new vertex v' , adjacent to v in O , by relaxing the inequality $u_{ij}y_i \leq p_j$, which is tight at v . Let T_u be the subtree of C rooted at C_u and J_u be the set of goods in the components of $(\mathcal{B}, \mathcal{G}, F(v))$ corresponding to the vertices of T_u . v' may also be obtained by increasing the prices of the goods in

¹ For the general (U, \mathbf{m}) , AdjacentVertex may be easily modified.

Table 3. AdjacentVertex Procedure

<p>AdjacentVertex(v)</p> $v' \leftarrow v$; if v is optimum then return v' ; endif $C \leftarrow$ component of $G(E(v), F(v))$ containing an unsaturated vertex; $C_s \leftarrow$ saturated vertex in C ; $C_u \leftarrow$ unsaturated vertex, adjacent to C_s , in C ; $e \leftarrow$ edge between C_s and C_u ; $(b_i, g_j) \leftarrow$ edge in $E(v) - F(v)$ corresponding to e ; if (b_i, g_j) is from C_s to C_u then $v' \leftarrow$ adjacent vertex obtained by relaxing $u_{ij}y_i \leq p_j$; else $v' \leftarrow$ adjacent vertex obtained by relaxing $z_{ij} \geq 0$; endif return v' ;

Table 4. Different cases for the new tight inequality

-
1. A non-zero edge (b_k, g_l) becomes zero, *i.e.*, $z_{kl} \geq 0$ becomes tight.
 2. A non-tight edge (b_k, g_l) becomes tight, *i.e.*, $u_{kl}y_k \leq p_l$ becomes tight.
 3. An unsaturated vertex in C becomes saturated, *i.e.*, $\sum_{l \in \mathcal{G}} z_{kl} \leq m_k$ becomes tight, where buyer k is a positive surplus buyer w.r.t. v .
-

J_u uniformly and by modifying y_i 's and z_{ij} 's accordingly till a new inequality becomes tight. Table 4 lists the three possible cases for the new inequality.

Case 2: We get a new vertex v' , adjacent to v in O , by relaxing the inequality $z_{ij} \geq 0$, which is tight at v . Let J be the set of goods in the components of $(\mathcal{B}, \mathcal{G}, F(v))$ corresponding to the vertices of C . v' may also be obtained by increasing the prices of the goods in J uniformly and by modifying the y_i 's and z_{ij} 's accordingly till a new inequality becomes tight. Table 4 lists the three possible cases for the new inequality.

Both the cases result in the new vertex v' adjacent to v in O , where \mathbf{p} as well as \mathbf{y} increase monotonically and $\sum_{j \in \mathcal{G}} p_j$ as well as $\sum_{i \in \mathcal{B}} y_i$ increase strictly going from v to v' . Hence the cost function value increases strictly going from v to v' . Note that v' is also a special vertex of O .

From the above discussion, the following lemma is straightforward.

Lemma 10. *If a special vertex v is not optimum, then there exists an adjacent special vertex v' such that the value of cost function is more at v' than v .*

There may be many simplex-like paths in O to reach at the optimum vertex using different pivoting rules. Algorithm 1 traces a particular simplex-like path in O , where the pivoting rule is such that there is at most one buyer with a positive surplus at every vertex on the path. In this algorithm, we do not consider the components, which contain only a single buyer.

Algorithm 1. A Simplex-like Pivoting Algorithm

```

 $U' \leftarrow \langle u_{11}, \dots, u_{1n} \rangle; \mathbf{m}' \leftarrow \langle m_1 \rangle;$ 
 $v \leftarrow$  special vertex corresponds to the solution of  $(U', \mathbf{m}')$ ;
 $i \leftarrow 2;$ 
while  $i \leq m$  do
  /* Note that the inequality  $y_i \geq 0$  is tight at  $v^*$ /
   $v \leftarrow$  vertex adjacent to  $v$  obtained by relaxing  $y_i \geq 0$ ;
  while surplus of buyer  $i$  w.r.t.  $v$  is non-zero do
     $v \leftarrow \text{AdjacentVertex}(v);$ 
  endwhile
   $i \leftarrow i + 1;$ 
endwhile

```

There are two types of iterations of the inner while loop, one in which we relax the inequality $z_{kl} \geq 0$ (Type 1) and the other in which we relax the inequality $u_{klyk} \leq p_l$ (Type 2) for some (b_k, g_l) .

Remark 11. Algorithm 1 provides a polyhedral interpretation to a sequential run of the so called *Basic Algorithm* in [5], where buyers are added one at a time.

Lemma 12. *Algorithm 1 takes at most $(m + n * 2^{m+n})$ iterations.*

Proof. Consider the iterations of Type 2 of the inner while loop, where we relax the tight inequality $u_{klyk} \leq p_l$ for some (b_k, g_l) . Let C_s^j be the component containing buyer k in the j^{th} such iteration. Note that C_s^j is a saturated component. Let B^j be the set of buyers and G^j be the set of goods in C_s^j , and $S^j = B^j \cup G^j$. Since prices monotonically increase, therefore all S^j 's are distinct. The total number of distinct S^j 's are clearly bounded by 2^{m+n} , and in every n iterations of inner while loop, one iteration has to be of Type 2, therefore the number of iterations of the algorithm is bounded by $(m + n * 2^{m+n})$. \square

Remark 13. A more refined bound is 2^{m+n+1} .

4 Analysis

In this section, we describe the main idea of Algorithm 1 and show that it is strongly polynomial for many special cases.

Main Idea of Algorithm 1. Consider the inner while loop for buyer i and let v be the current special vertex. The component C of $G(E(v), F(v))$ containing buyer i has exactly two vertices, one saturated (C_s) and one unsaturated (C_u), and an edge (b_k, g_l) between them. Note that buyer i belongs to C_u and $z_{kl} = 0$. Now, consider the tree T in $(\mathcal{B}, \mathcal{G}, E(v))$ rooted at buyer i . The edges are directed downwards, *i.e.*, away from the root. We increase the prices of the goods uniformly in T in order to decrease the surplus of buyer i . This increases the flow on the edges, which are from a buyer to a good (forward) and decreases the flow on the edges, which are from a good to a buyer (backward).

Therefore, when (b_k, g_l) be such that $g_l \in C_u$ and $b_k \in C_s$, we need to relax $u_{kl}y_k \leq p_l$, and when $g_l \in C_s$ and $b_k \in C_u$, we need to relax $z_{kl} \geq 0$ in order to increase the prices. It is also clear that during the price increase, only backward edges may be deleted. Moreover, since the prices of the goods in T increase, buyers in T may become interested in the goods outside T , and it implies that only forward edges may be added.

Theorem 14. *Algorithm 1 is strongly polynomial when either the number of buyers or goods is constant.*

Proof. W.l.o.g., we assume that (U, \mathbf{m}) are algebraically independent².

It is enough to show that the inner while loop takes a strongly polynomial number of iterations for every buyer i . Let C^j be the component of $G(E(v), F(v))$, which contains buyer i in the j^{th} iteration of the inner while loop for buyer i . If surplus of buyer i is not zero, then C^j contains exactly one saturated vertex, say C_s^j , and one unsaturated vertex, say C_u^j . Note that buyer i belongs to C_u^j .

Let (b_k, g_l) be the edge between C_u^j and C_s^j , and P_j be the path starting from buyer i and ending with the edge (b_k, g_l) in $(\mathcal{B}, \mathcal{G}, E(v))$.

Claim. All P_j 's are distinct.

Proof. Recall that when the edge (b_k, g_l) is such that buyer k belongs to C_s^j , we relax the inequality $u_{kl}y_k \leq p_l$, and when buyer k belongs to C_u^j , we relax the inequality $z_{kl} \geq 0$. In other words, we add the edge (b_k, g_l) when buyer k belongs to C_u^j and delete it when buyer k belongs to C_s^j .

We show that all P_j 's, which end in a good, are distinct, and a similar argument may be worked out for the case when they end in a buyer. A path P_j may repeat only when the last edge, say e , is deleted and added again, and this is possible only if some other edge more near to buyer i than e in P_j is deleted. The induction on the length of P_j proves the claim, because the edges between buyer i and the goods never break (buyer i always lies in C_u^j). \square

Since the length of any P_j is at most $2 * \min(m, n)$, therefore it is a constant when either m or n is constant. Hence the total number of distinct P_j 's are bounded by a polynomial in either m (if n is constant) or n (if m is constant). Hence the length of the simplex-like path in the Algorithm 1 is strongly polynomial when either the number of buyers or goods is constant. \square

Theorem 15. *Algorithm 1 is strongly polynomial when $\forall i \in \mathcal{B}, \forall j \in \mathcal{G}, u_{ij} = \alpha^{k_{ij}}$, where $0 \leq k_{ij} \leq \text{poly}(m, n)$ and $\alpha > 0$.*

Proof. We only need to show that for every buyer i , the inner while loop takes a strongly polynomial number of iterations. Consider the iterations of inner while loop for a buyer a . We monitor the values of $\frac{y_a}{p_b}, \forall b \in \mathcal{G}$. Note that $\frac{y_a}{p_b}$ for a good b remains same until both buyer a and good b are in the same component

² For the general (U, \mathbf{m}) , a similar proof may be worked out.

of $G(E(v), F(v))$, otherwise it strictly increases. Let C^j be the component of $G(E(v), F(v))$, which contains buyer a in the j^{th} iteration. If surplus of buyer a is not zero, then C^j contains exactly one saturated vertex, say C_s^j , and one unsaturated vertex, say C_u^j . Note that buyer a belongs to C_u^j .

Let (b_k, g_l) be the edge between C_u^j and C_s^j . There are two types of iterations, one in which we relax the inequality $z_{kl} \geq 0$ (Type 1) and the other in which we relax the inequality $u_{kl}y_k \leq p_l$ (Type 2). Let $z_{kl} \geq 0$ is relaxed in the j^{th} iteration, and $b_a, g_{j_1}, b_{i_1}, \dots, g_{j_k}, b_k, g_l$ be the path from b_a to g_l in C^j . Clearly, $\frac{y_a}{p_l} = \frac{u_{i_1 j_1} \dots u_{k j_k}}{u_{a j_1} \dots u_{i_{k-1} j_{k-1}} u_{kl}}$ (using the tight inequalities $u_{ij}y_i \leq p_j$), and the value of $\frac{y_a}{p_l}$ strictly increases when iteration of Type 1 occurs. Now, we consider the values of $\log_\alpha \frac{y_a}{p_j}$, $\forall j \in \mathcal{G}$. Clearly, these values monotonically increase when an iteration of Type 1 occurs. Since for every $j \in \mathcal{G}$, the value of $\log_\alpha \frac{y_a}{p_j}$ might be at most $n * \text{poly}(m, n)$, therefore for every buyer i , the number of iterations of inner while loop is bounded by $n^2 * \text{poly}(m, n)$. \square

Theorem 15 may be easily generalized to handle the case when some u_{ij} 's are zero. Many easy cases like all utilities are 0/1, non-zero utilities form a tree etc. may also be easily shown to be strongly polynomial in Algorithm 1.

5 Experimental Results

In this section, we report the experimental results of Algorithm 1. We ran Algorithm 1 on random instances of the Fisher market (*i.e.*, (U, \mathbf{m}) are generated uniformly at random), while keeping the number of buyers and goods same (*i.e.*, $m = n$). For each value of $m \in \{4, 8, 12, 16, 20, 24\}$, we ran 100 experiments. Table 5 summarizes the results in terms of the minimum (best), maximum (worst) and mean (average) number of pivoting steps taken by Algorithm 1.

Table 5. Number of Pivoting Steps Taken by Algorithm 1

# buyers/goods	4	8	12	16	20	24
min	6	31	84	136	245	223
max	24	80	168	235	320	514
mean	12.5	50.9	113.1	186.9	279.8	408.9

Clearly, the number of steps seem to increase quadratically with the size of instances, and even the worst case instance for each value of m requires fewer than $2m^2$ steps. Therefore, Algorithm 1 should have a much better bound.

6 Conclusion

We have presented a novel convex optimization formulation for the Fisher market problem whose feasible set is a polytope and it is guaranteed that there is a vertex

of this polytope which is an optimal solution. Exploiting this, we have developed a simplex-like vertex-marching algorithm which runs in strongly polynomial time for many special cases.

We feel that the strongly polynomial algorithm by Orlin [11] is neither polytopal nor very intuitive. The algorithms, which are polytopal and simplex-like are generally easier to understand, simpler to implement using standard math libraries, and run faster in practice. Therefore, an obvious open problem is to give a strongly polynomial, simplex-like algorithm; even a polynomial bound will be interesting. Another open problem is to give a linear programming formulation that captures the equilibrium prices for the Fisher market. Therefore, it will be interesting to construct a linear cost function on our polytope so that optimum vertex gives the equilibrium prices.

References

1. Adsul, B., Babu, C.S., Garg, J., Mehta, R., Sohoni, M.: Nash equilibria in Fisher market. arXiv:1002.4832 (2010)
2. Birnbaum, B., Devanur, N.R., Xiao, L.: New convex programs and distributed algorithms for Fisher markets with linear and spending constraint utilities. Working Paper (2010)
3. Chakraborty, S., Devanur, N.R., Karande, C.: Market equilibrium with transaction costs. arXiv:1001.0393 (2010)
4. Deng, X., Papadimitriou, C.H., Safra, S.: On the complexity of equilibria. In: STOC 2002 (2002)
5. Devanur, N.R., Papadimitriou, C.H., Saberi, A., Vazirani, V.V.: Market equilibrium via a primal-dual type algorithm. JACM 55(5) (2008)
6. Eisenberg, E., Gale, D.: Consensus of subjective probabilities: The pari-mutuel method. Annals Math. Stat. 30, 165–168 (1959)
7. Garg, J.: Algorithms for market equilibrium. First APS Report (2008), <http://www.cse.iitb.ac.in/~jugal/aps1.pdf>
8. Garg, J.: Fisher markets with transportation costs. Technical Report (2009), <http://www.cse.iitb.ac.in/~jugal/techRep.pdf>
9. Jain, K.: A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. In: FOCS 2004 (2004)
10. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, Cambridge (2007)
11. Orlin, J.B.: Improved algorithms for computing Fisher’s market clearing prices. In: STOC 2010 (2010)
12. Shmyrev, V.I.: An algorithm for finding equilibrium in the linear exchange model with fixed budgets. Journal of Applied and Industrial Mathematics 3(4), 505–518 (2009)

Nash Equilibria in Fisher Market

Bharat Adsul¹, Ch. Sobhan Babu², Jugal Garg¹,
Ruta Mehta¹, and Milind Sohoni¹

¹ Indian Institute of Technology, Bombay
{adsul,jugal,ruta,sohoni}@cse.iitb.ac.in
² Indian Institute of Technology, Hyderabad
sobhan@iith.ac.in

Abstract. Much work has been done on the computation of market equilibria. However due to strategic play by buyers, it is not clear whether these are actually observed in the market. Motivated by the observation that a buyer may derive a better payoff by feigning a different utility function and thereby manipulating the Fisher market equilibrium, we formulate the *Fisher market game* in which buyers strategize by posing different utility functions. We show that existence of a *conflict-free allocation* is a necessary condition for the Nash equilibria (NE) and also sufficient for the symmetric NE in this game. There are many NE with very different payoffs, and the Fisher equilibrium payoff is captured at a symmetric NE. We provide a complete polyhedral characterization of all the NE for the two-buyer market game. Surprisingly, all the NE of this game turn out to be symmetric and the corresponding payoffs constitute a piecewise linear concave curve. We also study the correlated equilibria of this game and show that third-party mediation does not help to achieve a better payoff than NE payoffs.

1 Introduction

A fundamental market model was proposed by Walras in 1874 [21]. Independently, Fisher proposed a special case of this model in 1891 [3], where a market comprises of a set of buyers and divisible goods. The money possessed by buyers and the amount of each good is specified. The utility function of every buyer is also given. The market equilibrium problem is to compute prices and allocation such that every buyer gets the highest utility bundle subject to her budget constraint and that the market clears. Recently, much work has been done on the computation of market equilibrium prices and allocation for various utility functions, for example [6,7,11,15].

The payoff (*i.e.*, happiness) of a buyer depends on the equilibrium allocation and in turn on the utility functions and initial endowments of the buyers. A natural question to ask is, can a buyer achieve a better payoff by feigning a different utility function? It turns out that a buyer may indeed gain by feigning! This observation motivates us to analyze the strategic behavior of buyers in the Fisher market. We analyze here the linear utility case described below.

Let \mathcal{B} be the set of buyers, and \mathcal{G} be the set of goods, and $|\mathcal{B}| = m, |\mathcal{G}| = n$. Let m_i be the money possessed by buyer i , and q_j be the total quantity of good j in the market. The utility function of buyer i is represented by the non-negative utility tuple $\langle u_{i1}, \dots, u_{in} \rangle$, where u_{ij} is the payoff, she derives from a unit amount of good j . Thus, if x_{ij} is the amount of good j allocated to buyer i , then the *payoff* she derives from her allocation is $\sum_{j \in \mathcal{G}} u_{ij} x_{ij}$. Market equilibrium or market clearing prices (p_1, \dots, p_n) , where p_j is the price of good j , and equilibrium allocation $[x_{ij}]_{i \in \mathcal{B}, j \in \mathcal{G}}$ satisfy the following constraints:

- **Market Clearing:** The demand equals the supply of each good, *i.e.*, $\forall j \in \mathcal{G}, \sum_{i \in \mathcal{B}} x_{ij} = q_j$, and $\forall i \in \mathcal{B}, \sum_{j \in \mathcal{G}} p_j x_{ij} = m_i$.
- **Optimal Goods:** Every buyer buys only those goods, which give her the maximum utility per unit of money, *i.e.*, if $x_{ij} > 0$ then $\frac{u_{ij}}{p_j} = \max_{k \in \mathcal{G}} \frac{u_{ik}}{p_k}$.

In this market model, by scaling u_{ij} 's appropriately, we may assume that the quantity of every good is one unit, *i.e.*, $q_j = 1, \forall j \in \mathcal{G}$. Equilibrium prices are unique and the set of equilibrium allocations is a convex set [14]. The following example illustrates a small market.

Example 1. Consider a 2 buyers, 2 goods market with $m_1 = m_2 = 10, q_1 = q_2 = 1, \langle u_{11}, u_{12} \rangle = \langle 10, 3 \rangle$ and $\langle u_{21}, u_{22} \rangle = \langle 3, 10 \rangle$. The equilibrium prices of this market are $\langle p_1, p_2 \rangle = \langle 10, 10 \rangle$ and the unique equilibrium allocation is $\langle x_{11}, x_{12}, x_{21}, x_{22} \rangle = \langle 1, 0, 0, 1 \rangle$. The payoff of both the buyers is 10.

In the above market, does a buyer have a strategy to achieve a better payoff? Yes indeed, buyer 1 can force price change by posing a different utility tuple, and in turn gain. Suppose buyer 1 feigns her utility tuple as $\langle 5, 15 \rangle$ instead of $\langle 10, 3 \rangle$, then coincidentally, the equilibrium prices $\langle p_1, p_2 \rangle$ are also $\langle 5, 15 \rangle$. The unique equilibrium allocation $\langle x_{11}, x_{12}, x_{21}, x_{22} \rangle$ is $\langle 1, \frac{1}{3}, 0, \frac{2}{3} \rangle$. Now, the payoff of buyer 1 is $u_{11} * 1 + u_{12} * \frac{1}{3} = 11$, and that of buyer 2 is $u_{22} * \frac{2}{3} = \frac{20}{3}$. Note that the payoffs are still calculated w.r.t. the true utility tuples.

This clearly shows that a buyer could gain by feigning a different utility tuple, hence the Fisher market is susceptible to gaming by strategic buyers. Therefore, the equilibrium prices w.r.t. the true utility tuples may not be the actual operating point of the market. The natural questions to investigate are: What are the possible operating points of this market model under strategic behavior? Can they be computed? Is there a preferred one? This motivates us to study the Nash equilibria of the *Fisher market game*, where buyers are the players and strategies are the utility tuples that they may pose.

Related work. Shapley and Shubik [18] consider a market game for the exchange economy, where every good has a trading post, and the strategy of a buyer is to bid (money) at each trading post. For each strategy profile, the prices are determined naturally so that market clears and goods are allocated accordingly, however agents may not get the optimal bundles. Many variants [2, 8] of this game have been extensively studied. Essentially, the goal is to design a mechanism to implement Walrasian equilibrium (WE), *i.e.*, to capture WE at a

NE of the game. The strategy space of this game is tied to the implementation of the market (in this case, trading posts). Our strategy space is the utility tuple itself, and is independent of the market implementation. It is not clear that bids of a buyer in the Shapley-Shubik game correspond to the feigned utility tuples.

In word auction markets as well, a similar study on strategic behavior of buyers (advertisers) has been done [49,19].

Our contributions. We formulate the Fisher market game, the strategy sets and the corresponding payoff function in Section 2. Every (pure) strategy profile defines a Fisher market, and therefore market equilibrium prices and a set of equilibrium allocations. The payoff of a buyer may not be same across all equilibrium allocations w.r.t. a strategy profile, as illustrated by Example 2 in Section 2. Furthermore, there may not exist an equilibrium allocation, which gives the maximum possible payoffs to all the buyers. This behavior causes a conflict of interest among buyers. A strategy profile is said to be *conflict-free*, if there is an equilibrium allocation which gives the maximum possible payoffs to all the buyers.

A strategy profile is called a Nash equilibrium strategy profile (*NESP*), if no buyer can unilaterally deviate and get a better payoff. In Section 3, we show that all NESPs are conflict-free. Using the equilibrium prices, we associate a bipartite graph to a strategy profile and show that this graph must satisfy certain conditions when the corresponding strategy profile is a NESP.

Next, we define *symmetric* strategy profiles, where all buyers play the same strategy. We show that a symmetric strategy profile is a NESP iff it is conflict-free. It is interesting to note that a symmetric NESP can be constructed for a given market game, whose payoff is the same as the Fisher payoff, *i.e.*, payoff when all buyers play truthfully. Example 11 shows that all NESPs need not be symmetric and the payoff w.r.t. a NESP need not be Pareto optimal (*i.e.*, efficient). However, the Fisher payoff is always Pareto optimal (see First Theorem of Welfare Economics [20]).

Characterization of all the NESPs seems difficult; even for markets with only three buyers. We study two-buyer markets in Section 4 and the main results are:

- All NESPs are symmetric and they are a union of at most $2n$ convex sets.
- The set of NESP payoffs constitute a piecewise linear concave curve and all these payoffs are Pareto optimal. The strategizing on utilities has the same effect as differing initial endowments (see Second Theorem of Welfare Economics [20]).
- The third-party mediation does not help in this game.

Some interesting observations about two-buyer markets are:

- The buyer i gets the maximum payoff among all Nash equilibrium payoffs when she imitates the other, *i.e.*, when they play $(\mathbf{u}_{-i}, \mathbf{u}_{-i})$, where \mathbf{u}_{-i} is the true utility tuple of the other buyer.
- There may exist NESPs, whose social welfare (*i.e.*, sum of the payoffs of both the buyers) is larger than that of the Fisher payoff (Example 17).

- For a particular payoff tuple, there is a convex set of NESPs and hence convex set of equilibrium prices. This motivates a seller to offer incentives to the buyers to choose a particular NESP from this convex set, which fetches the maximum price for her good. Example 18 illustrates this behavior.

Most qualitative features of these markets may carry over to oligopolies, which arise in numerous scenarios. For example, relationship between a few manufacturers of aircrafts or automobiles and many suppliers. Finally, we conclude in Section 5 that it is highly unlikely that buyers will act according to their true utility tuples in Fisher markets and discuss some directions for further research.

2 The Fisher Market Game

As defined in the previous section, a linear Fisher market is defined by the tuple $(\mathcal{B}, \mathcal{G}, (\mathbf{u}_i)_{i \in \mathcal{B}}, \mathbf{m})$, where \mathcal{B} is a set of buyers, \mathcal{G} is a set of goods, $\mathbf{u}_i = (u_{ij})_{j \in \mathcal{G}}$ is the true utility tuple of buyer i , and $\mathbf{m} = (m_i)_{i \in \mathcal{B}}$ is the endowment vector. We assume that $|\mathcal{B}| = m$, $|\mathcal{G}| = n$ and the quantity of every good is one unit.

The *Fisher market game* is a one-shot non-cooperative game, where the buyers are the players, and the *strategy set* is all possible utility tuples that they may pose, *i.e.*, $\mathbb{S}_i = \{ \langle s_{i1}, s_{i2}, \dots, s_{in} \rangle \mid s_{ij} \geq 0, \sum_{j \in \mathcal{G}} s_{ij} \neq 0 \}$, $\forall i \in \mathcal{B}$. Clearly, the set of all strategy profiles is $\mathbb{S} = \mathbb{S}_1 \times \dots \times \mathbb{S}_m$. When a strategy profile $S = (\mathbf{s}_1, \dots, \mathbf{s}_m)$ is played, where $\mathbf{s}_i \in \mathbb{S}_i$, we treat $\mathbf{s}_1, \dots, \mathbf{s}_m$ as utility tuples of buyers $1, \dots, m$ respectively, and compute the equilibrium prices and a set of equilibrium allocations w.r.t. S and \mathbf{m} .

Further, using the equilibrium prices (p_1, \dots, p_n) , we generate the corresponding *solution graph* G as follows: Let $V(G) = \mathcal{B} \cup \mathcal{G}$. Let b_i be the node corresponding to the buyer i , $\forall i \in \mathcal{B}$ and g_j be the node corresponding to the good j , $\forall j \in \mathcal{G}$ in G . We place an edge between b_i and g_j iff $\frac{s_{ij}}{p_j} = \max_{k \in \mathcal{G}} \frac{s_{ik}}{p_k}$, and call the edges of the solution graph as *tight edges*. Note that when the solution graph is a forest, there is exactly one equilibrium allocation, however this is not so, when it contains cycles. In the standard Fisher market (*i.e.*, strategy of every buyer is her true utility tuple), all equilibrium allocations give the same payoff to a buyer. However, this is not so when buyers strategize on their utility tuples: Different equilibrium allocations may not give the same payoff to a buyer. The following example illustrates this scenario.

Example 2. Consider the Fisher market of Example 1. Consider the strategy profile $S = (\langle 1, 19 \rangle, \langle 1, 19 \rangle)$. Then, the equilibrium prices $\langle p_1, p_2 \rangle$ are $\langle 1, 19 \rangle$ and the solution graph is a cycle. There are many equilibrium allocations and the allocations $[x_{11}, x_{12}, x_{13}, x_{14}]$ achieving the highest payoff for buyers 1 and 2 are $[1, \frac{9}{19}, 0, \frac{10}{19}]$ and $[0, \frac{10}{19}, 1, \frac{9}{19}]$ respectively. The payoffs corresponding to these allocations are $(11.42, 5.26)$ and $(1.58, 7.74)$ respectively. Note that there is no allocation, which gives the maximum possible payoff to both the buyers.

Let $\mathbf{p}(S) = (p_1, \dots, p_n)$ be the equilibrium prices, $G(S)$ be the solution graph, and $\mathbb{X}(S)$ be the set of equilibrium allocations w.r.t. a strategy profile S .

The payoff w.r.t. $X \in \mathbb{X}(S)$ is defined as $(u_1(X), \dots, u_m(X))$, where $u_i(X) = \sum_{j \in \mathcal{G}} u_{ij} x_{ij}$. Let $w_i(S) = \max_{X \in \mathbb{X}(S)} u_i(X), \forall i \in \mathcal{B}$.

Definition 3. A strategy profile S is said to be **conflict-free** if $\exists X \in \mathbb{X}(S)$, s.t. $u_i(X) = w_i(S), \forall i \in \mathcal{B}$. Such an X is called a **conflict-free allocation**.

When a strategy profile $S = (\mathbf{s}_1, \dots, \mathbf{s}_m)$ is not conflict-free, there is a conflict of interest in selecting a particular allocation for the play. If a buyer, say k , does not get the same payoff from all the equilibrium allocations, i.e., $\exists X \in \mathbb{X}(S), u_k(X) < w_k(S)$, then we show that for every $\delta > 0$, there exists a strategy profile $S' = (\mathbf{s}'_1, \dots, \mathbf{s}'_m)$, where $\mathbf{s}'_i = \mathbf{s}_i, \forall i \neq k$, such that $u_k(X') > w_k(S) - \delta, \forall X' \in \mathbb{X}(S')$ (Section 3.1). The following example illustrates the same.

Example 4. In Example 2 for $\delta = 0.1$, consider $S' = ((1.1, 18.9), (1, 19))$, i.e., buyer 1 deviates slightly from S . Then, $\mathbf{p}(S') = \langle 1.1, 18.9 \rangle$, and $G(S')$ is a tree; the cycle of Example 2 is broken. Hence there is a unique equilibrium allocation, and $w_1(S') = 11.41, w_2(S') = 5.29$.

Therefore, if a strategy profile S is not conflict-free, then for every choice of allocation $X \in \mathbb{X}(S)$ to decide the payoff, there is a buyer who may deviate and assure herself a better payoff. In other words, when S is not conflict-free, there is no way to choose an allocation X from $\mathbb{X}(S)$ acceptable to all the buyers. This suggests that only conflict-free strategies are interesting. Therefore, we may define the payoff function $\mathcal{P}_i : \mathbb{S} \rightarrow \mathbb{R}$ for each player $i \in \mathcal{B}$ as follows:

$$\forall S \in \mathbb{S}, \mathcal{P}_i(S) = u_i(X), \text{ where } X = \arg \max_{X' \in \mathbb{X}(S)} \prod_{i \in \mathcal{B}} u_i(X'). \quad (1)$$

Note that the payoff functions are well-defined and when S is conflict-free, $\mathcal{P}_i(S) = w_i(S), \forall i \in \mathcal{B}$.

3 Nash Equilibria: A Characterization

In this section, we prove some necessary conditions for a strategy profile to be a NESP of the Fisher market game defined in the previous section. Nash equilibrium [13] is a solution concept for games with two or more rational players. When a strategy profile is a NESP, no player benefits by changing her strategy unilaterally.

For technical convenience, we assume that $u_{ij} > 0$ and $s_{ij} > 0, \forall i \in \mathcal{B}, \forall j \in \mathcal{G}$. The boundary cases may be easily handled separately. Note that if $S = (\mathbf{s}_1, \dots, \mathbf{s}_m)$ is a NESP then $S' = (\alpha_1 \mathbf{s}_1, \dots, \alpha_m \mathbf{s}_m)$, where $\alpha_1, \dots, \alpha_m > 0$, is also a NESP. Therefore, w.l.o.g. we consider only the normalized strategies $\mathbf{s}_i = \langle s_{i1}, \dots, s_{in} \rangle$, where $\sum_{j \in \mathcal{G}} s_{ij} = 1$, $\forall i \in \mathcal{B}$. As mentioned in the previous section, the true utility tuple of buyer i is $\langle u_{i1}, \dots, u_{in} \rangle$. For convenience, we may assume that $\sum_{j \in \mathcal{G}} u_{ij} = 1$ and $\sum_{i \in \mathcal{B}} m_i = 1$ (w.l.o.g.).

¹ For simplicity, we do use non-normalized strategy profiles in the examples.

We show that all NESPs are conflict-free. However, not all conflict-free strategies are NESPs. A symmetric strategy profile, where all players play the same strategy (*i.e.*, $\forall i, j \in \mathcal{B}$, $\mathbf{s}_i = \mathbf{s}_j$), is a NESP iff it is conflict-free. If a strategy profile S is not conflict-free, then there is a buyer a such that $\mathcal{P}_a(S) < w_a(S)$. The ConflictRemoval procedure in the next section describes how she may deviate and assure herself payoff almost equal to $w_a(S)$.

3.1 Conflict Removal Procedure

Definition 5. Let S be a strategy profile, $X \in \mathbb{X}(S)$ be an allocation, and $P = v_1, v_2, v_3, \dots$ be a path in $G(S)$. P is called an **alternating path** w.r.t. X , if the allocation on the edges at odd positions is non-zero, *i.e.*, $x_{v_{2i-1}v_{2i}} > 0, \forall i \geq 1$. The edges with non-zero allocation are called **non-zero edges**.

Table 1. Conflict Removal Procedure

<p>ConflictRemoval(S, b_a, δ)</p> <p>while b_a belongs to a cycle in $G(S)$ do</p> <p style="padding-left: 20px;">$(p_1, \dots, p_n) \leftarrow \mathbf{p}(S)$;</p> <p style="padding-left: 20px;">$J \leftarrow \{j \in \mathcal{G} \mid \text{the edge } (b_a, g_j) \text{ belongs to a cycle in } G(S)\}$;</p> <p style="padding-left: 20px;">$g_b \leftarrow \arg \max_{j \in J} \frac{u_{aj}}{p_j}$;</p> <p style="padding-left: 20px;">$X \leftarrow$ an allocation in $\mathbb{X}(S)$ such that $u_a(X) = w_a(S)$ and x_{ab} is maximum;</p> <p style="padding-left: 20px;">$S \leftarrow$ Perturbation($S, X, b_a, g_b, \frac{\delta}{n}$);</p> <p>endwhile</p> <p>return S;</p>
<p>Perturbation(S, X, b_a, g_b, γ)</p> <p>$S' \leftarrow S$;</p> <p>if (b_a, g_b) does not belong to a cycle in $G(S)$ then</p> <p style="padding-left: 20px;">return S';</p> <p>endif</p> <p>$J_1 \leftarrow \{v \mid \text{there is an alternating path from } b_a \text{ to } v \text{ in } G(S) \setminus (b_a, g_b) \text{ w.r.t. } X\}$;</p> <p>$J_2 \leftarrow \{v \mid \text{there is an alternating path from } g_b \text{ to } v \text{ in } G(S) \setminus (b_a, g_b) \text{ w.r.t. } X\}$;</p> <p>$(p_1, \dots, p_n) \leftarrow \mathbf{p}(S)$; $l \leftarrow \sum_{g_j \in J_1} p_j$; $r \leftarrow \sum_{g_j \in J_2} p_j$;</p> <p>W.r.t. α, define prices of goods to be</p> <p style="padding-left: 20px;">$\forall g_j \in J_1 : (1 - \alpha)p_j$; $\forall g_j \in J_2 : (1 + \frac{l\alpha}{r})p_j$; $\forall g_j \in \mathcal{G} \setminus (J_1 \cup J_2) : p_j$;</p> <p>Raise α infinitesimally starting from 0 such that none of the three events occur:</p> <p style="padding-left: 20px;">Event 1: a new edge becomes tight;</p> <p style="padding-left: 20px;">Event 2: a non-zero edge becomes zero;</p> <p style="padding-left: 20px;">Event 3: payoff of buyer a becomes $u_a(X) - \gamma$;</p> <p>$s'_{ab} \leftarrow s_{ab} \frac{(1 + \frac{l\alpha}{r})}{(1 - \alpha)}$; $\mathbf{s}'_a \leftarrow \frac{\mathbf{s}'_a}{\sum_{j \in \mathcal{G}} s'_{aj}}$;</p> <p>return S';</p>

The ConflictRemoval procedure in Table 1 takes a strategy profile S , a buyer a and a positive number δ , and outputs another strategy profile S' , where $\mathbf{s}'_i = \mathbf{s}_i, \forall i \neq a$ such that $\forall X' \in \mathbb{X}(S'), u_a(X') > w_a(S) - \delta$. The idea is that if a buyer, say a , does not belong to any cycle in the solution graph of a strategy

profile S , then $u_a(X) = w_a(S)$, $\forall X \in \mathbb{X}(S)$. The procedure essentially breaks all the cycles containing b_a in $G(S)$ using the Perturbation procedure iteratively such that the payoff of buyer a does not decrease by more than δ .

The Perturbation procedure takes a strategy profile S , a buyer a , a good b , an allocation $X \in \mathbb{X}(S)$, where x_{ab} is maximum among all allocations in $\mathbb{X}(S)$ and a positive number γ , and outputs another strategy profile S' such that $s'_i = s_i$, $\forall i \neq a$ and $w_a(S') > u_a(X) - \gamma$. It essentially breaks all the cycles containing the edge (b_a, g_b) in $G(S)$.

A detailed explanation of both the procedures is given in [1]. In the next theorem, we use the ConflictRemoval procedure to show that all the NESPs in the Fisher market game are conflict-free.

Theorem 6. *If S is a NESP, then*

- (i) $\exists X \in \mathbb{X}(S)$ such that $u_i(X) = w_i(S)$, $\forall i \in \mathcal{B}$, i.e., S is conflict-free.
- (ii) the degree of every good in $G(S)$ is at least 2.
- (iii) for every buyer $i \in \mathcal{B}$, $\exists k_i \in K_i$ s.t. $x_{ik_i} > 0$, where $K_i = \{j \in \mathcal{G} \mid \frac{u_{ij}}{p_j} = \max_{k \in \mathcal{G}} \frac{u_{ik}}{p_k}\}$, $(p_1, \dots, p_n) = \mathbf{p}(S)$ and $[x_{ij}]$ is a conflict-free allocation.

Proof. Suppose there does not exist an allocation $X \in \mathbb{X}(S)$ such that $u_i(X) = w_i(S)$, $\forall i \in \mathcal{B}$, then there is a buyer $k \in \mathcal{B}$, such that $\mathcal{P}_k(S) < w_k(S)$. Clearly, buyer k has a deviating strategy (apply ConflictRemoval on the input tuple (S, k, δ) , where $0 < \delta < (w_k(S) - \mathcal{P}_k(S))$), which is a contradiction.

For part (ii), if a good b is connected to exactly one buyer, say a , in $G(S)$, then buyer a may gain by reducing s_{ab} , so that price of good b decreases and prices of all other goods increase by the same factor.

For part (iii), if there exists a buyer i such that $x_{ik_i} = 0$, $\forall k_i \in K_i$, then she may gain by increasing the utility for a good in K_i . \square

The following example shows that the above conditions are not sufficient.

Example 7. Consider a market with 3 buyers and 2 goods, where $\mathbf{m} = \langle 50, 100, 50 \rangle$, $\mathbf{u}_1 = \langle 2, 0.1 \rangle$, $\mathbf{u}_2 = \langle 4, 9 \rangle$, and $\mathbf{u}_3 = \langle 0.1, 2 \rangle$. Consider the strategy profile $S = (\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3)$ given by the true utility tuples. The payoff tuple w.r.t. S is $(1.63, 6.5, 0.72)$. It satisfies all the necessary conditions in the above theorem, however S is not a NESP because buyer 2 has a deviating strategy $\mathbf{s}'_2 = \langle 2, 3 \rangle$ and the payoff w.r.t. strategy profile $(\mathbf{s}_1, \mathbf{s}'_2, \mathbf{s}_3)$ is $(1.25, 6.75, 0.83)$.

3.2 Symmetric and Asymmetric NESPs

Recall that a strategy profile $S = (\mathbf{s}_1, \dots, \mathbf{s}_m)$ is said to be a *symmetric* strategy profile if $\mathbf{s}_1 = \dots = \mathbf{s}_m$, i.e., all buyers play the same strategy.

Proposition 8. *A symmetric strategy profile S is a NESP iff it is conflict-free.*

Proof. (\Rightarrow) is easy (Theorem 6). For (\Leftarrow) , suppose a buyer i may deviate and gain, then the prices have to be changed. In that case, all buyers except buyer i will be connected to only those goods, whose prices are decreased. This leads to a contradiction (refer to [1] for details). \square

Let $S^f = [s_{ij}]$ be a strategy profile, where $s_{ij} = u_{ij}, \forall i \in \mathcal{B}, \forall j \in \mathcal{G}$, i.e., true utility functions. All allocations in $\mathbb{X}(S^f)$ give the same payoff to the buyers (i.e., $\forall i \in \mathcal{B}, u_i(X) = w_i(S^f), \forall X \in \mathbb{X}(S^f)$), and we define *Fisher payoff* (u_1^f, \dots, u_m^f) to be the payoff derived when all buyers play truthfully.

Corollary 9. *A symmetric NESP can be constructed, whose payoff is the same as the Fisher payoff.*

Proof. Let $S = (\mathbf{s}, \dots, \mathbf{s})$ be a strategy profile, where $\mathbf{s} = \mathbf{p}(S^f)$. Clearly S is a symmetric NESP, whose payoff is the same as the Fisher payoff (refer to [1] for details). \square

Remark 10. The payoff w.r.t. a symmetric NESP is always Pareto optimal. For a Fisher market game, there is exactly one symmetric NESP iff the degree of every good in $G(S^f)$ is at least two [10].

The characterization of all the NESPs for the general market game seems hard; even for markets with only three buyers. The following example illustrates an asymmetric NESP, whose payoff is not Pareto optimal.

Example 11. Consider a market with 3 buyers and 2 goods, where $\mathbf{m} = \langle 50, 100, 50 \rangle$, $\mathbf{u}_1 = \langle 2, 3 \rangle$, $\mathbf{u}_2 = \langle 4, 9 \rangle$, and $\mathbf{u}_3 = \langle 2, 3 \rangle$. Consider the two strategy profiles given by $S_1 = (\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3)$ and $S_2 = (\mathbf{s}, \mathbf{s}, \mathbf{s})$, where $\mathbf{s}_1 = \langle 2, 0.1 \rangle$, $\mathbf{s}_2 = \langle 2, 3 \rangle$, $\mathbf{s}_3 = \langle 0.1, 3 \rangle$, and $\mathbf{s} = \langle 2, 3 \rangle$. The payoff tuples w.r.t. S_1 and S_2 are $(1.25, 6.75, 1.25)$ and $(1.25, 7.5, 1.25)$ respectively. Note that both S_1 and S_2 are NESPs for the above market (refer to [1] for details).

4 The Two-Buyer Markets

A two-buyer market consists of two buyers and a number of goods. These markets arise in numerous scenarios. The two firms in a duopoly may be considered as the two buyers with a similar requirements to fulfill from a large number of suppliers, for example, relationship between two big automotive companies with their suppliers.

In this section, we study two-buyer market game and provide a complete polyhedral characterization of NESPs, all of which turn out to be symmetric. Next, we study how the payoffs of the two buyers change with varying NESPs and show that these payoffs constitute a piecewise linear concave curve. For a particular payoff tuple on this curve, there is a convex set of NESPs, hence a convex set of equilibrium prices, which leads to a different class of non-market behavior such as incentives. Finally, we study the correlated equilibria of this game and show that third-party mediation does not help to achieve better payoffs than any of the NESPs.

Lemma 12. *All NESPs for a two-buyer market game are symmetric.*

Proof. If a NESP $S = (\mathbf{s}_1, \mathbf{s}_2)$ is not symmetric, then $G(S)$ is not a complete bipartite graph. Therefore there is a good, which is exclusively bought by a buyer, which is a contradiction (Theorem [6], part (ii)). \square

4.1 Polyhedral Characterization of NESPs

In this section, we compute all the NESPs of a Fisher market game with two buyers. Henceforth we assume that the goods are so ordered that $\frac{u_{1j}}{u_{2j}} \geq \frac{u_{1(j+1)}}{u_{2(j+1)}}$, for $j = 1, \dots, n - 1$. Chakrabarty et al. [5] also use such an ordering to design an algorithm for the linear Fisher market with two agents. Let $S = (\mathbf{s}, \mathbf{s})$ be a NESP, where $\mathbf{s} = (s_1, \dots, s_n)$ and $(p_1, \dots, p_n) = \mathbf{p}(S)$. The graph $G(S)$ is a complete bipartite graph. Since $m_1 + m_2 = 1$ and $\sum_{j=1}^n s_j = 1$, we have $p_j = s_j, \forall j \in \mathcal{G}$. In a conflict-free allocation $X \in \mathbb{X}(S)$, if $x_{1i} > 0$ and $x_{2j} > 0$, then clearly $\frac{u_{1i}}{p_i} \geq \frac{u_{1j}}{p_j}$ and $\frac{u_{2i}}{p_i} \leq \frac{u_{2j}}{p_j}$.

Definition 13. An allocation $X = [x_{ij}]$ is said to be a *nice allocation*, if it satisfies the property: $x_{1i} > 0$ and $x_{2j} > 0 \Rightarrow i \leq j$.

The main property of a *nice allocation* is that if we consider the goods in order, then from left to right, goods get allocated first to buyer 1 and then to buyer 2 exclusively, however they may share at most one good in between. Note that a symmetric strategy profile has a unique nice allocation.

Lemma 14. Every NESP has a unique conflict-free nice allocation.

Proof. The idea is to convert a conflict-free allocation into a nice allocation through an exchange s.t. payoff remains same (refer to [1] for details). \square

The non-zero edges in a nice allocation either form a tree or a forest containing two trees. We use the properties of nice allocations and NESPs to give the polyhedral characterization of all the NESPs. The convex sets B_k for all $1 \leq k \leq n$, as given in Table 2, correspond to all possible conflict-free nice allocations, where non-zero edges form a tree, and the convex sets B'_k for all $1 \leq k \leq n - 1$, as given in Table 3, correspond to all possible conflict-free nice allocations, where non-zero edges form a forest². Let $\mathbb{B} = \cup_{k=1}^n B_k \cup_{k=1}^{n-1} B'_k$ and $S^{NE} = \{(\alpha, \alpha) \mid \alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{B}\}$. Note that S^{NE} is a connected set.

Table 2. B_k

$\sum_{i=1}^{k-1} \alpha_i < m_1$	
$\sum_{i=k+1}^n \alpha_i < m_2$	
$\sum_{i=1}^n \alpha_i = m_1 + m_2$	
$u_{1j}\alpha_i - u_{1i}\alpha_j \leq 0$	$\forall i \leq k, \forall j \geq k$
$u_{2i}\alpha_j - u_{2j}\alpha_i \leq 0$	$\forall i \leq k, \forall j \geq k$
$\alpha_i \geq 0$	$\forall i \in \mathcal{G}$

Table 3. B'_k

$\sum_{i=1}^k \alpha_i = m_1$	
$\sum_{i=k+1}^n \alpha_i = m_2$	
$u_{1j}\alpha_i - u_{1i}\alpha_j \leq 0$	$\forall i \leq k, \forall j \geq k + 1$
$u_{2i}\alpha_j - u_{2j}\alpha_i \leq 0$	$\forall i \leq k, \forall j \geq k + 1$
$\alpha_i \geq 0$	$\forall i \in \mathcal{G}$

Lemma 15. A strategy profile S is a NESP iff $S \in S^{NE}$.

Proof. (\Leftarrow) is easy by the construction and Proposition 8. For the other direction, we know that every NESP has a conflict-free nice allocation (Lemma 14), and \mathbb{B} corresponds to all possible conflict-free nice allocations. \square

² In both the tables α_i 's may be treated as price variables.

4.2 The Payoff Curve

In this section, we consider the payoffs obtained by both the players at various NESPs. Recall that whenever a strategy profile S is a NESP, $\mathcal{P}_i(S) = w_i(S)$, $\forall i \in \mathcal{B}$. Henceforth, we use $w_i(S)$ as the payoff of buyer i for the NESP S . Let $\mathbb{F} = \{(w_1(S), w_2(S)) \mid S \in S^{NE}\}$ be the set of all possible NESP payoff tuples.

Let \mathcal{X} be the set of all nice allocations, and $\mathbb{H} = \{(u_1(X), u_2(X)) \mid X \in \mathcal{X}\}$. For $\alpha \in [0, 1]$, let $t(\alpha) = (\langle s_1, \dots, s_n \rangle, \langle s_1, \dots, s_n \rangle)$, where $s_i = u_{1i} + \alpha(u_{2i} - u_{1i})$, and $\mathbb{G} = \{(w_1(S), w_2(S)) \mid S = t(\alpha), \alpha \in [0, 1]\}$.


Proposition 16. \mathbb{F} is a piecewise linear concave (PLC) curve.

Proof. The proof is based on the following steps (refer to [1] for details).

1. \mathbb{H} is a PLC curve with $(0, 1)$ and $(1, 0)$ as the end points.
2. $\forall \alpha \in [0, 1]$, $t(\alpha) \in S^{NE}$, then clearly $\mathbb{G} \subset \mathbb{H}$. Since the nice allocation w.r.t. $t(\alpha)$ changes continuously as α moves from 0 to 1, so we may conclude that \mathbb{G} is a PLC curve with the end points $(w_1(S^1), w_2(S^1))$ and $(w_1(S^2), w_2(S^2))$, where $S^1 = t(0)$ and $S^2 = t(1)$.
3. $\mathbb{F} = \mathbb{G}$. □

The next example demonstrates the payoff curve for a small market game.

Example 17. Consider a market with 3 goods and 2 buyers, where $\mathbf{m} = \langle 7, 3 \rangle$, $\mathbf{u}_1 = \langle 6, 2, 2 \rangle$, and $\mathbf{u}_2 = \langle 0.5, 2.5, 7 \rangle$. The payoff curve for this game is shown in the figure. The first and the second line segment of the curve correspond to the sharing of good 2 and 3 respectively. The payoffs corresponding to the boundary NESPs $S^1 = t(0)$ and $S^2 = t(1)$ are $(7, 8.25)$ and $(9.14, 3)$ respectively.



Furthermore, the *Fisher payoff* $(8, 7)$ may be achieved by a NESP $t(0.2)$. Note that in this example the social welfare (*i.e.*, sum of the payoffs of both the buyers) from the Fisher payoff (15) is lower than that of the NESP S^1 (15.25).

4.3 Incentives

For a fixed payoff tuple on the curve \mathbb{F} , there is a convex set of NESPs and hence a convex set of prices, giving the same payoffs to the buyers, and these may be computed using the similar inequalities as defined in Tables 2 and 3. This leads to a different class of behavior, *i.e.*, motivation for a seller to offer incentives to the buyers to choose a particular NESP from this convex set, which fetches the maximum price for her good. The following example illustrates this possibility.

Example 18. Consider a market with 2 buyers and 4 goods, where $\mathbf{m} = \langle 10, 10 \rangle$, $\mathbf{u}_1 = \langle 4, 3, 2, 1 \rangle$, and $\mathbf{u}_2 = \langle 1, 2, 3, 4 \rangle$. Consider the two NESPs given by $S_1 = (\mathbf{s}_1, \mathbf{s}_1)$ and $S_2 = (\mathbf{s}_2, \mathbf{s}_2)$, where $\mathbf{s}_1 = \langle \frac{20}{3}, \frac{20}{3}, \frac{10}{3}, \frac{10}{3} \rangle$ and $\mathbf{s}_2 = \langle \frac{20}{3}, \frac{20}{3}, \frac{9}{3}, \frac{11}{3} \rangle$. Both S_1 and S_2 gives the payoff $(5.5, 8)$, however the prices are different, *i.e.*, $\mathbf{p}(S_1) = \langle \frac{20}{3}, \frac{20}{3}, \frac{10}{3}, \frac{10}{3} \rangle$ and $\mathbf{p}(S_2) = \langle \frac{20}{3}, \frac{20}{3}, \frac{9}{3}, \frac{11}{3} \rangle$. Clearly in S_2 , good 3 is penalized and good 4 is rewarded (compared to S_1).

4.4 Correlated Equilibria

We have seen in Section 4.2 that the two-buyer market game has a continuum of Nash Equilibria, with very different and conflicting payoffs. This makes it difficult to predict how a particular game will actually play out in practice, and if there is a different solution concept which may yield an outcome liked by both the players.

We examine the correlated equilibria framework as a possibility. Recall that according to the correlated equilibria, the mediator decides and declares a probability distribution π on all possible pure strategy profiles $(\mathbf{s}_1, \mathbf{s}_2) \in \mathbb{S}_1 \times \mathbb{S}_2$ beforehand. During the play, she suggests what strategy to play to each player privately, and no player benefits by deviating from the advised strategy. The question we ask: Is there a correlated equilibrium π such that the payoff w.r.t. π lies above the curve \mathbb{H} ? We continue with our assumption that $\frac{u_{1j}}{u_{2j}} \geq \frac{u_{1(j+1)}}{u_{2(j+1)}}$, $\forall j < n$.

Lemma 19. *For any strategy profile $S = (\mathbf{s}_1, \mathbf{s}_2)$, for every allocation $X \in \mathbb{X}(S)$, there exists a point (x_1, x_2) on \mathbb{H} such that $x_1 \geq u_1(X)$ and $x_2 \geq u_2(X)$.*

Proof. Any allocation X may be converted to a nice allocation through an exchange such that no buyer worse off (refer to [11] for details). \square

Corollary 20. *The correlated equilibrium does not give better payoff than any NE payoff to all the buyers.*

Remark 21. [10] extends this result for the general Fisher market game.

5 Conclusion

The main conclusion of the paper is that Fisher markets in practice will rarely be played with true utility functions. In fact, the utilities employed will usually be a mixture of a player's own utilities and her conjecture on the other player's true utilities. Moreover, there seems to be no third-party mediation which will induce players to play according to their true utilities so that the true Fisher market equilibrium may be observed. Further, any notion of market equilibrium should examine this aspect of players strategizing on their utilities. This poses two questions: (i) is there a mechanism which will induce players into revealing their true utilities? and (ii) how does this mechanism reconcile with the "invisible hand" of the market? The strategic behavior of agents and the question whether true preferences may ever be revealed, has been of intense study in economics [12, 17, 20]. The main point of departure for this paper is that buyers strategize directly on utilities rather than market implementation specifics, like trading posts and bundles. Hopefully, some of these analysis will lead us to a more effective computational model for markets.

On the technical side, the obvious next question is to completely characterize the NESPs for the general Fisher market game. We assumed the utility functions of the buyers to be linear, however Fisher market is gameable for the other class of utility functions as well. It will be interesting to do a similar analysis for more general utility functions.

References

1. Adsul, B., Babu, C.S., Garg, J., Mehta, R., Sohoni, M.: Nash equilibria in Fisher market. arXiv:1002.4832 (2010)
2. Amir, R., Sahi, S., Shubik, M., Yao, S.: A strategic market game with complete markets. *Journal of Economic Theory* 51, 126–143 (1990)
3. Brainard, W.C., Scarf, H.E.: How to compute equilibrium prices in 1891. Cowles Foundation, Discussion Paper-1272 (2000)
4. Bu, T., Deng, X., Qi, Q.: Forward looking Nash equilibrium for keyword auction. *Inf. Process. Lett.* 105(2), 41–46 (2008)
5. Chakrabarty, D., Devanur, N.R., Vazirani, V.V.: New results on rationality and strongly polynomial solvability in Eisenberg-Gale markets. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 239–250. Springer, Heidelberg (2006)
6. Codenotti, B., Pemmaraju, S., Varadarajan, K.: On the polynomial time computation of equilibria for certain exchange economies. In: SODA 2005 (2005)
7. Devanur, N.R., Papadimitriou, C.H., Saberi, A., Vazirani, V.V.: Market equilibrium via a primal-dual type algorithm. *JACM* 55(5) (2008)
8. Dubey, P., Geanakoplos, J.: From Nash to Walras via Shapley-Shubik. *Journal of Mathematical Economics* 39, 391–400 (2003)
9. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review* 97(1), 242–259 (2007)
10. Garg, J.: Nash equilibria in Fisher market. Working Manuscript (2010)
11. Jain, K.: A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. In: FOCS 2004 (2004)
12. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Microeconomic Theory*. Oxford University Press, Oxford (1995)
13. Nash, J.F.: Equilibrium points in n -person games. *Proc. of the National Academy of Sciences of the United States of America* 36(1), 48–49 (1950)
14. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
15. Orlin, J.B.: Improved algorithms for computing Fisher’s market clearing prices. In: STOC 2010 (2010)
16. Samuelson, P.A.: A note on the pure theory of consumers’ behaviour. *Economica* 5, 61–71 (1938)
17. Samuelson, P.A.: *Foundations of Economic Analysis*. Harvard University Press (1947)
18. Shapley, L., Shubik, M.: Trade using one commodity as a means of payment. *Journal of Political Economy* 85(5), 937–968 (1977)
19. Varian, H.: Position auctions. *International Journal of Industrial Organization* 25, 1163–1178 (2007)
20. Varian, H.: *Microeconomic Analysis*, 3rd edn (1992)
21. Walras, L.: *Elements of Pure Economics*, Translated by Jaffé, Allen & Urwin. London (1954)

Partition Equilibrium Always Exists in Resource Selection Games*

Elliot Anshelevich, Bugra Caskurlu, and Ameya Hate

Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY

Abstract. We consider the existence of Partition Equilibrium in Resource Selection Games. Super-strong equilibrium, where no subset of players has an incentive to change their strategies collectively, does not always exist in such games. We show, however, that partition equilibrium (introduced in [4] to model coalitions arising in a social context) always exists in general resource selection games, as well as how to compute it efficiently. In a partition equilibrium, the set of players has a fixed partition into coalitions, and the only deviations considered are by coalitions that are sets in this partition. Our algorithm to compute a partition equilibrium in any resource selection game (i.e., load balancing game) settles the open question from [4] about existence of partition equilibrium in general resource selection games. Moreover, we show how to always find a partition equilibrium which is also a Nash equilibrium. This implies that in resource selection games, we do not need to sacrifice the stability of individual players when forming solutions stable against coalitional deviations. In addition, while super-strong equilibrium may not exist in resource selection games, we show that its existence can be decided efficiently, and how to find one if it exists.

1 Introduction

In multi-agent systems, it is common to assume that the agents will change their existing behavior if they can reduce their cost by doing so. This assumption is at the heart of the study of Nash equilibrium in various settings. The concept of Nash equilibrium, however, becomes relevant only in scenarios where agents cannot form coalitions, and change their behavior as a group. The *Strong Equilibrium* [1] solution concept, where any subset of agents can form a coalition and deviate together if it is beneficial to all of them, addresses the weaknesses of the Nash equilibrium solution concept for the settings where players can form coalitions. A strong equilibrium represents the scenario where any group of players could form a coalition, and everyone has to strictly benefit from a deviation. In this paper, we relax these assumptions, and consider the cases where only some of the subsets of players could group themselves together into a deviating coalition, and where not everyone in a coalition has to strictly improve their utility in order to deviate.

* This work supported in part by NSF CCF-0914782.

We study these solution concepts in the context of Resource Selection Games (RSGs). RSGs model a wide range of scenarios, where a set of players are selecting exactly one of various resources, with the cost of using a resource depending on the type of the resource, as well as the number of players selecting this particular resource. They present a framework that can be used to model the problems of various communities like operations research, economics, computing systems, transportation, and communication networks. The atomic selfish routing game [2,6,13,14] on parallel link networks and selfish machine assignment [3,5,10,14] for identical jobs are among various problems modeled as RSGs in the algorithmic game theory community. RSGs fall into the class of potential games [11,12] for which existence of a pure Nash equilibrium is guaranteed. Holzman and Law-Yone [8,9] proved the existence of strong equilibrium in RSGs as well. However, *Super-Strong Equilibrium* (see below) is not guaranteed to exist in RSGs, which led Feldman and Tennenholtz [4] to define a concept of *Partition Equilibrium* and study its existence in the context of Resource Selection Games. In this paper, we greatly extend their results by showing the existence of partition equilibrium in *every* Resource Selection Game, as well as how to compute it.

1.1 Related Solution Concepts

The *Strong Equilibrium* (SE) solution concept assumes a coalition will deviate only if the deviation is strictly profitable to all members of the coalition. In a strong equilibrium, no subset of players is able to deviate with every player in the group strictly improving their utility.

Super-Strong Equilibrium (SSE) considers *weakly-profitable deviations*, where a coalition will deviate provided that no member of the coalition becomes worse off, and at least one member of the coalition strictly benefits. A super-strong equilibrium is a solution where no subset of players has such a deviation. This solution concept makes more sense in many settings, especially if agents will somewhat care about the utility of other agents (which perfectly make sense if the agents are friends, colleagues, family members). While strong equilibrium is guaranteed to exist in RSGs, there are RSG instances where super-strong equilibrium may not exist, even with 2 identical machines and 3 players [4]. Additionally, if we consider the formation of player coalitions as arising from a social context (i.e., a group of friends decide to form a coalition together), then the assumption that *any* subset of players can form a coalition is quite strong.

Partition Equilibrium was first defined in [4] as an attempt to model coalitions that arise from a social context. In this setting, the specification of the game contains a fixed partition T over the set of players. This partition divides the players into non-overlapping coalitions. In this solution concept, the only permissible deviations are the ones where a coalition is one of the sets in the fixed partition T . A solution is a stable solution if no coalition has an *weakly-profitable deviation*, i.e., a deviation where at least one member of the coalition strictly benefits and no member of the coalition becomes worse off. [4] called such a stable solution a T -SSE, since a partition equilibrium is a super-strong equilibrium, but with the only coalitions that are allowed to deviate being the

sets of partition T . Also observe that unlike strong equilibrium or super-strong equilibrium, partition equilibrium solutions are not a subset of Nash equilibrium solutions.

Feldman and Tennenholtz [4] studied the existence of partition equilibrium in the context of resource selection games and proved that partition equilibrium exists in the following special cases:

- All the resources are identical, i.e., they share the same latency function, or
- There are only 2 resources in the system, or
- Each coalition is composed of 1 or 2 players.

Note that partition equilibrium is a solution concept in a non-transferable utility game, i.e., money transfers among the players are not allowed. The *Collusion Equilibrium* solution concept [7,6] is the analogue of partition equilibrium in transferable utility games. In this solution concept, there is also a fixed partition over the players which forms the non-overlapping coalitions. The only difference is that money transfers among the players are permitted, and therefore a deviation is an improving deviation if it reduces the total cost of the players in the coalition. Observe that collusion equilibrium is a stronger solution concept in the sense that an allocation of players to resources that constitutes a collusion equilibrium (no coalition can reduce its total cost by deviating) is also a partition equilibrium allocation but not vice versa. Hayrapetyan, Tardos and Wexler [7] studied the existence and computation of collusion equilibrium in the context of resource selection games. They proved the existence of collusion equilibrium (and therefore, partition equilibrium) in the special case where the latency functions of the resources are convex. Their proof is constructive, i.e., they give an algorithm that produces a collusion equilibrium solution which may *not* be a Nash equilibrium solution.

1.2 Our Results

Our main result is the proof of existence (and efficient computation) of an allocation A of players to resources such that A is *both* a partition equilibrium and a Nash equilibrium allocation. This result holds for general resource allocation games, with no assumptions about the latency functions of different resources (except them being increasing), on the size of the coalitions, or on the number of machines. This resolves an open question from [4] about the existence of partition equilibrium for general RSGs. Moreover, our results provide the interesting insight that for every partition T there exists a solution where no coalition of T would gain by deviating (i.e., it is a T -SSE), *and* no single player would gain by deviating (i.e., it is a Nash equilibrium). This implies that we do not need to sacrifice the stability of individual players when forming solutions stable against coalitional deviations.

In Section 2, we present a formal definition of resource selection games and give a complete characterization of Nash equilibrium solutions for these games. In Section 3, we give a set of sufficient conditions for coalitions such that if a

coalition satisfies the given conditions on a Nash equilibrium allocation, then that coalition does not have an improving deviation. In Section 4, we give an algorithm that produces a Nash equilibrium allocation of players to resources such that all coalitions satisfy the sufficient conditions given in Section 3. In Section 5, we show that for any resource selection game instance, the existence of super-strong equilibrium is efficiently decidable, and if super-strong equilibrium exists, then we can compute one efficiently.

In summary, this paper shows that we can always find a SSE if one exists, but even for games which do not admit a SSE, we can find a solution that is stable for any set in a given partition T , as well as for any individual player.

2 Model and Preliminaries

We now formally define the resource selection game. We have n players (jobs) and m resources (machines). The strategy of each player is to select exactly one of the m machines. Each machine i has a strictly increasing latency function $f_i(n_i)$ which only depends on the number of players n_i that select machine i . The cost of each player that selected machine i is $f_i(n_i)$.

In this paper we will consider partition equilibrium and super-strong equilibrium (SSE), both of which are solution concepts involving stability against coalitional deviations. Specifically, by an *improving deviation* by a coalition of players C , we will mean a *weakly-profitable deviation*, i.e., a deviation where no player in C increases their cost, and at least one player of C strictly decreases their cost.

A SSE is an allocation of jobs to machines, so that no subset of jobs has an improving deviation. As shown in 4, a SSE does not always exist, although a strong equilibrium (where a deviation will only occur if every member of a coalition strictly profits) always exists in resource selection games [8,9].

Now suppose that we have a fixed partition $T = T_1, \dots, T_k$ over the set of players such that $T_i \cap T_j = \emptyset$, i.e., the sets are not overlapping. Each set T_i represents a coalition of players that are willing to deviate as a group. Then, a *partition equilibrium* or T -SSE is an allocation of jobs to machines such that no set of jobs in partition T has an improving deviation. Then, it is clear that a SSE is also a T -SSE for every partition T , as well as a Nash equilibrium. A T -SSE, on the other hand, is not necessarily a Nash equilibrium.

2.1 Nash Equilibrium

Since we are going to show that the existence of an allocation that is a T -SSE and a Nash equilibrium, we first give a complete characterization of Nash equilibrium solutions.

Let u be the minimum makespan of our system, i.e., the minimum value of $\max_i f_i(n_i)$ that can be achieved for any allocation of jobs to machines. Notice that since the latency of a machine depends only on the number of jobs assigned to this machine, then u is easily computable using a greedy algorithm. We classify

the machines into two groups. A resource i is called a 'type 1' resource if there exists a positive integer z such that $f_i(z) = u$. In other words, a resource is a 'type 1' resource if it can attain a latency of u . We say that a resource i is a 'type 2' resource if it cannot attain a latency of u , i.e., there is no positive integer z such that $f_i(z) = u$.

For each machine i , define m_i as the maximum number of jobs a machine can accept while i attains a latency at most u , i.e., $m_i = \max_z \{f_i(z) \leq u\}$.

Proposition 1. *An allocation A of jobs onto machines is a Nash equilibrium if and only if each type 2 machine i is allocated exactly m_i jobs and each type 1 machine i is allocated either m_i or $m_i - 1$ jobs, with at least one type 1 machine i allocated exactly m_i jobs.*

Proof. *if:* Note that when a job deviates it has to move to another machine, thereby increasing the number of jobs on that machine. If the number of jobs on any machine increases, then that machine will experience a latency of at least u . Since all jobs are currently experiencing a cost of at most u , the latency of any job after moving to a different machine will not decrease. This proves that if all the above conditions are satisfied then the allocation is a Nash equilibrium.

only if: If the makespan of a solution is more than u , say α , then this means that some machine i has more than m_i jobs on it. This implies that there exists a machine j that has less than m_j jobs on it. Then by transferring a job from machine i to j we can reduce the latency faced by that job from α to at most u . Hence a Nash equilibrium will always have a makespan of u . Also if any type 2 machine has less than m_i jobs, or a type 1 machine has less than $m_i - 1$ jobs on it, then by moving a job that faces a latency of u to this machine, we can reduce its latency. It is trivial to see that any type 2 machine will not have more than m_i jobs on it since that will increase the makespan to more than u . Hence any such allocation will not form a Nash equilibrium. This proves that in order for an allocation to be a Nash equilibrium, all the above conditions must be fulfilled. ■

By Proposition [1](#), some type 1 machines i are allocated m_i jobs and therefore the jobs on them are experiencing a cost of u , and some type 1 machines are (possibly) allocated $m_i - 1$ jobs and the jobs on those machines are experiencing a cost strictly less than u . Given a Nash equilibrium solution, we use the term *high machine* to refer to a type 1 machine i that has m_i jobs and use H to denote this set of machines. We use the term *low machine* to refer to all other type 1 machines and use L to denote this set of machines throughout the paper. We use R to denote the set of type 2 machines.

Given a game instance (i.e., the set of machines with their latency functions, the set of players, and the partition specified on it), the set of type 1 and type 2 machines can be readily decided, i.e., the same set of machines will be type 1 machines and the same set of machines will be type 2 machines in any Nash equilibrium allocation A . However, the splitting of type 1 machines into high machines H and low machines L depends on the Nash equilibrium solution selected. Let A and A' be two different Nash equilibrium allocations and let

H, H' and L, L' be the corresponding high and low machines for these Nash equilibrium allocations. Observe that $|H| = |H'|$ (and therefore $|L| = |L'|$) even though H and H' (and therefore L and L') may be different sets of machines. The number of high machines in any Nash equilibrium will be same.

3 Sufficient Conditions for Stability

In this paper, we want to construct a Nash equilibrium solution that is also a partition equilibrium for any given partition of the players. So, we want to construct a Nash equilibrium allocation such that none of the coalitions has an improving deviation. Given a Nash equilibrium allocation A , whether a coalition T_k has an improving deviation or not depends on the number of jobs of this coalition allocated to each machine. In this section, we will give a set of sufficient conditions for a coalition T_k not to have an improving deviation. Observe that if all the coalitions satisfy these sufficient conditions then none of the coalitions will have an improving deviation, which implies that the allocation A is also a partition equilibrium. For a type 1 machine i , we use l_i to denote $f_i(m_i - 1)$, i.e., the latency that it would experience if it were a low machine.

Following lemmas, the proofs which are not included due to lack of space, will help in finding the sufficient conditions for stability:

Lemma 1. *If the number of jobs on a machine k is the same before and after an improving deviation, then there exists an equivalent improving deviation (i.e., with the same number of jobs on each machine) where no jobs move to or from machine k .*

Lemma 2. *If a coalition T_k , that has 0 or 1 jobs on a high machine i in a Nash equilibrium allocation A , has an improving deviation D , then T_k has another improving deviation D' , where no jobs move to or from i .*

Theorem 1. *Given a Nash equilibrium allocation A and a coalition T_k , let x_i denote the number of jobs of the coalition T_k allocated to machine i in A . Then coalition T_k does not have an improving deviation if for every high machine i such that $x_i \geq 2$ the following conditions are satisfied:*

- for every low machine j such that $l_j > l_i$, we have that $x_j \geq x_i$ and
- for every low machine j such that $l_j \leq l_i$, we have that $x_j \geq x_i - 1$.

Proof. For the purpose of contradiction, assume there exists a coalition T_k that satisfies all the conditions and yet has an improving deviation D . Let A' denote the allocation of jobs to machines if coalition T_k takes its improving deviation D , and x'_i be the number of jobs T_k has on machine i in allocation A' . Since the allocation of the jobs of all coalitions except T_k are the same in both A and A' , the change in the number of jobs on any machine i is as much as the change in the number of jobs coalition T_k has on i . No machine i can have more than m_i jobs allocated to it in allocation A' since otherwise, the jobs on i (at least one of which is a member of T_k) will experience a latency more than u , which will imply that the deviation is not an improving deviation.

If coalition T_k has 0 or 1 jobs on a high machine i in allocation A then there exists another improving deviation D' , where no jobs move to or from i by Lemma 2. We will assume that D has this property. Notice that if $x_i < 2$ for a high machine i , then i is also a high machine in allocation A' . Let $x_h = \max_{i \in H} \{x_i\}$. We will first show that $x_h \geq 2$. Otherwise, all machines in H remain high after the deviation. If any other machine $j \notin H$ became high after deviation D , then jobs on j would experience a cost of u . However, all jobs with cost of u in A are on machines of H after the deviation, which means that the jobs on j have strictly increased their cost due to deviation D , and therefore D could not be an improving deviation. Thus, if $x_h < 2$, then the set of high machines is the same before and after D . In addition, if any machine $j \notin H$ has less jobs in A' than it did in A , then another machine must have more jobs, which would cause those jobs to experience a cost of at least u . By the argument above, this cannot happen, and so $x_h \geq 2$.

Lemma 3. *Let H' be the set of machines with latency of exactly u in allocation A' . Then, $|H'| \leq |H|$.*

Proof. In allocation A , coalition T_k has $\sum_{i \in H} x_i$ jobs experiencing a latency of u , whereas in allocation A' , coalition T_k has $\sum_{i \in H'} x'_i$ jobs experiencing a latency of u . Let $x_h = \max_{i \in H} \{x_i\}$ and let $x_l = \min_{i \in L} \{x_i\}$. The sufficient conditions state that $x_l \geq x_h - 1$, since $x_h \geq 2$ as shown above. If $i \in H'$ was a low machine before the deviation, then $x'_i = x_i + 1$, and so it has at least as many jobs of T_k in A' as any high machine of allocation A . If $i \in H'$ was a high machine before the deviation, then $x'_i = x_i$. Thus $|H'| > |H|$ would imply that $\sum_{i \in H'} x'_i > \sum_{i \in H} x_i$, which means that coalition T_k has more jobs that are experiencing a latency of u in allocation A' than allocation A . However, that would contradict with D being an improving deviation, and so it has to be that $|H'| \leq |H|$. ■

Note that the total number of jobs in any Nash equilibrium allocation A can be expressed as $\sum_{i \in R} m_i + \sum_{i \in H} m_i + \sum_{i \in L} (m_i - 1)$. If a type 2 machine $i \in R$ has less than m_i jobs in A' then the number of machines that has latency of u would be strictly more than $|H|$, i.e., $|H'| > |H|$. Therefore, the number of jobs coalition T_k has on any type 2 machine in allocation A' is exactly x_i . Since deviation D does not change the number of jobs T_k has on any type 2 machine, there exists an equivalent improving deviation where the jobs of type 2 machines do not change by Lemma 1, and we will assume that D has this property.

Observe that if a type 1 machine i has less than $m_i - 1$ jobs in allocation A' , then $|H'| > |H|$, thus violating Lemma 3. Therefore, every type 1 machine has either m_i or $m_i - 1$ jobs in allocation A' . In other words, by Proposition 1 A' is also a Nash equilibrium allocation.

Since A' is a Nash equilibrium, we can assume without loss of generality that deviation D made a certain number of high machines become low, and the same number of low machines become high. Using Lemma 1 we can assume that the machines on which the number of jobs did not change also did not take part in deviation D . Let the set of machines the become low after the deviation be

H^- and the set of machines that become high after the deviation be L^+ . We know that $|H^-| = |L^+|$. Now consider the total latency faced by jobs on machines belonging to $H^- \cup L^+$ before deviation, say α , and after the deviation, say β .

$$\begin{aligned}\alpha &= \sum_{i \in H^-} ux_i + \sum_{j \in L^+} l_j x_j \\ \beta &= \sum_{i \in H^-} (x_i - 1)l_i + \sum_{j \in L^+} (x_j + 1)u\end{aligned}$$

We now prove the following lemma:

Lemma 4. *For every perfect matching P between the machines of H^- and L^+ that pairs $i \in H^-$ with $j \in L^+$, it must be true that $l_j \leq l_i$ and $x_j = x_i - 1$.*

Proof. Let P be any perfect matching between the machines of H^- and L^+ (note that $|H^-| = |L^+|$). Consider a pair of machines $(i, j) \in P$ such that $i \in H^-$ and $j \in L^+$. If $l_j > l_i$ then we know that $x_j \geq x_i$. This means that the total number of jobs facing a latency u on machines i, j after deviation: $(x_j + 1)$ is strictly more than before: (x_i) . If $l_j \leq l_i$ then we know that $x_j \geq x_i - 1$. This would mean that the total number of jobs facing a latency u on machines i, j after deviation: $(x_j + 1)$ is at least as much as before: (x_i) .

This implies that if there exists even one pair of machines (i, j) such that $l_j > l_i$, then the total number of jobs facing a latency of u after deviation will strictly increase. On the other hand if $l_j \leq l_i$ but $x_j > x_i - 1$ then too it is easy to see that the number of jobs facing a latency of u after deviation strictly increases.

Hence D is a valid deviation only if for every $(i, j) \in P$, $l_j \leq l_i$ and $x_j = x_i - 1$. ■

Consider any perfect matching P between the machines of H^- and L^+ . We can now compare the values of α and β :

$$\begin{aligned}\alpha &= \sum_{i \in H^-} ux_i + \sum_{j \in L^+} l_j x_j \\ &= \sum_{j \in L^+} u(x_j + 1) + \sum_{j \in L^+} l_j x_j \quad \dots \text{(Lemma 4)} \\ &= \sum_{j \in L^+} u(x_j + 1) + \sum_{(i,j) \in P} l_j (x_i - 1) \quad \dots \text{(Lemma 4)} \\ &\leq \sum_{j \in L^+} u(x_j + 1) + \sum_{i \in H^-} l_i (x_i - 1) \quad \dots \text{(For every } (i, j) \in P, l_j \leq l_i) \\ &= \beta\end{aligned}$$

This means that the total cost faced by jobs of machines of $H^- \cup L^+$ will at best remain the same. If the latency of some job decreases then the latency of some other has to increase in order to keep the sum constant. This means that

the latency faced by every job can at best remain the same. But an improving deviation requires that D must strictly improve the latency of some job. Hence an improving deviation D does not exist. ■

4 Partition Equilibrium

We now present an algorithm that constructs an allocation of jobs such that all the sufficient conditions of Theorem 1 are satisfied. Thus we will create a Nash equilibrium that is also a partition equilibrium. For this purpose we will use the properties of Nash equilibrium as described in Section 2. Particularly we will use the fact that, given the total number of jobs, every Nash equilibrium will have the same number of high machines, which we denote by q . The algorithm gives an allocation of jobs over all type-1 machines. Since the sufficient conditions do not have restrictions on jobs of the type-2 machines, remaining jobs can be arbitrarily allocated to them so that each machine has m_i jobs. We also define a set of *active machines* as all machines i that have less than m_i jobs on them. Let $q = |H|$ be the number of high machines in any Nash equilibrium allocation. The algorithm is as follows:

- Begin with an empty allocation. Note that all machines are active at this time.
- Obtain an ordering on the set of all active machines based on non-increasing values of their l_i .
- For every coalition, place jobs sequentially starting from the first active machine according to the above ordering. If the number of jobs in this coalition exceeds the number of active machines then rollover and continue placing jobs from the first active machine in the ordering.
- If at any step a machine i has m_i jobs placed on it, i.e. i becomes high, then remove it from the set of active machines.
- When q machines become high, place remaining jobs on the active machines arbitrarily such that they have $m_i - 1$ jobs on them.

Example. Consider an example with 4 machines and 2 coalitions in Figure 1. Coalition 1 has 6 jobs and coalition 2 has 4 jobs. All the machines in this example are of type-1. Also they have been sorted in non-increasing order of their l_i -values. The blocks represent the jobs and height of the j 'th block on machine i is given by $f_i(j) - f_i(j - 1)$. Figure 1 illustrates various stages during the implementation of the algorithm. Observe that in any Nash equilibrium for this input exactly $q = 2$ machines will be high.

Notice that making sure that our algorithm places no more than m_i jobs on any machine is crucial not only to create a Nash equilibrium, but also to create a partition equilibrium. For example, in Figure 1, if we did not stop adding jobs to machines once they have m_i jobs, then we would end up with $3 = m_2 + 1$ jobs on machine 2. If $f_2(3) > f_4(3)$, then this would not be a partition equilibrium, since Coalition 1 would have an improving deviation by moving two of its jobs to machine 4 from machine 2, and one job to machine 2 from machine 4.

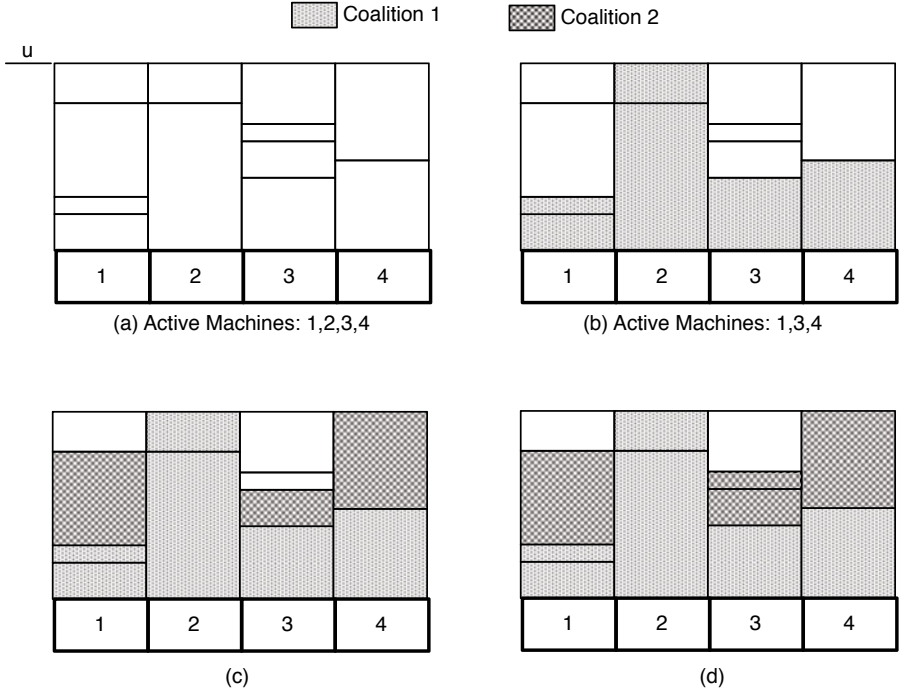


Fig. 1. (a) In the beginning all machines are active. (b) Jobs of coalition 1 are placed and machine 2 becomes inactive. (c) 3 out of 4 jobs of coalition 2 have been placed and $q = 2$ machines have become high. (d) The remaining job is placed on machine 3 making it low. Sufficiency conditions of Theorem [1](#) are now satisfied.

Theorem 2. *The above algorithm produces a partition equilibrium and a Nash equilibrium.*

Proof. The algorithm makes exactly q machines high hence due to the property of NE we know that there are sufficient jobs to make rest of the machines low, i.e., put $m_i - 1$ jobs on them. Consider a coalition C . If this coalition has only 0 or 1 jobs on every high machine then the conditions of Theorem [1](#) are fulfilled, and so C has no improving deviation. Consider a high machine i on which coalition C has more than one jobs. Let us look at the time-step when the algorithm has put α jobs of coalition C on i . Now before putting the $(\alpha + 1)$ 'th job on i the algorithm puts one job on every low machine. This is true because the low machines are exactly the ones that do not run out of space for jobs until the high machines are completely filled. This implies that for every low machine j , $x_j \geq x_i - 1$.

Also if a low machine j is such that $l_j > l_i$ then the algorithm puts one job on every such machine j before i . This follows from the ordering obtained on the machines on the basis of the l_i -values. This means that if $l_j > l_i$ then $x_j \geq x_i$.

This proves that both sufficient conditions of Theorem [1](#) are fulfilled by the final allocation, which is also a Nash equilibrium by Proposition [1](#). Hence the allocation obtained by the algorithm is a partition equilibrium. ■

5 Existence and Computation of SSE

For a resource selection game, SSE may or may not exist (see [4](#) for an example where it does not exist). In this section, we show that for a given instance of a resource selection game, we can efficiently determine whether there exists a SSE or not. We also give an algorithm that finds a SSE if it exists.

Theorem 3. *Given a resource selection game G , there is a polynomial time algorithm that returns a SSE allocation if it exists, and returns “no” if G does not have a SSE.*

Proof. Since every SSE is also a Nash equilibrium, then each type 2 machine i has to have exactly m_i jobs in any SSE allocation.

Recall that the total number of jobs in the system is exactly as much as $\sum_{i \in R} m_i + \sum_{i \in H} (m_i - 1) + \sum_{i \in L} (m_i - 1) + q$ in any Nash equilibrium allocation, with q being the number of high machines in any Nash equilibrium. If all type 1 machines are high machines, i.e., $L = \emptyset$, then all the machines in the system will have exactly m_i jobs and no coalition can have an improving deviation. This is because any non-trivial deviation would require moving a job so that its resulting latency is strictly more than u , and so cannot be an improving deviation. Therefore, if $L = \emptyset$ then any Nash equilibrium allocation is also a SSE allocation. So, a SSE allocation can be obtained simply by assigning m_i jobs to all machines.

Consider the case, where $L \neq \emptyset$. Assume that a high machine i has 2 or more jobs. Consider a coalition composed of 2 jobs such that both are allocated to i . If one of the jobs of this coalition moves to a low machine, then the cost of the moving job will not change, while the other member of the coalition strictly benefits. Therefore, an allocation where a high machine i has 2 or more jobs is not a SSE if $L \neq \emptyset$. Thus, G does not have a SSE if $L \neq \emptyset$ and G does not have at least q type 1 machines for which $m_i = 1$.

If there are at least q type 1 machines for which $m_i = 1$ then any Nash equilibrium allocation A where q of the type 1 machines, for which $m_i = 1$ are high machines, is a SSE. This is because no subset of players has more than 1 job on any high machine in A , and therefore no coalition has an improving deviation by Theorem [1](#). SSE allocation then can simply be obtained by placing 1 job on q machines for which $m_i = 1$ and assigning the remaining jobs to all other machines in a way that every type 2 machine has exactly m_i jobs and every remaining type 1 machine has exactly $m_i - 1$ jobs allocated to it. ■

References

1. Aumann, R.: Acceptable Points in General Cooperative n -person Games. In: Contributions to Theory of Games IV. Princeton Univ. Press, Princeton (1959)
2. Awerbuch, B., Azar, Y., Epstein, L.: The Price of Routing Unsplittable Flow. In: Proceedings of The Thirty-Seventh Annual ACM Symposium on Theory of Computing, pp. 57–66. ACM, New York (2005)
3. Durr, C., Kim Thang, N.: Non-clairvoyant Scheduling Games. In: Mavronicolas, M., Papadopoulos, V.G. (eds.) SAGT 2009. LNCS, vol. 5814, pp. 135–146. Springer, Heidelberg (2009)
4. Feldman, M., Tenneholtz, M.: Partition Equilibrium. In: Mavronicolas, M., Papadopoulos, V.G. (eds.) SAGT 2009. LNCS, vol. 5814, pp. 48–59. Springer, Heidelberg (2009)
5. Fiat, A., Kaplan, H., Levi, 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)
6. Fotakis, D., Kontogiannis, S., Spirakis, P.: Atomic Congestion Games Among Coalitions. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 572–583. Springer, Heidelberg (2006)
7. Hayrapetyan, A., Tardos, É., Wexler, T.: The Effect of Collusion in Congestion Games. In: Proceedings of The Thirty-Eighth Annual ACM Symposium on Theory of Computing, pp. 89–98. ACM, New York (2006)
8. Holzman, R., Law-Yone, N.: Strong Equilibrium in Congestion Games. *Games and Economic Behavior* 21, 85–101 (1997)
9. Holzman, R., Law-Yone, N.: Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences* 46, 193–205 (2003)
10. Immorlica, N., Li, L., Mirrokni, V., Schulz, A.: Coordination Mechanisms for Selfish Scheduling. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 55–69. Springer, Heidelberg (2005)
11. Monderer, D., Shapley, L.: Potential Games. *Games and Economic Behavior* 14, 124–143 (1996)
12. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
13. Roughgarden, T.: Selfish Routing with Atomic Players. In: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1184–1185. Society for Industrial and Applied Mathematics, Philadelphia (2005)
14. Suri, S., Toth, C., Zhou, Y.: Selfish Load Balancing and Atomic Congestion Games. *Algorithmica* 47(1), 79–96 (2007)

Mixing Time and Stationary Expected Social Welfare of Logit Dynamics*

Vincenzo Auletta, Diodato Ferraioli, Francesco Pasquale,
and Giuseppe Persiano

Dipartimento di Informatica e Applicazioni “Renato M. Capocelli”
{auletta,ferraioli,pasquale,giuper}@dia.unisa.it

Abstract. We study *logit dynamics* [Blume, *Games and Economic Behavior*, 1993] for strategic games. At every stage of the game a player is selected uniformly at random and she plays according to a *noisy* best-response dynamics where the noise level is tuned by a parameter β . Such a dynamics defines a family of ergodic Markov chains, indexed by β , over the set of strategy profiles. Our aim is twofold: On the one hand, we are interested in the expected social welfare when the strategy profiles are random according to the stationary distribution of the Markov chain, because we believe it gives a meaningful description of the long-term behavior of the system. On the other hand, we want to estimate how long it takes, for a system starting at an arbitrary profile and running the logit dynamics, to get close to the stationary distribution; i.e., the *mixing time* of the chain.

In this paper we study the stationary expected social welfare for the 3-player congestion game that exhibits the worst Price of Anarchy [Christodoulou and Koutsoupias, *STOC'05*], for 2-player coordination games (the same class of games studied by Blume), and for a simple n -player game. For all these games, we give almost-tight upper and lower bounds on the mixing time of logit dynamics.

1 Introduction

The evolution of a system is determined by its dynamics and complex systems are often described by looking at the equilibrium states induced by their dynamics. Once the system enters an equilibrium state, it stays there and thus it can be rightly said that an equilibrium state describes the long-term behavior of the system. In this paper we are mainly interested in *selfish* systems whose individual components are selfish agents. The state of a selfish system is fully described by a vector of *strategies*, each controlled by one agent, and each state assigns a payoff to each agent. The agents are selfish in the sense that they pick their strategy so to maximize their payoff, given the strategies of the other agents. The notion of a Nash equilibrium is the classical notion of equilibrium for selfish systems and it corresponds to the equilibrium induced by the *best-response* dynamics. The

* Work partially supported by EU project FRONTS and by MIUR PRIN project COGENT.

observation that selfish systems are described by their equilibrium states (that is, by the Nash equilibria) has motivated the notion of Price of Anarchy [15] (and Price of Stability [1]) and the efficiency analysis of selfish systems based on such notions.

The analysis based on Nash equilibria inherits some of the shortcomings of the concept of a Nash equilibrium. First of all, the best-response dynamics assumes that the selfish agents have complete knowledge of the current state of the system; that is, of the payoff associated with each possible choice and of the strategies chosen by the other agents. Instead, in most cases, agents have only an approximate knowledge of the system state. Moreover, in presence of multiple equilibria, it is not clear which equilibrium will be reached by the system as it may depend on the initial state of the system. The notion of Price of Anarchy solves this problem by considering the worst case equilibrium whereas Price of Stability focuses on the best case equilibrium. Finally, Nash equilibria are hard to compute [7,5] and thus for some system it might take very long to enter a Nash equilibrium. In this case using equilibrium states to describe the system performance is not well justified. Rather, one would like to analyze the performance of a system by using a dynamics (and its related equilibrium notion) that has the following three properties: the dynamics takes into account the fact that the system components might have a perturbed or noisy knowledge of the system; the equilibrium state exists and is unique for every system; independently from the starting state; the system enters the equilibrium very quickly.

In this paper, we consider *noisy best-response dynamics* in which the behavior of the agents is described by a parameter $\beta \geq 0$ (β is sometimes called the *inverse temperature*). The case $\beta = 0$ corresponds to agents picking their strategies completely at random (that is, the agents have no knowledge of the system) and the case $\beta = \infty$ corresponds to the best-response dynamics (in which the agents have full and complete knowledge of the system). The intermediate values of β correspond to agents that are roughly guided by the best-response dynamics but can make a sub-optimal response with some probability that depends on β (and on the associated payoff). We will study a specific noisy best-response dynamics for which the system evolves according to an ergodic Markov chain for all $\beta \geq 0$. For these systems, it is natural to look at the stationary distribution (which is the equilibrium state of the Markov chain) and to analyze the performance of the system at the stationary distribution. We stress that the noisy best-response dynamics well models agents that only have approximate or noisy knowledge of the system and that for ergodic Markov chains (such as the ones arising in our study) the stationary distribution is known to exist and to be unique. Moreover, to justify the use of the stationary distribution for analyzing the performance of the system, we will study how fast the Markov chain converges to the stationary distribution.

Related Works and Our Results. Several dynamics, besides the best response dynamics, and several notions of equilibrium, besides Nash equilibria, have been considered to describe the evolution of a selfish system and to analyze its performance. See, for example, [11,20,19].

Equilibrium Concepts Based on the Best-Response. In case the game does not possess a Pure Nash equilibrium, the best-response dynamics will eventually cycle over a set of states (in a Nash equilibrium the set is a singleton). These states are called *sink equilibria* [12]. Sink equilibria exist for all games and, in some context, they seem a better approximation of the real setting than *mixed* Nash equilibria. Unfortunately, sink equilibria share two undesirable properties with Nash equilibria: a game can have more than one sink equilibrium and sink equilibria seem hard to compute [9]. Other notions of equilibrium state associated with best-response dynamics are the *unit-recall equilibria* and *component-wise unit-recall equilibria* (see [9]). We point out though that the former does not always exist and that the latter imposes too strict limitations on the players.

No-Regret Dynamics. Another broadly explored set of dynamics are the no-regret dynamics (see, for example, [11]). The regret of a user is the difference between the long term average cost and average cost of the best strategy in hindsight. In the no-regret dynamics the regret of every player after t step is $o(t)$ (sublinear with time). In [10,14] it is showed that the no-regret dynamics converges to the set of Correlated Equilibria. Note that the convergence is to the *set* of Correlated Equilibria and not to a specific correlated equilibrium.

Our Work. In this paper we consider a specific noisy best-response dynamics called the *logit* dynamics (see [4]) and we study its mixing time (that is, the time it takes to converge to the stationary distribution) for various games. Specifically,

- We start by analyzing the logit dynamics for a simple 3-player linear congestion game (the CK game [6]) which exhibits the worst Price of Anarchy among linear congestion games. We show that the convergence time to stationarity of the logit dynamics is upper bounded by a constant independent of β . Moreover, we show that the expected social cost at stationarity is smaller than the cost of the worst Nash equilibrium for all β .
- We then study the 2×2 coordination games studied by [4]. Here we show that, under some conditions, the expected social welfare at stationarity is better than the social welfare of the worst Nash equilibrium. We give exponential in β upper and lower bounds on the convergence time to stationarity for all values of β .
- Finally, we apply our analysis to a simple n -player game, the OR-game, and give upper and lower bound on the convergence time to stationarity. In particular, we prove that for $\beta = \mathcal{O}(\log n)$ the convergence time is polynomial in n .

The *logit* dynamics has been first studied by Blume [4] who showed that, for 2×2 coordination games, the long-term behaviour of the Markov chain is concentrated in the risk dominant equilibrium (see [13]) for sufficiently large β . Ellison [8] studied different noisy best-response dynamics for 2×2 games and assumed that interaction among players were described by a graph; that is, the utility of a player is determined only by the strategies of the adjacent players. Specifically, Ellison [8] studied interaction modeled by rings and showed that some large fraction of the players will eventually choose the risk dominant strategy. Similar

results were obtained by Peyton Young [21] for the logit dynamics and for more general families of graphs. Montanari and Saberi [17] gave bounds on the hitting time of the risk dominant equilibrium states for the logit dynamics in terms of some graph theoretic properties of the underlying interaction network. Asadpour and Saberi [2] studied the hitting time for a class of congestion games. We notice that none of [4,8,21] gave any bound on the convergence time to the risk dominant equilibrium. Montanari and Saberi [17] were the first to do so but their study focuses on the hitting time of a specific configuration.

From a technical point of view, our work follows the lead of [4,8,21] and extends their technical findings by giving bounds on the mixing time of the Markov chain of the logit dynamics. We stress that previous results only proved that, for sufficiently large β , *eventually* the system concentrates around certain states without further quantifying the rate of convergence nor the asymptotic behaviour of the system for small values of β . Instead, we identify the stationary distribution of the logit dynamics as the *global* equilibrium and we evaluate the social welfare at stationarity and the time it takes the system to reach it (the mixing time) as explicit functions of the inverse temperature β of the system. For $\beta \rightarrow \infty$, the *logit* dynamics tends to the best response dynamics. It should come to no surprise that for large β the mixing time could be super-polynomial.

We choose to start our study from the class of coordination games considered in [4] for which we give tight upper and lower bound on the mixing time and then look also at other 2-player games and a simple n -player game (the OR-game). Despite its game-theoretic simplicity, the analytical study of the mixing time of the Markov chain associated with the OR-game as a function of β is far from trivial. Also we notice that the results of [17] cannot be used to derive upper bounds on the mixing time as in [17] the authors give a tight estimation of the hitting time only for a specific state of the Markov chain. The mixing time instead is upper bounded by the *maximum* hitting time.

From a more conceptual point of view, our work tries (similarly to [12,9,18]) to introduce a solution concept that well models the behaviour of selfish agents, is uniquely defined for any game and is quickly reached by the game. We propose the stationary distribution induced by the logit dynamics as a possible solution concept and exemplify its use in the analysis of the performance of some 2×2 games (as the ones considered in [4,8,21]), in games used to obtain tight bounds on the Price of Anarchy and on a simple multiplayer game.

Organization of the Paper. In Section 2 we formally describe the logit dynamics Markov chain for a strategic game. In Sections 3, 4, and 5 we study the stationary expected social welfare and the mixing time of the logit dynamics for CK game, coordination games, and the OR-game respectively. Due to lack of space, the proofs are omitted and are available in the full version [3]. Finally, in Section 6 we present conclusions and some open problems.

Notation. We write \bar{S} for the complementary set of a set S and $|S|$ for its size. We use bold symbols for vectors, when $\mathbf{x} = (x_1, \dots, x_n) \in \{0, 1\}^n$ we write $|\mathbf{x}|$ for the number of 1s in \mathbf{x} ; i.e., $|\mathbf{x}| = |\{i \in [n] : x_i = 1\}|$. We use the standard game theoretic notation (\mathbf{x}_{-i}, y) to mean the vector obtained from \mathbf{x} by replacing the

i -th entry with y , i.e. $(\mathbf{x}_{-i}, y) = (x_1, \dots, x_{i-1}, y, x_{i+1}, \dots, x_n)$. We use standard Markov chain terminology (see [16]).

2 The Model and the Problem

A *strategic game* is a triple $([n], \mathcal{S}, \mathcal{U})$, where $[n] = \{1, \dots, n\}$ is a finite set of *players*, $\mathcal{S} = \{S_1, \dots, S_n\}$ is a family of non-empty finite sets (S_i is the set of *strategies* for player i), and $\mathcal{U} = \{u_1, \dots, u_n\}$ is a family of *utility functions* (or *payoffs*), where $u_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$ is the utility function of player i .

Consider the following *noisy* best-response dynamics, introduced in [4] and known as *logit dynamics*: At every time step

1. Select one player $i \in [n]$ uniformly at random;
2. Update the strategy of player i according to the following probability distribution over the set S_i of her strategies. For every $y \in S_i$

$$\sigma_i(y | \mathbf{x}) = \frac{1}{T_i(\mathbf{x})} e^{\beta u_i(\mathbf{x}_{-i}, y)} \quad (1)$$

where $\mathbf{x} \in S_1 \times \dots \times S_n$ is the strategy profile played at the current time step, $T_i(\mathbf{x}) = \sum_{z \in S_i} e^{\beta u_i(\mathbf{x}_{-i}, z)}$ is the normalizing factor, and $\beta \geq 0$ is the *inverse noise*.

From (1) it is easy to see that, for $\beta = 0$ player i selects her strategy uniformly at random, for $\beta > 0$ the probability is biased toward strategies promising higher payoffs, and for $\beta \rightarrow \infty$ player i chooses her best response strategy (if more than one best response is available, she chooses uniformly at random one of them). Moreover observe that probability $\sigma_i(y | \mathbf{x})$ does not depend on the strategy x_i currently adopted by player i .

The above dynamics defines an ergodic finite Markov chain with the set of strategy profiles as state space, and where the transition probability from profile $\mathbf{x} = (x_1, \dots, x_n)$ to profile $\mathbf{y} = (y_1, \dots, y_n)$ is zero if the two profiles differ at more than one player and it is $\frac{1}{n} \sigma_i(y_i | \mathbf{x})$ if the two profiles differ exactly at player i . More formally, we have the following definition.

Definition 1 (Logit dynamics [4]). Let $\mathcal{G} = ([n], \mathcal{S}, \mathcal{U})$ be a strategic game and let $\beta \geq 0$ be the inverse noise. The logit dynamics for \mathcal{G} is the Markov chain $\mathcal{M}_\beta = \{X_t : t \in \mathbb{N}\}$ with state space $\Omega = S_1 \times \dots \times S_n$ and transition matrix

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i=1}^n \frac{e^{\beta u_i(\mathbf{x}_{-i}, y_i)}}{T_i(\mathbf{x})} \mathbb{I}_{\{y_j = x_j \text{ for every } j \neq i\}}. \quad (2)$$

It is easy to see that, if $([n], \mathcal{S}, \mathcal{U})$ is a potential game with exact potential Φ , then the Markov chain given by (2) is reversible and its stationary distribution is the Gibbs measure

$$\pi(\mathbf{x}) = \frac{1}{Z} e^{\beta \Phi(\mathbf{x})} \quad (3)$$

where $Z = \sum_{\mathbf{y} \in S_1 \times \dots \times S_n} e^{\beta \Phi(\mathbf{y})}$ is the normalizing constant (the *partition function* in physicists' language). Except for the Matching Pennies example in Subsection 2.1, all the games we analyse in this paper are potential games.

Let $W : S_1 \times \dots \times S_n \rightarrow \mathbb{R}$ be a *social welfare function* (in this paper we assume that W is simply the sum of all the utility functions $W(\mathbf{x}) = \sum_{i=1}^n u_i(\mathbf{x})$, but clearly any other function of interest can be analysed). We study the *stationary expected social welfare*, i.e. the expectation of W when the strategy profiles are random according to the stationary distribution π of the Markov chain,

$$\mathbf{E}_\pi [W] = \sum_{\mathbf{x} \in S_1 \times \dots \times S_n} W(\mathbf{x}) \pi(\mathbf{x}).$$

Since the Markov chain defined in (2) is irreducible and aperiodic, from every initial profile \mathbf{x} the distribution $P^t(\mathbf{x}, \cdot)$ of chain X_t starting at \mathbf{x} will eventually converge to π as t tends to infinity. We will be interested in the *mixing time* t_{mix} of the chain, i.e. the time needed to have that $P^t(\mathbf{x}, \cdot)$ is *close* to π for every initial configuration \mathbf{x} . More formally, we define

$$t_{\text{mix}}(\varepsilon) = \min_{t \in \mathbb{N}} \max_{\mathbf{x} \in \Omega} \{ \|P^t(\mathbf{x}, \cdot) - \pi\|_{\text{TV}} \leq \varepsilon \}$$

where $\|P^t(\mathbf{x}, \cdot) - \pi\|_{\text{TV}} = \frac{1}{2} \sum_{\mathbf{y} \in \Omega} |P^t(\mathbf{x}, \mathbf{y}) - \pi(\mathbf{y})|$ is the *total variation distance*, and we set $t_{\text{mix}} = t_{\text{mix}}(1/4)$.

2.1 An Example: Matching Pennies

As an example consider the classical *Matching Pennies* game:

	H	T	
H	+1, -1	-1, +1	(4)
T	-1, +1	+1, -1	

The update probabilities (1) for the logit dynamics are, for every $x \in \{H, T\}$

$$\sigma_1(H | (x, H)) = \sigma_1(T | (x, T)) = \frac{1}{1+e^{-2\beta}} = \sigma_2(T | (H, x)) = \sigma_2(H | (T, x))$$

$$\sigma_1(T | (x, H)) = \sigma_1(H | (x, T)) = \frac{1}{1+e^{2\beta}} = \sigma_2(H | (H, x)) = \sigma_2(T | (T, x)).$$

So the transition matrix (2) is

$$P = \left(\begin{array}{c|cccc} & HH & HT & TH & TT \\ \hline HH & 1/2 & b/2 & (1-b)/2 & 0 \\ HT & (1-b)/2 & 1/2 & 0 & b/2 \\ TH & b/2 & 0 & 1/2 & (1-b)/2 \\ TT & 0 & (1-b)/2 & b/2 & 1/2 \end{array} \right)$$

where we named $b = \frac{1}{1+e^{-2\beta}}$ for readability sake.

Since every column of the matrix adds up to 1, the uniform distribution π over the set of strategy profiles is the stationary distribution for the logit dynamics. The expected stationary social welfare is thus 0 for every inverse noise β .

As for the mixing time, it is easy to see that it is upper bounded by a constant independent of β . Indeed, a direct calculation shows that, for every $\mathbf{x} \in \{HH, HT, TH, TT\}$ and for every $\beta \geq 0$, it holds that

$$\|P^3(\mathbf{x}, \cdot) - \pi\|_{\text{TV}} \leq \frac{7}{16} < \frac{1}{2}.$$

3 Warm Up: A 3-Player Congestion Game

In this section we study the CK game, a simple 3-player linear congestion game introduced in [6] that exhibits the worst Price of Anarchy of the average social welfare among linear congestion games with 3 or more players. This game has two equilibria: one with social welfare -6 (which is also optimal) and one with social welfare -15 . As we shall see briefly, the stationary expected social welfare of the logit dynamics is always larger than the social welfare of the worst Nash equilibrium and, for large enough β , players spend most of the time in the best Nash equilibrium. Moreover, we will show that the mixing time of the logit dynamics is bounded by a constant independent of β ; that is, the stationary distribution guarantees a good social welfare and it is quickly reached by the system.

Let us now describe the CK game. We have 3 players and 6 facilities divided into two sets: $G = \{g_0, g_1, g_2\}$ and $H = \{h_0, h_1, h_2\}$. Player $i \in \{0, 1, 2\}$ has two strategies: Strategy “0” consists in selecting facilities (g_i, h_i) ; Strategy “1” consists in selecting facilities $(g_{i+1}, h_{i-1}, h_{i+1})$ (index arithmetic is modulo 3). The cost of a facility is the number of players choosing such facility, and the cost of a player is the sum of the costs of the facilities she selected. It is easy to see that this game has two pure Nash equilibria: when every player plays strategy 0 (each player pays 2, which is optimal), and when every player plays strategy 1 (each player pays 5). The game is a congestion game, and thus a potential game with following potential function:

$$\Phi(\mathbf{x}) = \sum_{j \in GUH} \sum_{i=1}^{L_{\mathbf{x}}(j)} i$$

where $L_{\mathbf{x}}(j)$ is the number of players using facility j in configuration \mathbf{x} .

Stationary Expected Social Welfare and Mixing Time. The logit dynamics for the CK game gives the following update probabilities (see Equation (II))

$$\begin{aligned} \sigma_i(0 \mid \mathbf{x}_{-i} = 0) &= \frac{1}{1+e^{-4\beta}} & \sigma_i(1 \mid \mathbf{x}_{-i} = 0) &= \frac{1}{1+e^{4\beta}} \\ \sigma_i(0 \mid \mathbf{x}_{-i} = 1) &= \frac{1}{1+e^{-2\beta}} & \sigma_i(1 \mid \mathbf{x}_{-i} = 1) &= \frac{1}{1+e^{2\beta}} \\ \sigma_i(0 \mid \mathbf{x}_{-i} = 2) &= \frac{1}{2} & \sigma_i(1 \mid \mathbf{x}_{-i} = 2) &= \frac{1}{2}. \end{aligned}$$

It is easy to check that the following distribution is stationary for the logit dynamics:

$$\begin{aligned}\pi[(0, 0, 0)] &= \frac{e^{-6\beta}}{Z(\beta)} \\ \pi[(0, 0, 1)] &= \pi[(0, 1, 0)] = \pi[(1, 0, 0)] = \frac{e^{-10\beta}}{Z(\beta)} \\ \pi[(0, 1, 1)] &= \pi[(1, 1, 0)] = \pi[(1, 0, 1)] = \pi[(1, 1, 1)] = \frac{e^{-12\beta}}{Z(\beta)}\end{aligned}$$

where $Z(\beta) = e^{-6\beta} + 3e^{-10\beta} + 4e^{-12\beta}$. Let k be the number of players playing strategy 1; the social welfare is -6 when $k = 0$, it is -13 if $k = 1$, it is -16 if $k = 2$, and -15 when $k = 3$. Thus the stationary expected social welfare is

$$\mathbf{E}_\pi[W] = -\frac{6e^{-6\beta} + 39e^{-10\beta} + (48 + 15)e^{-12\beta}}{e^{-6\beta} + 3e^{-10\beta} + 4e^{-12\beta}} = -\frac{3[2 + 13e^{-4\beta} + 21e^{-6\beta}]}{1 + 3e^{-4\beta} + 4e^{-6\beta}}.$$

For $\beta = 0$, we have $\mathbf{E}_\pi[W] = -27/2$ which is better than the social welfare of the worst Nash equilibrium. As β tends to ∞ , $\mathbf{E}_\pi[W]$ approaches the optimal social welfare. Furthermore, we observe that $\mathbf{E}_\pi[W]$ increases with β and thus we can conclude that the long-term behavior of the logit dynamics gives a better social welfare than the worst Nash equilibrium for any $\beta \geq 0$.

Theorem 1 (Mixing time of CK game). *There exists a constant τ such that the mixing time t_{mix} of the logit dynamics of the CK game is upper bounded by τ for every $\beta \geq 0$.*

4 Coordination Games

Coordination Games are two-player games in which the players have an advantage in selecting the same strategy. They are often used to model the spread of a new technology [21]: two players have to decide whether to adopt or not a new technology. We assume that the players would prefer choosing the same technology and that choosing the new technology is risk dominant.

We analyse the mixing time of the logit dynamics for 2×2 coordination games and compute the stationary expected social welfare of the game as a function of β . We show that, for large enough β , players will spend most of the time in the risk dominant equilibrium and the expected utility is better than the one associated with the worst Nash equilibrium. Similar results can be obtained for anti coordination games (see [3]).

We denote by 0 the NEW strategy and by 1 the OLD strategy. The game is formally described by the following payoff matrix

	0	1	
0	(a, a)	(c, d)	(5)
1	(d, c)	(b, b)	

We assume that $a > d$ and $b > c$ (meaning that they prefer to coordinate) and that $a - d > b - c$ (meaning that strategy 0 is the risk dominant strategy [11])

for each player). Notice that we do not make any assumption on the relation between a and b . It is easy to see that this game is a potential game and the following function is an exact potential for it:

$$\Phi(0,0) = a - d \quad \Phi(0,1) = \Phi(1,0) = 0 \quad \Phi(1,1) = b - c.$$

This game has two pure Nash equilibria: $(0,0)$, where each player has utility a , and $(1,1)$, where each player has utility b . As $d + c < a + b$, the social welfare is maximized in correspondence of one of the two equilibria and the Price of Anarchy is equal to $\max\{b/a, a/b\}$.

Stationary Expected Social Welfare and Mixing Time. The logit dynamics for the game defined by the payoffs in Table 5 gives the following update probabilities for any strategy $x \in \{0,1\}$ (see Equation (11))

$$\begin{aligned} \sigma_1(0 | (x,0)) = \sigma_2(0 | (0,x)) &= \frac{1}{1+e^{-(a-d)\beta}} & \sigma_1(1 | (x,0)) = \sigma_2(1 | (0,x)) &= \frac{1}{1+e^{(a-d)\beta}} \\ \sigma_1(0 | (x,1)) = \sigma_2(0 | (1,x)) &= \frac{1}{1+e^{(b-c)\beta}} & \sigma_1(1 | (x,1)) = \sigma_2(1 | (1,x)) &= \frac{1}{1+e^{-(b-c)\beta}}. \end{aligned}$$

Theorem 2 (Expected social welfare). *The stationary expected social welfare $\mathbf{E}_\pi[W]$ of the logit dynamics for the coordination game is*

$$\mathbf{E}_\pi[W] = 2 \cdot \frac{a + be^{-((a-d)-(b-c))\beta} + (c+d)e^{-(a-d)\beta}}{1 + e^{-((a-d)-(b-c))\beta} + 2e^{-(a-d)\beta}}.$$

The following observation gives conditions on β and the players' utility for which the expected social welfare $\mathbf{E}_\pi[W]$ obtained by the logit dynamics is better than the social welfare SW_N of the worst Nash Equilibrium.

Observation 2. *For the coordination game described in Table 5, we have*

- if $a > b$ and $b \leq \max\{\frac{a+c+d}{3}, \frac{c+d}{2}\}$ then $\mathbf{E}_\pi[W] > \text{SW}_N$ for all β ;
- if $a > b$ and $b > \max\{\frac{a+c+d}{3}, \frac{c+d}{2}\}$ then $\mathbf{E}_\pi[W] > \text{SW}_N$ for all sufficiently large β ;
- if $a < b$ and $a \leq \max\{\frac{b+c+d}{3}, \frac{c+d}{2}\}$ then $\mathbf{E}_\pi[W] > \text{SW}_N$ for all β ;
- if $a < b$ and $a > \max\{\frac{b+c+d}{3}, \frac{c+d}{2}\}$ then $\mathbf{E}_\pi[W] > \text{SW}_N$ for all sufficiently large β ;
- if $a = b$ then $\mathbf{E}_\pi[W] < \text{SW}_N$ for any β, a, c and d .

Theorem 3 (Mixing Time of Coordination Games). *The mixing time of the logit dynamics with parameter β for the coordination game of Table 5 is $\Theta(e^{(b-c)\beta})$.*

5 A Simple n -Player Game: OR-Game

In this section we consider the following simple n -player potential game that we here call *OR-game*. For the upper bound we use the path coupling technique on the Hamming graph with carefully chosen edge weights. Every player has two

strategies, say $\{0, 1\}$, and each player pays the OR of the strategies of all players (including herself). More formally, the utility function of player $i \in [n]$ is

$$u_i(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{x} = \mathbf{0}; \\ -1, & \text{otherwise.} \end{cases}$$

Notice that the OR-game has $2^n - n$ Nash equilibria. The only profiles that are not Nash equilibria are the n profiles with exactly one player playing 1. Nash equilibrium $\mathbf{0}$ has social welfare 0, while all the others have social welfare $-n$. Despite its simplicity, the analysis of the mixing time is far from trivial (see full version [3]).

In Theorem 4 we show that the stationary expected social welfare is always better than the social welfare of the worst Nash equilibrium, and it is *significantly* better for large β . Unfortunately, in Theorem 5 we show that, if β is large enough to guarantee a *good* stationary expected social welfare, then the time needed to get close to the stationary distribution is exponential in n . Finally, in Theorem 6 we give upper bounds on the mixing time showing that, if β is relatively small then the mixing time is polynomial in n , while for large β the upper bound is exponential in n and it is almost-tight with the lower bound.

Theorem 4 (Expected social welfare). *The stationary expected social welfare of the logit dynamics for the OR-game is $\mathbf{E}_\pi[W] = -\alpha n$ where $\alpha = \alpha(n, \beta) = \frac{(2^n - 1)e^{-\beta}}{1 + (2^n - 1)e^{-\beta}}$.*

In the next theorem we show that the mixing time can be polynomial in n only if $\beta \leq c \log n$ for some constant c .

Theorem 5 (Lower bound on mixing time). *The mixing time of the logit dynamics for the OR-game is*

1. $\Omega(e^\beta)$ if $\beta < \log(2^n - 1)$;
2. $\Omega(2^n)$ if $\beta > \log(2^n - 1)$.

In the next theorem we give upper bounds on the mixing time depending on the value of β . The theorem shows that, if $\beta \leq c \log n$ for some constant c , the mixing time is effectively polynomial in n with degree depending on c . The use of the path coupling technique in the proof of the theorem requires a careful choice of the edge-weights.

Theorem 6 (Upper bound on mixing time). *The mixing time of the logit dynamics for the OR-game is*

1. $\mathcal{O}(n \log n)$ if $\beta < (1 - \varepsilon) \log n$, for an arbitrary small constant $\varepsilon > 0$;
2. $\mathcal{O}(n^{c+3} \log n)$ if $\beta \leq c \log n$, where $c \geq 1$ is an arbitrary constant.

Moreover the mixing time is $\mathcal{O}(n^{5/2} 2^n)$ for every β .

6 Conclusions and Open Problems

In this paper we studied strategic games where at every run a player is selected uniformly at random and she is assumed to choose her strategy for the next run according to a *noisy best-response*, where the noise level is tuned by a parameter β . Such dynamics defines a family of ergodic Markov chains, indexed by β , over the set of strategy profiles. We study the long-term behavior of the system by analysing the expected social welfare when the strategy profiles are random according to the stationary distribution of such chains, and we compare it with the social welfare at Nash equilibria.

In order for such analysis to be meaningful we are also interested in the *mixing time* of the chains, i.e. how long it takes, for a chain starting at an arbitrary profile, to get close to its stationary distribution. The analysis of the mixing time is usually far from trivial even for very simple games.

We study several examples of applications of this approach to games with two and three players and to a simple n -players game. We started by showing that the social welfare at stationarity for the 3-player linear congestion game that attains the maximum Price of Anarchy is larger than the social welfare of the worst Nash equilibrium. This result is made significant by the fact that, for all β , the logit dynamics converges at the stationary distribution in constant time. For 2-player coordination games the mixing time turns out to be exponential in β and we give conditions for the expected social welfare at stationarity to be smaller than the social welfare of the worst Nash equilibrium. In the n -player OR-game, the mixing time is $\mathcal{O}(n \log n)$ for β up to $\log n$; if $\beta < c \log n$ with $c > 1$ constant, the mixing time is polynomial in n with the degree depending on the constant c ; finally, for large β the mixing time is exponential in n .

We leave several questions for further investigation. For example, we would like to close gaps between upper and lower bounds for the mixing time of the OR-game. Moreover, we would like to investigate logit dynamics for notable classes of n -player games.

Acknowledgements

We thank Paolo Penna and Carmine Ventre for helpful discussions and pointers to the literature.

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. *SIAM Journal on Computing* 38(4), 1602–1623 (2008)
2. Asadpour, A., Saberi, A.: On the inefficiency ratio of stable equilibria in congestion games. In: Leonardi, S. (ed.) *WINE 2009*. LNCS, vol. 5929, pp. 545–552. Springer, Heidelberg (2009)
3. Auletta, V., Ferraioli, D., Pasquale, F., Persiano, G.: Mixing time and stationary expected social welfare of logit dynamics. *CoRR* abs/1002.3474 (2010)

4. Blume, L.E.: The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5, 387–424 (1993)
5. Chen, X., Deng, X., Teng, S.-H.: Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM* 56(3) (2009)
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. ACM, New York (2005)
7. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. *SIAM Journal on Computing* 39(1), 195–259 (2009)
8. Ellison, G.: Learning, local interaction, and coordination. *Econometrica* 61(5), 1047–1071 (1993)
9. Fabrikant, A., Papadimitriou, C.H.: The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond. In: *Proc. of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008)*, pp. 844–853. ACM, New York (2008)
10. Foster, D.P., Vohra, R.V.: Calibrated learning and correlated equilibrium. *Games and Economic Behavior* 21(1-2), 40–55 (1997)
11. Fudenberg, D., Levine, D.K.: *The Theory of Learning in Games*. MIT, Cambridge (1998)
12. Goemans, M., Mirrokni, V., Vetta, A.: Sink equilibria and convergence. In: *Proc. of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005)*, pp. 142–154. IEEE, Los Alamitos (2005)
13. Harsanyi, J.C., Selten, R.: *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge (1988)
14. Hart, S., Mas-Colell, A.: A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5), 1127–1150 (2000)
15. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. *Computer Science Review* 3(2), 65–69 (2009); Preliminary version in STACS 1999
16. Levin, D., Peres, Y., Wilmer, E.L.: *Markov Chains and Mixing Times*. American Mathematical Society, Providence (2008)
17. Montanari, A., Saberi, A.: Convergence to equilibrium in local interaction games. In: *Proc. of the 50th Annual Symposium on Foundations of Computer Science (FOCS 2009)*. IEEE, Los Alamitos (2009)
18. Roth, A., Balcan, N., Kalai, A., Mansour, Y.: On the equilibria of alternating move games. In: *Proc. of ACM-SIAM Symposium on Discrete Algorithms, SODA 2010* (2010)
19. Sandholm, W.H.: *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge (2010)
20. Peyton Young, H.: *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, Princeton (1998)
21. Peyton Young, H.: The diffusion of innovations in social networks. *Economics Working Paper Archive* number 437, Johns Hopkins University, Department of Economics (2000)

Pareto Efficiency and Approximate Pareto Efficiency in Routing and Load Balancing Games

Yonatan Aumann and Yair Dombb

Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel
aumann@cs.biu.ac.il, yair.biu@gmail.com

Abstract. We analyze the Pareto efficiency, or inefficiency, of solutions to routing games and load balancing games, focusing on Nash equilibria and greedy solutions to these games. For some settings, we show that the solutions are necessarily Pareto optimal. When this is not the case, we provide a measure to *quantify* the distance of the solution from Pareto efficiency. Using this measure, we provide upper and lower bounds on the “Pareto inefficiency” of the different solutions. The settings we consider include load balancing games on identical, uniformly-related, and unrelated machines, both using pure and mixed strategies, and nonatomic routing in general and some specific networks.

1 Introduction

Efficiency, and the efficient utilization of resources, is a key interest in economics. Efficiency can be defined in many ways, depending on the situation and goals, but perhaps one of the most rudimentary and basic efficiency notions is that of Pareto Efficiency. Pareto efficiency captures the idea that an outcome is clearly inefficient if it is possible to achieve an improvement “on all fronts” simultaneously; for example, in game theory an outcome of a game is (weakly) Pareto optimal if there is no other outcome in which *all players* are (strictly) better off. Unfortunately, it is well known that strategic behavior by players can frequently lead to Pareto inefficient outcomes, such as in the famous Prisoner’s Dilemma. Thus, Nash equilibrium may be Pareto inefficient.

In this work, we study the Pareto efficiency, or inefficiency, of two well known games: routing games and load balancing games (also known as *job scheduling games*). These games have received a lot of attention in the past decade, mainly in the context of the Price of Anarchy and the Price of Stability, measures that quantify the loss in social welfare due to selfishness and inability of players to coordinate. We analyze these games with respect to the Pareto efficiency of solutions to the games. Specifically, we focus on Nash equilibria and greedy solutions, and analyze their Pareto efficiency. In some cases we can show that the solutions are necessarily Pareto optimal. When this is not the case, we wish to *quantify* how far the solution is from Pareto efficiency, since it would be different if all players can improve their outcome ten-fold or just by 10%. Thus, we introduce the notion of *approximate Pareto efficiency*, defined shortly. With

this definition in hand, we show that while Pareto optimality is not always guaranteed, the inefficiency in the settings we consider can frequently be bounded by a constant.

Approximate-Pareto-Efficiency. We now present the formal definition for quantifying the distance of an outcome from Pareto efficiency. Conceptually, an outcome is α -Pareto-deficient if there is a different outcome which improves all players by at least an α factor.

Definition 1. Let G be a game with a set P of players, and A a possible outcome of G . We denote by $\text{cost}(i, A)$ the cost of player i in the outcome A ¹. For outcomes A, A' , we say that A' α -Pareto-dominates A if it holds that

$$\forall i \in P : \quad \alpha \cdot \text{cost}(i, A') \leq \text{cost}(i, A) .$$

We say that A is α -Pareto-deficient if there exists an alternative outcome A' of G that α -Pareto-dominates A .

We say that outcome A is α -Pareto-efficient (α -PE) if it is not β -Pareto-deficient for any $\beta > \alpha$.

Thus, in an α -Pareto-deficient outcome, all players can simultaneously improve their outcome by a factor of at least α . In an α -Pareto-efficient outcome, it is impossible to improve all players simultaneously by more than α . Note that for $\alpha = 1$, 1-Pareto-efficient coincides with Pareto optimality.

This Work. As mentioned, in this work we consider routing and load balancing games, with several flavors of each. For each class of games, we consider the following issues:

1. **Bounding the Pareto inefficiency of *any* Nash equilibrium:** we seek the smallest possible α such that *every* Nash equilibrium in any game of the class is α -Pareto-efficient.
2. **Bounding the Pareto inefficiency of the “best” Nash equilibrium:** we seek the smallest possible α such that for any game in the class there *exists* a Nash equilibrium that is α -Pareto-efficient.
3. **Bounding the Pareto inefficiency of a greedy assignment process:** The greedy solution is defined as follows. Assume that the players are (arbitrarily) ordered, and each player, in its turn, chooses a strategy that minimize her cost at the time of choosing (ties are broken arbitrarily). We seek the *smallest* α such that *every* outcome achieved by a greedy solution is α -Pareto-efficient.

Results. We consider selfish load balancing and selfish routing games. For load balancing games we consider the settings of identical machines, uniformly-related machines, and unrelated machines. In addition, we consider both the case where only pure strategies are permitted and the case that mixed strategies are also allowed. We obtain:

¹ Due to the nature of the routing and load balancing games we consider, we use a cost formulation of the notions. Analogous definitions can be defined for value/utility.

- Pure strategies only: If only pure strategies are allowed, any Nash equilibrium is necessarily Pareto optimal for both identical and uniformly-related machines. For unrelated machines, the Pareto-deficiency of a Nash equilibrium can be arbitrarily large, but there necessarily exists a Nash equilibrium which is Pareto optimal.

The greedy solution is Pareto optimal for identical machines, and necessarily 2-Pareto-efficient for uniformly-related machines. We were unable to show a bound on the Pareto-deficiency of the greedy solution for unrelated machines, but it can be shown that the upper bound of 2 does not hold for this case (i.e. there are cases in which the Pareto-deficiency of the greedy solution is strictly larger than 2).

- Mixed strategies: If mixed strategies are allowed, then on identical machines any Nash equilibrium is necessarily $(2 - \frac{1}{m})$ -Pareto-efficient, where m is the number of machines. This bound is tight, in the sense that for any m , there exists a setting with m machines that exhibits a Nash equilibrium which is $(2 - \frac{1}{m})$ -Pareto-deficient. For uniformly-related machines with mixed strategies, we show that any Nash equilibrium is necessarily 4-Pareto-efficient. We do not know to say if this bound is tight, and suspect that it is not. For unrelated machines, the worst Nash can be arbitrarily Pareto-deficient. For the best Nash equilibrium in mixed strategies, we do not have any tight bounds (of course, the upper bounds for the worst Nash apply for the best Nash as well). The greedy process is not well defined for such strategies.

For selfish routing games we consider the case of nonatomic games with monotone cost functions. We show:

- For general networks, for any family of cost functions, the Pareto efficiency of any Nash equilibrium is necessarily bounded by the Price of Anarchy for this class of functions. This bound is tight, in the sense that there exists a game for which the only Nash equilibrium exhibits this level of Pareto-deficiency. Hence, the same bound also holds for the best Nash.
- For the special case of networks with only parallel edges between a single source and a single sink (which we call *parallel-edge networks*), we show that any equilibrium is Pareto Optimal. Also, any greedy solution is necessarily Pareto optimal, as is any solution that uses all edges.

The results are summarized in Table [1](#). Unfortunately, due to the strict page limit in these proceedings, most of the proofs are omitted from this extended abstract. They all appear in the full version of the paper.

1.1 Related Work

Pareto efficiency is a desirable property for solutions of games. In cooperative games, such as in Nash's bargaining game [\[13\]](#), it is usually *required* that solutions be Pareto optimal. In non-cooperative game theory, it is well known that Nash equilibria are frequently Pareto inefficient, as illustrated by the famous prisoner's dilemma. Several works aimed at developing a deeper understanding

Table 1. Summary of results (PO stands for Pareto Optimal, and PE for Pareto Efficient)

Setting	Any Nash	Best Nash	Greedy	Section
Routing – General Networks	Equals the POA for the class		<i>open</i>	2.1
Routing – Parallel-Edge Networks	PO		PO	2.2
Load Balancing (Pure) – Identical Machines	PO		PO	3.1
Load Balancing (Pure) – Uniformly-Related Machines	PO		2-PE	3.1
Load Balancing (Pure) – Unrelated Machines	∞	PO	<i>open</i>	3
Load Balancing (Mixed) – Identical Machines	$(2 - \frac{1}{m})$ -PE	<i>open</i>	<i>N/A</i>	4
Load Balancing (Mixed) – Uniformly-Related Machines	4-PE (<i>not tight</i>)	<i>open</i>	<i>N/A</i>	4
Load Balancing (Mixed) – Unrelated Machines	∞	<i>open</i>	<i>N/A</i>	3

of this phenomenon. Examples include [\[5\]](#), which gives sufficient conditions for inefficiency of equilibria, [\[3\]](#), which computes the probability of inefficient (pure) Nash equilibria in finite random games, and [\[8,12\]](#), which consider the Pareto optimality of different social choice rules.

Pareto efficient solutions are also sought in multi-objective optimization problems. In this case, the *Pareto front* is defined as the set of solutions from which not all objectives can be improved simultaneously. Several works (for example [\[11,15,4,10\]](#)) have considered various approximation notions of the Pareto front, by additive or multiplicative terms, and provided algorithms for finding such solution sets.

Also related is the line of research on the Price of Anarchy [\[9\]](#) and the Price of Stability [\[1\]](#). The Price of Anarchy bounds the distance of any Nash equilibrium from an optimal outcome, defined using a social welfare function. Likewise, the Price of Stability bounds the distance of the “best” Nash equilibrium from the optimal social welfare. Some of the issues we consider in this work (namely the Pareto inefficiency of any/the “best” equilibrium) resemble these concepts, although our Paretian efficiency concept is distinct from social welfare efficiency, and cannot be expressed using any real-valued social welfare (or cost) function. It is worth noting that if the utilitarian social welfare function is considered, it can be shown that the Pareto-deficiency of the worst and best equilibria provide lower bounds for the POA and POS (resp.). The same holds for the egalitarian social welfare function, in the case that only pure strategies are allowed.

Finally, while our “worst/best Nash” questions cannot be expressed as special cases of the classical POA/POS, they can be formulated using the IR_{min} measure, presented by Feldman and Tamir in [\[7\]](#). Using the notation therein, the worst Nash is α -Pareto-deficient iff $\alpha = \sup_{E \in \mathcal{E}} IR_{min}(E, P)$, and the best

Nash is β -Pareto-deficient iff $\beta = \inf_{E \in \mathcal{E}} IR_{min}(E, P)$, where P is the set of all players and \mathcal{E} the set of equilibria. However, while the results in [7] aim at bounding the simultaneous improvement of the players in *any* possible coalition, we focus on the special case of the “grand coalition”, involving all the players. In addition, we consider routing games, and various settings in load balancing games (including mixed strategies, uniformly-related machines and unrelated machines), whereas [7] focuses on load balancing games with pure strategies on identical machines.

2 Selfish Routing Games

A *multi-commodity network* is a directed multigraph $N = (V, E)$ (possibly containing parallel edges) together with a collection $\{(s_1, t_1), \dots, (s_k, t_k)\} \subseteq V \times V$ of source-sink vertex pairs, called *commodities*. We denote the set of edges E by $[m]$ (where $[m] = \{1, \dots, m\}$), and with each edge $j \in [m]$ we associate a *cost function* $c_j(\cdot)$, and denote $c = (c_1(\cdot), \dots, c_m(\cdot))$. (We assume throughout that $c_j(\cdot)$ is continuous for all j ; for the results of Section 2.2 we additionally assume that $c_j(\cdot)$ is nondecreasing.) Finally, for each commodity i there is some amount r_i of traffic that needs to be routed from s_i to t_i . Thus, a *multi-commodity selfish routing game* is simply a triple (N, r, c) .

The players in a selfish routing game are infinitesimally-small “traffic units” that make independent routing decisions, possibly using different paths to go from the commodity’s source to its sink. A *flow* f in (N, r, c) is a vector, indexed by all the $s_i - t_i$ paths for all i , indicating the amount of traffic using each path. We denote by f_j the total amount of traffic traversing edge j . We say that a flow f is *feasible* if for every i , it routes an amount r_i of traffic from s_i to t_i . The cost incurred to a player p using a path P in the flow f is simply $\text{cost}(p, f) = \sum_{j \in P} c_j(f_j)$; an *equilibrium flow* (sometimes termed *Wardrop equilibrium*, first presented in [17]) is defined naturally as a flow in which no unit of traffic can decrease its cost by unilaterally changing its path. A useful characterization of a Wardrop equilibrium is that all paths with nonzero flow of the same commodity i have the same cost γ , and all other paths from s_i to t_i have cost of at least γ . It is also well known that equilibrium flows exist for every network, and that all equilibrium flows on a network have the exact same cost (see [2] and Chapter 18 in [14]). Since equilibrium is unique in this sense, there is no distinction between “worst Nash” and “best Nash” in routing games.

2.1 General Networks

Our bounds for Pareto efficiency on general networks relate to the Price of Anarchy for such networks. Let \mathcal{P} be the set of all $s_i - t_i$ paths for every commodity i , and let f_P be the amount of traffic using the path $P \in \mathcal{P}$. The utilitarian social cost of flow f is $C(f) = \sum_{P \in \mathcal{P}} \sum_{j \in P} c_j(f_j) f_P = \sum_{j \in [m]} c_j(f_j) f_j$. The Price of Anarchy for a game (N, r, c) is defined as $POA(N, r, c) = \max_f \frac{C(f^E)}{C(f)}$, where f^E is an equilibrium flow, and the maximum is taken over all feasible flows f . We

now show that the POA of a class of games tightly bounds the Pareto-efficiency for the class. We start with the single commodity case.

Theorem 1. *Let \mathcal{C} be a class of continuous cost functions, and let $\mathcal{G}(\mathcal{C})$ be the set of single commodity selfish routing games with cost functions from \mathcal{C} . If there exists $\rho(\mathcal{C}) = \sup_{G \in \mathcal{G}(\mathcal{C})} POA(G)$, then in any $G \in \mathcal{G}(\mathcal{C})$ any equilibrium is $\rho(\mathcal{C})$ -Pareto-efficient, and for every $\epsilon > 0$ there exists $G \in \mathcal{G}(\mathcal{C})$ in which equilibria are $(\rho(\mathcal{C}) - \epsilon)$ -Pareto-deficient. If $POA(\mathcal{G}(\mathcal{C}))$ is unbounded, then equilibria of games in $\mathcal{G}(\mathcal{C})$ may be arbitrarily Pareto-deficient.*

Proof. We provide only a sketch of the proof, and the full details appear in the full version of this paper. First, if $\rho(\mathcal{C})$ exists, then for any game $G \in \mathcal{G}(\mathcal{C})$, there is no feasible flow with average cost smaller by factor $> \rho(\mathcal{C})$ than the average cost in equilibrium. This implies that in G , any equilibrium is $\rho(\mathcal{C})$ -Pareto-efficient.

For the lower bound (whether the Price of Anarchy of the class is bounded or unbounded) it suffices to show that for any game $G \in \mathcal{G}(\mathcal{C})$ with $POA(G) = \rho$ and every $\epsilon > 0$ there exists $G' \in \mathcal{G}(\mathcal{C})$ in which equilibria are $(\rho - \epsilon)$ -Pareto-deficient. Let P_1, \dots, P_M be an enumeration of the $s - t$ paths in N (a flow f is thus a non-negative vector in \mathbb{R}^M) and w.l.o.g. assume that the amount to that needs to be routed is 1, and so in an equilibrium flow f^E all players pay $C(f^E)$.

The idea behind the construction of $G' = (N', r'c')$ is as follows. We create a new network N' by “concatenating” q copies of N (for some large enough q), connecting every two adjacent copies by placing a zero-cost edge going from the sink of the first to the source of the second. A flow that routes all the traffic exactly as in equilibrium in each copy is an equilibrium in N' , and its cost is q -times that of the original equilibrium flow. We now look at the optimal flow in N ; since the total cost function $C(\cdot)$ is continuous, there is a flow $(\frac{p_1}{q}, \dots, \frac{p_M}{q})$ (for large enough q) routing rational amounts on the paths in N , and having a total cost larger than that of the optimal flow by at most ϵ . We can now use that latter flow in every copy of N in N' , keeping the *amounts* routed on each path the same in every copy, but changing the sets of *players* routed on these paths. This can be done to achieve a flow in N' with total cost of q -times the optimum (up to an additive factor of ϵ) in which all players are incurred the same cost. It then follows that this flow $(\rho - \epsilon)$ -Pareto-dominates the equilibrium flow.

For the multi-commodity case, Roughgarden, in [16], proves that under some additional conditions on the class of allowable cost functions, the worst POA for multi-commodity instances can be achieved (up to an arbitrarily small additive factor) on single-commodity “Pigou network” instances. We therefore immediately get that under the same conditions (namely that the class of allowable cost functions is both standard and diverse, and that all cost functions are monotone) the Pareto-deficiency of a Nash equilibrium in multi-commodity instances cannot be significantly worse than that of a single-commodity instance with cost functions from the same class.

2.2 Parallel-Edge Networks

Consider the special case of single-commodity networks with no nodes except a single source and a single sink, and only parallel edges connecting the two. We call such networks *parallel-edge networks*, and further assume that all cost functions in these networks are nondecreasing. Interestingly, while such networks exhibit the worst case examples of POA (as proven in [16]), the next theorem shows that all equilibrium flows in such networks are Pareto-optimal.

Theorem 2. *Let (N, r, c) be a selfish routing game on a parallel-edge network, where the cost functions of all edges are nondecreasing. Then if f is an equilibrium flow for (N, r, c) , f is Pareto optimal.*

The following lemma is straightforward:

Lemma 1. *Let (N, r, c) be a selfish routing game on a parallel-edge network, where the cost functions of all edges are nondecreasing. Then if f is a flow obtained by a greedy online process, f is an equilibrium flow.*

Corollary 1. *Let (N, r, c) be a selfish routing game on a parallel-edge network, where the cost functions of all edges are nondecreasing. Then if f is a flow obtained by a greedy online process, f is Pareto optimal.*

For the case of parallel-edge networks with linear cost functions, we can show an even stronger result, that will also be used in the analysis of load balancing games with pure strategies.

Theorem 3. *Let N be a parallel-edge network with linear cost functions. Let f be a flow on N , and let I be the set of edges with positive flow in f . Let f^* be another flow obtained from f by shifting at most an α fraction of each edge in I to edges not in I . Then, if f^* γ -Pareto-dominates f , then $\gamma \leq \frac{1}{1-\alpha}$.*

When $I = [m]$ it is immediate to observe that $\alpha = 0$ and we obtain:

Corollary 2. *Let G be a parallel-edge network with linear cost functions, and f a flow such that $f_i > 0$ for all $i \in [m]$. Then f is Pareto optimal.*

Thus, every flow on a parallel-edge network that uses all the edges is Pareto optimal, even if it is not an equilibrium flow.

3 Load Balancing Games – Pure Strategies

A *load balancing game* is defined by a set $[m]$ of machines and a set $[n]$ of tasks, where each task is associated with a weight function $w_i : [m] \rightarrow \mathbb{R}$ such that $w_i(j)$ is the weight of task i on machine j . We say that the machines are *uniformly-related* if there are constants $\{w_i\}_{i \in [n]}$ and $\{s_j\}_{j \in [m]}$ such that for all i, j it holds that $w_i(j) = \frac{w_i}{s_j}$. The machines are *identical* if this holds with all

$s_j = 1$. If the machines are not uniformly-related, we say that the game is played on *unrelated machines*.

A pure strategy profile is a function $A : [n] \rightarrow [m]$ assigning every task i to a single machine $j = A(i)$. The cost incurred to a task k assigned by A to machine j is $\text{cost}(k, A) = \sum_{i:A(i)=j} w_i(j)$ (i.e. we assume that on every machine, the tasks are executed in parallel). An assignment A is thus in Nash equilibrium if no player can benefit by unilaterally moving to another machine, i.e. if for every task $k \in [n]$ and machine $j \neq A(k)$ it holds that $\text{cost}(k, A) \leq w_k(j) + \sum_{i:A(i)=j} w_i(j)$.

Theorem 4. *Let G be a load balancing game, then G has a Pareto optimal Nash equilibrium in pure strategies.*

The theorem is a direct consequence of the fact that the (egalitarian) POS in such games is 1 (as shown in [6]), and the detailed argument is given in the full version of this paper. However, while the best Nash is always Pareto optimal, the worst Nash on unrelated machines may be arbitrarily Pareto-deficient, as the following example shows: Let $\epsilon > 0$ be arbitrarily small and consider an instance with m machines and $n = m$ tasks, such that for all i , $w_i(i) = 1$ and for all $i \neq j$, $w_i(j) = \epsilon$. It is easy to observe that the identity assignment $A(i) = i$ is a Nash equilibrium that is $\frac{1}{\epsilon}$ -Pareto-deficient.

3.1 Uniformly-Related and Identical Machines

Load balancing games on uniformly-related machines can be viewed as *atomic routing games* (where each player controls a non-negligible amount of traffic) on parallel-edge networks with linear cost functions. Interestingly, we can use the results for nonatomic selfish routing games to derive bounds for load balancing games.

Theorem 5. *Let G be a load balancing game on uniformly-related machines, and let $A : [n] \rightarrow [m]$. If either*

1. *A is an equilibrium assignment, or,*
2. *the machines of $[m]$ are identical and A is the result of a greedy online assignment process,*

then A is Pareto optimal.

Proof. Let $\{w_i\}_{i \in [n]}$ be the set of job weights and $s = \{s_j\}_{j \in [m]}$ be the machine speeds. Denote $W = \sum_{i \in [n]} w_i$. We define a selfish routing game $G'(W, s) = (N, r, c)$ on a parallel-edge network by creating a set of edges $[m]$ with cost function $c_j(x) = \frac{x}{s_j}$ for every $j \in [m]$, and $r = W$. Every assignment $A : [n] \rightarrow [m]$ for G induces a feasible flow f^A on $G'(W, s)$ in which the flow on an edge j is $\sum_{i:A(i)=j} w_i$; furthermore, every player in the routing game originates from a single task $i \in [n]$ in the load balancing game, and pays $\text{cost}(i, A)$ in f^A . Therefore, if an assignment A^* Pareto dominates A , then f^{A^*} Pareto dominates f^A by the same factor.

Assume that A is an equilibrium assignment, and define $I = \{j \in [m] \mid \exists i : A(i) = j\}$ as the set of machines j that some task uses. Then in $G'(W, s)$, I is the set of edges with nonzero flow in f^A . Assume by contradiction that A is not Pareto optimal, so there exists another assignment A^* that Pareto dominates A . Define $I^* = \{j \in [m] \mid \exists i : A^*(i) = j\}$, then clearly, $I^* \subseteq I$; otherwise let $j \in I^* \setminus I$ and let i be such that $A^*(i) = j$. Since A is an equilibrium, it holds that $\text{cost}(i, A) \leq \frac{w_i}{s_j} \leq \text{cost}(i, A^*)$, and thus A^* does not Pareto dominate A because player i pays in it at least as much as it payed in A . However, if $I^* \subseteq I$ then the flow f^{A^*} routes all the traffic on the edges of I , and thus by applying Theorem 3 with $\alpha = 0$ we get that A^* cannot Pareto dominate A ; a contradiction.

Now, assume that A is a result of a greedy online assignment process on identical machines, and define I as above. There are two cases: If $I = [m]$, then by Corollary 2 we are done. Otherwise, there are machines that are not used by any of the tasks; however, since A was obtained by a greedy process and the machines have identical speeds, it has to be that on every machine $j \in I$ there is only a *single task* (or the second task that arrived to j would have preferred to use some vacant machine $\ell \in [m] \setminus I$). Thus, every task pays the minimum possible cost (of its weight divided by the uniform speed) and there is no way to reduce the cost of any of the tasks, so again A is Pareto optimal.

Unlike with identical machines, if A is the result of a greedy online assignment process on non-identical machines, A may be Pareto dominated by another assignment. For example, assume that we have three machines with speeds 2, 1 and 1, and three tasks with weights 1, 1 and 2. Consider the following scenario: A task of weight 1 arrives first, and chooses the fast (speed 2) machine. Then arrives the other unit weighted task, and (being indifferent about which machine to choose) chooses the fast machine as well. Finally, the heavy (weight 2) task arrives, and again chooses the fast machine (as it too would have the same cost on all the machines). In this assignment all the tasks pay a cost of $\frac{1+1+2}{2} = 2$; however, if we assign each of the light tasks to a (distinct) slow machine and the heavy task to the fast machine we get that every task pays only 1. Thus, the online greedy assignment is 2-Pareto-deficient. The following theorem, whose proof again utilizes Theorem 3, establishes that this is the worst possible case.

Theorem 6. *Let G be a load balancing game on uniformly-related machines, and let A be the result of a greedy online assignment process. Then A is 2-Pareto-efficient.*

4 Load Balancing Games – Mixed Strategies

A mixed strategy of a player $i \in [n]$ in a load balancing game is a distribution $p_i = (p_i^1, \dots, p_i^m)$ over the set of machines, so that i chooses to use machine j with probability p_i^j . The expected cost for player i is thus $\text{cost}(i, p) = \sum_{j \in [m]} p_i^j \cdot (w_i(j) + \sum_{h \in [n] - i} p_h^j w_h(j))$. As one would expect, a profile p is in equilibrium if no player can benefit by unilaterally switching to a different distribution p'_i . Note that mixed strategies are a superset of pure strategies; therefore,

we immediately obtain that for unrelated machines the worst Nash in mixed strategies may be arbitrarily Pareto-deficient. We were unable to give any tight bounds for the best Nash equilibria in mixed strategies.

We will thus focus on bounding the distance of the worst Nash from Pareto optimality for uniformly-related and identical machines. We will denote the expected weight on machine j in the strategy profile p by $E_p[W_j] = \sum_{h \in [n]} p_h^j w_h$, and thus we can also write:

$$\text{cost}(i, p) = \sum_{j \in [m]} p_i^j \cdot \frac{\sum_{h \in [n]} p_h^j w_h + (1 - p_i^j) w_i}{s_j} = \sum_{j \in [m]} p_i^j \cdot \frac{E_p[W_j] + (1 - p_i^j) w_i}{s_j}$$

For a load balancing game with identical machines we show that every equilibrium profile is $(2 - \frac{1}{m})$ -Pareto-efficient, and that this is tight, i.e. the worst equilibrium in a game may indeed be $(2 - \frac{1}{m})$ -Pareto-deficient.

Theorem 7. *In load balancing games on identical machines with mixed strategies, all Nash equilibria are $(2 - \frac{1}{m})$ -Pareto-efficient, and this bound is tight.*

For the case of uniformly-related machines with mixed strategies, we show that every equilibrium is 4-Pareto-efficient; however, we do not know to show that this bound is tight.

Theorem 8. *In load balancing games on uniformly-related machines with mixed strategies, all Nash equilibria are 4-Pareto-efficient.*

Proof. Assume that there exist a set $[m]$ of machines, a set $[n]$ of tasks, and strategy profiles p, q such that for every player $i \in [n]$, $\text{cost}(i, q) \leq \frac{1}{4} \text{cost}(i, p)$. The idea is to show that there exists a task k that can unilaterally improve its cost from the profile p , implying that p cannot be an equilibrium. To that end, we first define another strategy profile r , which is a variation of the profile q .

For a task i in the profile q , let B_i be the set of “bad” machines to which i gives nonzero probability and on which it pays over twice its expected cost, i.e.,

$$B_i = \left\{ j \mid q_i^j > 0 \wedge \frac{E_q[W_j] + (1 - q_i^j) w_i}{s_j} > 2 \cdot \text{cost}(i, q) \right\}.$$

We also denote the remaining (“good”) machines to which i gives nonzero probability by G_i , so

$$G_i = \left\{ j \mid q_i^j > 0 \wedge \frac{E_q[W_j] + (1 - q_i^j) w_i}{s_j} \leq 2 \cdot \text{cost}(i, q) \right\}.$$

Since all q_i^j are non-negative, it holds that the total probability every i gives to bad machines is $b_i = \sum_{j \in B_i} q_i^j < \frac{1}{2}$ (or the expected cost for i would have exceeded $\text{cost}(i, q)$). We create the new strategy profile r as follows. First, for every i, j with $q_i^j = 0$ we set $r_i^j = 0$ as well. For every i, j such that $j \in B_i$, we

also set $r_i^j = 0$. However, in order to keep the vector r_i a distribution, we add the total “missing” probability b_i to the machines in G_i ; specifically, for every i, j with $j \in G_i$ we set $r_i^j = \frac{q_i^j}{1-b_i}$.

We now have

$$\sum_{j \in [m]} E_p[W_j] = \sum_{i \in [n]} w_i \sum_{j \in [m]} p_i^j = \sum_{i \in [n]} w_i = \sum_{i \in [n]} w_i \sum_{j \in [m]} r_i^j = \sum_{j \in [m]} E_r[W_j],$$

so there is bound to be a machine ℓ with $E_r[W_\ell] > 0$ and $E_p[W_\ell] \leq E_r[W_\ell]$. Let k be a task with $r_k^\ell > 0$.

We now show that in the strategy profile p , player k can reduce its cost to less than $4 \cdot \text{cost}(k, q)$, by choosing to use machine ℓ with probability 1. The cost incurred to k when doing so is

$$\frac{E_p[W_\ell] + (1 - p_k^\ell)w_k}{s_\ell} \leq \frac{E_r[W_\ell] + w_k}{s_\ell} \leq \frac{\sum_{i \in [n]-k} r_i^\ell w_i + 2w_k}{s_\ell}.$$

Recall that for every i, j it holds that $r_i^j \leq \frac{q_i^j}{1-b_i}$ and that $b_i < \frac{1}{2}$; this implies that $r_i^j < 2q_i^j$. Thus,

$$\frac{\sum_{i \in [n]-k} r_i^\ell w_i + 2w_k}{s_\ell} < 2 \cdot \frac{\sum_{i \in [n]-k} q_i^\ell w_i + w_k}{s_\ell} \leq 2 \cdot 2 \cdot \text{cost}(k, q),$$

where the last inequality holds since $\frac{\sum_{h \in [n]-k} q_h^\ell w_h + w_k}{s_\ell}$ is exactly the cost k pays on machine ℓ in the profile q . We chose k such that $r_k^\ell > 0$ and so it must be that $\ell \in G_k$; this implies that in the profile q , k pays on ℓ at most $2 \cdot \text{cost}(i, q)$. Combining the two inequalities above we get that $\frac{E_p[W_\ell] + (1 - p_k^\ell)w_k}{s_\ell} < 4 \cdot \text{cost}(k, q)$; however, we assumed that the original cost of k in the profile p was $\text{cost}(k, p) \geq 4 \cdot \text{cost}(k, q)$, so unilaterally moving to ℓ is beneficial for k in the profile p . We thus conclude that if some profile q 4-Pareto-dominates another profile p , then p cannot be an equilibrium.

5 Open Problems

A natural direction for further research is the analysis of the Pareto efficiency/deficiency of solutions in other games, as well as other solution concepts in these and other games. In addition, there are a few cases left open in this work, including:

- *Flows obtained by greedy online processes on general routing networks.* Unlike with parallel-edge networks, such flows in general networks need not be equilibrium flows, even in single-commodity instances. However, the lower bound of the POA value ρ still holds for such flows.
- *Online greedy assignments for load balancing on unrelated machines.* It can be shown that there are instances in which greedy assignments are not 2-Pareto-efficient. What is the Pareto efficiency/deficiency this case?

- (Worst) mixed equilibrium for load balancing on uniformly-related machines. We have shown that such equilibria are always 4-PE, but suspect that the real bound may be lower.

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: FOCS, pp. 295–304 (2004)
2. Beckmann, M., McGuire, C.B., Winsten, C.B.: Studies in the economics of transportation. Yale University Press, New Haven (1956)
3. Cohen, J.E.: Cooperation and self-interest: Pareto-inefficiency of Nash equilibria in finite random games. Proceedings of the National Academy of Sciences of the United States of America 95(17), 9724–9731 (1998)
4. Diakonikolas, I., Yannakakis, M.: Small approximate pareto sets for biobjective shortest paths and other problems. SIAM Journal on Computing 39(4), 1340–1371 (2009), <http://link.aip.org/link/?SMJ/39/1340/1>
5. Dubey, P.: Inefficiency of nash equilibria. Mathematics of Operations Research 11(1) (1986)
6. 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)
7. Feldman, M., Tamir, T.: Approximate strong equilibrium in job scheduling games. In: Monien, B., Schroeder, U.-P. (eds.) SAGT 2008. LNCS, vol. 4997, pp. 58–69. Springer, Heidelberg (2008)
8. Hurwicz, L., Schmeidler, D.: Construction of outcome functions guaranteeing existence and pareto optimality of nash equilibria. Econometrica 46(6) (1978)
9. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: STACS, pp. 404–413 (1999)
10. Legriel, J., Guernic, C.L., Cotton, S., Maler, O.: Approximating the pareto front of multi-criteria optimization problems. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010. LNCS, vol. 6015, pp. 69–83. Springer, Heidelberg (2010)
11. Loridan, P.: ϵ -solutions in vector minimization problems. Journal of Optimization Theory and Applications 43(2) (1984)
12. Maskin, E.: Nash equilibrium and welfare optimality. The Review of Economic Studies, Special Issue: Contracts 66(1) (1999)
13. Nash, J.F.: The bargaining problem. Econometrica 18(2) (1950)
14. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, New York (2007)
15. Papadimitriou, C.H., Yannakakis, M.: On the approximability of trade-offs and optimal access of web sources (extended abstract). In: Proceedings 41st Annual Symposium on Foundations of Computer Science, pp. 86–92 (2000)
16. Roughgarden, T.: The price of anarchy is independent of the network topology. Journal of Computer and System Sciences, 428–437 (2002)
17. Wardrop, J.G.: Some theoretical aspects of road traffic research. Proceedings of the Institution of Civil Engineers, Part II 1(36), 352–362 (1952)

On Nash-Equilibria of Approximation-Stable Games^{*}

Pranjal Awasthi¹, Maria-Florina Balcan², Avrim Blum¹,
Or Sheffet¹, and Santosh Vempala²

¹ Carnegie Mellon University, Pittsburgh, Pennsylvania

² Georgia Institute of Technology, Atlanta, Georgia

Abstract. One reason for wanting to compute an (approximate) Nash equilibrium of a game is to predict how players will play. However, if the game has multiple equilibria that are far apart, or ϵ -equilibria that are far in variation distance from the true Nash equilibrium strategies, then this prediction may not be possible even in principle. Motivated by this consideration, in this paper we define the notion of games that are *approximation stable*, meaning that all ϵ -approximate equilibria are contained inside a small ball of radius Δ around a true equilibrium, and investigate a number of their properties. Many natural small games such as matching pennies and rock-paper-scissors are indeed approximation stable. We show furthermore there exist 2-player n -by- n approximation-stable games in which the Nash equilibrium and all approximate equilibria have support $\Omega(\log n)$. On the other hand, we show all (ϵ, Δ) approximation-stable games must have an ϵ -equilibrium of support $O(\frac{\Delta^{2-o(1)}}{\epsilon^2} \log n)$, yielding an immediate $n^{O(\frac{\Delta^{2-o(1)}}{\epsilon^2} \log n)}$ -time algorithm, improving over the bound of [11] for games satisfying this condition. We in addition give a polynomial-time algorithm for the case that Δ and ϵ are sufficiently close together. We also consider an inverse property, namely that all *non*-approximate equilibria are *far* from some true equilibrium, and give an efficient algorithm for games satisfying that condition.

1 Introduction

One reason for wanting to compute a Nash equilibrium or approximate equilibrium of a game is to predict how players will play. However, if the game has multiple equilibria that are far apart, or ϵ -equilibria that are far from the true Nash equilibrium strategies, then this prediction may not be possible even in principle. Motivated by this consideration, in this paper we define the notion of games that are (ϵ, Δ) -*approximation stable*, meaning that all ϵ -approximate equilibria are contained inside a small ball of radius Δ (in variation distance) around a true equilibrium, and investigate a number of their properties. If a game is approximation-stable for small Δ , then this means that even if players are only approximately best-responding, or even if the game matrix is not a perfect description of players' true payoffs, stationary play should in principle be predictable. Many natural small 2-player games such as matching pennies and rock-paper-scissors are indeed approximation-stable for Δ close to ϵ . In this paper we analyze fundamental properties of approximation-stable games.

^{*} This work was supported in part by NSF grants CCF-0830540 and CCF-0953192, ONR grant N00014-09-1-0751, and AFOSR grant FA9550-09-1-0538.

We show first that all (ϵ, Δ) approximation-stable games must have an ϵ -equilibrium of support at most $O(\frac{\Delta^2 \log(1+1/\Delta)}{\epsilon^2} \log n)$, yielding an immediate $n^{O(\frac{\Delta^2 \log(1+1/\Delta)}{\epsilon^2} \log n)}$ -time algorithm for finding an ϵ -equilibrium, improving by a factor $O(\Delta^2 \log(1+1/\Delta))$ in the exponent over the bound of [11] for games satisfying this condition (and reducing to the bound of [11] in the worst case when $\Delta = 1$). Note that by assumption, this approximate equilibrium is also Δ -close to a true Nash equilibrium. We in addition give improved bounds yielding polynomial-time algorithms for the case that Δ and ϵ are sufficiently close together. Specifically, for $\Delta \leq 2\epsilon - 6\epsilon^2$ we give an algorithm for finding $O(\epsilon)$ -equilibria in time $n^{O(1/\epsilon)}$. On the other hand, we show that for $\Delta = O(\sqrt{\epsilon})$, there exist n -action approximation-stable games in which the Nash equilibrium and all approximate equilibria have support $\Omega(\log n)$, extending results of Feder et al. [10]. We also consider an inverse property, namely that all *non*-approximate equilibria are *far* from some true equilibrium, and give an efficient algorithm for finding approximate equilibria in games satisfying that condition.

Note that the classic notion of a stable Nash equilibrium is substantially more restrictive than the condition we consider here: it requires that (1) any infinitesimal deviation from the equilibrium by any player should make the deviating player strictly worse off (a *strict* equilibrium, implying that the equilibrium must be in pure strategies) and (2) such a deviation should not give the *other* player any incentive to deviate. Our condition can be viewed in a sense as a weaker, approximation version of requirement (1), namely any deviation by distance Δ from the equilibrium should make *at least one* of the two players have at least ϵ incentive to deviate.

Related Work: There has been substantial work exploring the computation of Nash equilibria in 2-player $n \times n$ general-sum games. Unfortunately, the complexity results in this area have been almost uniformly negative. A series of papers has shown that it is PPAD complete to compute Nash equilibria, even in 2 player games, even when payoffs are restricted to lie in $\{0, 1\}$ [7][16].

A structural result of Lipton et al. [11] shows that there always exist ϵ -approximate equilibria with support over at most $O((\log n)/\epsilon^2)$ strategies: this gives an immediate $n^{O(\log n/\epsilon^2)}$ -time algorithm for computing ϵ -approximate equilibria and has also been shown to be essentially tight [10]. There has also been a series of results [8][25] on polynomial-time algorithms for computing approximate equilibria for larger values of ϵ . The best polynomial-time approximation guarantee known is 0.3393 [12].

For special classes of games, better results are known. For example, Barany et al. considered two player games with randomly chosen payoff matrices, and showed that with high probability, such games have Nash equilibria with small support [4].

Our work is also motivated by that of Balcan et al. [2] who consider *clustering* problems under approximation stability – meaning that all near-optimal solutions to the objective function should be close in the space of solutions – and give efficient algorithms for stable instances for several common objectives. Results relating incentive to deviate and distance to equilibria in general games appear in [9].

2 Definitions and Preliminaries

We consider 2-player n -action general-sum games. Let R denote the payoff matrix to the row player and C denote the payoff matrix of the column player. We assume all

payoffs are scaled to the range $[0, 1]$. We say that a pair of mixed strategies (p, q) is an ϵ -equilibrium if for all rows i , we have $e_i^T Rq \leq p^T Rq + \epsilon$, and for all columns j , we have $p^T C e_j \leq p^T Cq + \epsilon$. We will typically use (p^*, q^*) to denote a Nash equilibrium, which is an ϵ -equilibrium for $\epsilon = 0$. Note that in a Nash equilibrium (p^*, q^*) , all rows i in the support of p^* satisfy $e_i^T Rq^* = p^* Rq^*$ and similarly all columns j in the support of q^* satisfy $p^{*T} C e_j = p^{*T} Cq^*$.

We also are interested in the distance between mixed strategies. For probability distributions in this context, the most natural notion is variation distance, which we use here. Specifically we define:

$$d(q, q') = \frac{1}{2} \sum_i |q_i - q'_i| = \sum_i \max(q_i - q'_i, 0). \quad (1)$$

We then define the distance between two strategy pairs as the maximum of the row-player's and column-player's distances, that is:

$$d((p, q), (p', q')) = \max[d(p, p'), d(q, q')]. \quad (2)$$

We now present our main definition, namely that of a game being *approximation stable*.

Definition 1. A game satisfies (ϵ, Δ) -approximation stability if there exists a Nash equilibrium (p^*, q^*) such that any (p, q) that is an ϵ -equilibrium is Δ -close to (p^*, q^*) , i.e. $d((p, q), (p^*, q^*)) \leq \Delta$.

So, fixing ϵ , a smaller Δ means a stronger condition and a larger Δ means a weaker condition. Every game is $(\epsilon, 1)$ -approximation stable, and as Δ gets smaller, we might expect for the game to exhibit more useful structure. Many natural games such as matching pennies and rock-paper-scissors satisfy (ϵ, Δ) -approximation stability for $\Delta = O(\epsilon)$; see Section 2.2 for analysis of a few simple examples. We note that this definition is very similar to a condition used in Balcan et al. [2] in the context of clustering problems.

All our results also apply to a weaker notion of approximation stability that allows for multiple equilibria, so long as moving distance Δ from any equilibrium produces a solution in which at least one player has ϵ incentive to deviate. Specifically,

Definition 2. A game satisfies (ϵ, Δ) -weak approximation stability if, for any Nash equilibrium (p^*, q^*) and any (p, q) such that $d((p, q), (p^*, q^*)) = \Delta$, (p, q) is not an ϵ' -equilibrium for any $\epsilon' < \epsilon$.

Organization of This Paper: We now begin with a few useful facts about the region around Nash equilibria and the relation between ϵ and Δ in any game, as well as a few simple examples of games satisfying (ϵ, Δ) -approximation stability for $\Delta \approx \epsilon$. We then in Section 3 analyze properties of approximation-stable games, showing that every (ϵ, Δ) -approximation stable game must have an ϵ -equilibrium of support $O(\frac{\Delta^2 \log(1+1/\Delta) \log(n)}{\epsilon^2})$, yielding an immediate $n^{O(\frac{\Delta^2 \log(1+1/\Delta) \log(n)}{\epsilon^2})}$ -time algorithm.

Note that for large Δ this exponent simply reduces to the $O(\frac{\log(n)}{\epsilon^2})$ bound of [11], but improves as Δ approaches ϵ . In Section 5 we give a near-matching lower bound, showing that there exist approximation-stable games with all approximate equilibria having support $\Omega(\log n)$. In Section 4 we analyze games where Δ is especially close

to ϵ , and give polynomial-time algorithms for finding approximate equilibria when $\Delta \leq 2\epsilon - O(\epsilon^2)$. Finally, in Section 6 we consider the inverse condition that all strategies within distance Δ of some Nash equilibrium are ϵ -equilibria, and give an efficient algorithm for computing (ϵ/Δ) -approximate equilibria in this case.

2.1 Preliminaries

We begin with a few preliminary facts that apply to any 2-player general-sum game.

Claim 1. *If (p, q) is α -close to a Nash equilibrium (p^*, q^*) (i.e., if $d((p, q), (p^*, q^*)) \leq \alpha$), then (p, q) is a 3α -Nash equilibrium.*

Proof. (omitted)

Claim 1 is useful because while it may be hard to determine how close some pair (p, q) is to a true equilibrium, it is easy to check how much incentive players have to deviate. Say that a Nash equilibrium (p^*, q^*) is *non-trivial* if at least one of p^* or q^* does not have full support over all the rows or columns. Notice trivial Nash equilibria, if they exist, can be computed in polynomial-time using Linear programming. We then have:

Claim 2. *For any nontrivial Nash equilibrium (p^*, q^*) and any $\alpha > 0$, there exists (p, q) such that $d((p, q), (p^*, q^*)) \geq \alpha$ and (p, q) is an α -approximate equilibrium.*

Proof. Without loss of generality, assume that p^* does not have full support. Let e_i be a row not in the support of p^* . Consider a pair of distributions (p, q^*) where $p = (1 - \alpha)p^* + \alpha e_i$. Since i was not in the support of p^* , (p, q^*) has variation distance α from (p^*, q^*) . Yet, in (p, q^*) , with probability $(1 - \alpha)$ both the players are playing best responses to each other. Hence, no player has more than α incentive to deviate. \square

Corollary 1. *Assume that the game \mathcal{G} satisfies (ϵ, Δ) -approximation stability and has a non-trivial Nash equilibrium. Then we must have $\Delta \geq \epsilon$.*

2.2 Some Simple Examples

A number of natural small games satisfy (ϵ, Δ) -approximation stability for every $\epsilon > 0$ and for $\Delta = O(\epsilon)$. Here, we give a few simple examples.

Game 1: The row and the column matrices are 2×2 as follows:

$$R = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} \quad C = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}$$

Here, the only Nash equilibrium (p^*, q^*) is for the row player to play row 1 and the column player to play column 1, which are dominant strategies. Any deviation by distance Δ from p^* will give the row player Δ incentive to deviate, regardless of the strategy of the column player. Similarly, any deviation of Δ from q^* will give the column player a Δ incentive to deviate regardless of the strategy of the row player. Hence, for every $\epsilon \in [0, 1]$, this game is (ϵ, Δ) -stable for $\Delta = \epsilon$.

Game 2: This game is simply matching pennies:

$$R = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

Denoting the indicator vectors as e_1 and e_2 , the Nash equilibrium (p^*, q^*) is equal to $(\frac{1}{2}(e_1 + e_2), \frac{1}{2}(e_1 + e_2))$. We now show that for any strategy which is Δ far from (p^*, q^*) , at least one player must have ϵ incentive to deviate for $\epsilon = \Delta \frac{(1+2\Delta)}{(1+4\Delta)}$.

Specifically, let (p, q) be Δ -far from (p^*, q^*) , and without loss of generality assume $d(p, p^*) = \Delta$. We may further assume without loss of generality (by symmetry) that $p = (\frac{1}{2} + \Delta)e_1 + (\frac{1}{2} - \Delta)e_2$. Let $q = (\frac{1}{2} - \Delta')e_1 + (\frac{1}{2} + \Delta')e_2$ for $\Delta' \in [-\Delta, \Delta]$. In this case the row player is getting a payoff $p^T Rq = (\frac{1}{2} - 2\Delta\Delta')$. Furthermore, he can move to row 2 and get payoff $e_2^T Rq = (\frac{1}{2} + \Delta')$. Hence, the incentive to deviate $(e_2 - p)^T Rq \geq \Delta'(1 + 2\Delta)$. Similarly, the column player has payoff $p^T Cq = (\frac{1}{2} + 2\Delta\Delta')$, whereas $p^T Ce_2 = (\frac{1}{2} + \Delta)$, and hence has at least $\Delta(1 - 2\Delta')$ incentive to deviate. The maximum of these two is at least $\Delta \frac{(1+2\Delta)}{(1+4\Delta)}$ (with this value occurring at $\Delta' = \frac{\Delta}{1+4\Delta}$). Therefore, the incentive to deviate in any (p, q) that is Δ -far from (p^*, q^*) is at least this large. Solving for Δ as a function of ϵ , this game is (ϵ, Δ) -approximation stable for $\Delta = \epsilon + O(\epsilon^2)$.

Game 3: Rock, Paper, Scissors.

$$R = \begin{bmatrix} 0.5 & 0 & 1 \\ 1 & 0.5 & 0 \\ 0 & 1 & 0.5 \end{bmatrix} \quad C = \begin{bmatrix} 0.5 & 1 & 0 \\ 0 & 0.5 & 1 \\ 1 & 0 & 0.5 \end{bmatrix}$$

A case analysis (omitted) shows that this game is (ϵ, Δ) -approximation stable for $\Delta = 4\epsilon$, for any $\epsilon \leq \frac{1}{6}$.

3 The Support of Equilibria in Stable Games

We now show that approximation-stable games have structure that can be used to improve the efficiency of algorithms for computing approximate equilibria.

Theorem 1. *For any game satisfying (ϵ, Δ) -approximation stability, there exists an ϵ -equilibrium where each player’s strategy has support $O((\Delta/\epsilon)^2 \log(1 + 1/\Delta) \log n)$.*

Corollary 2. *There is an algorithm to find ϵ -equilibria in games satisfying (ϵ, Δ) -approximation stability, running in time $n^{O((\Delta/\epsilon)^2 \log(1+1/\Delta) \log n)}$.*

Let $S = c(\Delta/\epsilon)^2 \log n$ for some absolute constant c , and let (p^*, q^*) denote the Nash equilibrium such that all ϵ -equilibria lie within distance Δ of (p^*, q^*) . Theorem 1 is proven in stages. First, in Lemma 1 we show that given a pair of distributions (p, q) , if p is near-uniform over a large support then p can be written as a convex combination $p = xp_1 + (1 - x)p_2$ where p_1 and p_2 have disjoint supports, and for every column j , j ’s performance against p_1 is close to its performance against p_2 . This implies p^* itself cannot be near-uniform over a large sized support, since otherwise we could write it in this way and then shift Δ probability mass from p_2 to p_1 , producing a new distribution p' such that under (p', q^*) , the column player has less than ϵ incentive to deviate (and the row player has zero incentive to deviate since $\text{supp}(p') \subseteq \text{supp}(p^*)$). This contradicts the fact that the game is (ϵ, Δ) -approximation stable. We then build on this to show that if p^* is not near-uniform and does have a large support, it must be well-approximated by

a distribution of small support (roughly $O(S \log \frac{1}{\Delta})$). This analysis combines Lemma [1](#) together with the sampling idea of Lipton et al. [\[11\]](#). The same, of course, applies to q^* . For the rest of this section we assume that $\Delta \leq 1/4$.

Lemma 1. *For any distributions p and q , if p satisfies $\|p\|_2^2 \leq \frac{1}{S}$ where $S = c(\Delta/\epsilon)^2 \log n$ for sufficiently large constant c , then p can be written as a convex combination $p = xp_1 + (1-x)p_2$ of two distributions p_1 and p_2 over disjoint supports such that:*

- (i) $x \leq 3/4 \leq 1 - \Delta$.
- (ii) $\forall j, (p_1 - p)^T C(e_j - q) < \frac{\epsilon}{4\Delta}$

The point of Lemma [1](#) is that by (i) and (ii), modifying p by moving Δ probability mass from p_2 to p_1 can improve the performance of e_j relative to q for the column player by at most ϵ . The proof of Lemma [1](#) makes extensive use of the Hoeffding Bound:

Theorem 2 (Hoeffding Bound). *Let $X_i, i = 1, 2, \dots, n$, be n random variables, s.t. $\forall i, X_i \in [a_i, b_i]$. Let $\mu_i = \mathbf{E}[X_i]$. Then for every $t > 0$ we have that:*

$$\Pr \left[\sum_i X_i > t + \sum_i \mu_i \right] \leq \exp \left(-\frac{t^2}{\sum_i (b_i - a_i)^2} \right) \quad (3)$$

Proof (Lemma [1](#)). Let r be a random subset of the support of p ; that is, for every element in $\text{supp}(p)$, add it to r with probability $1/2$. Also, let C_i denote the i th entry of Cq . The idea of the proof is just to argue that for any column j , by the Hoeffding bound, with high probability over the choice of r , the distribution p_1 induced by p restricted to r satisfies the desired condition that $p_1^T C(e_j - q)$ is within $\frac{\epsilon}{4\Delta}$ of $p^T C(e_j - q)$. We then simply perform a union bound over j .

Fix column e_j . Let Y_{ij} be the random variable defined as $2p_i(C_{ij} - C_i)$ if element i was added to r , and 0 otherwise. Observe that $\mathbf{E}[\sum_i Y_{ij}] = \frac{1}{2} \sum_i 2p_i(C_{ij} - C_i) = p^T C(e_j - q)$. Let Z_i be the random variable defined as $2p_i$ with probability $1/2$ (if element i was added to r), and 0 otherwise. Observe $\mathbf{E}[\sum_i Z_i] = 1$. Observe also that for every i we have that $Z_i, Y_{ij} \in [-2p_i, 2p_i]$.

The obvious reason for defining Y_{ij} and Z_i is that by denoting the distribution p restricted to r (renormalized to have L_1 norm equal to 1) as p_r , we have:

$$p_r^T C(e_j - q) = \frac{\sum_{i \in r} p_i (C_{ij} - C_i)}{\sum_{i \in r} p_i} = \frac{\sum_i Y_{ij}}{\sum_i Z_i} \quad (4)$$

so by bounding the numerator from above and the denominator from below, we can hope to find r for which $p_r^T C(e_j - q) < \mathbf{E}[\sum_i Y_{ij}] + (\epsilon/4\Delta)$, thus decomposing p into the desired $p_1 = p_r$ and $p_2 = p_{\bar{r}}$. We can do this using the Hoeffding bound and plugging the value of S :

$$\Pr \left[\sum_i Y_{ij} > p^T C(e_j - q) + \frac{\epsilon}{10\Delta} \right] < \exp \left(\frac{-(\epsilon/10\Delta)^2}{\sum_i (4p_i)^2} \right) \leq \exp \left(\frac{-S\epsilon^2}{(40\Delta)^2} \right) < \frac{1}{2n},$$

where the last inequality is by definition of S . Thus, $\Pr[\exists j, \sum_i Y_{ij} > p^T C(e_j - q) + \frac{\epsilon}{10\Delta}] < 1/2$. Similarly (and even simpler), we have that $\Pr[\sum_i Z_i < 1 - \frac{\epsilon}{10\Delta}] < 1/2$, and so the existence of r for which both events do not hold is proven. Observe that for this specific r we have that

$$\frac{\sum_i Y_{ij}}{\sum_i Z_i} \leq \frac{p^T C(e_j - q) + \epsilon/10\Delta}{1 - \epsilon/10\Delta} \leq p^T C(e_j - q) + \frac{\epsilon/5\Delta}{1 - \epsilon/10\Delta} \leq p^T C(e_j - q) + \frac{\epsilon}{4\Delta},$$

using the fact that $p^T C(e_j - q) \leq 1$. Thus, we have the desired decomposition of p . \square

Proof (Theorem [1](#)). We begin by partitioning p^* into its *heavy* and *light* parts. Specifically, greedily remove the largest entries of p^* and place them into a set H (the heavy elements) until either (a) $\Pr[H] \geq 1 - 4\Delta$, or (b) the remaining entries L (the light elements) satisfy the condition that $\forall i \in L, \Pr[i] \leq \frac{1}{5}\Pr[L]$ for S as in Lemma [1](#), whichever comes first. We analyze each case in turn.

If case (a) occurs first, then clearly H has at most $S \log(1/4\Delta)$ elements. We now simply apply the sampling argument of Lipton et al [\[11\]](#) to L and union the result with H . Specifically, decompose p^* as $p^* = \beta p_H + (1 - \beta)p_L$, where β denotes the total probability mass over H . Applying the sampling argument of [\[11\]](#) to p_L , we have that by sampling a multiset \mathcal{X} of S elements from $\text{supp}(p_L) = L$, we are guaranteed, by definition of S , that for any column e_j , $|(U_{\mathcal{X}})^T C e_j - p_L^T C e_j| \leq (\epsilon/8\Delta)$, where $U_{\mathcal{X}}$ is the uniform distribution over \mathcal{X} . This means that for $\tilde{p} = \beta p_H + (1 - \beta)U_{\mathcal{X}}$, all columns e_j satisfy $|p^{*T} C e_j - \tilde{p}^T C e_j| \leq \epsilon/2$. We have thus found (the row portion of) an ϵ -equilibrium with support of size $S(1 + \log(1/4\Delta))$ as desired, and now simply apply the same argument to q^* .

If (b) occurs first, we show that the game cannot satisfy (ϵ, Δ) -approximation stability. Specifically, let p_L denote the induced distribution produced by restricting p^* to L and renormalizing so that $\sum_i (p_L)_i = 1$, then $\sum_i (p_L)_i^2 \leq \frac{1}{S} \sum_i (p_L)_i = \frac{1}{S}$. Using Lemma [1](#), we deduce we can write p_L as a convex combination $p_L = x p_1 + (1 - x)p_2$ of p_1 and p_2 satisfying the properties of Lemma [1](#). Again, by denoting β as the total probability mass over H , we have:

$$p^* = \beta p_H + (1 - \beta)x p_1 + (1 - \beta)(1 - x)p_2 \quad (5)$$

where p_H is the induced distribution over H . We now consider the transition from p^* to p' defined as

$$p' = \beta p_H + ((1 - \beta)x + \Delta)p_1 + ((1 - \beta)(1 - x) - \Delta)p_2 \quad (6)$$

Notice that by Lemma [1](#), $x \leq \frac{3}{4}$ and hence $(1 - \beta)(1 - x) - \Delta \geq (1 - \beta)/4 - \Delta \geq 0$, so p' is a valid probability distribution. Also, since p_1 and p_2 are distributions over disjoint support, p' is Δ far from p^* . Note that since p' is obtained from an internal deviation within the support of p^* , the row player has no incentive to deviate when playing p' against q^* . So, if the game is (ϵ, Δ) -approximation stable, then playing p' against q^* must cause the column player to have more than ϵ incentive to deviate. However, by transitioning from p^* to p' the expected gain of switching from q^* to any e_j is

$$\begin{aligned} p'^T C(e_j - q) &= (p^* + \Delta(p_1 - p_2))^T C(e_j - q^*) \\ &\leq \Delta(p_1 - p_2)^T C(e_j - q^*) \quad (\text{since } p^{*T} C q^* \geq p^{*T} C e_j) \end{aligned}$$

From Lemma [1](#) we know that for every column j , $(p_1 - p_L)^T C(e_j - q^*) < \frac{\epsilon}{4\Delta}$. Also we have that $p_2 = \frac{1}{1-x}(p_L - x p_1)$. Using this we can write $\Delta(p_1 - p_2)^T C(e_j - q^*) = \frac{\Delta}{1-x}(p_1 - p_L)^T C(e_j - q^*) < \frac{\Delta}{1-x}(\frac{\epsilon}{4\Delta}) \leq \epsilon$ where the last step follows from $x \leq 3/4$. So the column player has less than ϵ incentive to deviate which contradicts the fact that the game is (ϵ, Δ) -approximation stable. \square

4 Polynomial-Time Algorithms When Δ and ϵ Are Close

We now show that if $\Delta \leq 2\epsilon - 6\epsilon^2$, then there must exist an $O(\epsilon)$ -equilibrium where each player's strategy has support $O(1/\epsilon)$. Thus, in this case, for constant ϵ , we have a polynomial-time algorithm for computing $O(\epsilon)$ -equilibria.

Theorem 3. *For any game satisfying (ϵ, Δ) -approximation stability for $\Delta \leq 2\epsilon - 6\epsilon^2$, there exists an $O(\epsilon)$ -equilibrium where each player's strategy has support $O(1/\epsilon)$. Thus, $O(\epsilon)$ -equilibria can be computed in time $n^{O(1/\epsilon)}$.*

Proof. Let (p^*, q^*) be a Nash equilibrium of the game. First, if there is no set S of rows having a combined total probability mass $x \in [\Delta, \Delta + \epsilon]$ in p^* , then this implies that except for rows of total probability mass less than Δ , all rows in the support of p^* have probability greater than ϵ . Therefore, p^* is Δ -close to a distribution of support at most $1/\epsilon$. If this is true for q^* as well, then this implies (p^*, q^*) is Δ -close to a pair of strategies (p, q) each of support $\leq 1/\epsilon$, which by Claim [1](#) and the assumption $\Delta < 2\epsilon$, is an $O(\epsilon)$ -equilibrium as desired. So, to prove the theorem, it suffices to show that if such a set S exists, then the game cannot satisfy (ϵ, Δ) -approximation stability for $\Delta \leq 2\epsilon - 6\epsilon^2$.

Therefore, assume for contradiction that p^* can be written as a convex combination

$$p^* = xp_1 + (1-x)p_2, \quad (7)$$

where p_1, p_2 have disjoint supports and $x \in [\Delta, \Delta + \epsilon]$. Let $\gamma = p_1^T C q^* - p_2^T C q^*$ and let $V_C = p^{*T} C q^*$. We now consider two methods for moving distance Δ from p^* : moving probability from p_1 to p_2 , and moving probability from p_2 to p_1 . Let

$$p' = (x - \Delta)p_1 + (1 - x + \Delta)p_2 \quad (8)$$

$$= (1 + \frac{\Delta}{1-x})p^* - (\frac{\Delta}{1-x})p_1. \quad (9)$$

Since p' has distance Δ from p^* and its support is contained in the support of p^* , by approximation-stability, there must exist some column e_j such that $p'^T C e_j \geq p^{*T} C e_j + \epsilon$. By [\(8\)](#) we have $p'^T C e_j = V_C - \Delta(p_1 - p_2)^T C e_j = V_C - \Delta\gamma$. By [\(9\)](#) and the fact that $p^{*T} C e_j \leq V_C$ we have that $p'^T C e_j \leq V_C(1 + \frac{\Delta}{1-x})$. Therefore we have the constraint

$$V_C(1 + \frac{\Delta}{1-x}) \geq V_C - \Delta\gamma + \epsilon. \quad (10)$$

Now, consider moving Δ probability mass from p_2 to p_1 . Specifically, let

$$p'' = (x + \Delta)p_1 + (1 - x - \Delta)p_2 \quad (11)$$

$$= (1 - \frac{\Delta}{1-x})p^* + (\frac{\Delta}{1-x})p_1. \quad (12)$$

Again, there must exist some column e_k such that $p''^T C e_k \geq p^{*T} C e_k + \epsilon$. By [\(11\)](#) we have $p''^T C e_k = V_C + \Delta(p_1 - p_2)^T C e_k = V_C + \Delta\gamma$. By [\(12\)](#) and the fact that $p^{*T} C e_k \leq V_C$ we have that $p''^T C e_k \leq V_C(1 - \frac{\Delta}{1-x}) + \frac{\Delta}{1-x}$. Therefore we have the constraint

$$V_C(1 - \frac{\Delta}{1-x}) + \frac{\Delta}{1-x} \geq V_C + \Delta\gamma + \epsilon. \quad (13)$$

From constraint [\(10\)](#) we have $V_C(\frac{\Delta}{1-x}) \geq \epsilon - \Delta\gamma$ and from constraint [\(13\)](#) we have $V_C(\frac{\Delta}{1-x}) \leq \frac{\Delta}{1-x} - \Delta\gamma - \epsilon$. Therefore, $\frac{\Delta}{1-x} \geq 2\epsilon$, contradicting $\Delta \leq 2\epsilon - 6\epsilon^2$. \square

5 Stable Games of Large Support

We now give a near-matching lower bound to the results of Section 3 showing that there exist stable games in which the Nash equilibrium and all approximate equilibria have support $\Omega(\log n)$.

Theorem 4. *For any $\Delta \leq 1/2$, there exist n -by- n games satisfying (ϵ, Δ) -approximation stability for $\epsilon = \Delta^2/32$ such that all ϵ -Nash equilibria have supports of size at least $(1 - \Delta) \lg(n)$.*

Thus, Theorem 4 implies the following near-matching lower bound to Theorem 1.

Corollary 3. *For any $\Delta \leq 1/2$ there exists an (ϵ, Δ) -approximation stable game \mathcal{G} for some $\epsilon > 0$ such that all ϵ -equilibria have support $\Omega(\frac{\Delta^4}{\epsilon^2} \log n)$.*

Proof. The proof builds on a construction in Feder et al. [10] exhibiting a game in which all approximate equilibria have support of size $\Omega(\log n)$. However, the game in [10] does not satisfy stability and so a more involved construction and argument is needed. We now present the construction of the matrix R . The game will be constant sum with $C = \mathbf{1} - R$. Let $k = \log_2(n)$ and let $\alpha = \Delta/4$. The matrix R looks like:

$$\begin{bmatrix} X & Y \\ Z & W \end{bmatrix}$$

Where:

- X is k by k with all entries equal to 0.5.
- W is $n - k$ by $n - k$ with all entries equal to 0.5.
- Z is $n - k$ by k where each row has $(0.5 - \alpha)k$ entries equal to 1 and $(0.5 + \alpha)k$ entries equal to 0. Specifically, all $\binom{k}{(0.5 - \alpha)k}$ different such rows appear. We can add multiple copies of these rows if needed to fill out the matrix.
- Y is k by $n - k$ where each column has $(0.5 - \alpha)k$ entries equal to 0 and $(0.5 + \alpha)k$ entries equal to 1. Specifically, all $\binom{k}{(0.5 - \alpha)k}$ different such columns appear. We can add multiple copies of these columns if needed to fill out the matrix.

We begin with two observations about the above construction:

Observation 1: This game has a Nash equilibrium (p^*, q^*) which is uniform over the first k rows and columns.

Observation 2: The minimax value of this game is $1/2$ to each player. So any (p, q) in which one player gets less than $1/2 - \epsilon$ is not ϵ -Nash.

We now prove that this game satisfies (ϵ, Δ) approximation-stability for $\epsilon = \Delta^2/32$. Let (p, q) be some pair of distributions such that $d((p, q), (p^*, q^*)) > \Delta$. Recall that $d((p, q), (p^*, q^*)) = \max[d(p, p^*), d(q, q^*)]$ and assume without loss of generality that $d(q, q^*) > \Delta$. We want to show that this is *not* an ϵ -Nash equilibrium. It will be convenient to write $q = q' + q''$ where q' is nonzero only over the first k columns and q'' is nonzero only over the remaining $n - k$ columns.

Case 1: Suppose that $|q''| > \beta$ for $\beta = \Delta/4$. Then, one possible response of the row player is to play p^* , achieving a payoff $p^{*T} R q$ greater than:

$$0.5(1 - \beta) + (0.5 + \alpha)\beta = 0.5 + \alpha\beta. \tag{14}$$

Thus, if $p^T Rq \leq 0.5 + \frac{\alpha\beta}{2}$ then this is not an $\frac{\alpha\beta}{2}$ -equilibrium (since the row player would have more than $\frac{\alpha\beta}{2}$ incentive to deviate to p^*) and if $p^T Rq > 0.5 + \frac{\alpha\beta}{2}$ then this is also not an $\frac{\alpha\beta}{2}$ -equilibrium (since $p^T Cq = 1 - p^T Rq < 0.5 - \frac{\alpha\beta}{2}$ and yet $p^T Cq^* \geq 0.5$ by Observation 2, so now the column player has more than $\frac{\alpha\beta}{2}$ incentive to deviate). Plugging in $\alpha = \beta = \Delta/4$, we get $\epsilon = \alpha\beta/2 = \Delta^2/32$ as desired.

Case 2: $|q''| \leq \beta$. Define $d'(q, q^*) = \sum_{i=1}^k \max(q_i - q_i^*, 0)$. So, $d'(q, q^*) > \Delta - \beta$. For conceptual convenience, let us sort the entries of q' (i.e., the first k entries of q) in decreasing order. We now claim that

$$\sum_{i=1}^{(0.5-\alpha)k} q_i > 1/2 + \alpha\beta. \quad (15)$$

This will imply at least one player has more than ϵ incentive to deviate since one possible response of the row player is to play the row in matrix Z with 1's in the first $(0.5 - \alpha)k$ entries, gaining a value greater than $1/2 + \alpha\beta$. Thus, if $p^T Rq \leq 0.5 + \alpha\beta/2$ then the row-player has more than $\alpha\beta/2$ incentive to deviate to that row in Z , and if $p^T Rq > 0.5 + \alpha\beta/2$ then the column player has more than $\alpha\beta/2$ incentive to deviate to q^* . So, all that remains is to prove inequality (15). Let $c = q_{(0.5-\alpha)k}$.

Case 2a: $c \geq 1/k$. In this case we simply use the fact that since the columns are sorted in decreasing order of q_i , at least an $(0.5 - \alpha)$ fraction of the quantity $d'(q, q^*) = \sum_{i=1}^k \max(q_i - q_i^*, 0)$ (think of this as the ‘‘excess’’ of q' over q^*) must be in the first $(0.5 - \alpha)k$ columns. In addition, we have the remaining ‘‘non-excess’’ $\sum_{i=1}^{(0.5-\alpha)k} \min(q_i, q_i^*) = [(0.5 - \alpha)k](1/k) = 0.5 - \alpha$. So, summing these two and using $d'(q, q^*) > \Delta - \beta$ we get: $\sum_{i=1}^{(0.5-\alpha)k} q_i > (0.5 - \alpha)(1 + \Delta - \beta) = 0.5 + \alpha\beta + (0.5\Delta - 0.5\beta - \alpha - \alpha\Delta) \geq 0.5 + \alpha\beta$, where the last inequality comes from our choice of $\alpha = \beta = \Delta/4$ and assumption that $\Delta \leq 1/2$.

Case 2b: $c \leq 1/k$. This implies that *all* the $d(q, q^*) - \beta$ ‘‘excess’’ of q' over q^* must be in the first $(0.5 - \alpha)k$ columns. In addition, these columns must contain at least a $(0.5 - \alpha)$ fraction of the ‘‘non-excess’’ $\sum_{i=1}^k \min(q_i, q_i^*)$. This latter quantity in turn equals $1 - d(q, q^*)$, by using the fact $d(q, q^*) = \sum_{i=1}^k \max[q_i^* - q_i, 0]$. Putting this together we have: $\sum_{i=1}^{(0.5-\alpha)k} q_i > (\Delta - \beta) + (0.5 - \alpha)(1 - \Delta) = 0.5 - \alpha + \alpha\Delta - \beta + \Delta/2 \geq 0.5 + \alpha\Delta$, where the last inequality comes from our choice of $\alpha = \beta = \Delta/4$.

This completes Case 2 and the proof. \square

This example can be extended if desired to make the game be non-constant sum and also so that the sum $R + C$ of the two matrices does not have a constant rank.

6 Inverse Conditions

In this section we consider an inverse condition to approximation-stability, namely that for some true equilibrium (p^*, q^*) , all *non*-approximate equilibria are *far* from (p^*, q^*) . In particular,

Definition 3. A game is (ϵ, Δ) -smooth if for some equilibrium (p^*, q^*) , all strategy pairs (p, q) such that $d((p, q), (p^*, q^*)) \leq \Delta$ are ϵ -equilibria.

We now show that games that are (ϵ, Δ) -smooth for Δ large compared to ϵ have the property that good approximate equilibria can be computed efficiently. (Recall by Claim 1 that all games are (ϵ, Δ) -smooth for $\Delta \leq \epsilon/3$.)

Theorem 5. There is a polynomial-time algorithm to find an (ϵ/Δ) -approximate equilibrium in any game that is (ϵ, Δ) -smooth.

We prove Theorem 5 through a series of claims as follows.

Claim. Let \mathcal{G} be (ϵ, Δ) -smooth for equilibrium (p^*, q^*) . Then for every row i we have $e_i^T Rq^* \geq p^{*T} Rq^* - \epsilon/\Delta$.

Proof. Let $V_R = p^{*T} Rq^*$. Since (p^*, q^*) is a Nash equilibrium, any row $e_i \in \text{supp}(p^*)$ will get an expected payoff of V_R against q^* as well. Now consider a row $e_i \notin \text{supp}(p^*)$. Let $p = (1 - \Delta)p^* + \Delta e_i$ and consider the pair (p, q^*) . This pair is Δ -close to (p^*, q^*) and hence, by the assumption that the game is (ϵ, Δ) -smooth, must be an ϵ -equilibrium. This means that $p^T Rq^* \geq V_R - \epsilon$. So we get $(1 - \Delta)p^{*T} Rq^* + \Delta e_i^T Rq^* \geq V_R - \epsilon$, and using the fact that $p^{*T} Rq^* = V_R$, this implies that $e_i^T Rq^* \geq V_R - \frac{\epsilon}{\Delta}$.

Similarly, we have:

Claim. Let \mathcal{G} be (ϵ, Δ) -smooth for equilibrium (p^*, q^*) . Then for every column j we have $p^{*T} C e_j \geq p^{*T} C q^* - \epsilon/\Delta$.

Using these claims, we can efficiently compute an $\frac{\epsilon}{\Delta}$ -approximate equilibrium in smooth games.

Proof (Theorem 5): Solve the following linear program for a pair of strategies p, q and values V_R, V_C :

$$e_i^T Rq \geq V_R - \frac{\epsilon}{\Delta}, \quad \forall i \quad (16)$$

$$e_i^T Rq \leq V_R, \quad \forall i \quad (17)$$

$$p^T C e_j \geq V_C - \frac{\epsilon}{\Delta}, \quad \forall j \quad (18)$$

$$p^T C e_j \leq V_C, \quad \forall j \quad (19)$$

From the previous claims we have that $(p^*, q^*, V_R = p^{*T} Rq^*, V_C = p^{*T} Cq^*)$ is a feasible solution to the above LP. Also, when playing (p, q) , the row and the column players are getting expected payoff at least $V_R - \frac{\epsilon}{\Delta}$ and $V_C - \frac{\epsilon}{\Delta}$ respectively. Furthermore, by deviating from p , the row player can get a payoff of at most V_R and by deviating from q , the column player cannot get more than V_C . Hence, (p, q) is an $\frac{\epsilon}{\Delta}$ -approximate Nash equilibrium. \square

7 Open Questions and Conclusions

In this work we define and analyze a natural notion of approximation-stability for 2-player general-sum games, motivated by the goal of finding approximate equilibria for

predictive purposes. We show that one can improve over the general Lipton et al. [11] bound based on the extent to which the given game satisfies this condition. Furthermore, if $\Delta < 2\epsilon - O(\epsilon^2)$ we show there must exist approximate equilibria of small support, yielding an algorithm to find them in time $n^{O(1/\epsilon)}$. On the other hand, we show that approximation-stable games with $\Delta = O(\sqrt{\epsilon})$ can have all approximate equilibria of support $\Omega(\log n)$. We also analyze an inverse condition for which we show finding (ϵ/Δ) -approximate equilibria can be done efficiently. One open problem is to better understand for what values of Δ (as a function of ϵ) one can find $O(\epsilon)$ -approximate equilibria efficiently under the assumption of (ϵ, Δ) -approximation-stability. For instance, can one extend the $n^{O(1/\epsilon)}$ -time algorithm from $\Delta < 2\epsilon - O(\epsilon^2)$ to $\Delta = \text{poly}(\epsilon)$? Recently Balcan and Braverman [3] have shown this may be intrinsically hard: specifically, for $\Delta = \epsilon^{1/4}$, they show an $n^{\text{poly}(1/\epsilon)}$ algorithm to find ϵ -equilibria in such games would imply a PTAS in general games. In fact, [3] motivates the following interesting question: could there be an algorithm that for every (ϵ, Δ) finds a Δ -equilibrium in time $O(n^{\text{poly}(1/\epsilon)})$? This may be solvable even if a PTAS is hard for general games, which itself still remains an open question.

References

1. Abbott, T.G., Kane, D.M., Valiant, P.: On the complexity of two-player win-lose games. In: Proc. 46th FOCS, pp. 113–122 (2005)
2. Balcan, M.F., Blum, A., Gupta, A.: Approximate clustering without the approximation. In: Proc. 20th SODA, pp. 1068–1077 (2009)
3. Balcan, M.F., Braverman, M.: Approximate Nash equilibria under stability conditions (2010) (manuscript)
4. Bárány, I., Vempala, S., Vetta, A.: Nash equilibria in random games. *Random Struct. Algorithms* 31(4), 391–405 (2007)
5. Bosse, H., Byrka, J., Markakis, E.: New algorithms for approximate Nash equilibria in bimatrix games. *Theor. Comput. Sci.* 411(1), 164–173 (2010)
6. Chen, X., Deng, X.: Settling the complexity of two-player Nash equilibrium. In: Proc. 47th FOCS, pp. 261–272 (2006)
7. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39(1), 195–259 (2009)
8. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: Progress in approximate Nash equilibria. In: Proc. 8th ACM-EC (2007)
9. Etesami, K., Yannakakis, M.: On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.* 39(6), 2531–2597 (2010)
10. Feder, T., Nazerzadeh, H., Saberi, A.: Approximating nash equilibria using small-support strategies. In: Proc. 8th ACM-EC, pp. 352–354 (2007)
11. Lipton, R.J., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: Proc. 4th ACM-EC, pp. 36–41 (2003)
12. Tsaknakis, H., Spirakis, P.: An optimization approach for approximate Nash equilibria. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 42–56. Springer, Heidelberg (2007)

Improved Lower Bounds on the Price of Stability of Undirected Network Design Games*

Vittorio Bilò¹, Ioannis Caragiannis², Angelo Fanelli³, and Gianpiero Monaco⁴

¹ Department of Mathematics, University of Salento, Lecce, Italy

² RACTI & Department of Computer Engineering and Informatics,
University of Patras, Greece

³ Division of Mathematical Sciences, School of Physical and Mathematical Sciences,
Nanyang Technological University, Singapore

⁴ Mascotte Project, I3S (CNRS/UNSA) INRIA, Sophia Antipolis, France

Abstract. Bounding the price of stability of undirected network design games with fair cost allocation is a challenging open problem in the Algorithmic Game Theory research agenda. Even though the generalization of such games in directed networks is well understood in terms of the price of stability (it is exactly H_n , the n -th harmonic number, for games with n players), far less is known for network design games in undirected networks. The upper bound carries over to this case as well while the best known lower bound is $42/23 \approx 1.826$. For more restricted but interesting variants of such games such as broadcast and multicast games, sublogarithmic upper bounds are known while the best known lower bound is $12/7 \approx 1.714$. In the current paper, we improve the lower bounds as follows. We break the psychological barrier of 2 by showing that the price of stability of undirected network design games is at least $348/155 \approx 2.245$. Our proof uses a recursive construction of a network design game with a simple gadget as the main building block. For broadcast and multicast games, we present new lower bounds of $20/11 \approx 1.818$ and 1.862, respectively.

1 Introduction

Network design is among the most well-studied problems in the combinatorial optimization literature. A natural definition is as follows. We are given a graph consisting of a set of nodes and edges among them representing potential links. Each edge has an associated cost which corresponds to the cost for establishing the corresponding link. We are also given connectivity requirements as pairs of source-destination nodes. The objective is to compute a subgraph of the original graph of minimum total cost that satisfies the connectivity requirements. In other words, we seek to establish a network that satisfies the connectivity requirements

* This work was partially supported by the grant NRF-RF2009-08 “Algorithmic aspects of coalitional games” and the PRIN 2008 research project COGENT (COmputational and GamE-theoretic aspects of uncoordinated NeTworks) funded by the Italian Ministry of University and Research.

at the minimum cost. This optimization problem is known as Minimum Steiner Forest and generalizes well-studied problems such as the Minimum Spanning Tree and Minimum Steiner Tree.

In this paper, we consider a game-theoretic variant of network design that was first considered in [2]. Instead of considering the connectivity requirements as a global goal, we assume that each connectivity requirement is desired by a different player. The players participate in a non-cooperative game; each of them selects as her strategy a path from her source to the destination and is charged for part of the cost of the edges she uses. According to the fair cost sharing scheme we consider in the current paper, the cost of an edge is shared equally among the players using the edge. The *social cost* of an *assignment* (i.e., a snapshot of players' strategies) is the cost of the edges contained in at least one path. An optimal assignment would contain a set of edges of minimum cost so that the connectivity requirements of the players are satisfied. Unfortunately, this does not necessarily mean that all players are satisfied with this assignment since a player may have an incentive to deviate from its path to another one so that her individual cost is smaller. Eventually, the players will reach a set of strategies (and a corresponding network) that satisfies their connectivity requirements and in which no player has any incentive to deviate to another path; such outcomes are known as *Nash equilibria*. Interestingly, even though the optimal solution is always a forest, Nash equilibria may contain cycles.

The non-optimality of the outcomes of *network design games* (which is typical when selfish behavior comes into play) leads to the following question that has been a main line of research in Algorithmic Game Theory: How is the system performance affected by selfish behavior? The notion of the *price of anarchy* (introduced in [8]; see also [10]) quantifies the deterioration of performance. In general terms, it is defined as the ratio of the social cost of the worst possible Nash equilibrium over the optimal cost. Hence, it is pessimistic in nature and (as its name suggests) provides a worst-case guarantee for conditions of total anarchy. Instead, the notion of the *price of stability* (introduced in [2]) is optimistic in nature. It is defined as the ratio of the social cost of the best equilibrium over the optimal cost and essentially asks: What is the best one can hope for the system performance given that the players are selfish?

The aim of the current paper is to determine better lower bounds on the price of stability for network design games in an attempt to understand the effect of selfishness on the efficiency of outcomes in such games. We usually refer to network design games as multi-source network design games in order to capture the most general case in which players may have different sources. An interesting variant is when each player wishes to connect a particular common node, which we call the *root*, with her destination node; we refer to such network design games as multicast games. An interesting special case of multicast games is the class of broadcast games: in such games, there is a player for each non-root node of the network that has this node as her destination.

The existence of Nash equilibria in network design games is guaranteed by a potential function argument. Rosenthal [11] defined a potential function over all

assignments of a network design game so that the difference in the potential of two assignments that differ in the strategy of a single player equals the difference of the cost of that player in these assignments; hence, an assignment that locally minimizes the potential function is a Nash equilibrium. So, the price of stability is well-defined in network design games. Anshelevich *et al.* [2] considered network design games in directed graphs and proved that the price of stability is at most H_n . Their proof considers a Nash equilibrium that can be reached from an optimal assignment when the players make arbitrary selfish moves. The main argument used is that the potential of the Nash equilibrium is strictly smaller than that of the optimal assignment and the proof follows due to the fact that the potential function of Rosenthal approximates the social cost of an assignment within a factor of at most H_n . This approach suggests a general technique for bounding the price of stability and has been extended to other games as well; see [3,5]. For directed graphs, the bound of H_n was also proved to be tight [2]. Although the upper bound proof carries over to undirected network design games, the lower bound does not. The bound of H_n is the only known upper bound for multi-source network design games in undirected graphs. Better upper bounds are known for single-source games. For broadcast games, Fiat *et al.* [7] proved an upper bound of $O(\log \log n)$ while Li [9] presented an upper bound of $O(\log n / \log \log n)$ for multicast games. These bounds are not known to be tight either and, actually, the gap with the corresponding lower bounds is large. For single-source games, in the full version of [7] Fiat *et al.* present a lower bound of $12/7 \approx 1.714$; their construction uses a broadcast game. This was the best lower bound known for the multi-source case as well until the recent work of Christodoulou *et al.* [6] who presented an improved lower bound of $42/23 \approx 1.826$. Higher (i.e., super-constant) lower bounds are only known for weighted variants of network design games (see [1,4]).

In this paper, we present better lower bounds for general undirected network design games, as well as for the restricted variants of broadcast and multicast games. For the general case, we present a game that has price of stability at least $348/155 \approx 2.245$, improving the previously best known lower bound of [6]. Our proof uses a simple gadget as the main building block which is augmented by a recursive construction to our lower bound instance. The particular recursive construction of the game has two advantages. Essentially, the recursive construction blows up the price of stability of the gadget used as the main building block. Furthermore, recursion allows to handle successfully the technical difficulties in the analysis. We believe that our construction could be extended to use more complicated gadgets as building blocks that would probably lead to better lower bounds on the price of stability. For multicast games, we present a lower bound of 1.862. Our proof uses a game on a graph with a particular structure. For this game, we prove sufficient conditions on the edge costs of the graph so that a particular assignment is the unique Nash equilibrium. Then, the construction that yields the lower bound is the solution of a linear program which has the edge costs as variables, the sufficient conditions as constraints, an additional constraint that upper-bounds the optimal cost by 1, and its objective

is to maximize the cost of the unique Nash equilibrium. The particular lower bound was obtained in a game with 100 players using the linear programming solver of Matlab. A slight variation of this construction yields our lower bound for broadcast games. In this case, we are able to obtain a more compact set of sufficient conditions so that there is a unique Nash equilibrium. As a result, we have a formal proof that the price of stability approaches $20/11 \approx 1.818$ when the number of players is large.

2 Preliminaries

In an undirected network design game, we are given an undirected graph $G = (V, E)$ in which each edge $e \in E$ has a non-negative cost c_e . There are n players; player i wishes to establish a connection between two nodes $s_i, t_i \in V$ called the source and destination node of player i , respectively. The set of strategies available to player i consists of all paths connecting nodes s_i and t_i in G . We call an *assignment* any set of strategies σ , with one strategy per player. Given an assignment σ , let $n_e(\sigma)$ be the number of players using edge e in σ . Then, the cost of player i in σ is defined as $\text{cost}_i(\sigma) = \sum_{e \in \sigma_i} \frac{c_e}{n_e(\sigma)}$. Let $G(\sigma)$ be the subgraph of G which contains the edges of G that are used by at least one player in assignment σ . The social cost of the assignment σ is simply the total cost of the edges in $G(\sigma)$ which coincides with the sum of the costs of the players.

An assignment σ is called a *Nash equilibrium* if for any player i and for any other assignment σ' that differs from σ only in the strategy of player i , it holds $\text{cost}_i(\sigma) \leq \text{cost}_i(\sigma')$. It can be easily seen that any Nash equilibrium is a *proper* assignment, in the sense that the edges used by any pair of players do not form any cycle. The *price of stability* of a network design game is defined as the ratio of the minimum social cost among all Nash equilibria over the optimal cost.

Network design games with $s_i = s$ for any player i are called *multicast* games. We refer to node s as the *root* node. Multicast games in which there is one player for any non-root node that has this node as destination are called *broadcast* games. We also use the term *multi-source* games to refer to the general class of undirected network design games and the term *single-source* games in order to refer to multicast and broadcast games.

3 The Lower Bound for Multi-source Games

In this section, we prove the following theorem.

Theorem 1. *For any $\delta > 0$, there exists an undirected network design game with price of stability at least $348/155 - \delta$.*

We will construct a network design game on a connected undirected graph so that there is a distinct player associated with each edge of the graph that wishes to connect the endpoints of the edge. The construction uses integer parameters $k \geq 3$ and $t \geq 2$. We start with the *gadget* construction depicted in Figure [11a](#).

We use the terms *left gadget player* and *right gadget player* for the players associated with the left and right gadget edge of a gadget, respectively. We also use the term *floor players* for the players associated with floor edges. Given an edge e , we build a *block under this edge* by putting k gadgets so that the leftmost node of the first gadget coincides with the left endpoint of e , the rightmost node of i -th gadget coincides with the leftmost node of the $(i + 1)$ -th gadget for $i = 1, \dots, k - 1$, and the rightmost node of the k -th gadget coincides with the right endpoint of e (see Figure 1(b)). We refer to e as the *ceiling edge* of the block.




Fig. 1. (a) The gadget used in the proof of Theorem 1 (b) The construction of a block under a ceiling edge (with $k = 3$).

We set $x = 28/109$, $y = 33/109$, $z = 30/109 - \epsilon$, $w = 35/109 - \epsilon$, and $\alpha = 63/218 - \epsilon$, where ϵ is a negligibly small but strictly positive number. If g denotes the cost of the ceiling edge, then the cost of the edges in each gadget of the block under it are defined as follows: $\frac{xg}{\alpha k^2}$ for each of the left floor edges, $\frac{(1-x-y)g}{\alpha k^2}$ for each of the middle floor edges, $\frac{yg}{\alpha k^2}$ for each of the right floor edges, $\frac{zg}{\alpha k}$ for the left gadget edge, and $\frac{wg}{\alpha k}$ for the right gadget edge. So, the total cost of the floor edges of the block is g/α while the total cost of all edges of the block is $g(1 + z + w)/\alpha$.

Now, our construction starts with a *roof edge* of cost 1 (and an associated *roof player*) and a block under it. The roof edge has level t and the block under it has level $t - 1$. We build blocks of level $t - 2$ by building a block under each of the floor edges of the block of level $t - 1$. We continue recursively and define all blocks down to level 1. Clearly, for $j = 1, \dots, t - 1$, the total cost of the floor edges of level j is $g\alpha^{j-t}$ while the total cost of all edges of level j is $g(1 + z + w)\alpha^{j-t}$.

Hence, the total cost of the edges in the graph is

$$1 + \sum_{i=1}^{t-1} (1 + z + w)\alpha^{-i} = \frac{348 - 436\epsilon}{155 + 218\epsilon}\alpha^{1-t} - \frac{193 - 654\epsilon}{155 + 218\epsilon}$$

while the cost of the floor edges of level 1 is α^{1-t} and upper-bounds the optimal cost (since the floor edges of level one constitute a spanning tree of the whole graph). For any $\delta > 0$, we can set t and ϵ appropriately so that the ratio of the total cost of edges over the optimal cost is at least $348/155 - \delta$.

In order to complete the proof of the theorem, it suffices to prove that the assignment in which each player uses her direct edge is the unique Nash equilibrium; the rest of this section is devoted to proving this claim. We will refer to the players associated to floor edges (respectively, gadget edges) at blocks of level j as the floor players of level j (respectively, the gadget players of level j). A floor player of level j follows a *non-increasing* strategy if she uses neither a gadget edge of her gadget nor any edge of level $j' > j$. A gadget player of level j follows a non-increasing strategy if she does not use any edge of level $j' > j$. In the opposite case, we say that the player follows an *increasing* strategy.

In an assignment, a player may use a floor edge or connect its endpoints by being routed through the block under the edge. In the latter case, we say that the player *crosses* the floor edge. We also say that a player is *external* to a gadget (respectively, external to a block) if she does not correspond to any edge of the gadget (respectively, block) and uses or crosses its edges.

In a proper assignment, the sets of non-increasing strategies of the gadget players of a gadget can belong to one of the following types (Figure 2); any other set of non-increasing strategies violates the fact that the assignment is proper.

- Type A: Both gadget players use their direct edges.
- Type B: The left gadget player uses her direct edge and the right gadget player uses or crosses the middle and right floor edges.
- Type C: Both gadget players use the left gadget edge. The right gadget player uses or crosses the left and right floor edges as well.
- Type D: The right gadget player uses her direct edge and the left gadget player uses or crosses the left and middle floor edges.
- Type E: Both gadget players use the right gadget edge. The left gadget player uses or crosses the left and right floor edges as well.
- Type F: The left gadget player uses or crosses the left and middle floor edges and the right gadget player uses or crosses the middle and right floor edges.

We are ready to significantly restrict the structure of assignments we have to consider as candidates to be Nash equilibria.

Lemma 1. *At any Nash equilibrium, all players besides the roof player follow non-increasing strategies. Furthermore, at each block: either there are no external players and the gadget players have strategies of type A or there are $h > 0$ external players and each of them experiences cost more than g/h , where g is the cost of the ceiling edge of the block.*

Proof. Consider a Nash equilibrium. We will prove the claim inductively (on the block level). We will first prove it for the blocks of level 1. In this case, there is no block under any floor edge and players do not cross the floor edges.




Fig. 2. The six possible types for the players of a gadget that follow non-increasing strategies. The dashed lines denote the paths used by the left and the right gadget player. Only the gadget edges that are used by some player are shown.

Consider a block of level 1 and assume that a floor player p follows an increasing strategy. Then, she should connect the endpoints of her floor edge to the two closest gadget edge endpoints by using $k - 1$ floor edges. Furthermore, observe that neither a gadget player of the same gadget nor an external player to this gadget uses these floor edges (since this would imply that they also use the direct edge of player p and the assignment would not be proper). Similarly, the players associated to the $k - 1$ floor edges use their direct edges. Hence, player p uses each of the $k - 1$ floor edges together with one floor player. Since $k \geq 3$, this means that the cost she experiences at the $k - 1 \geq 2$ floor edges plus the non-zero cost she experiences at the other edges she uses is strictly larger than the cost of her direct edge and she would have an incentive to move to its direct edge. So, all floor players of the block follow non-increasing strategies.

Now, assume that a gadget player p follows an increasing strategy, i.e., her path contains the endpoints of her gadget. This means that there are no external players to the current block nor other gadget players within the current block that follow increasing strategies (any such player should connect the endpoints of the gadget of p and the assignment would not be proper). So, there are at least $k - 1$ gadgets whose gadget (and floor) players follow non-increasing strategies.

We focus on such a gadget of the current block and assume that there are $h \geq 0$ external players; these can be players that are external to the block or a player from another gadget of the same block that follows an increasing strategy. In the inequalities below, we use the following claim.

Claim. Let ζ, η be positive integers. Then, $\frac{1}{\zeta+h} \geq \frac{\eta}{(\zeta+\eta)h}$ for any integer $h \geq \eta$.

We consider the six different cases for the strategies of the gadget players. If the strategies of the gadget players are of type A, then all the external players (if any) are routed either through the left gadget edge and the right floor edges of the gadget, or through the left floor edges and the right gadget edge, or through the left gadget edge, the middle floor edges, and the right gadget edge (any other case violates the fact that the assignment is proper). In the first subcase, the cost of each external player at the edges of the gadget is

$$\frac{g}{\alpha k} \left(\frac{z}{1+h} + \frac{y}{1+h} \right) \geq \frac{g}{\alpha k h} \left(\frac{z}{2} + \frac{y}{2} \right) = \frac{g}{\alpha k h} \left(\frac{63}{218} - \frac{\epsilon}{2} \right) > \frac{g}{k h}.$$

In the second subcase, the cost of each external player is

$$\frac{g}{\alpha k} \left(\frac{x}{1+h} + \frac{w}{1+h} \right) \geq \frac{g}{\alpha kh} \left(\frac{x}{2} + \frac{w}{2} \right) = \frac{g}{\alpha kh} \left(\frac{63}{218} - \frac{\epsilon}{2} \right) > \frac{g}{kh}.$$

In the third subcase, the cost of each external player is again

$$\frac{g}{\alpha k} \left(\frac{z}{1+h} + \frac{1-x-y}{1+h} + \frac{w}{1+h} \right) \geq \frac{g}{\alpha kh} \left(\frac{w}{2} + \frac{1-x-y}{2} + \frac{w}{2} \right) > \frac{g}{kh}.$$

If the strategies of the gadget players are of type B, all the external players are routed through the left gadget edge and the right floor edges. We will first show that $h \geq 2$. Indeed, if at most one external player is routed through the gadget, the cost of the right gadget player would be at least

$$\frac{g}{\alpha k} \left(\frac{1-x-y}{2} + \frac{y}{3} \right) = \frac{g}{\alpha k} \cdot \frac{35}{109} > \frac{gw}{\alpha k},$$

i.e., this player would have an incentive to move and use her direct edge. So, since $h \geq 2$, the cost of each external player at the edges of the gadget is

$$\frac{g}{\alpha k} \left(\frac{z}{1+h} + \frac{y}{2+h} \right) \geq \frac{g}{\alpha kh} \left(\frac{2z}{3} + \frac{y}{2} \right) = \frac{g}{\alpha kh} \left(\frac{73}{218} - \frac{2\epsilon}{3} \right) > \frac{g}{kh}.$$

If the strategies of the gadget players are of type C, all the external players are routed through the left gadget edge and the right floor edges. We will show again that $h \geq 2$. Indeed, if at most one external player is routed through the gadget, the cost of the right gadget player would be at least

$$\frac{g}{\alpha k} \left(\frac{x}{2} + \frac{z}{3} + \frac{y}{3} \right) = \frac{g}{\alpha k} \left(\frac{35}{109} - \frac{\epsilon}{3} \right) > \frac{gw}{\alpha k},$$

i.e., this player would have an incentive to move. So, since $h \geq 2$, the cost of each external player at the edges of the gadget is

$$\frac{g}{\alpha k} \left(\frac{z}{2+h} + \frac{y}{2+h} \right) \geq \frac{g}{\alpha kh} \left(\frac{z}{2} + \frac{y}{2} \right) = \frac{g}{\alpha kh} \left(\frac{63}{218} - \frac{\epsilon}{2} \right) > \frac{g}{kh}.$$

If the strategies of the gadget players are of type D, all the external players are routed through the left floor edges and the right gadget edge. We will show again that $h \geq 2$. Indeed, if at most one external player is routed through the gadget, the cost of the left gadget player would be at least

$$\frac{g}{\alpha k} \left(\frac{x}{3} + \frac{1-x-y}{2} \right) = \frac{g}{\alpha k} \cdot \frac{100}{327} > \frac{gz}{\alpha k},$$

i.e., this player would have an incentive to move. So, since $h \geq 2$, the cost of each external player at the edges of the gadget is

$$\frac{g}{\alpha k} \left(\frac{x}{2+h} + \frac{w}{1+h} \right) \geq \frac{g}{\alpha kh} \left(\frac{x}{2} + \frac{2w}{3} \right) = \frac{g}{\alpha kh} \left(\frac{112}{327} - \frac{2\epsilon}{3} \right) > \frac{g}{kh}.$$

If the strategies of the gadget players are of type E, then all the external players are routed through the left floor edges and the right gadget edge. We will show again that $h \geq 2$. Indeed, if at most one external player is routed through the gadget, the cost of the left gadget player would be at least

$$\frac{g}{\alpha k} \left(\frac{x}{3} + \frac{w}{3} + \frac{y}{2} \right) = \frac{g}{\alpha k} \left(\frac{75}{218} - \frac{\epsilon}{3} \right) > \frac{gz}{\alpha k},$$

i.e., this player would have an incentive to move. So, since $h \geq 2$, the cost of each external player at the edges of the gadget is

$$\frac{g}{\alpha k} \left(\frac{x}{2+h} + \frac{w}{2+h} \right) \geq \frac{g}{\alpha kh} \left(\frac{x}{2} + \frac{w}{2} \right) = \frac{g}{\alpha kh} \left(\frac{63}{218} - \frac{\epsilon}{2} \right) > \frac{g}{kh}.$$

If the strategies of the gadget players are of type F, then all the external players are routed through the floor edges. We will show that $h > 0$ in this case. Indeed, if there were no external players that are routed through the gadget, the cost of the left gadget player would be

$$\frac{g}{\alpha k} \left(\frac{x}{2} + \frac{1-x-y}{3} + \frac{y}{2} \right) = \frac{g}{\alpha k} \cdot \frac{93}{218} > \frac{gz}{\alpha k},$$

i.e., this player would have an incentive to move. So, the cost of each external player at the edges of the gadget is

$$\frac{g}{\alpha k} \left(\frac{x}{2+h} + \frac{1-x-y}{3+h} + \frac{y}{2+h} \right) \geq \frac{g}{\alpha kh} \left(\frac{x}{3} + \frac{1-x-y}{4} + \frac{y}{3} \right) > \frac{g}{kh}.$$

Now, consider again the gadget player p which follows an increasing strategy. In each of the other $k - 1$ gadgets of the same block, the gadget players have strategies of types A or F and the cost player p experiences at the edges of the gadget is more than $\frac{g}{k}$. Her total cost through the edges of the $k - 1 \geq 2$ gadgets different than her own one would be $\frac{g(k-1)}{k} \geq \frac{g}{\alpha k} \max\{z, w\}$, i.e., she would have an incentive to move and use her direct edge instead. So, all gadget players of the block follow non-increasing strategies as well.

Now, assume that no external player is routed through the block. Then, by the above discussion, the only case in which the gadget players of a gadget do not have an incentive to move is when they follow strategies of type A. If one external player is routed through the block, then the gadget players follow strategies of type A or F and the cost experienced by the external player at each gadget is more than g/k , i.e., more than g in total. If $h \geq 2$ external players are routed through the block, then each of them experiences cost more than $\frac{g}{kh}$ at each gadget, i.e., more than g/h in total.

We have completed the proof of the base of the induction. Now, assuming that the statement is true for blocks of levels up to j , we have to prove it for blocks of level $j + 1$. The proof of the induction step is almost identical to the proof of the induction base. The only difference is that, now, a player may cross a floor edge in order to connect its endpoints. Then, when h players cross a floor edge,

they are external to the block under the edge and (by the induction hypothesis) the cost they experience when crossing the edge is more than its cost over h (as opposed to exactly its cost over h which we had in the induction base). This inequality (instead of equality) does not affect any of the inequalities above and the proof of the induction step completes in the very same way. \square

Lemma 2. *At any Nash equilibrium, there are no external players at any block.*

Proof. Assume that this is not the case and consider a Nash equilibrium with external players at some block. Consider the block of highest level that has some external player routed through it. Then, it is either the block of level $t - 1$ or (by Lemma 1) some block under a floor edge of a gadget of the higher-level block whose gadget players follows strategies of type A. In both cases, exactly one player is routed through the block (i.e., the player corresponding to its ceiling edge) and, by Lemma 1, her cost at the edges of the block is more than the cost of the ceiling edge of the block. Hence, this player has an incentive to move and use the ceiling edge instead. The lemma follows. \square

Now, Theorem 1 follows by Lemmas 1 and 2 since they imply that the assignment in which every player uses her direct edge is the unique Nash equilibrium.

4 Lower Bounds for Single-Source Games

In this section, we present our lower bounds for multicast and broadcast games. We note that since all players have a common source node in such games, in any proper assignment the set of edges that are used by at least one player is a tree that is rooted at the source node and spans the destinations of all players. Also, any such tree defines in a unique way the strategies of the players in a proper assignment. So, when considering Nash equilibria in multicast or broadcast games, it suffices to restrict our attention to assignments defined by trees spanning the root node and the destination nodes of all players. We refer to them as *multicast* or *broadcast trees* depending on whether the game is a multicast or a broadcast game.

Our lower bound for multicast games uses the graph M_n depicted in Figure 3. There are n players; player i wishes to connect node s to node t_i . The cost of the edges is defined by the tuple $C = (x_2, \dots, x_n, y_1, \dots, y_n, z_1, \dots, z_n)$. We denote by τ the multicast tree formed by the edges (s, t_i) for $i = 1, \dots, n$. The next lemma provides a sufficient condition so that the assignment defined by tree τ is the unique Nash equilibrium of the multicast game on M_n ; its formal proof is omitted due to lack of space.

Lemma 3. *The assignment defined by tree τ is the unique Nash equilibrium of the multicast game on graph M_n if C is such that for $i = 2, \dots, n$ and for $k = 1, \dots, i - 1$ it holds*

$$z_k < \frac{z_i}{\min\{2i - 2k, n - k\} + 1} + \frac{y_i}{\min\{2i - 2k, n - k\}} + \sum_{p=0}^{i-k-1} \frac{x_{i-p}}{i - k - p} + y_k,$$




Fig. 3. The graphs M_n (left) and B_n (right)

and for $i = 1, \dots, n - 1$ and for $j = i + 1, \dots, n$ it holds

$$z_j < \frac{z_i}{\min\{2j - 2i, j\}} + \frac{y_i}{\min\{2j - 2i, j\} - 1} + \sum_{p=1}^{j-i} \frac{x_{i+p}}{j - i - p + 1} + y_j.$$

Now, we can use Lemma 3 to obtain lower bounds on the price of stability of multicast games by solving the following linear program. The variables of the linear program are the edge costs of the tuple C . The objective is to maximize the cost $\sum_{i=1}^n z_i$ of tree τ subject to the two sets of constraints in the statement of Lemma 3 and the additional constraint $z_1 + \sum_{i=2}^n x_i + \sum_{i=1}^n y_i \leq 1$ which upper-bounds the optimal cost by 1 (observe that the left-hand side of this constraint is the cost of the multicast tree containing all edges of M_n besides (s, t_i) for $i = 2, \dots, n$). Then, the objective value of this linear program denotes the price of stability of the multicast game on M_n for the particular values of the edge costs that correspond to the solution of the linear program. We obtained our lower bound on the price of stability using the linear programming solver of Matlab. Note that we have used $n = 100$ and have simulated the strict inequalities in the conditions of Lemma 3 by using standard inequalities and adding a constant of 10^{-6} on their left-hand side. The following statement summarizes our best observed lower bound.

Theorem 2. *There exists a multicast game with price of stability at least 1.862.*

Our lower bound for broadcast games uses the graph B_n depicted at the right part of Figure 3. In this case, the cost of the edges is defined by the tuple $C = (x_2, \dots, x_n, z_1, \dots, z_n)$. Again, there are n players; player i wishes to connect node s to node t_i . Denote by τ the broadcast tree formed by the edges (s, t_i) for $i = 1, \dots, n$. Observe that the graph B_n is obtained from M_n by contracting the edges (t_i, v_i) . Hence, any Nash equilibrium of the multicast game on graph M_n with $y_i = 0$ for $i = 1, \dots, n$ corresponds to a Nash equilibrium of the broadcast game on graph B_n of the same cost (and vice versa) while the cost of the optimal assignment is the same in both cases. So, we can apply the same technique we

used above by further constraining the variable y_i to be zero for $i = 1, \dots, n$. Fortunately, we are able to define a much more compact set of conditions for C in order to guarantee that the assignment defined by τ is the unique Nash equilibrium of the broadcast game on B_n . Our related result is the following; due to lack of space, the formal proof is omitted.

Theorem 3. *For any $\delta > 0$, there exists a broadcast game with price of stability at least $20/11 - \delta$.*

We remark that the graph B_n has the same structure with the lower bound construction of [7] albeit with a different definition of the edge costs that yields the improved lower bound on the price of stability.

References

1. Albers, S.: On the value of coordination in network design. In: Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 294–303 (2008)
2. 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 on Computing* 38(4), 1602–1623 (2008)
3. 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)
4. Chen, H.-L., Roughgarden, T.: Network design with weighted players. *Theory of Computing Systems* 45, 302–324 (2009)
5. Christodoulou, G., Koutsoupias, E.: 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., Chung, C., Ligett, K., Pyrga, E., van Stee, R.: On the price of stability for undirected network design. In: Bampis, E., Jansen, K. (eds.) Approximation and Online Algorithms. LNCS, vol. 5893, pp. 86–97. Springer, Heidelberg (2010)
7. Fiat, A., Kaplan, H., Levy, M., Olonetsky, S., Shabo, R.: On the price of stability for designing undirected networks with fair cost allocations. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 608–618. Springer, Heidelberg (2006)
8. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
9. Li, J.: An $O(\log n / \log \log n)$ upper bound on the price of stability for undirected Shapley network design games. *Information Processing Letters* 109(15), 876–878 (2009)
10. Papadimitriou, C.H.: Algorithms, games and the internet. In: Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC), pp. 749–753 (2001)
11. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)

On the Rate of Convergence of Fictitious Play

Felix Brandt¹, Felix Fischer², and Paul Harrenstein¹

¹ Institut für Informatik, Technische Universität München, 80538 München, Germany
{brandtf,harrenst}@in.tum.de

² Harvard School of Engineering and Applied Sciences, Cambridge, MA 02138, USA
fischerf@seas.harvard.edu

Abstract. Fictitious play is a simple learning algorithm for strategic games that proceeds in rounds. In each round, the players play a best response to a mixed strategy that is given by the empirical frequencies of actions played in previous rounds. There is a close relationship between fictitious play and the Nash equilibria of a game: if the empirical frequencies of fictitious play converge to a strategy profile, this strategy profile is a Nash equilibrium. While fictitious play does not converge in general, it is known to do so for certain restricted classes of games, such as constant-sum games, non-degenerate $2 \times n$ games, and potential games. We study the rate of convergence of fictitious play and show that, in all the classes of games mentioned above, fictitious play may require an exponential number of rounds (in the size of the representation of the game) before *some* equilibrium action is eventually played. In particular, we show the above statement for symmetric constant-sum win-lose-tie games.

1 Introduction

A common criticism of Nash equilibrium, the most prominent solution concept of the theory of strategic games, is that it fails to capture how players' deliberation processes actually reach a steady state. When considering a set of human or artificial agents engaged in a parlor game or a more austere decision-making situation, it is somewhat hard to imagine that they would after some deliberation arrive at a Nash equilibrium, a carefully chosen probability distribution over all possible courses of action. One reason why this behavior is so hard to imagine is that Nash equilibrium rests on rather strong assumptions concerning the rationality of players and the ability to reliably carry out randomizations. Another concern is that in many settings finding a Nash equilibrium is computationally intractable.

A more reasonable scenario would be that agents face a strategic situation by playing the game in their heads, going through several rounds of speculation and counterspeculation as to how their opponents might react and how they would react in turn. This is the idea underlying *fictitious play* (FP). FP proceeds in rounds. In the first round, each player arbitrarily chooses one of his actions. In subsequent rounds, each player looks at the empirical frequency of play of their respective opponents in previous rounds, interprets it as a probability distribution, and myopically plays a pure best response against this distribution. FP

can also be seen as a *learning algorithm* for games that are played repeatedly, such that the intermediate best responses are actually played. This interpretation rests on the simplifying assumption that the other players follow a fixed strategy.

FP was originally introduced by [7] as an algorithm to approximate the value of constant-sum games, or equivalently compute approximate solutions to linear programs [10]. Shortly after, it was shown that FP does indeed converge to the desired solution [24]. While convergence does not extend to arbitrary games, as illustrated by [25], it does so for quite a few interesting classes of games, and much research has focussed—and still focusses—on identifying such classes ([3], and the references therein). Both as a linear program solver and as a learning algorithm, FP is easily outperformed by more sophisticated algorithms. However, FP is of captivating simplicity and therefore is considered as one of the most convincing explanations of Nash equilibrium play. As I put it: “Brown’s results are not only computationally valuable but also quite illuminating from a substantive point of view. Imagine a pair of players repeating a game over and over again. It is plausible that at every stage a player attempts to exploit his knowledge of his opponent’s past moves. Even though the game may be too complicated or too nebulous to be subjected to an adequate analysis, experience in repeated plays may tend to a statistical equilibrium whose (time) average return is approximately equal to the value of the game” [16, p. 443].

In this paper, we show that in virtually all classes of games where FP is known to converge to a Nash equilibrium, it may take an exponential number of rounds (in the representation of the game) before any equilibrium action is played at all. While it was widely known that FP does not converge rapidly, the strength of our results is still somewhat surprising. They do not depend on the choice of a metric for comparing probability distributions. Rather, we show that the empirical frequency of FP after an exponential number of rounds can be *arbitrarily far* from any Nash equilibrium for *any* reasonable metric. This casts doubt on the plausibility of FP as an explanation of Nash equilibrium play.

2 Related Work

As mentioned above, FP does not converge in general. I showed this using a variant of Rock-Paper-Scissors and argued further that “if fictitious play is to fail, the game must contain elements of both coordination and competition” [23, p. 24]. This statement is perfectly consistent with the fact that FP is guaranteed to converge for both constant-sum games [24] and identical interest games, i.e., games that are best-response equivalent (in mixed strategies) to a common payoff game [20]. Other classes of games where FP is known to converge include two-player games solvable by iterated elimination of strictly dominated strategies [21] and non-degenerate 2×2 games [17]. While the proof of I was initially thought to apply to the class of all 2×2 games, this was later shown to be false [18]. The result was recently extended to non-degenerate $2 \times n$ games [2]. Since every non-degenerate 2×2 game is best-response equivalent to either a constant-sum game

or a common payoff game [20], the result of [1] follows more easily by combining those of [24] and [20].

To our knowledge, the *rate* of convergence of FP has so far only been studied in 2×2 games. For this class of games, FP converges at a rate of $O(T^{-1})$, where T is the number of rounds, as soon as both players have played an equilibrium action at least once [13]. We will see, however, that even in 2×2 games the latter may only happen after an exponential number of rounds.

Von Neumann [27] proposed a variant of FP and compared it to Dantzig's Simplex method. Indeed, there are some interesting similarities between the two. [8] recently studied the ability of FP to find approximate Nash equilibria. In addition to worst-case guarantees on the approximation ratio—which are rather weak—[1] showed that in random games a good approximation is typically achieved after a relatively small number of rounds. Similarly, the Simplex method is known to work very well in practice. As we show in this paper, FP also shares one of the major shortcomings of the Simplex method—its exponential worst-case running time.

Since FP is one of the earliest and simplest algorithms for learning in games, it inspired many of the algorithms that followed: the variant due to [1], a similar procedure suggested by [11], improvements like smooth FP [11], the regret minimization paradigm [15], and a large number of specialized algorithms put forward by the artificial intelligence community (e.g., [22, 9]).

Despite its conceptual simplicity and the existence of much more sophisticated learning algorithms, FP continues to be employed successfully in the area of artificial intelligence. Recent examples include equilibrium computation in Poker [12] and in anonymous games with continuous player types [23], and learning in sequential auctions [28].

3 Preliminaries

An accepted way to model situations of strategic interaction is by means of a *normal-form game* (see, e.g., [16]). We will focus on games with two players.

A *two-player game* $\Gamma = (P, Q)$ is given by two matrices $P, Q \in \mathbb{R}^{m \times n}$ for positive integers m and n . Player 1, or the row player, has a set $A = \{1, \dots, m\}$ of actions corresponding to the rows of these matrices, player 2, the column player, a set $B = \{1, \dots, n\}$ of actions corresponding to the columns. To distinguish between them, we usually denote actions of the row player by a^1, \dots, a^m and actions of the column player by b^1, \dots, b^n . Both players are assumed to simultaneously choose one of their actions. For the resulting action profile $(i, j) \in A \times B$, they respectively obtain payoffs p_{ij} and q_{ij} .

A *strategy* of a player is a probability distribution $s \in \Delta(A)$ or $t \in \Delta(B)$ over his actions, i.e., a nonnegative vector $s \in \mathbb{R}^m$ or $t \in \mathbb{R}^n$ such that $\sum_i s_i = 1$ or $\sum_j t_j = 1$, respectively. In a slight abuse of notation, we write p_{st} and q_{st} for the expected payoff of players 1 and 2 given a strategy profile $(s, t) \in \Delta(A) \times \Delta(B)$. A strategy is called *pure* if it chooses some action with probability one, and the set of pure strategies can be identified in a natural way with the set of actions.

A two-player game is called a *constant-sum game* if $p_{ij} + q_{ij} = p_{i'j'} + q_{i'j'}$ for all $i, i' \in A$ and $j, j' \in B$. Since all results in this paper hold invariably under positive affine transformations of the payoffs, such games can conveniently be represented by a single matrix P containing the payoffs of player 1; player 2 is then assumed to minimize the values in P . A constant-sum game is further called *symmetric* if P is a skew-symmetric matrix. In symmetric games, both players have the same set of actions, and we usually denote these actions by a^1, a^2, \dots, a^m . A game is a *common payoff game* if $p_{ij} = q_{ij}$ for all $i \in A$ and $j \in B$. Finally, a game is *non-degenerate* if for each strategy, the number of best responses of the other player is at most the support size of that strategy, i.e., the number of actions played with positive probability.

An action $i \in A$ of player 1 is said to *strictly dominate* another action $i' \in A$ if it provides a higher payoff for every action of player 2, i.e., if for all $j \in B$, $p_{ij} > p_{i'j}$. Dominance among actions of player 2 is defined analogously. A game is then called *solvable via iterated strict dominance* if strictly dominated actions can be removed iteratively such that exactly one action remains for each player.

A pair (s, t) of strategies is called a *Nash equilibrium* if the two strategies are best responses to each other, i.e., if $p_{st} \geq p_{it}$ for every $i \in A$ and $q_{st} \geq q_{sj}$ for every $j \in B$. A Nash equilibrium is *quasi-strict* if actions played with positive probability yield strictly more payoff than actions played with probability zero. By the minimax theorem [26], every Nash equilibrium (s, t) of a constant-sum game satisfies $\min_j \sum_i p_{ij}s_i = \max_i \sum_j p_{ij}t_j = \omega$ for some $\omega \in \mathbb{R}$, also called the value of the game.

Fictitious play (FP) was originally introduced to approximate the value of constant-sum games, and has subsequently been studied in terms of its convergence to Nash equilibrium in more general classes of games. It proceeds in rounds. In the first round, each player arbitrarily chooses one of his actions. In subsequent rounds, each player looks at the empirical frequency of play of his respective opponents in previous rounds, interprets it as a probability distribution, and myopically plays a pure best response against this distribution. Fix a game $\Gamma = (P, Q)$ with $P, Q \in \mathbb{R}^{m \times n}$. Denote by u_i and v_i the i th unit vector in \mathbb{R}^m and \mathbb{R}^n , respectively. Then, a *learning sequence* of Γ is a sequence $(x^0, y^0), (x^1, y^1), (x^2, y^2), \dots$ of pairs of non-negative vectors $(x^i, y^i) \in \mathbb{R}^m \times \mathbb{R}^n$ such that $x^0 = \mathbf{0}, y^0 = \mathbf{0}$, and for all $k \geq 0$,

$$\begin{aligned} x^{k+1} &= x^k + u_i && \text{where } i \text{ is the index of a maximum component of } Py^k \text{ and} \\ y^{k+1} &= y^k + v_j && \text{where } j \text{ is the index of a maximum component of } x^kQ. \end{aligned}$$

A learning sequence $(x^0, y^0), (x^1, y^1), (x^2, y^2), \dots$ of a game Γ is said to converge if for some Nash equilibrium s of Γ ,

$$\lim_{k \rightarrow \infty} \left(\frac{x^k}{k}, \frac{y^k}{k} \right) = s,$$

where both division and limit are to be interpreted component-wise. We then say that FP converges for Γ if every learning sequence of Γ converges to a Nash equilibrium.

An alternative definition of a learning sequence, in which players update their beliefs alternatingly instead of simultaneously, can be obtained by replacing $x^k Q$ by $x^{k+1} Q$ in the last condition above. [3] distinguishes between simultaneous and alternating FP, and points out that [1] actually introduced the latter variant, while almost all subsequent work routinely uses the former. We henceforth concentrate on simultaneous FP, or simply FP, but note that with some additional work all of our results can be shown to hold for alternating FP as well.

4 Results

We now present several results concerning the convergence rate of FP. Taken together, they cover virtually all classes of games for which FP is known to converge.

4.1 Symmetric Constant-Sum Games and Games Solvable by Iterated Strict Dominance

Let us first consider games with arbitrary payoffs. Our first result concerns two large classes of games where FP is guaranteed to converge: constant-sum games and games solvable by iterated strict dominance.

Theorem 1. *In symmetric two-player constant-sum games, FP may require exponentially many rounds (in the size of the representation of the game) before an equilibrium action is eventually played. This holds even for games solvable via iterated strict dominance.*

Proof. Consider the symmetric two-player constant-sum game $\Gamma = (P, Q)$ with payoff matrix P for player 1 as shown in Figure 1, where $0 < \epsilon < 1$. It is readily appreciated that (a^3, a^3) is the only Nash equilibrium of this game, as it is the only action profile that remains after iterated elimination of strictly dominated actions. Consider an arbitrary integer $k > 1$. We show that for $\epsilon = 2^{-k}$, FP may take 2^k rounds before either player plays action a^3 . Since the game can clearly be encoded using $O(k)$ bits in this case, the theorem follows.

Let FP start with both players choosing action a^1 . Since the game is symmetric, we can assume the actions for each step of the learning sequence to be identical for both players. After the first round $Py^1 = (0, 1, 2^{-k})$, and both players will play a^2 in round 2. We claim that they will continue to do so at least until round 2^k . To see this, observe that for all i with $1 \leq i < 2^k$, we have $Py^i = (-i + 1, 1, 2^{-k}i)$. As $2^{-k}i < 1$, both players will choose a^2 round $i + 1$.

Table 1 summarizes this development. It follows that the action sequence

$$(a^1, a^1) \underbrace{(a^2, a^2), \dots, (a^2, a^2)}_{2^k - 1 \text{ times}}$$

gives rise to a learning sequence that is exponentially long in k and in which no equilibrium action is played. □

$$\begin{array}{c}
 \begin{array}{ccc}
 & a^1 & a^2 & a^3 \\
 a^1 & \boxed{\begin{array}{ccc} 0 & -1 & -\epsilon \end{array}} \\
 a^2 & \boxed{\begin{array}{ccc} 1 & 0 & -\epsilon \end{array}} \\
 a^3 & \boxed{\begin{array}{ccc} \epsilon & \epsilon & 0 \end{array}}
 \end{array}
 \end{array}$$

Fig. 1. Symmetric constant-sum game used in the proof of Theorem 1. Player 1 chooses rows, player 2 chooses columns. Outcomes are denoted by the payoff of player 1.

Table 1. A learning sequence of the game depicted in Figure 1, where $\epsilon = 2^{-k}$

Round i	(a^i, a^i)	Py^i
0	—	$(0, 0, 0)$
1	(a^1, a^1)	$(0, 1, 2^{-k})$
2	(a^2, a^2)	$(-1, 1, 2^{-k}2)$
3	(a^2, a^2)	$(-2, 1, 2^{-k}3)$
	\vdots	\vdots
2^k	(a^2, a^2)	$(-2^k + 1, 1, 1)$

This result is tight in the sense that FP converges very quickly in symmetric 2×2 games. Up to renaming of actions, every such game can be described by a matrix

$$\begin{array}{c}
 \begin{array}{cc}
 & a^1 & a^2 \\
 a^1 & \boxed{\begin{array}{cc} 0 & -\alpha \end{array}} \\
 a^2 & \boxed{\begin{array}{cc} \alpha & 0 \end{array}}
 \end{array}
 \end{array}$$

for some $\alpha \geq 0$. If $\alpha = 0$, every strategy profile is a Nash equilibrium. Otherwise, action a^1 is strictly dominated for both players, and both players will play the equilibrium action a^2 from round 2 onwards.

4.2 Non-degenerate $2 \times n$ Games and Identical Interest Games

Another class of games where FP is guaranteed to converge are non-degenerate $2 \times n$ games. We again obtain a strong negative result concerning the convergence rate of FP, which also applies to games with identical interests.

Theorem 2. *In non-degenerate 2×3 games, FP may require exponentially many rounds (in the size of the representation of the game) before an equilibrium action is eventually played. This holds even for games with identical interests.*

Proof. Consider the 2×3 game $\Gamma = (P, Q)$ shown in Figure 2, where $0 < \epsilon < 1$. It is easily verified that Γ is non-degenerate and that the players have identical interests. The action profile (a^2, b^3) is the only action profile that remains

	b^1	b^2	b^3
a^1	(1, 1)	(2, 2)	(0, 0)
a^2	(0, 0)	$(2 + \epsilon, 2 + \epsilon)$	(3, 3)

Fig. 2. Non-degenerate two-player game with identical interests used in the proof of Theorem 2. Outcomes are denoted by a pair of payoffs for the two players.

Table 2. A learning sequence of the game shown in Figure 2, where $\epsilon = 2^{-k}$

Round i	(a^i, b^i)	Py^i	x^iQ
0	—	(0, 0)	(0, 0, 0)
1	(a^1, b^1)	(1, 0)	(1, 2, 0)
2	(a^1, b^2)	$(3, 2 + 2^{-k})$	(2, 4, 0)
3	(a^1, b^2)	$(5, 4 + 2^{-k}2)$	(3, 6, 0)
\vdots	\vdots	\vdots	\vdots
2^k	(a^1, b^2)	$(2^{k+1} - 1, 2^{k+1} - 1 - 2^{-k})$	$(2^k, 2^{k+1}, 0)$

after iterated elimination of strictly dominated actions, and thus the only Nash equilibrium of the game.

Now consider an integer $k > 1$. We show that for $\epsilon = 2^{-k}$, FP may take 2^k rounds before actions a^2 or b^3 are played. Since in this case the game can clearly be encoded using $O(k)$ bits, the theorem follows.

Let FP start with both players choosing action a^1 . Then, $Py^1 = (1, 0)$ and $x^1Q = (1, 2, 0)$. Accordingly, in the second round, the row player will choose a^1 , and the column player b^2 . Hence, $Py^2 = (3, 2 + 2^{-k})$ and $x^2Q = (2, 4, 0)$. Hereafter, for at least another $2^k - 1$ rounds, the players will choose the same actions as in round 2, because for all i with $2 \leq i \leq 2^k$, $x^iQ = (i, 2i, 0)$, $Py^i = (2i - 1, 2i - 1 + 2^{-k}(i - 1))$, and $2i - 1 > 2i - 1 + 2^{-k}(i - 1)$. Accordingly, the sequence of pairs of actions

$$(a^1, b^1) \underbrace{(a^1, b^2), \dots, (a^1, b^2)}_{2^k \text{ times}},$$

which contains no equilibrium actions, gives rise to a learning sequence that is exponentially long in k . Figure 2 illustrates both sequences. □

This result is again tight: in any 2×2 game, one of the players must always play an equilibrium action almost immediately. Indeed, given that the initial action profile is not itself an equilibrium, one of the players plays his second action in the following round. But what about the other player? By looking at the subgame of the game in Figure 2 induced by actions $\{a^1, a^2\}$ and $\{b^1, b^2\}$, and at the learning sequence used to obtain Theorem 2, we find that it might still take exponentially many rounds for *one* of the two players until he plays an equilibrium action for the first time.

Theorem 2 also applies to potential games [19], which form a superclass of games with identical interests. For the given ordering of its actions, the game of Figure 2 further has strategic complementarities and diminishing returns¹ which implies results analogous to Theorem 2 for classes of games in which convergence of FP was respectively claimed by [14]² and shown by [4].

4.3 Games with Constant Payoffs

The proofs of the previous two theorems crucially rely on exponentially small payoffs, so one may wonder if similar results can still be obtained if additional constraints are imposed on the payoffs. While this is certainly not the case for games where both the payoffs and the number of actions are constant, we find that a somewhat weaker version of Theorem 1 holds for games with constant payoffs, and in particular for symmetric constant-sum win-lose-tie games, i.e., symmetric constant-sum games with payoffs in $\{-1, 0, 1\}$.

For each integer k we define a symmetric constant-sum game Γ^k with a unique (mixed) Nash equilibrium and show that FP may take a number of rounds exponential in k before an equilibrium action is played. In contrast to the previous result, however, this result not only assumes a worst-case initial action profile, but also a worst-case learning sequence.

Theorem 3. *In symmetric constant-sum win-lose-tie games, FP may require exponentially many rounds (in the size of the game) before an equilibrium action is eventually played.*

Proof. Fix an integer $k > 1$. We construct a symmetric constant-sum win-lose-tie game $\Gamma^k = (P^k, Q^k)$ with a $(2k + 1) \times (2k + 1)$ payoff matrix $P^k = (p_{ij}^k)$ for player 1 such that for all i, j with $1 \leq j \leq i \leq 2k + 1$,

$$p_{ij}^k = \begin{cases} 1 & \text{if } j = 1 \text{ and } 2 \leq i \leq k + 1, \text{ or} \\ & \text{if } j = 1 \text{ and } i = 2k + 1, \text{ or} \\ & \text{if } j \neq 1 \text{ and } i = j + k, \\ -1 & \text{if } j \neq 1 \text{ and } i > j + k, \\ 0 & \text{otherwise.} \end{cases}$$

For $i < j$, let $p_{ij}^k = -p_{ji}^k$. Thus Γ^k clearly is a symmetric constant-sum game. To illustrate the definition, Γ^4 is shown in Figure 3.

Further define, for each k , a strategy profile (s^k, s^k) of Γ^k such that for all i with $1 \leq i \leq 2k + 1$,

$$s_i^k = \begin{cases} 2^{2k+1-i}/(2^k - 1) & \text{if } i > k + 1, \\ 0 & \text{otherwise.} \end{cases}$$

¹ A two-player game with totally ordered sets of actions is said to have strategic complementarities if the advantage of switching to a higher action, according to the ordering, increases when the opponent chooses a higher action, and diminishing returns if the advantage of increasing one's action is decreasing.

² The proof of this claim later turned out to be flawed [5].

	a^1	a^2	a^3	a^4	a^5	a^6	a^7	a^8	a^9
a^1	0	-1	-1	-1	-1	0	0	0	-1
a^2	1	0	0	0	0	-1	1	1	1
a^3	1	0	0	0	0	0	-1	1	1
a^4	1	0	0	0	0	0	0	-1	1
a^5	1	0	0	0	0	0	0	0	-1
a^6	0	1	0	0	0	0	0	0	0
a^7	0	-1	1	0	0	0	0	0	0
a^8	0	-1	-1	1	0	0	0	0	0
a^9	1	-1	-1	-1	1	0	0	0	0

Fig. 3. Symmetric constant-sum game Γ^4 used in the proof of Theorem 3. The game possesses a quasi-strict equilibrium (s^4, s^4) with $s^4 = (0, 0, 0, 0, 0, \frac{8}{15}, \frac{4}{15}, \frac{2}{15}, \frac{1}{15})$.

It is not hard to see that (s^k, s^k) is a quasi-strict equilibrium of Γ^k . Moreover, since Γ^k is both a symmetric and a constant-sum game, the support of any equilibrium strategy of Γ^k is contained in that of s^k (cf. [6]). We will now show that, when starting with (a^1, a^1) , FP in Γ^k may take at least 2^k rounds before an equilibrium action is played for the first time.

Consider the sequence a_1, \dots, a_{2^k} with $a_j = a^{1+\lceil \log_2 j \rceil}$ for all j with $1 \leq j \leq 2^k$, i.e., the sequence

$$a^1, a^2, a^3, a^3, \dots, \underbrace{a^i, \dots, a^i}_{2^{i-2} \text{ times}}, \dots, \underbrace{a^{k+1}, \dots, a^{k+1}}_{2^{k-1} \text{ times}}.$$

The length of this sequence is clearly exponential in k . Further define vectors x^0, \dots, x^{2^k} of dimension $2k + 1$ such that $x^0 = \mathbf{0}$, and for i with $1 \leq j \leq 2k + 1$, $x^{j+1} = x^j + u_i$ when $a_{j+1} = i$.

We now claim that $(x^0, x^0), \dots, (x^{2^k}, x^{2^k})$ is a learning sequence of Γ^k , i.e., that $j + 1$ is the index of a maximal component of both $P^k y^j$ and $x^j Q^k$. Table 3 shows the development of this sequence for $k = 4$.

By symmetry of Γ^k it suffices to prove the claim for $P^k y^j$. After the first round, we have for all i with $1 \leq i \leq 2k + 1$,

$$(P^k y^1)_i = \begin{cases} 1 & \text{if } 1 < i \leq k + 1, \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, since $\{a_2, \dots, a_{2^k}\} \subseteq \{a^2, \dots, a^{k+1}\}$, we have that $(P^k y^j)_i = 1$ for all i with $1 < i \leq k + 1$ and all j with $1 < j \leq 2^k$. It, therefore, suffices to show

Table 3. A learning sequence of the game Γ^4 shown in Figure 3

Round i	(a^j, a^j)	$P^4 y^i$
0	—	(0, 0, 0, 0, 0, 0, 0, 0)
1	(a^1, a^1)	(0, 1, 1, 1, 1, 0, 0, 0, 1)
2	(a^2, a^2)	(-1, 1, 1, 1, 1, 1, -1, -1, 0)
3	(a^3, a^3)	(-2, 1, 1, 1, 1, 1, 0, -2, -1)
4	(a^3, a^3)	(-3, 1, 1, 1, 1, 1, 1, -3, -2)
5	(a^4, a^4)	(-4, 1, 1, 1, 1, 1, 1, -2, -3)
6	(a^4, a^4)	(-5, 1, 1, 1, 1, 1, 1, -1, -4)
7	(a^4, a^4)	(-6, 1, 1, 1, 1, 1, 1, 0, -5)
8	(a^4, a^4)	(-7, 1, 1, 1, 1, 1, 1, 1, -6)
9	(a^5, a^5)	(-8, 1, 1, 1, 1, 1, 1, 1, -5)
10	(a^5, a^5)	(-9, 1, 1, 1, 1, 1, 1, 1, -4)
11	(a^5, a^5)	(-10, 1, 1, 1, 1, 1, 1, 1, -3)
12	(a^5, a^5)	(-11, 1, 1, 1, 1, 1, 1, 1, -2)
13	(a^5, a^5)	(-12, 1, 1, 1, 1, 1, 1, 1, -1)
14	(a^5, a^5)	(-13, 1, 1, 1, 1, 1, 1, 1, 0)
15	(a^5, a^5)	(-14, 1, 1, 1, 1, 1, 1, 1, 1)
16	(a^5, a^5)	(-15, 1, 1, 1, 1, 1, 1, 1, 2)

that $(P^k y^j)_i$ for all i with $i = 1$ or $k + 1 < i < 2k + 1$ and all j with $1 < j \leq 2^k$. Since, $p_{1i} = -1$ for all i with $1 < i \leq k + 1$, the former is obvious. For the latter, it can be shown by a straightforward if somewhat tedious induction on j that for all i with $1 \leq i < k$ and all j with $1 < j \leq 2^k$,

$$(P^k y^j)_{i+k+1} = \begin{cases} 1 - j & \text{if } j \leq 2^{i-1}, \\ 1 + j - 2^i & \text{if } 2^{i-1} < j \leq 2^i, \\ 1 & \text{otherwise, and} \end{cases}$$

$$(P^k y^j)_{2k+1} = \begin{cases} 2 - j & \text{if } j \leq 2^{k-1}, \\ 2 + j - 2^k & \text{otherwise.} \end{cases}$$

It follows that $(P^k y^j)_i \leq 1$ for all i with $1 \leq i \leq 2k + 1$ and all j with $1 \leq j < 2^k$, thus proving the claim. □

5 Conclusion

We have studied the rate of convergence of fictitious play, and obtained mostly negative results: for almost all of the classes of games where FP is known to converge, it may take an exponential number of rounds before some equilibrium action is eventually played. These results hold already for games with very few actions, given that one of the payoffs is exponentially small compared to the

others. Slightly weaker results can still be salvaged for symmetric constant-sum games and games solvable by iterated strict dominance, even if payoffs are in the set $\{-1, 0, 1\}$. It is an open question whether this result can be strengthened to match that for games with arbitrary payoffs, and whether a similar result can be obtained for the classes of games covered by Theorem 2, i.e., for potential games and identical interest games.

While it was known that fictitious play does not converge rapidly, the strength of our results is still somewhat surprising. They do not depend on the choice of a metric for comparing probability distributions. Rather, the empirical frequency of FP after an exponential number of rounds can be *arbitrarily far* from any Nash equilibrium for *any* reasonable metric. This casts doubt on the plausibility of fictitious play as an explanation of Nash equilibrium play.

Acknowledgements. This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/3-2, BR 2312/3-3, and FI 1664/1-1. We would like to thank Vincent Conitzer and Troels Bjerre Sørensen for valuable discussions.

References

1. Bellman, R.: On a new iterative algorithm for finding the solutions of games and linear programming problems. In: Research Memorandum P-473, The RAND Corporation (1953)
2. Berger, U.: Fictitious play in $2 \times n$ games. *Journal of Economic Theory* 120, 139–154 (2005)
3. Berger, U.: Brown’s original fictitious play. *Journal of Economic Theory* 135, 572–578 (2007)
4. Berger, U.: Learning in games with strategic complementarities revisited. *Journal of Economic Theory* 143, 292–301 (2008)
5. Berger, U.: The convergence of fictitious play in games with strategic complementarities: A comment. MPRA Paper No. 20241, Munich Personal RePEc Archive (2009)
6. Brandt, F., Fischer, F.: On the hardness and existence of quasi-strict equilibria. In: Monien, B., Schroeder, U.-P. (eds.) SAGT 2008. LNCS, vol. 4997, pp. 291–302. Springer, Heidelberg (2008)
7. Brown, G.W.: Iterative solutions of games by fictitious play. In: Koopmans, T.C. (ed.) *Activity Analysis of Production and Allocation*, pp. 374–376. John Wiley and Sons, Inc., Chichester (1951)
8. Conitzer, V.: Approximation guarantees for fictitious play. In: Proc. of 47th Annual Allerton Conference on Communication, Control and Computing (2009)
9. Conitzer, V., Sandholm, T.: AWESOME: A general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Machine Learning* 67(1-2), 23–43 (2007)
10. Dantzig, G.B.: A proof of the equivalence of the programming problem and the game problem. In: Koopmans, T.C. (ed.) *Activity Analysis of Production and Allocation*, pp. 330–335. John Wiley & Sons Inc., Chichester (1951)
11. Fudenberg, D., Levine, D.: Universal consistency and cautious fictitious play. *Journal of Economic Dynamics and Control* 19, 1065–1089 (1995)

12. Ganzfried, S., Sandholm, T.: Computing an approximate jam/fold equilibrium for 3-player no-limit texas hold'em tournaments. In: Proc. of 7th AAMAS Conference, pp. 919–925 (2008)
13. Gjerstad, S.: The rate of convergence of continuous fictitious play. *Journal of Economic Theory* 7, 161–178 (1996)
14. Hahn, S.: The convergence of fictitious play in 3×3 games with strategic complementarities. *Economic Letters* 64, 57–60 (1999)
15. Hannan, J.: Approximation to Bayes risk in repeated plays. In: Drescher, M., Tucker, A.W., Wolfe, P. (eds.) *Contributions to the Theory of Games*, vol. 3, pp. 97–139. Princeton University Press, Princeton (1957)
16. Luce, R.D., Raiffa, H.: *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc., Chichester (1957)
17. Miyasawa, K.: On the convergence of the learning process in a 2×2 nonzero sum two-person game. In: *Research Memorandum 33*, Econometric Research Program. Princeton University, Princeton (1961)
18. Monderer, D., Sela, A.: A 2×2 game without the fictitious play property. *Games and Economic Behavior* 14, 144–148 (1996)
19. Monderer, D., Shapley, L.S.: Potential games. *Games and Economic Behavior* 14(1), 124–143 (1996)
20. Monderer, D., Shapley, L.S.: Fictitious play property for games with identical interests. *Journal of Economic Theory* 68, 258–265 (1996)
21. Nachbar, J.H.: “Evolutionary” selection dynamics in games: Convergence and limit properties. *International Journal of Game Theory* 19, 59–89 (1990)
22. Powers, R., Shoham, Y.: New criteria and a new algorithm for learning in multi-agent systems. In: *Advances in Neural Information Processing Systems*, vol. 17, pp. 1089–1096. MIT Press, Cambridge (2004)
23. Rabinovich, Z., Gerding, E., Polukarov, M., Jennings, N.R.: Generalised fictitious play for a continuum of anonymous players. In: Proc. of 21st IJCAI, pp. 245–250 (2009)
24. Robinson, J.: An iterative method of solving a game. *Annals of Mathematics* 54(2), 296–301 (1951)
25. Shapley, L.: Some topics in two-person games. In: Drescher, M., Shapley, L.S., Tucker, A.W. (eds.) *Advances in Game Theory*. *Annals of Mathematics Studies*, vol. 52, pp. 1–29. Princeton University Press, Princeton (1964)
26. von Neumann, J.: Zur Theorie der Gesellschaftspiele. *Mathematische Annalen* 100, 295–320 (1928)
27. von Neumann, J.: A numerical method to determine optimum strategy. *Naval Research Logistics Quarterly* 1(2), 109–115 (1954)
28. Zhu, W., Wurman, P.R.: Structural leverage and fictitious play in sequential auctions. In: Proc. of 18th AAAI Conference, pp. 385–390 (2002)

On Learning Algorithms for Nash Equilibria

Constantinos Daskalakis¹, Rafael Frongillo², Christos H. Papadimitriou²,
George Pierrakos², and Gregory Valiant²

¹ MIT

costis@csail.mit.edu

² UC Berkeley

{raf, christos, georgios, gvaliant}@cs.berkeley.edu

Abstract. Can learning algorithms find a Nash equilibrium? This is a natural question for several reasons. Learning algorithms resemble the behavior of players in many naturally arising games, and thus results on the convergence or non-convergence properties of such dynamics may inform our understanding of the applicability of Nash equilibria as a plausible solution concept in some settings. A second reason for asking this question is in the hope of being able to prove an impossibility result, not dependent on complexity assumptions, for computing Nash equilibria via a restricted class of reasonable algorithms. In this work, we begin to answer this question by considering the dynamics of the standard multiplicative weights update learning algorithms (which are known to converge to a Nash equilibrium for zero-sum games). We revisit a 3×3 game defined by Shapley [10] in the 1950s in order to establish that fictitious play does not converge in general games. For this simple game, we show via a potential function argument that in a variety of settings the multiplicative updates algorithm impressively fails to find the unique Nash equilibrium, in that the cumulative distributions of players produced by learning dynamics actually drift *away* from the equilibrium.

1 Introduction

In complexity, once a problem is shown intractable, research shifts towards two directions: (a) polynomial algorithms for more modest goals such as special cases and approximation, and (b) exponential lower bounds for restricted classes of algorithms. In other words, we weaken either the problem or the algorithmic model. For the problem of finding Nash equilibria in games, the first avenue has been followed extensively and productively, but, to our knowledge, not yet the second. It *has* been shown that a general and natural class of algorithms fails to solve *multiplayer* games in polynomial time in the number of players [4] — but such games have an exponential input anyway, and the point of that proof is to show, via communication complexity arguments, that, if the players do not know the input, they have to communicate large parts of it, at least for some games, in order to reach equilibrium.

We conjecture that a very strong lower bound result, of sweeping generality, is possible even for bimatrix games. In particular, we suspect that a broad class of algorithms that maintains and updates mixed distributions in essentially arbitrary ways can be shown

¹ In addition, of course, to the perennial challenge of collapsing complexity classes. . .

to fail to efficiently find Nash equilibria in bimatrix games, as long as these algorithms cannot identify the matrices — since our ambition here falls short of proving that $P \neq NP$, such restriction needs to be in place. In this paper we start on this path of research.

In targeting restricted classes of algorithms, it is often most meaningful to focus on algorithmic ideas which are known to perform well under certain circumstances or in related tasks. For games, *learning* is the undisputed champion among algorithmic styles. By learning we mean a large variety of algorithmic ways of playing games which maintain weights for the strategies (unnormalized probabilities of the current mixed strategy) and update them based on the performance of the current mixed strategy, or single strategy sampled from it, against the opponent's mixed strategy (or, again, sampled strategy). Learning algorithms are known to converge to the Nash equilibrium in zero-sum games [2], essentially because they can be shown to have diminishing regret. Furthermore, in general games, a variant in which regret is minimized explicitly [5] is known to always converge to a correlated equilibrium. Learning is of such central importance in games that it is broadly discussed as a loosely defined equilibrium concept — for example, it has been recently investigated viz. the price of anarchy [17,9].

There are three distinct variants of the learning algorithmic style with respect to games: In the first, which we call the *distribution payoff* setting, the players get feedback on the expected utility of the opponent's mixed strategy on all of their strategies — in other words, in a bimatrix game (R, C) , if the row player plays mixed strategy x and the column player y , then the row player sees at each stage the vector Cy^T while the column player sees $x^T R$. In the second variant which we call the *stochastic setting*, we sample from the two mixed strategies and both players learn the payoffs of all of their strategies against the one chosen by the opponent — that is, the row player learns the C_j , the whole column corresponding to the column player's choice, and vice-versa. A third variant is the *multi-armed setting*, introduced in [2], in which the players sample the distributions and update them according to the payoff of the combined choices. In all three cases we are interesting in studying the behavior of the cumulative distributions of the players, and see if they converge to the Nash equilibrium (as is the case for zero-sum games).

An early fourth kind of learning algorithm called *fictitious play* does not fall into our framework. In fictitious play both players maintain the opponent's histogram of past plays, adopt the belief that this histogram is the mixed strategy being played by the opponent, and keep best-responding to it. In 1951 Julia Robinson proved that fictitious play (or more accurately, the cumulative distributions of players resulting from fictitious play) converges to the Nash equilibrium in zero-sum games. Incidentally, Robinson's inductive proof implies a convergence that is exponentially slow in the number of strategies, but Karlin [6] conjectured in the 1960s *quadratic* convergence; this conjecture remains open. Shapley [10] showed that fictitious play fails to converge in a particular simple 3×3 nonzero-sum game (it does converge in all $2 \times n$ games).

But how about learning dynamics? Is there a proof that this class of algorithms fails to solve the general case of the Nash equilibrium problem? This question has been discussed in the past, and has in fact been treated extensively in Zinkevich's thesis [14]. Zinkevich presents extensive experimental results showing that, for the same 3×3 game considered by Shapley in [10] (and which is the object of our investigation), as well as

in a variant of the same game, the cumulative distributions do not converge to a Nash equilibrium (we come back to Zinkevich’s work later in the last section). However, to our knowledge there is no actual proof in the literature establishing that learning algorithms fail to converge to a Nash equilibrium.

Our main result is such a non-convergence proof; in fact, we establish this for each of the variants of learning algorithms. For each of the three styles, we consider the standard learning algorithm in which the weight updates are *multiplicative*, that is, the weights are multiplied by an exponential in the observed utility, hence the name multiplicative experts weight update algorithms. (In the multi-armed setting, we analyze the variant of the multiplicative weights algorithm that applies in this setting, in which payoffs are scaled so as to boost low-probability strategies). In all three settings, our results are negative: for Shapley’s 3×3 game the learning algorithms fail, in general, to converge to the unique Nash equilibrium. In fact, we prove the much more striking result that in all settings, the dynamics lead the players’ cumulative distributions *away* from the equilibrium exponentially quickly. The precise statements of the theorems differ, reflecting the different dynamics and the analytical difficulties they entail.

At this point it is important to emphasize that most of the work on the field focuses on proving the non-convergence of private distributions of the players, i.e. the distribution over strategies of each player at each time-step. In general, this is easy to do. In sharp contrast, we prove the non-convergence of the cumulative distributions of the players; the cumulative distribution is essentially the time-average of the private distributions played up to some time-step. This is a huge difference, because this weaker definition of convergence (corresponding to a realistic sense of what it means to play a mixed strategy in a repeated game) yields a much stronger result. Only Shapley in his original paper [10] (and Benaim and Hirsch [15] for a more elaborate setting) prove non-convergence results for the cumulative distributions, but for fictitious play dynamics. We show this for multiplicative weight updates, arguably (on the evidence of its many other successes, see the survey [12]) a much stronger class of algorithms.

2 The Model

We start by describing the characteristics of game-play; to do that we need to specify the type of information that the players receive at each time step. In this section we briefly describe the three “learning environments” which we consider, and then for each environment describe the types of learning algorithms which we consider.

2.1 Learning Environments

The first setting we consider is the *distribution payoff* setting, in which each player receives a vector of the expected payoffs that each of his strategies would receive, given the distribution of the other player. Formally, we have the following definition:

Definition 1. [Distribution payoff setting] *Given mixed strategy profiles $\mathbf{c}_t = (c_1, \dots, c_n)$, and $\mathbf{r}_t = (r_1, \dots, r_n)^T$ with $\sum r_i = \sum c_i = 1$ for the column and row player, respectively, and payoff matrices C, R of the underlying game,*

$$\mathbf{r}_{t+1} = f(R\mathbf{c}_t^T, \mathbf{r}_t), \quad \mathbf{c}_{t+1} = g(\mathbf{r}_t^T C, \mathbf{c}_t),$$

where f, g are update functions of the row, and column player, respectively, with the condition that $\mathbf{r}_{t+1}, \mathbf{c}_{t+1}$ are distributions.

It may seem that this setting gives too much information to the players, to the point of being unrealistic. We consider this setting for two reasons; first, intuitively, if learning algorithms can find Nash equilibria in any setting, then they should in this setting. Since we will provide largely negative results, it is natural to consider this setting that furnishes the players with the most power. The second reason for considering this setting is that in this setting, provided f, g are deterministic functions, the entire dynamics is deterministic, simplifying the analysis. Our results and proof approaches for this setting provide the guiding intuition for our results in the more realistic learning settings.

The second setting we consider, is the *stochastic* setting, in which each player selects a single strategy to play, according to their private strategy distributions, \mathbf{r}_t and \mathbf{c}_t , and each player may update his strategy distribution based on the entire vector of payoffs that his different strategies would have received given the single strategy choice of the opponent. Formally, we have:

Definition 2. [Stochastic setting] Given mixed strategy profiles \mathbf{r}_t , and \mathbf{c}_t for the row and column player, respectively, at some time t , and payoff matrices R, C of the underlying game, the row and column players select strategies i , and j according to \mathbf{r}_t and \mathbf{c}_t , respectively, and

$$\mathbf{r}_{t+1} = f(R_{\cdot,j}, \mathbf{r}_t), \quad \mathbf{c}_{t+1} = g(C_{i,\cdot}, \mathbf{c}_t),$$

where f, g are update functions of the row and column player, respectively, and $\mathbf{r}_{t+1}, \mathbf{c}_{t+1}$ are required to be distributions, and $M_{i,\cdot}, M_{\cdot,i}$, respectively, denote the i^{th} row and column of matrix M .

Finally, we will consider the *multi-armed* setting, in which both players select strategies according to their private distributions, knowing only the single payoff value given by their combined choices of strategies.

Definition 3. [Multi-armed setting] Given mixed strategy profiles \mathbf{r}_t , and \mathbf{c}_t for the row and column player, respectively, at some time t , and payoff matrices R, C of the underlying game, the row and column players select strategies i , and j according to \mathbf{r}_t and \mathbf{c}_t , respectively, and

$$\mathbf{r}_{t+1} = f(R_{i,j}, \mathbf{r}_t), \quad \mathbf{c}_{t+1} = g(C_{i,j}, \mathbf{c}_t),$$

where f, g are update functions of the row, and column player, respectively, and $\mathbf{r}_{t+1}, \mathbf{c}_{t+1}$ are distributions.

While the multi-armed setting is clearly the weakest setting to learn in, it is also, arguably, the most realistic and closely resembles the type of setting in which many everyday games are played.

Almost all of the results in this paper refer to the non-convergence of the cumulative distributions of the players, defined as:

$$R_{i,t} = \frac{\sum_{j=0}^t r_{i,j}}{t}, \quad C_{i,t} = \frac{\sum_{j=0}^t c_{i,j}}{t}$$

2.2 Learning Algorithms

For each game-play setting, the hope is to characterize which types of learning algorithms are capable of efficiently converging to an equilibrium. In this paper, we tackle the much more modest goal of analyzing the behavior of standard learning models that are known to perform well in each setting. For the distribution payoff setting, and the stochastic setting, we consider the dynamics induced by multiplicative weight updates. Specifically, for a given update parameter $\epsilon > 0$, at each timestep t , a player's distribution $\mathbf{w}_t = (w_{1,t}, \dots, w_{n,t})$ is updated according to

$$w_{i,t+1} = \frac{w_{i,t}(1 + \epsilon)^{P_i}}{\sum_i w_{i,t}(1 + \epsilon)^{P_i}},$$

where P_i is the payoff that the i^{th} strategy would receive at time t . We focus on this learning algorithm as it is extraordinarily successful, both practically and theoretically, and is known to have vanishing regret (which, by the min-max theorem, guarantees that cumulative distributions $\sum_{t=1}^T \frac{\mathbf{w}_t}{T}$ converge to the Nash equilibrium for zero-sum games [12]).

For the multi-armed setting, the above weight update algorithm is not known to perform well, as low-probability strategies are driven down by the dynamics. There is a simple fix, first suggested in [11]; one scales the payoffs by the inverse of the probability with which the given strategy was played, then applies multiplicative weights as above with the scaled payoffs in place of the raw payoff. Intuitively, this modification gives the low-weight strategies the extra boost that is needed in this setting. Formally, given update parameter ϵ , and distribution \mathbf{w}_t , if strategy s is chosen at time t , and payoff P is received, we update according to the following:

$$\begin{aligned} w_s^* &= w_{s,t}(1 + \epsilon)^{P/w_{s,t}} \\ w_{i \neq s}^* &= w_{i,t} \\ w_{j,t+1} &= \frac{w_j^*}{\sum_k w_k^*}. \end{aligned}$$

We note that this update scheme differs slightly from the originally proposed scheme in [11], in which a small drift towards the uniform distribution is explicitly added. We omit this drift as it greatly simplifies the analysis; additionally, arguments from [13] can be used to show that our update scheme also has the guarantee that the algorithm will have low-regret in expectation (and thus the dynamics converge for zero-sum games).

2.3 The Game

For all of our results, we will make use of Shapley's 3×3 bimatrix game with row and column payoffs given by

$$R = \begin{pmatrix} 0 & 1 & 2 \\ 2 & 0 & 1 \\ 1 & 2 & 0 \end{pmatrix}, \quad C = \begin{pmatrix} 0 & 2 & 1 \\ 1 & 0 & 2 \\ 2 & 1 & 0 \end{pmatrix}.$$

This game has a single Nash equilibrium in which both players play each strategy with equal probabilities. It was originally used by Shapley to show that fictitious play does not converge for general games.

3 Distribution Payoff Setting

In this section we consider the deterministic dynamics of running the experts weights algorithm in the distribution payoff setting. We show that under these dynamics, provided that the initial distributions satisfy $\mathbf{r} \neq \mathbf{c}$, the cumulative distributions R_t, C_t tend away from the Nash equilibrium. The proof splits into three main pieces; first, we define a potential function, which we show is strictly increasing throughout the dynamics, and argue that the value of the potential cannot be bounded by any constant. Next, we argue that given a sufficiently large value of the potential function, eventually the private row and column distributions $\mathbf{r}_t, \mathbf{c}_t$ must become unbalanced in the sense that for some $i \in \{1, 2, 3\}$, $r_i > .999$ and $c_i < .001$ (or $r_i < .001, c_i > .999$). Finally, given this imbalance, we argue that the dynamics consists of each player switching between essentially pure strategies, with the amount of time spent playing each strategy increasing in a geometric progression, from which it follows that the cumulative distributions will not converge.

Each of the three components of the proof, including the potential function argument, will also apply in the stochastic, and multi-armed settings, although the details will differ.

Before stating our main non-convergence results, we start by observing that in the case that both players perform multiplicative experts weight updates with parameters $\epsilon_R = \epsilon_C$, and start with identical initial distributions $\mathbf{r} = \mathbf{c}$, the dynamics *do* converge to the equilibrium. In fact, not only do the cumulative distributions R_t, C_t converge, but so do the private distributions $\mathbf{r}_t, \mathbf{c}_t$.

Proposition 1. *If both players start with a common distribution $\mathbf{r} = \mathbf{c}$ and perform their weight updates with $\epsilon_R = \epsilon_C = \epsilon \leq 3/5$, then the dynamics of $\mathbf{r}_t, \mathbf{c}_t$ converge to the Nash equilibrium exponentially fast.*

The proof is simple and is delegated to the full version of this paper. We now turn our attention to the main non-convergence result of this section—if the initial distributions are not equal, then the dynamics diverge.

Theorem 1. *In the distribution payoff setting, with a row player performing experts weight updates with parameter $1 + \epsilon_R$, and column player performing updates with parameter $1 + \epsilon_C$, the cumulative distributions $R_t = \sum_{i=0}^t \frac{r_i}{t}, C_t = \sum_{i=0}^t \frac{c_i}{t}$ diverge, provided that the initial weights do not satisfy $\mathbf{r}_i = \mathbf{c}_i^\alpha$, with $\alpha = \frac{\log(1+\epsilon_R)}{\log(1+\epsilon_C)}$.*

The first component of the proof will hinge upon the following potential function for the dynamics:

$$\Phi(\mathbf{r}, \mathbf{c}) := \log \left(\max_i \left(\frac{r_i}{c_i^\alpha} \right) \right) - \log \left(\min_i \left(\frac{r_i}{c_i^\alpha} \right) \right), \tag{1}$$

with $\alpha = \frac{\log(1+\epsilon_R)}{\log(1+\epsilon_C)}$. We are going to use the same potential function for the other two learning settings as well. The following lemma argues that $\Phi(\mathbf{r}_t, \mathbf{c}_t)$ increases unboundedly.

Lemma 1. *Given initial private distributions $\mathbf{r}_0, \mathbf{c}_0$ such that $\Phi(\mathbf{r}_0, \mathbf{c}_0) \neq 0$, then $\Phi(\mathbf{r}_t, \mathbf{c}_t)$ is strictly increasing, and for any constant k , there exists some t_0 such that $\Phi(\mathbf{r}_{t_0}, \mathbf{c}_{t_0}) > k$.*

Proof. We consider the change in Φ after one step of the dynamics. For convenience, we give the proof in the case that $\epsilon_R = \epsilon_C = \epsilon$; without this assumption identical arguments yield the desired general result. Also note that without loss of generality, by the symmetry of the game, it suffices to consider the case when $r_{1,t} \geq c_{1,t}$. The dynamics define the following updates:

$$\left(\frac{r_{1,t+1}}{c_{1,t+1}}, \frac{r_{2,t+1}}{c_{2,t+1}}, \frac{r_{3,t+1}}{c_{3,t+1}} \right) = \frac{n_1}{n_2} \left(\frac{r_{1,t}(1+\epsilon)^{c_2+2c_3}}{c_{1,t}(1+\epsilon)^{r_2+2r_3}}, \frac{r_{2,t}(1+\epsilon)^{2c_1+c_3}}{c_{2,t}(1+\epsilon)^{2r_1+r_3}}, \frac{r_{3,t}(1+\epsilon)^{c_1+2c_2}}{c_{3,t}(1+\epsilon)^{r_1+2r_2}} \right),$$

for some positive normalizing constants n_1, n_2 . By the symmetry of the game, it suffices to consider the following two cases: when $\operatorname{argmax}_i(r_i/c_i) = 1$ and $\operatorname{argmin}_i(r_i/c_i) = 2$, and the case when $\operatorname{argmax}_i(r_i/c_i) = 1$ and $\operatorname{argmin}_i(r_i/c_i) = 3$. We start by considering the first case:

$$\begin{aligned} \Phi(\mathbf{r}_{t+1}, \mathbf{c}_{t+1}) &= \log \left(\max_i \left(\frac{r_i}{c_i} \right) \right) - \log \left(\min_i \left(\frac{r_i}{c_i} \right) \right) \\ &\geq \log \left(\frac{r_1}{c_1} \right) - \log \left(\frac{r_2}{c_2} \right) \\ &= \log(n_1/n_2) + \log \left(\frac{r_{1,t}}{c_{1,t}} \right) + (c_{2,t} + 2c_{3,t} - r_{2,t} - 2r_{3,t}) \log(1 + \epsilon) \\ &\quad - \log(n_1/n_2) - \left(\log \left(\frac{r_{2,t}}{c_{2,t}} \right) + (c_{3,t} + 2c_{1,t} - r_{3,t} - 2r_{1,t}) \log(1 + \epsilon) \right) \\ &= \Phi(\mathbf{r}_t, \mathbf{c}_t) + (-2c_{1,t} + c_{2,t} + c_{3,t} - r_{2,t} - r_{3,t} + 2r_{1,t}) \log(1 + \epsilon) \\ &= \Phi(\mathbf{r}_t, \mathbf{c}_t) + 3(r_{1,t} - c_{1,t}) \log(1 + \epsilon) \end{aligned}$$

In the case second case, where $\operatorname{argmax}_i(r_i/c_i) = 1$ and $\operatorname{argmin}_i(r_i/c_i) = 3$, a similar calculation yields that

$$\Phi(\mathbf{r}_{t+1}, \mathbf{c}_{t+1}) \geq \Phi(\mathbf{r}_t, \mathbf{c}_t) + 3(c_{3,t} - r_{3,t}) \log(1 + \epsilon).$$

In either case, note that Φ is strictly increasing unless $r_i/c_i = 1$ for each i , which can only happen when $\Phi(\mathbf{r}_t, \mathbf{c}_t) = 0$.

To see that Φ is unbounded, we first argue that if the private distributions \mathbf{r}, \mathbf{c} are both sufficiently far from the boundary of the unit cube, then the value of the potential function will be increasing at a rate proportionate to its value. If \mathbf{r} or \mathbf{c} is near the boundary of the unit cube, and $\max_i |r_i - c_i|$ is small, then we argue that the dynamics will drive the private distributions towards the interior of the unit cube. Thus it will follow that the value of the potential function is unbounded.

Specifically, if $\mathbf{r}, \mathbf{c} \in [1, 1]^3$, then from the derivative of the logarithm, we have

$$30 \max_i |r_i - c_i| \geq \Phi(\mathbf{r}, \mathbf{c})$$

and thus provided $\mathbf{r}_t, \mathbf{c}_t$ are in this range $\Phi(\mathbf{r}_{t+1}, \mathbf{c}_{t+1}) \geq \Phi(\mathbf{r}_t, \mathbf{c}_t) \left(1 + \frac{\log(1+\epsilon)}{30} \right)$. If $\mathbf{r}, \mathbf{c} \notin [1, 1]^3$, then arguments from the proof of Proposition [11](#) can be used to show that

after some time t_0 , either $\mathbf{r}_{t_0}, \mathbf{c}_{t_0} \in [.2, 1]^3$, or for some time $t' < t_0$, $\max_i |r_i - c_i| \geq .01$, in which case by the above arguments the value of the potential function must have increased by at least $.01 \log(1 + \epsilon)$, and thus our lemma holds. \square

The above lemma guarantees that the potential function will get arbitrarily large. We now leverage this result to argue that there is some time t_0 and a coordinate i such that r_{i,t_0} is very close to 1, whereas c_{i,t_0} is very close to zero. The proof consists of first considering some time at which the potential function is quite large. Then, we argue that there must be some future time at which for some i, j with $i \neq j$, the contributions of coordinates i and j to the value of the potential function are both significant. Given that $|\log(r_i/c_i)|$ and $|\log(r_j/c_j)|$ are both large, we then argue that after some more time, we get the desired imbalance in some coordinate k , namely that $r_k > .999$ and $c_k < .001$ (or vice versa).

Lemma 2. *Given initial distributions $\mathbf{r}_0 = (r_{1,0}, r_{2,0}, r_{3,0})$, $\mathbf{c}_0 = (c_{1,0}, c_{2,0}, c_{3,0})$, with $\Phi(\mathbf{r}_0, \mathbf{c}_0) \geq 40 \log_{1+\epsilon_R}(2000)$, assuming that the cumulative distributions converge to the equilibrium, then there exists $t_0 > 0$ and i such that either $r_{i,t_0} > .999$ and $c_{i,t_0} < .001$, or $r_{i,t_0} < .001$, and $c_{i,t_0} > .999$.*

Proof. For convenience, we will assume all logarithms are to the base $1 + \epsilon_R$, unless otherwise specified. For ease of notation, let $k = \lceil \log_{1+\epsilon_R}(2000) \rceil$. Also, for simplicity, we give the proof in the case that $\epsilon_R = \epsilon_C = \epsilon$; as above, the proof of the general case is nearly identical.

Assuming for the sake of contradiction that the cumulative distributions converge to the equilibrium of the game, it must be the case that there exists some time $t > 0$ for which $\arg \max_i |\log(r_{i,t}/c_{i,t})| \neq \arg \max_i |\log(r_{i,0}/c_{i,0})|$, and thus, without loss of generality, we may assume that at time 0, for some i, j with $i \neq j$,

$$\left| \log \left(\frac{r_{i,0}}{c_{i,0}} \right) \right| > 13k, \text{ and } \left| \log \left(\frac{r_{j,0}}{c_{j,0}} \right) \right| > 13k.$$

Without loss of generality, we may assume that $r_i > c_i$. We will first consider the cases in which $\log(r_i/c_i) > 13k$ and $\log(r_j/c_j) > 13k$, and then will consider the cases when $\log(r_i/c_i) > 13k$ and $\log(r_i/c_i) < -13k$.

Consider the case when $\log(r_1/c_1) > 13k$ and $\log(r_2/c_2) > 13k$. Observe that $c_3 > r_3$ and that $k = \ln(2000)/\ln(1 + \epsilon_R) \geq \ln(2000)/\epsilon_R$. Let t_0 be the smallest time at which $\log(r_{3,t_0}) - \max(\log(r_{1,t_0}), \log(r_{2,t_0})) \leq k$. We argue by induction, that

$$\log(c_{3,t}) - \max(\log(c_{1,t}), \log(c_{2,t})) - (\log(r_{3,t}) - \max(\log(r_{1,t}), \log(r_{2,t}))) \geq 12k,$$

for any $t \in \{0, \dots, t_0 - 1\}$. When $t = 0$, this quantity is at least $13k$. Assuming the claim holds for all $t < t'$, for some fixed $t' < t_0 - 1$, we have that $\sum_{t=0}^{t'+1} r_{1,t} \leq \frac{2}{2\epsilon_R} \frac{1}{2000}$, where the factor of 2 in the numerator takes into account the fact that the payoffs are slightly different than 2, 1, 0, for the three row strategies. Similarly, $\sum_{t=0}^{t'+1} r_{2,t} \leq \frac{2}{\epsilon_R} \frac{1}{2000}$. Thus we have that

$$\begin{aligned} \log(c_{3,t'+1}) - \log(c_{1,t'+1}) &\geq \log(c_{3,0}) - \log(c_{1,0}) - 2(t' + 1) - \frac{4}{2\epsilon_R} \frac{1}{2000} \\ &\geq \log(c_{3,0}) - \log(c_{1,0}) - 2(t' + 1) - k \end{aligned}$$

Similarly, we can write a corresponding expression for $\log(c_{3,t'+1}) - \log(c_{2,t'+1})$, from which our claim follows.

Thus we have that $\log(c_{3,t_0}) - \max(\log(c_{1,t_0}), \log(c_{2,t_0})) \geq 12k$, and $\log(r_{3,t_0}) - \max(\log(r_{1,t_0}), \log(r_{2,t_0})) \leq k$. After another $2.1k$ timesteps, we have that $\log(r_{3,t_0}) - \max(\log(r_{1,t_0}), \log(r_{2,t_0})) \leq -k$, and $\log(c_{3,t_0}) - \max(\log(c_{1,t_0}), \log(c_{2,t_0})) \geq 7k$. If $\log(r_{1,t_0+2.1k}) - \log(r_{2,t_0+2.1k}) < -k$, then we are done, since $r_{2,t_0+2.1k} > .999$, $c_{2,t_0+2.1k} < .001$. If $\log(r_{1,t_0+2.1k}) - \log(r_{2,t_0+2.1k}) > -k$, then it must be the case that $\log(r_{1,t_0+4.2k}) - \log(r_{2,t_0+4.2k}) > k$, at which point we still have $\log(c_{3,t_0+4.2k}) - \max(\log(c_{1,t_0+4.2k}), \log(c_{2,t_0+4.2k})) > 2k$, so we have $r_{1,t_0+4.2k} > .999$, $c_{1,t_0+4.2k} < .001$. The case when $\log(r_1/c_1) > 13k$ and $\log(r_3/c_3) > 13k$ is identical.

In the case when $\log(r_1/c_1) > 13k$ and $\log(r_2/c_2) < -13k$, we let t_0 be the first time at which either $\log(r_{1,t_0}) - \log(r_{3,t_0}) > -k$ or $\log(c_{2,t_0}) - \log(c_{3,t_0}) > -k$. As above, we can show by induction that $\log(r_{2,t_0} - \max(\log(r_{1,t_0}), \log(r_{3,t_0}))) < -12k$, and $\log(c_{1,t_0} - \max(\log(c_{2,t_0}), \log(c_{3,t_0}))) < -12k$. After another $2.1k$ timesteps, either $r_1 > .999$, and $c_1 < .001$ or $c_{2,t_0+2.1k} > .1$, in which case after an additional $2.1k$ timesteps, $c_2 > .999$ and $r_2 < .001$.

The remaining case when $\log(r_1/c_1) > 13k$ and $\log(r_3/c_3) < -13k$, is identical, as can be seen by switching the players and permuting the rows and columns of the matrix. \square

The following lemma completes our proof of Theorem [II](#).

Lemma 3. *Given initial distributions $\mathbf{r}_0 = (r_{1,0}, r_{2,0}, r_{3,0})$, $\mathbf{c}_0 = (c_{1,0}, c_{2,0}, c_{3,0})$, such that for some i , $r_{i,0} > .999$ and $c_{i,0} < .001$, the cumulative distributions defined by*

$$R_{i,t} = \frac{\sum_{j=0}^t r_{i,j}}{t}, C_{i,t} = \frac{\sum_{j=0}^t c_{i,j}}{t}$$

do not converge, as $t \rightarrow \infty$.

Proof. As above, for the sake of clarity we present the proof in the case that $\epsilon_R = \epsilon_C = \epsilon$. Throughout the following proof, all logarithms will be taken with base $1 + \epsilon$.

Assume without loss of generality that $r_{1,0} > .999$ and $c_{1,0} < .001$. First note that if $c_{2,t} < 1/2$ then r_1 will must increase and c_1 will decrease, and thus without loss of generality, we may assume that $r_{1,0} \geq .999$, $c_{1,0} < .001$, and $c_{2,0} \geq 1/2$. For some $k \leq \log 10$, it must be the case that after k timesteps we have $c_{2,k} \geq .9$, and $\log(r_{1,k}) - \log(r_{i,k}) \geq \log 999 - k$, for $i = 2, 3$. At this point $\log(c_2/c_3)$, $\log(c_3/c_1)$, and $\log(r_1/r_2)$, $\log(r_3/r_2)$ will all continue to increase until $r_3 \geq 1/3 - .001$. Let t_1 denote the number of steps before $r_1 < .9$, and note that

$$t_1 \geq \log 999 - k - \log 10.$$

At this point, we must have

$$\log(r_1/r_2) \geq .9t_1, \log(c_2/c_3) \geq .9t_1, \log(c_3/c_1) \geq .9t_1.$$

After another at most $\log 10$ steps, $r_3 > .9$, and r_3 will continue to increase until $c_2 < .9$. Let t_2 denote the time until $c_2 \leq .9$, which must mean that $c_1 \approx .1$ since c_3 is decreasing, and note that

$$t_2 \geq 1.8t_1 - 2 \log 10,$$

where the last term is due to the steps that occur when neither r_1 nor r_3 were at least $.9$. At this time point, we must have that

$$\log(c_2/c_3) \geq .9t_2, \log(r_3/r_1) \geq .9t_2, \log(r_1/r_2) \geq .9t_2.$$

After another at most k^3 steps, $c_1 > .9$, and we can continue arguing as above, to yield that after another $t_3 \geq 1.8t_2 - 2 \log 10$ steps, $r_3 < .9$, $r_2 \approx .1$, and $\log(c_1/c_2) \geq .9t_3, \log(c_2/c_3) \geq .9t_3$. Inductively applying these arguments shows that the amount of time during which the weight of a single strategy is held above $.9$, increases by a factor of at least 1.8 in each iteration, and thus the cumulative distributions $\sum_{j=1}^t r_i/t$ cannot converge. \square

4 Stochastic Setting

In this section we prove an analog of Theorem [1](#) for the multiplicative weights learning algorithm in the stochastic setting. We show that in this setting, no matter the initial configuration, with probability tending towards 1 , the cumulative distributions of the row and column player will be far from the Nash equilibrium. To show this, we will make use of the same potential function [\(1\)](#) as in the proof of Theorem [1](#), and analyze its *expected* drift. Although the expectation operator doesn't commute with the application of the potential function (and thus we cannot explicitly use the monotonicity of the potential function as calculated above), unsurprisingly, in expectation the potential function increases. While the drift in the potential function vanished at the equilibrium in the distribution payoff setting, in this setting, the randomness, together with the non-negativity of the potential function allow us to bound the expected drift by a positive constant when the distributions are not near the boundary of the unit cube. Given this, as in the previous section we will then be able to show that for any constant, with probability 1 after a sufficiently long time the value of the potential function will be at least that constant. Given this, analogs of Lemmas [2](#) and [3](#) then show that the cumulative distributions tend away from the equilibrium with all but inverse exponential probability. Our main theorem in this setting is the following.

Theorem 2. *If the row player uses multiplicative updates with update parameter $(1 + \epsilon_R)$, and the column player uses multiplicative updates with update parameter $(1 + \epsilon_C)$, then from any initial pair of distributions, after t time steps, either the dynamics have left the simplex $r_i, c_i \in (1/3 - .2, 1/3 + .2)$ at some time step $t_0 \leq t$, or with all but inverse exponential probability will be at distance $\exp(-\Omega(t))$ from the equilibrium.*

To prove the theorem, we need the following lemma –whose proof is deferred to the full version– that establishes the desired drift of potential [\(1\)](#).

Lemma 4. *If $r_i, c_i \in (1/3 - .2, 1/3 + .2)$, then*

$$\mathbb{E}[\Phi(\mathbf{r}_{t+1}, \mathbf{c}_{t+1}) | \mathbf{r}_t, \mathbf{c}_t] \geq \Phi(\mathbf{r}_t, \mathbf{c}_t) + \max \left(\frac{\Phi(\mathbf{r}_t, \mathbf{c}_t) \log(1 + \epsilon_R)}{240}, \frac{(\log(1 + \epsilon_R))^2}{24000} \right).$$

We are now prepared to finish our proof of Theorem 2. We do so by analyzing the one-dimensional random walk defined by the value of the potential function over time. As long as our pair of distributions has probability values in $(1/3 - .2, 1/3 + .2)$, there is a constant (a function of ϵ_R) drift pushing us away from the equilibrium (which corresponds to the minimum of the potential function). Using martingale arguments we can show then that with all but inverse exponential probability the value of the potential function will be γt for some constant γ , independent of t , unless we have exited the ball of radius 0.2 around the equilibrium.

Proof of theorem 2 We wish to analyze the random walk $(\mathbf{r}_0, \mathbf{c}_0), (\mathbf{r}_1, \mathbf{c}_1), \dots$, where the evolution is according to the stochastic dynamics. To do this analysis, we'll consider the one dimensional random walk X_0, X_1, \dots , where $X_i = \Phi(\mathbf{r}_t, \mathbf{c}_t)$, assuming that the walk starts within the ball $r_i, c_i \in (1/3 - .2, 1/3 + .2)$. Note first that $|X_{t+1} - X_t| \leq 4 \log(1 + \epsilon_R)$. Next, from the X_i 's, we can define a martingale sequence Y_0, Y_1, \dots where $Y_0 = X_0$, and for $i \geq 1, Y_{i+1} := Y_i + X_{i+1} - \mathbb{E}[X_{i+1}|X_i]$.

Clearly the sequence Y_i has the bounded difference property, specifically $|Y_{i+1} - Y_i| \leq 8 \log(1 + \epsilon_R)$, and thus we can apply Azuma's inequality² to yield that with probability at least $1 - 2 \exp(-t^{2/3}/2), Y_t \geq Y_0 - t^{5/6} 8 \log(1 + \epsilon_R)$.

Notice next that, from our definition of the martingale sequence $\{Y_t\}_t$ and Lemma 4, it follows that, as long as the distributions are contained within the ball $r_i, c_i \in (1/3 - .2, 1/3 + .2), X_t \geq Y_t + t \cdot \frac{(\log(1+\epsilon_R))^2}{24000}$.

Let us then define T to be the random time where the distributions exit the ball for the first time, and consider the sequence of random variables $\{Y_{t \wedge T}\}_t$. Clearly, the new sequence is also a martingale, and from the above we get $X_{t \wedge T} \geq Y_{t \wedge T} + (t \wedge T) \cdot \frac{(\log(1+\epsilon_R))^2}{24000}$, and, with probability at least $1 - 2 \exp(-t^{2/3}/2), Y_{t \wedge T} \geq Y_0 - t^{5/6} 8 \log(1 + \epsilon_R)$. Hence, with probability at least $1 - 2 \exp(-t^{2/3}/2), X_{t \wedge T} \geq Y_0 - t^{5/6} 8 \log(1 + \epsilon_R) + (t \wedge T) \cdot \frac{(\log(1+\epsilon_R))^2}{24000}$ and the theorem follows. ■

5 Multi-armed Setting

Perhaps unsurprising in light of the inability of multiplicative weight updates to converge to the Nash equilibrium in the stochastic setting, we show the analogous result for the multi-armed setting. The proof very closely mirrors that of Theorem 2 and, in fact the only notable difference is in the calculation of the expected drift of the potential function. The analogous of Lemma 4 can be easily shown to hold and the rest of the proof follows easily; we defer details to the full version.

6 Conclusions and Open Problems

We showed that simple learning approaches which are known to solve zero-sum games cannot work for Nash equilibria in general bimatrix games; we did so by considering the simplest possible game. Some of our non-convergence proofs are rather daunting; it

² **Azuma's inequality:** Let X_1, X_2, \dots be a martingale sequence with the property that for all $t, |X_t - X_{t+1}| \leq c$; then for all positive t , and any $\gamma > 0, \Pr[|X_t - X_1| \geq c\gamma\sqrt{t}] \leq 2e^{-\gamma^2/2}$.

would be interesting to investigate whether considering more complicated games results in simpler (and easier to generalize to larger classes of algorithms) proofs. In particular, Shapley's game has a unique Nash equilibrium; intuitively, one algorithmically nasty aspect of Nash equilibria in nonzero-sum games is their non-convexity: there may be multiple discrete equilibria. Zinkevich [14] has taken an interesting step in this direction, defining a variant of Shapley's game with an extra pure Nash equilibrium. However, after quite a bit of effort, it seems to us that a non-convergence proof in Zinkevich's game may not be ultimately much easier than the ones presented here.

Despite the apparent difficulties, however, we feel that a very strong lower bound, valid for a very large class of algorithms, may ultimately be proved.

References

1. Blum, A., Hajiaghayi, M., Ligett, K., Roth, A.: Regret minimization and price of total anarchy. In: STOC 2008 (2008)
2. Friedman, E., Shenker, S.: Learning and implementation on the Internet. Working paper (1997)
3. Freund, Y., Schapire, R.E.: Adaptive game playing using multiplicative weights. *Games and Economic Behavior* 29, 79–103 (1999)
4. Hart, S., Mansour, Y.: The communication complexity of uncoupled Nash equilibrium procedures. In: STOC 2007 (2007)
5. Hart, S., Mas-Colell, A.: A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5), 1127–1150 (2000)
6. Karlin, S.: *Mathematical Methods and Theory in Games, Programming and Economics*. Dover, New York (1994)
7. Kleinberg, R., Piliouras, G., Tardos, É.: Multiplicative updates outperform generic no-regret learning in congestion games. In: STOC 2009 (2009)
8. Robinson, J.: An iterative method of solving a game. In: *Annals of Mathematics* (1951)
9. Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: STOC 2009 (2009)
10. Shapley, L.S.: Some topics in two-person games. In: Dresher, M., Aumann, R.J., Shapley, L.S., Tucker, A.W. (eds.) *Advances in Game Theory* (1964)
11. Auer, P., Cesa-Bianchi, N., Freund, Y., Schapire, R.E.: The Nonstochastic Multiarmed Bandit Problem. *SIAM J. Comput.* 32(48), 77 (2002)
12. Arora, S., Hazan, E., Kale, S.: *The Multiplicative Weights Update Method: a Meta Algorithm and Applications* (2005) (Manuscript)
13. Abernethy, J., Rakhlin, A.: Beating the adaptive bandit with high probability. In: COLT 2009 (2009)
14. Zinkevich, M.: Theoretical guarantees for algorithms in multi-agent settings. Tech. Rep. CMU-CS-04-161, Carnegie Mellon University (2004)
15. Benaim, M., Hirsch, M.W.: Mixed Equilibria and Dynamical Systems Arising from Fictitious Play in Perturbed Games. *Games and Economic Behavior* 29, 36–72 (1999)

On the Structure of Weakly Acyclic Games^{*}

Alex Fabrikant^{1,**}, Aaron D. Jaggard^{2,***}, and Michael Schapira^{3,†}

¹ Princeton University, Princeton, NJ, USA
alex@cal.berkeley.edu

² Colgate University, Hamilton, NY, USA & Rutgers University,
Piscataway, NJ, USA
adj@dimacs.rutgers.edu

³ Yale University, New Haven, CT, USA & UC Berkeley, Berkeley, CA, USA
michael.schapira@yale.edu

Abstract. The class of *weakly acyclic games*, which includes potential games and dominance-solvable games, captures many practical application domains. Informally, a weakly acyclic game is one where natural distributed dynamics, such as better-response dynamics, cannot enter *inescapable oscillations*. We establish a novel link between such games and the existence of pure Nash equilibria in subgames. Specifically, we show that the existence of a *unique* pure Nash equilibrium in every *subgame* implies the weak acyclicity of a game. In contrast, the possible existence of *multiple* pure Nash equilibria in every subgame is *insufficient* for weak acyclicity.

1 Introduction

In many domains, convergence to a pure Nash equilibrium is a fundamental problem. In many engineered agent-driven systems that fare best when steady at a pure Nash equilibrium, convergence to equilibrium is expected [7,9] to happen via *better-response (best-response) dynamics*: Start at some strategy profile. Players take turns, in some arbitrary order, with each player making a better response (best response) to the strategies of the other players, i.e., choosing a strategy that increases (maximizes) their utility, given the current strategies of the other players. Repeat this process until no player wants to switch to a different strategy, at which point we reach a pure Nash equilibrium.

For better-response dynamics to converge to a pure Nash equilibrium regardless of the initial strategy profile, a *necessary* condition is that, from every strategy profile, there exist *some* better-response improvement path (that is, a sequence of players' better responses) leading from that strategy profile to a pure Nash equilibrium. Games for which this property holds are called “weakly acyclic games”

^{*} Partially supported by the DIMACS Special Focus on Communication Security and Information Privacy.

^{**} Supported by a Cisco URP grant and a Princeton University postdoctoral fellowship.

^{***} Partially supported by NSF grants 0751674 and 0753492.

[†] Supported by NSF grant 0331548.

[10,17]. Both potential games [12,15] and dominance-solvable games [13] are special cases of weakly acyclic games.

In games that are not weakly acyclic, under better-/best-response dynamics, there are starting states from where the game is guaranteed to oscillate indefinitely. Moreover, the weak acyclicity of a game implies that natural decentralized dynamics (e.g., randomized better-/best-response, or no-regret dynamics) are stochastically guaranteed to reach a pure Nash equilibrium [8,17]. Thus, weakly acyclic games capture the possibility of reaching pure Nash equilibria via simple, local, globally-asynchronous interactions between strategic agents, independently of the starting state. We assert this is *the* realistic notion of “convergence” in most distributed systems.

1.1 A Motivating Example

We now look at an example inspired by interdomain routing that has this natural form of convergence despite it being, formally, possible that the network will never converge. In keeping with results that we study here, we consider best-response dynamics of a routing model in which each node can see each other node’s current strategy, i.e., its “next hop” (the node to which it forwards its data en route to the destination), as contrasted with models where nodes depending on path announcements to learn this information. (Levin et al. [7] formalized routing dynamics in which nodes learn about forwarding through path announcements.)

Consider the network on four nodes shown in Fig. 1. Each of the nodes 1, 2, and 3 is trying to get a path for network traffic to the destination node d . A strategy of a node i is a choice of a neighbor to whom i will forward traffic; the strategy space of node i , S_i , is its neighborhood in the graph.. The utility of d is independent of the outcome, and the utility u_i of node $i \neq d$ depends only on the path that i ’s traffic takes to the destination (and is $-\infty$ if there is no path). We only need to consider the relationships between the values of u_i on all possible paths; the actual values of the utilities do not make a difference. Using $132d$ to denote the path from 1 to 2 to 3 to d , and similarly for other paths, here we assume the following: $u_1(132d) > u_1(1d) > u_1(13d) > -\infty$; $u_2(213d) > u_2(2d) > u_2(21d) > -\infty$; $u_3(321d) > u_3(3d) > u_3(32d) > -\infty$; and $u_i(P) = -\infty$ for all other paths P , e.g., $u_1(12d) = -\infty$. These preferences are indicated by the lists of paths in order of decreasing preference next to the nodes in Fig. 1. (d, d, d) is a the unique pure




Fig. 1. Instance of the interdomain routing game that is weakly acyclic and has a best-response cycle

¹ In some of the economics literature, the terms “weak finite-improvement path property” (weak FIP) and “weak finite best-response path property” (weak FBRP) are also used, for weak acyclicity under better- and best-response dynamics, respectively.

Nash equilibrium in this game, and, ideally, the dynamics would always converge to it. However, there exists a best-response cycle:

$$\begin{array}{cccccccccccc}
 1d & & 132d & & 13d & & 13d & & 1d & & 1d & & 1d \\
 2d & \xrightarrow{1} & 2d & \xrightarrow{3} & 2d & \xrightarrow{2} & 213d & \xrightarrow{1} & 21d & \xrightarrow{3} & 21d & \xrightarrow{2} & 2d \xrightarrow{1} \dots \\
 32d & & 32d & & 3d & & 3d & & 3d & & 321d & & 32d
 \end{array}$$

Here, each triple lists the paths that nodes 1, 2, and 3 get; the nodes’ strategies correspond to the second node in their respective paths. The node above the arrow between two triples is the one that makes a best response to get from one triple to the next.

Once the network is in one of these states², there is a fair activation sequence (i.e., in which every node is activated infinitely often) such that each activated node best responds to the then-current choices of the other nodes and such that the network never converges to a stable routing tree (a pure Nash equilibrium).

Although this cycle seems to suggest that the network in Fig. 1 would be operationally troublesome, it is not as problematic as we might fear. From every point in the state space, there is a sequence of best-response moves that leads to the unique pure Nash equilibrium. We may see this by inspection in this case, but this example also satisfies the hypotheses of our main theorem below. So long as each node has some positive probability of being the next activated node, then, with probability 1, the network will *eventually* converge to the unique stable routing tree, regardless of the initial configuration of the network.

1.2 Our Results

Weak acyclicity is connected to the study of the computational properties of *sink equilibria* [2,4], minimal collections of states from which best-response dynamics cannot escape: a game is weakly acyclic if and only if all sinks are “singletons”, that is, pure Nash equilibria. Unfortunately, Mirrokni and Skopalik [11] found that reliably checking weak acyclicity is extremely computationally intractable in the worst case (PSPACE-complete) even in succinctly-described games. This means, inter alia, that not only can we not hope to consistently check games in these categories for weak acyclicity, but we cannot even hope to have general short “proofs” of weak acyclicity, which, once somehow found, could be tractably checked.

With little hope of finding robust, effective ways to consistently check weak acyclicity, we instead set out to find *sufficient* conditions for weak acyclicity: finding usable properties that imply weak acyclicity may yield better insights into at least *some* cases where we need weak acyclicity for the application.

In this work, we focus on general normal-form games. Potential games, the much better understood subcategory of weakly acyclic games, are known to have

² For example, this might happen if the link between 2 and d temporarily fails. 2 would always choose to send traffic to 1 (if anywhere); 1 would eventually converge to sending traffic directly to d (with 2 sending its traffic to 1), and 3 would then be able to send its traffic along $321d$. Once the failed link between 2 and d is restored, 2’s best response to the choices of the other nodes is to send its traffic directly to d , resulting in the first configuration of the cycle above.

the following property, which we will refer to as *subgame stability*, abbreviated **SS**: not only does a pure Nash equilibrium exist in the game, but a pure Nash equilibrium exists in each of its *subgames*, i.e., in each game obtained from the original game by the removal of players' strategies. Subgame stability is a useful property in many contexts. For example, in network routing games, subgame stability corresponds to the important requirement that there be a stable routing state even in the presence of arbitrary network malfunctions [5]. We ask the following natural question: When is the strong property of subgame stability *sufficient* for weak acyclicity?

Yamamori and Takahashi [16] prove the following two results³:

Theorem: [16] *In 2-player games, subgame stability implies weak acyclicity, even under best response.*

Theorem: [16] *There exist $3 \times 3 \times 3$ games for which subgame stability holds that are not weakly acyclic under best response.*

Thus, subgame stability is sufficient for weak acyclicity in 2-player games, yet is not always sufficient for weak acyclicity in games with $n > 2$ players. Our goal in this work is to (1) identify sufficient conditions for weak acyclicity in the general n -player case; and (2) pursue a detailed characterization of the boundary between games for which subgame stability does imply weak acyclicity and games for which it does not.

Our main result for n -player games shows that a constraint stronger than **SS**, that we term “*unique subgame stability*” (**USS**), is sufficient for weak acyclicity:

Theorem: *If every subgame of a game Γ has a unique pure Nash equilibrium then Γ is weakly acyclic, even under best response.*

This result casts an interesting contrast against the negative result in [16]: *unique equilibria* in subgames guarantee weak acyclicity, but the existence of *more* pure Nash equilibria in subgames can lead to violations of weak acyclicity. Hence, perhaps counter-intuitively, too many stable states can potentially result in persistent instability of local dynamics.

We consider **SS** games, **USS** games, and also the class of *strict and subgame stable* games **SSS**, i.e., subgame stable games which have no ties in the utility functions. We observe that these three classes of games form the hierarchy $\text{USS} \subset \text{SSS} \subset \text{SS}$. We examine the number of players, number of strategies, and the *strictness* of the game (the constraint that there are no ties in the utility function), and give a complete characterization of the weak acyclicity implications of each of these. Our contributions are summarized in Table [1].

1.3 Other Related Work

Weak acyclicity has been specifically addressed in a handful of specially-structured games: in an applied setting, BGP with backup routing [1], in a game-theoretical

³ Yamamori and Takahashi use the terms *quasi-acyclicity* for weak acyclicity under best response, and *Pure Nash Equilibrium Property (PNEP)* for subgame stability.

Table 1. Results summary: The impact of USS/SSS/SS on weak acyclicity: \checkmark marks classes with guaranteed weak acyclicity, even under best response; \nexists marks classes which admit counter-examples which are not weakly acyclic even under better response. *: only for strict games.

	2 players		3 players				4+ players	
	$2 \times M$	$3 \times M$	$2 \times 2 \times 2$	$\frac{2 \times 2 \times M}{2 \times 3 \times M}$	$2 \times 4 \times 4$	$3 \times 3 \times 3$	$2 \times 2 \times 2 \times 2$	
\exists pNE	\checkmark (Lma 4)	\nexists (easy)	\checkmark^* (Lma 5)	\nexists (easy)				
SS	\checkmark [16]		\nexists (Thm 2)			\nexists (Thm 2 & [16])	\nexists (Thm 2)	
SSS			\checkmark (Lma 5)	\checkmark (Thm 3)	\nexists (Thm 4)	\nexists (Thm 4 & [16])	\nexists (Thm 5)	
USS			\checkmark (Thm 1)					

setting, games with “strategic complementarities” [3,6] (a supermodularity condition on lattice-structured strategy sets), and in an algorithmic setting, in several kinds of succinct games [11]. Milchtaich [10] studied Rosenthal’s congestion games [15] and proved that, in interesting cases, such games are weakly acyclic even if the payoff functions (utilities) are not universal but player-specific. Marden et al. [9] formulated the cooperative-control-theoretic consensus problem as a potential game (implying that it is weakly acyclic); they also defined and investigated a time-varying version of weak acyclicity.

1.4 Outline of Paper

In the following, we recall the relevant concepts and definitions in Section 2, present our sufficient condition for weak acyclicity in Section 3, and our characterization of weak acyclicity implications in Section 4.

2 Weakly Acyclic Games and Subgame Stability

We use standard game-theoretic notation. Let Γ be a *normal-form game* with n players $1, \dots, n$. We denote by S_i be the *strategy space* of the i^{th} player. Let $S = S_1 \times \dots \times S_n$, and let $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$ be the cartesian product of all strategy spaces but S_i . Each player i has a *utility function* u_i that specifies i ’s *payoff* in every strategy-profile of the players. For each strategy $s_i \in S_i$, and every $(n - 1)$ -tuple of strategies $s_{-i} \in S_{-i}$, we denote by $u_i(s_i, s_{-i})$ the utility of the strategy profile in which player i plays s_i and all other players play their strategies in s_{-i} . We will make use of the following definitions.

Definition 1 (better-response strategies). A strategy $s'_i \in S_i$ is a better-response of player i to a strategy profile (s_i, s_{-i}) if $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$.

Definition 2 (best-response strategies). A strategy $s_i \in S_i$ is a best response of player i to a strategy profile $s_{-i} \in S_{-i}$ of the other players if $s_i \in \operatorname{argmax}_{s'_i \in S_i} u_i(s'_i, s_{-i})$

Definition 3 (pure Nash equilibria). A strategy profile s is a pure Nash equilibrium if, for every player i , s_i is a best response of i to s_{-i}

Definition 4 (better- and best-response improvement paths). A better-response (best-response) improvement path in a game Γ is a sequence of strategy profiles s^1, \dots, s^k such that for every $j \in [k - 1]$ (1) s^j and s^{j+1} only differ in the strategy of a single player i and (2) i 's strategy in s^{j+1} is a better response to s_{-i}^j (best response to s_{-i}^j and $u_i(s_i^{j+1}, s_{-i}^j) > u_i(s_i^j, s_{-i}^j)$). The better-response dynamics (best-response dynamics) graph for Γ is the graph on the strategy profiles in Γ whose edges are the better-response (best-response) improvement paths of length 1.

We will use $\Delta R_\Gamma(s)$ and $BR_\Gamma(s)$ to denote the set of all states reachable by, respectively, better and best responses when starting from s in Γ .

We are now ready to define weakly acyclic games [17]. Informally, a game is weakly acyclic if a pure Nash equilibrium can be reached from any initial strategy profile via a better-response improvement path.

Definition 5 (weakly acyclic games). A game Γ is weakly acyclic if, from every strategy profile s , there is a better-response improvement path s^1, \dots, s^k such that $s^1 = s$, and s^k is a pure Nash equilibrium in Γ . (I.e., for each s , there's a pure Nash equilibrium in $\Delta R_\Gamma(s)$.)

We also coin a parallel definition based on best-response dynamics.

Definition 6 (weak acyclicity under best response). A game Γ is weakly acyclic under best response if, from every strategy profile s , there is a best-response improvement path s^1, \dots, s^k such that $s^1 = s$ and s^k is a pure Nash equilibrium in Γ . (I.e., for each s , there's a pure Nash equilibrium in $BR_\Gamma(s)$.)

Weak acyclicity of either kind is equivalent to requiring that, under the respective dynamics, the game has no “non-trivial” sink equilibria [4,2], i.e., sink equilibria containing more than one strategy profile. Conventionally, sink equilibria are defined with respect to best-response dynamics, but the original definition by Goemans et al. [4] takes into account better-response dynamics as well.

The following follows easily from definitions:

Claim. If a game is weakly acyclic under best response then it is weakly acyclic.

On the other hand, the game in Figure 2, mentioned, e.g., in [8], is weakly acyclic, but not weakly acyclic under best response.

Curiously, all of our results apply both to weak acyclicity in its conventional better-response sense and to weak acyclicity under best response. Thus, unlike weak acyclicity itself, the conditions presented in this paper are “agnostic” to the better-/best-response distinction (like the notion of pure Nash equilibria itself).

We now present the notion of subgame stability.

Definition 7 (subgames). A subgame of a game Γ is a game Γ' obtained from Γ via the removal of players' strategies.

	H	T	X
H	2,0	0,2	0,0
T	0,2	2,0	0,0
X	0,0	1,0	3,3

Fig. 2. Matching pennies with a “better-response” escape route, but a best response persistent cycle

	c_0		c_1	
	b_0	b_1	b_0	b_1
a_0	2,2,2	1,2,2	2,1,2	2,2,1
a_1	2,2,1	2,1,2	1,2,2	0,0,0

Fig. 3. $2 \times 2 \times 2$ subgame-stable game with a non-trivial sink

Definition 8 (subgame stability). Subgame stability is said to hold for a game Γ if every subgame of Γ has a pure Nash equilibrium. We use **SS** to denote the class of subgame stable games.

Definition 9 (unique subgame stability). Unique subgame stability is said to hold for a game Γ if every subgame of Γ has a unique pure Nash equilibrium. We use **USS** to denote the class of such games.

We will also consider games in which no player has two or more equally good responses to any fixed set of strategies played by the other players. Following, e.g., [14], we define *strict games* as follows.

Definition 10 (strict game). A game Γ is strict if, for any two distinct strategy profiles $s = (s_1, \dots, s_n)$ and $s' = (s'_1, \dots, s'_n)$ such that there is some $j \in [n]$ for which $s' = (s'_j, s_{-j})$ (i.e., s and s' differ only in j 's strategy), then $u_j(s) \neq u_j(s')$.

Definition 11 (SSS). We use **SSS** to denote the class of games that are both strict and subgame stable.

It's easy to connect unique subgame stability and strictness. To do so, we use the next definition, which will also play a role in our main proofs.

Definition 12 (subgame spanned by profiles). For game Γ with n players and profiles s^1, \dots, s^k in Γ , the subgame spanned by s^1, \dots, s^k is the subgame Γ' of Γ in which the strategy space for player i is $S'_i = \{s_i^j | 1 \leq j \leq k\}$.

Claim. The categories **USS**, **SSS**, and **SS** form a hierarchy: $\text{USS} \subset \text{SSS} \subset \text{SS}$

Proof. $\text{SSS} \subset \text{SS}$ by definition. To see that $\text{USS} \subset \text{SSS}$ observe the following. If a game is not strict, there are $s_j, s'_j \in S_j$ and s_{-j} such that $u_j(s_j, s_{-j}) = u_j(s'_j, s_{-j})$. Both strategy profiles in the subgame spanned by (s_j, s_{-j}) and (s'_j, s_{-j}) are pure Nash equilibria, violating unique subgame stability. \square

3 Sufficient Condition for Weak Acyclicity with n Players

When is weak acyclicity guaranteed in n -player games for $n \geq 3$? We prove that the existence of a *unique* pure Nash equilibrium in every subgame implies weak acyclicity. We note that this is not true when subgames can contain multiple pure Nash equilibria [16]. Thus, while at first glance, introducing extra equilibria

might seem like it would make it harder to get “stuck” in a non-trivial component of the state space with no “escape path” to an equilibrium, this intuition is false; allowing extra pure Nash equilibria in subgames actually enables the existence of non-trivial sinks.

Theorem 1. *Every game Γ that has a unique pure Nash equilibrium in every subgame $\Gamma' \subseteq \Gamma$ is weakly acyclic under best-response (as are all of its subgames).*

We shall need the following technical lemma:

Lemma 1. *If s is a strategy profile in Γ , and Γ' is the subgame of Γ spanned by $BR_\Gamma(s)$, then any best-response improvement path s, s^1, \dots, s^k in Γ' that starts at s is also a best-response improvement path in Γ . x*

Proof. We proceed by induction on the length of the path. The base case is tautological. Inductively, assume s, \dots, s_i is a best-response improvement path in Γ . The strategy s^{i+1} is a best response to s^i in Γ' by some player j . This guarantees that s^i is not a best response by j to s^i_{-j} in Γ' , let alone in Γ , so $\Gamma' \supseteq BR_\Gamma(s) \supseteq BR_\Gamma(s^i)$ must contain a best-response \hat{s}^i_j to s^i_{-j} in Γ , and since s^{i+1}_j is a best-response in Γ' , we are guaranteed that $u_j(\hat{s}^i_j, s^i_{-j}) = u_j(s^{i+1}_j, s^i_{-j})$, so s^{i+1} must be a best-response in Γ . \square

We may now prove Theorem \square

Proof (Proof of Theorem \square). To prove Theorem \square , assume that Γ is a game satisfying the hypotheses of the theorem, and for a subgame $\Delta \subseteq \Gamma$, denote by s_Δ the unique pure Nash equilibrium in Δ . We will proceed by induction up the semilattice of subgames of Γ . The base cases are trivial: any $1 \times \dots \times 1$ subgame is weakly acyclic for lack of any transitions. Suppose that for some subgame Γ' of game Γ we know that every strict subgame $\Gamma'' \subsetneq \Gamma'$ is weakly acyclic.

Suppose that Γ' is not weakly acyclic: it has a state s from which its unique pure Nash equilibrium $s_{\Gamma'}$ cannot be reached by best responses. Let Γ'' be the game spanned by $BR(s)$. Consider the cases of (i) $s_{\Gamma'} \in \Gamma''$ and (ii) $s_{\Gamma'} \notin \Gamma''$:

Case (i): $s_{\Gamma'} \in \Gamma''$. This requires that, for an arbitrary player j with more than 1 strategy in Γ' , there be a best-response improvement path from s to some profile \hat{s} where j plays the same strategy as it does in $s_{\Gamma''}$. Take one such j , and let Γ^j be the subgame of Γ' where j is restricted to playing \hat{s}_j only. Since $s_{\Gamma'}$ is in Γ^j , the inductive hypothesis guarantees a best-response improvement path in Γ^j from \hat{s} to $s_{\Gamma'}$. By construction, that path must only involve best responses by players other than j , who have the same strategy options in Γ^j as they did in Γ' , so that path is also a best-response improvement path in Γ' , assuring a best-response improvement path in Γ' from s to $s_{\Gamma'}$ via \hat{s} .

Case (ii): $s_{\Gamma'} \notin \Gamma''$. Then, Γ'' 's unique pure equilibrium $s_{\Gamma''}$ must be distinct from $s_{\Gamma'}$. Since $s_{\Gamma'}$ is the only pure equilibrium in Γ' , $s_{\Gamma''}$ must have an outgoing best-response edge to some profile \hat{s} in Γ' . But the inductive hypothesis ensures that $s_{\Gamma''} \in BR_{\Gamma''}(s)$; by Lemma \square , $s_{\Gamma''} \in BR_{\Gamma'}(s)$, which then ensures that \hat{s} must also be in $BR_{\Gamma'}(s)$, and hence in Γ'' , so $s_{\Gamma''}$ isn't an equilibrium in Γ'' . \square

4 Characterizing the Implications of Subgame Stability

[16] establishes that in 2-player games, subgame stability implies weak acyclicity, even under best response, yet this is not true in 3x3x3 games. We now present a complete characterization of when subgame stability is sufficient for weak acyclicity, as a function of game size and strictness. Our next result shows that the two-player theorem of [16] is maximal:

Theorem 2. *Subgame stability is not sufficient for weak acyclicity even in non-strict $2 \times 2 \times 2$ games.*

Proof. The game in Fig. 3 can be seen to provide the needed counterexample.

However, if we require the games to be strict, subgame stability turns out to be somewhat useful in 3-player games:

Theorem 3. *In any strict $2 \times 2 \times M$ or $2 \times 3 \times M$ game, subgame stability implies weak acyclicity, even under best response.*

The proof of the theorem rests on several technical lemmas:

Lemma 2. *In strict games, neither a pure Nash equilibrium and strategy profiles differing from it in only one player's action can be part of a non-trivial sink of the best-response dynamics.*

Proof. A pure Nash equilibrium always forms a 1-node sink. If the game is strict, profiles differing by one player's action have to give that one player a strictly lower payoff, requiring a best-response transition to the equilibrium's sink. Any node connected to either cannot be in a sink.

Lemma 3. *The profiles of a game that constitute a non-trivial sink of the best-response dynamics cannot be all contained within a subgame which is weakly acyclic under best-response.*

Proof (sketch). The lemma comes from considering, for a sink of Γ contained in a weakly acyclic subgame Γ' , a best response path in Γ' from the sink to an equilibrium. The first transition on that path that is not a best response in Γ (or, in absence of such, the transition from the equilibrium of Γ' that makes it a non-equilibrium in Γ), will have to lead out of Γ' but remain in the sink.

We then consider the corner cases of 3-player, $2 \times 2 \times 2$ strict games, and 2-player, $2 \times m$ games, where weak acyclicity requires even less than subgame stability. The former result forms the base case for Theorem 3, and both might also be of independent interest.

Lemma 4 (proof in the full version). *In any $2 \times m$ game, and if there is a pure Nash equilibrium, the game is weakly acyclic, even under best response.*

Lemma 5. *In any strict $2 \times 2 \times 2$ game, if there is a pure Nash equilibrium, the game is weakly acyclic, even under best response.*

Proof. In strict $2 \times 2 \times 2$ games, Lemma 2 leaves 4 other strategy profiles, with the possible best-response transitions forming a star in the underlying undirected graph. Since best-response links are antisymmetric ($s \rightarrow s'$ and $s' \rightarrow s$ cannot both be best-response moves), there can be no cycle among those 4 profiles, and thus no non-trivial sink components. \square

Proof (sketch of Theorem 3). The full proof is long and technical, and is relegated to the full version of the paper.

We treat the $2 \times 2 \times M$ case first. Naming the equilibrium of the game (a_0, b_0, c_0) , Lemmas 2 and 3 guarantee that the sink must contain a profile where player 3 plays c_0 , yet the only such profile that can be in the sink is (a_1, b_1, c_0) , the total degree of which in the best-response directed graph is at most 1 (also by Lemma 2), which cannot happen for a node in a non-trivial sink.

The $2 \times 3 \times M$ case is much more complex. The proof operates inductively on M . From the inductive hypothesis, the $2 \times 2 \times M$ result, and Lemma 3, we get that the *smallest* $2 \times 3 \times M$ game Γ that is not weakly acyclic under best response must have a non-trivial sink spanning Γ . Given such a sink, we then use Lemma 2 and a similar result that excludes from the non-trivial sink any profiles adjacent to the equilibrium of the $2 \times 2 \times M$ subgame that does not contain the global equilibrium. The proof concludes by a detailed examination of the possible structures of such a sink under all those constraints, which yield a contradiction in every case.

Theorem 3 is maximal. All bigger sizes of 3-player games admit subgame-stable counter-examples that are not weakly acyclic:

Theorem 4. *In non-degenerate⁴ strict 3-player games, the existence of pure Nash equilibria in every subgame is insufficient to guarantee weak acyclicity, for any game with at least 3 strategies for each player, and any game with at least 4 strategies for 2 of the players.*

Proof (sketch). The first half of the theorem follows directly from a specific counterexample game in [16]. There, the strict 3-player, $3 \times 3 \times 3$ game in question is stated to demonstrate that SSS does not imply weak acyclicity under best response. However, their very same counter-example is not even weakly acyclic under better response. Here, we give a $2 \times 4 \times 4$ counter-example to establish the second half of the theorem, and a $3 \times 3 \times 3$ counterexample slightly cleaner than the one in [16]. Close inspection of the games $\Gamma_{3,3,3}$ and $\Gamma_{4,4,2}$ shown in Figure 4 reveals that these are not weakly acyclic but are strict and subgame stable.

With 4 or more players, a more mechanistic approach produces analogous examples even with just 2 strategies per player:

Theorem 5. *In a strict n -player game for an arbitrary $n \geq 4$, the existence of pure Nash equilibria in every subgame is insufficient to guarantee weak acyclicity, even with only 2 strategies per player.*

⁴ Each player has 2 or more strategies.

	c_0			c_1			c_2		
	b_0	b_1	b_2	b_0	b_1	b_2	b_0	b_1	b_2
a_0	0, 0, 0	5, 5, 4	5, 4, 5	4, 5, 5	0, 1, 1	0, 2, 1	5, 5, 4	5, 4, 5	0, 2, 2
a_1	5, 4, 5	1, 1, 0	4, 5, 5	1, 0, 1	5, 5, 5	1, 2, 1	1, 0, 2	1, 1, 2	1, 2, 2
a_2	4, 5, 5	2, 1, 0	2, 2, 0	5, 5, 4	2, 1, 1	2, 2, 1	2, 0, 2	2, 1, 2	2, 2, 2

	c_0				c_1			
	b_0	b_1	b_2	b_3	b_0	b_1	b_2	b_3
a_0	5, 5, 5	0, 1, 0	0, 2, 0	0, 3, 0	5, 5, 4	0, 1, 1	5, 4, 5	0, 3, 1
a_1	1, 0, 0	1, 1, 0	5, 5, 4	5, 4, 5	1, 0, 1	1, 1, 1	4, 5, 5	1, 3, 1
a_2	2, 0, 0	5, 4, 5	2, 2, 0	4, 5, 5	2, 0, 1	2, 1, 1	2, 2, 1	2, 3, 1
a_3	5, 4, 5	4, 5, 5	3, 2, 0	3, 3, 0	3, 0, 1	3, 1, 1	3, 2, 1	5, 5, 5

Fig. 4. 3-player strict subgame stable games that are not weakly acyclic, even under better-response dynamics

Proof. For strategy profiles in $\{0, 1\}^n$, using indices mod n , set the utilities to:

$$u(\mathbf{s}) = \begin{cases} (4, \dots, 4) & \text{at } \mathbf{s} = (1, \dots, 1) \\ (3, \dots, 3, \underset{i^{\text{th}}}{2}, 3, \dots, 3) & \text{when } s_{i-1} = s_i = 1, s_{-(i-1,i)} = 0 \\ (3, \dots, 3, \underset{i+1^{\text{th}}}{2}, 3, \dots, 3) & \text{when } s_i = 1, s_{-i} = 0 \\ \mathbf{s} & \text{else (for the “sheath”).} \end{cases}$$

Similarly to Theorem 4, this plants a global pure Nash equilibrium at $(1, \dots, 1)$, and creates a “fragile” better-response cycle. Here, the cycle alternates between profiles with edit distance $n - 1$ and $n - 2$ from the global pure Nash equilibrium. At every point of the cycle, the only non-sheath profiles 1 step away are its predecessor and successor on the cycle, so the cycle is persistent. Since each profile with edit distance $n - 1$ from the equilibrium is covered, removing any player’s 1 strategy breaks the cycle, thus guaranteeing a pure Nash equilibrium in every subgame by the same reasoning as above. □

We note that, in the counter-example results — Theorems 2, 4, and 5 — the counter-example games of fixed size easily extend to games with extra strategies for some or all players, or with extra players, by “padding” the added part of the payoff table with negative, unique values that, for the added profiles, make payoffs independent of the other players, such as, e.g., $u_i(s) = -s_i$. This preserves SS, SSS, and USS properties without changing weak acyclicity. Thus, this completes our classification of weak acyclicity under the three subgame-based properties, as shown in Table 1.

5 Concluding Remarks

The connection between weak acyclicity and unique subgame stability that we present is surprising, but not immediately practicable: in most succinct game

representations, there is no reason to believe that checking unique subgame stability will be tractable in many general settings. In a complexity-theoretic sense, USS is *closer* to tractability than weak acyclicity: Any reasonable game representation will have some “reasonable” representation of subgames, i.e., one in which *checking* whether a state is a pure Nash equilibrium is tractable, which puts unique subgame stability in a substantially easier complexity class, Π_3P , than the class PSPACE for which weak acyclicity is complete in many games.

We leave open the important question of finding efficient algorithms for checking unique subgame stability, which may well be feasible in particular classes of games. Also open and relevant, of course, is the question of more broadly applicable and tractable conditions for weak acyclicity. In particular, there may well be other levels of the subgame stability hierarchy between SSS and USS that could give us weak acyclicity in broader classes of games.

References

1. Engelberg, R., Schapira, M.: Hierarchy of Internet routing games, 12 p. (in submission 2010)
2. Fabrikant, A., Papadimitriou, C.H.: The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond. In: Proc. of SODA, pp. 844–853 (2008)
3. Friedman, J.W., Mezzetti, C.: Learning in games by random sampling. Journal of Economic Theory 98, 55–84 (2001)
4. Goemans, M., Mirrokni, V., Vetta, A.: Sink equilibria and convergence. In: Proceedings of IEEE FOCS, pp. 142–151 (2005)
5. Griffin, T.G., Shepherd, F.B., Wilfong, G.: The stable paths problem and interdomain routing. IEEE/ACM Trans. Netw. 10(2), 232–243 (2002)
6. Kukushkin, N.S., Takahashi, S., Yamamori, T.: Improvement dynamics in games with strategic complementarities. Intl. J. Game Theory 33(2), 229–238 (2005)
7. Levin, H., Schapira, M., Zohar, A.: Interdomain routing and games. In: Proceedings of ACM STOC, pp. 57–66 (2008)
8. Marden, J.R., Young, H.P., Arslan, G., Shamma, J.S.: Payoff-based dynamics in multi-player weakly acyclic games. SIAM J. on Control and Optimization 48, 373–396 (2009)
9. Marden, J., Arslan, G., Shamma, J.: Connections between cooperative control and potential games illustrated on the consensus problem. In: Proceedings of the European Control Conference (July 2007)
10. Milchtaich, I.: Congestion games with player-specific payoff functions. GEB 13, 111–124 (1996)
11. Mirrokni, V.S., Skopalik, A.: On the complexity of nash dynamics and sink equilibria. In: ACM Conference on Electronic Commerce, pp. 1–10 (2009)
12. Monderer, D., Shapley, L.S.: Potential games. GEB 14, 124–143 (1996)
13. Moulin, H.: Dominance solvable voting schemes. Econometrica 47, 1337–1351 (1979)
14. Nisan, N., Schapira, M., Zohar, A.: Asynchronous best-reply dynamics. In: Proceedings of the Workshop on Internet Economics, pp. 531–538 (2008)
15. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory 2, 65–67 (1973)
16. Yamamori, T., Takahashi, S.: The pure Nash equilibrium property and the quasi-acyclic condition. Economics Bulletin 3(22), 1–6 (2002)
17. Young, H.P.: The evolution of conventions. Econometrica 61(1), 57–84 (1993)

A Direct Reduction from k -Player to 2-Player Approximate Nash Equilibrium*

Uriel Feige** and Inbal Talgam-Cohen

Weizmann Institute of Science, PO Box 26, Rehovot 76100, Israel
{uriel.feige,inbal.talgam}@weizmann.ac.il

Abstract. We present a direct reduction from k -player games to 2-player games that preserves approximate Nash equilibrium. Previously, the computational equivalence of computing approximate Nash equilibrium in k -player and 2-player games was established via an indirect reduction. This included a sequence of works defining the complexity class PPAD, identifying complete problems for this class, showing that computing approximate Nash equilibrium for k -player games is in PPAD, and reducing a PPAD-complete problem to computing approximate Nash equilibrium for 2-player games. Our direct reduction makes no use of the concept of PPAD, eliminating some of the difficulties involved in following the known indirect reduction.

1 Introduction

This manuscript addresses the computation of Nash equilibrium for games represented in normal form. It is known that for 2-player games this problem is PPAD-complete [5], and for k players it is in PSPACE [11]. Moreover, for sufficiently small ϵ , computing ϵ -well-supported Nash equilibrium for 2-player games remains PPAD-complete [6], and for k players it is in PPAD [8]. It follows that, for appropriate choices of ϵ , ϵ -well-supported Nash in k -player games reduces to ϵ -well-supported Nash in 2-player games. However, this chain of reductions is indirect, passing through intermediate notions other than games, and also rather complicated.

In this manuscript we present a direct, "game theoretic" polynomial-time reduction from k -player to 2-player games. In our reduction, every pure strategy of each of the k players is represented by a corresponding pure strategy of one of the 2 players. Previously, a direct reduction preserving exact Nash equilibrium was known from k -player to 3-player games [3]. Such a reduction cannot exist to 2-player games due to issues of irrationality [21], hence the need to consider the notion of ϵ -well-supported Nash in this context. Our reduction guarantees that

* A full version of this paper is available at <http://arxiv.org/abs/1007.3886>. Work supported in part by The Israel Science Foundation (grant No. 873/08).

** The author holds the Lawrence G. Horowitz Professorial Chair at the Weizmann Institute.

for appropriate choices of ϵ_2 and ϵ_k , given any ϵ_2 -well-supported Nash for the 2-player game, normalizing its probabilities according to the above correspondence gives an ϵ_k -well-supported Nash for the k -player game.

The direct reduction makes no use of the concept of PPAD. This eliminates some of the difficulties involved in following the known indirect reduction. It is inevitable that unlike the indirect reduction, our reduction by itself does not establish the PPAD-completeness of computing (or approximating) Nash equilibria. Nevertheless, the new gadgets we introduce are relevant to the notion of PPAD-completeness, as they can be used in other reductions among PPAD problems. Moreover, our reduction provides an alternative to the proof of [8] that finding an approximate Nash equilibrium in k -player games is in PPAD.

The first step of our reduction "linearizes" a k -player game by replacing the multilateral interactions among the k players with bilateral interactions among pairs of players. In the next step, two representative "super-players" replace the multiple players, resulting in a 2-player game. In terms of techniques, the first step of the reduction uses and extends the machinery of gadget games developed by Goldberg and Papadimitriou [13]. We introduce a new gadget for performing approximate multiplication using linear operations, in order to bridge the gap between multiplicative and linear games. The second step of the reduction uses similar methods to [13] and [20] in order to replace multiple players by 2 players, resulting in a combination of a generalized Matching Pennies game and an imitation game.

1.1 Preliminaries

Let $[n] = \{1, \dots, n\}$. Let $\|\mathbf{v}\| = \sum_i |v_i|$, and let \mathbf{v}^{-i} be the vector obtained from \mathbf{v} by removing the i 'th entry. For vectors \mathbf{u} and \mathbf{v} of length n , let $\mathbf{u} \otimes \mathbf{v}$ denote their tensor product written as a vector of length n^2 , where entry $(i - 1)n + j$ is $u_i v_j$. We write $x = y \pm z$ to denote $y - z \leq x \leq y + z$. For vectors, $\mathbf{x} = \mathbf{y} \pm z$ denotes $y_i - z \leq x_i \leq y_i + z$ for every i .

Normal Form Games. G_k is a normal form game with player set $[k]$, where each player's pure strategy set is $[n]$. A *pure strategy profile* $\mathbf{s} \in [n] \times \dots \times [n]$ contains one pure strategy per player, and a *mixed strategy profile* $\mathbf{p} = (\mathbf{p}^1, \dots, \mathbf{p}^k)$ is defined analogously. Let $\tilde{\mathbf{p}} = \mathbf{p}^1 \otimes \dots \otimes \mathbf{p}^k$ be the corresponding *joint mixed strategy* distribution, such that for every pure strategy profile \mathbf{s} , entry $\tilde{\mathbf{p}}[\mathbf{s}] = \prod_i p_{s_i}^i$ is the probability that for every i player i plays pure strategy s_i . Unless stated otherwise, payoffs are rationals in $[0, 1]$. Let M^i (sometimes denoted $M_{G_k}^i$) be player i 's $n \times n^{k-1}$ payoff matrix, where $M^i[j, \mathbf{s}^{-i}]$ is the payoff for playing pure strategy j against \mathbf{s}^{-i} . Given a mixed strategy profile \mathbf{p}^{-i} , the *expected payoff vector* $\mathbf{u}_{G_k}^i = M^i \tilde{\mathbf{p}}^{-i}$ contains player i 's expected payoffs $\mathbf{u}_{G_k}^i[j]$ for playing pure strategy j against \mathbf{p}^{-i} . The expected payoffs are algebraic functions in the probabilities played by the others.

Polymatrix (Linear) Games. Games in which every player plays bilaterally against others, and receives the sum of payoffs obtained from the bilateral interactions. Thus, polymatrix games are actually collections of 2-player games in which every player plays the same strategy in every game she participates

in. Unlike normal form games, the size of polymatrix games is polynomial in n even when the number of players is non-constant. G_m is a polymatrix game with player set $[m]$, where player i has $2 \leq n_i \leq n$ pure strategies and $m - 1$ payoff matrices $M^{i,i'}$ of size $n_i \times n_{i'}$. Entry $M^{i,i'}[j, j']$ is the payoff to player i for playing j against player i' who plays j' . If the interaction between i and i' does not exist or does not influence i 's payoff, $M^{i,i'}$ is set to be all-zeros. Given a pure strategy profile \mathbf{s}^{-i} , the total payoff for playing j is $\sum_{i' \neq i} M^{i,i'}[j, \mathbf{s}^{-i}[i']]$. Given a mixed strategy profile \mathbf{p}^{-i} , the expected payoff vector of player i is $\mathbf{u}_{G_m}^i = \sum_{i' \neq i} M^{i,i'} \mathbf{p}^{i'}$. Equivalently, let $M^i = (M^{i,1} \dots M^{i,m})$ be a matrix containing all of i 's payoff matrices as submatrices, then $\mathbf{u}_{G_m}^i = M^i \mathbf{p}^{-i}$. The expected payoffs are thus linear functions in the probabilities of the others.

Nash Equilibrium. A mixed strategy profile whose supports include only *best response* pure strategies. Given a mixed strategy profile \mathbf{p}^{-i} , j is a best response for player i if it maximizes the expected payoff, i.e., $\mathbf{u}_{G_m}^i[j] = \max_{j' \in [n]} \{\mathbf{u}_{G_m}^i[j']\}$; j is an ϵ -best response if it maximizes the expected payoff up to an additive factor of ϵ . In the context of reductions from k -player to 2-player games, the following fact motivates consideration of approximate rather than exact Nash equilibrium: 2-player games always have a rational Nash equilibrium, while k -player games do not [21]. Out of several possible notions of approximation, we focus on ϵ -well-supported Nash equilibrium, whose supports contain only ϵ -best responses. We shall primarily be interested in small, non-constant values of ϵ , namely $\epsilon = 1/\text{poly}(n)$ and $\epsilon = 1/\exp(n)$. A related weaker notion is ϵ -Nash equilibrium, from which deviating unilaterally cannot improve the expected payoff by more than ϵ (see [11] for a discussion).

Definition 1 (ϵ_k - k NASH and ϵ_m -LINEAR-NASH). *Given a pair of normal form game G_k and accuracy parameter ϵ_k , the problem ϵ_k - k NASH is to find an ϵ_k -well-supported Nash equilibrium of G_k . The problem ϵ_m -LINEAR-NASH is the same for polymatrix game G_m and accuracy parameter ϵ_m .*

1.2 Our Results

Let $(G_{m_1}, \epsilon_{m_1}), (G_{m_2}, \epsilon_{m_2})$ be two pairs of games and accuracy parameters. The games have m_1, m_2 players respectively; player i has n_i^1, n_i^2 pure strategies respectively. The following is based on Bubelis's notion of *reduction scheme* [3].

Definition 2 (Mapping between Games). *A mapping includes:*

- A function $g : [m_1] \rightarrow [m_2]$ mapping players of G_{m_1} to players of G_{m_2} ;
- For every $i \in [m_1]$, an injective function $h_i : [n_i^1] \rightarrow [n_{g(i)}^2]$ mapping pure strategies of player i to distinct pure strategies of player $g(i)$.

Definition 3 (Direct Reduction). *A direct reduction from $(G_{m_1}, \epsilon_{m_1})$ to $(G_{m_2}, \epsilon_{m_2})$ is a mapping from G_{m_1} to G_{m_2} , such that for every ϵ_{m_2} -well-supported Nash equilibrium $(\mathbf{q}^1, \dots, \mathbf{q}^{m_2})$ of G_{m_2} , an ϵ_{m_1} -well-supported Nash equilibrium $(\mathbf{p}^1, \dots, \mathbf{p}^{m_1})$ of G_{m_1} can be obtained by renormalizing probabilities as follows: $\mathbf{p}^i[j] = (1/z) \mathbf{q}^{g(i)}[h_i(j)]$ (where z is a normalization factor).*

Theorem 4 (Main). *For every $\epsilon_k < 1$, there exists a direct reduction from ϵ_k - k NASH to ϵ_2 -2NASH, where $\epsilon_2 = \text{poly}(\epsilon_k/|G_k|)$. The reduction runs in polynomial time in $|G_k|$ and in $\log(1/\epsilon_k)$.*

Corollary 5. *There is a direct, polynomial time reduction from $(1/\exp(n))$ - k NASH to $(1/\exp(n))$ -2NASH, and from $(1/\text{poly}(n))$ - k NASH to $(1/\text{poly}(n))$ -2NASH.*

Proof of Theorem 4: By combining Theorem 15 (linearizing reduction) with Theorem 18 (reduction from linear to bimatrix games), and plugging in the parameters of Lemma 8 (logarithmic-sized linear multiplication gadget). \square

The proof of Lemma 8 appears in the full version. For simplicity of presentation we prove here the slightly weaker Lemma 7 (polynomial-sized gadget), resulting in a reduction that's polynomial time in $1/\epsilon_k$ instead of $\log(1/\epsilon_k)$.

1.3 Related Work

Bubelis [3] shows a direct reduction from k -player to 3-player games. This reduction relies heavily on the multiplicative nature of 3-player games. Examples of direct reductions involving 2-player games include *symmetrization* [12], and reduction to *imitation games* [20]. We use imitation games in Section 4.

PPAD-completeness. PPAD is the class of total search problems polynomial-time reducible to the abstract path-following problem END OF THE LINE [22]. The known results can be summarized by the two following chains of reductions, each forming an indirect reduction (according to Definition 3 of directness) from k -player games to 2-player games:

- $1/\exp(n)$ - k NASH \leq END OF THE LINE \leq 3D-BROUWER \leq ADDITIVE GRAPHICAL NASH \leq $1/\exp(n)$ -2NASH
- $1/\exp(n)$ - k NASH \leq END OF THE LINE \leq 2D-BROUWER \leq n D-BROUWER \leq $1/\text{poly}(n)$ -2NASH

Reductions in chain 1 are by [17,8], [22,8], [5,8] and [8], respectively, and in chain 2 they are by [17,8], [4], [6] and [6], respectively. For an overview of these celebrated results see [23]. In comparison, our reduction can be written as: ϵ_k - k NASH \leq ϵ_m -LINEAR-NASH \leq ϵ_2 -2NASH, where $\epsilon_k, \epsilon_m, \epsilon_2$ can either all be $1/\exp(n)$ or $1/\text{poly}(n)$. Note there is gap between the second chain of reductions and our results - the second chain achieves a stronger reduction from $1/\exp(n)$ - k NASH to $1/\text{poly}(n)$ -2NASH. Achieving a direct version of this result by [6] is an interesting open problem. Note also that our reduction from ϵ_m -LINEAR-NASH to ϵ_2 -2NASH is somewhat similar to the reduction from ADDITIVE GRAPHICAL NASH to $1/\exp(n)$ -2NASH, however our reduction does not require the input game to be bipartite nor does it limit the number of interactions per player.

Constant Approximations. Another open question is the complexity of ϵ_k - k NASH and ϵ_2 -2NASH for *constant* values of ϵ_k, ϵ_2 . As a quasi-polynomial algorithm is known [11,9], these problems are not believed to be PPAD-complete. The current state-of-the-art is a polynomial-time algorithm for ϵ_2 -2NASH where $\epsilon_2 \approx 0.667$ [16]. For finding ϵ_2 -Nash equilibrium rather than ϵ_2 -well-supported

Nash equilibrium, there is an algorithm where $\epsilon_2 \approx 0.339$ [24] (see also [9,2,25]). On the negative side, several algorithmic techniques have been ruled out [15,10].

Reductions to 2 players and linearization. The empirical success of the Lemke-Howson algorithm [18] for finding Nash equilibrium in 2-player games has motivated research on extending it to a more general class of games. Daskalakis et al. show an indirect reduction from *succinct* games to 2-player games [7]. Govindan and Wilson present a non-polynomial linearizing reduction, which reduces multiplayer games to polymatrix games while preserving approximate Nash equilibrium [14]. Linearization is also related to the formulation of PPAD as the class of fixed-point problems for piecewise-linear functions [11].

2 A Linear Multiplication Gadget

Theorem 6 (Linear Multiplication Gadget). *There exist constants $\epsilon_0 < 1, c, d$ and an increasing polynomial function f such that the following holds. For every $\epsilon < \epsilon_0$, there exists a linear multiplication gadget $G_* = G_*(\epsilon)$ of size $O(m \cdot f(\frac{1}{\epsilon}))$, such that in an ϵ -well-supported Nash equilibrium, the output of G_* equals the product of its m inputs up to an additive error of $\pm dm\epsilon^c$.*

Lemma 7 (Polynomial-Sized Construction). *Theorem 6 holds with the following parameters:¹ $\epsilon_0 = \frac{1}{4}$, $c = 1$, $d = 19$ and $f(x) = x^2$.*

Lemma 8 (Logarithmic-Sized Construction). *Theorem 6 holds with the following parameters: $\epsilon_0 = \frac{1}{10^5}$, $c = \frac{1}{2}$, $d = 3$ and $f(x) = \log x$.*

Both constructions use standard gadgets as building blocks. The second construction gives a smaller gadget with size $O(m \log \frac{1}{\epsilon})$ instead of $O(\frac{m}{\epsilon})$, but is also more complicated. Its details appear in the full version. The rest of this section describes the first construction and proves Lemma 7.

2.1 Linear Gadgets

Goldberg and Papadimitriou developed the framework of gadgets [13], carefully-engineered games that simulate arithmetic calculations and are useful in many PPAD-completeness results. Gadget players are typically *binary*.

Definition 9 (Binary Player). *A binary player P is a player that has exactly two pure strategies 0 and 1. We say P represents the numerical value $p \in [0, 1]$ if her mixed strategy is p , i.e. she plays pure strategy 1 with probability p .*

Gadget games have three kinds of binary players - one or more input players,² one output player, and one or more auxiliary players. The *size* of a gadget is the number of its auxiliary and output players. The values represented by the

¹ The choice of $d = 19$ simplifies the proof, but can be replaced with a smaller value.


² Gadgets can also have non-binary input players, in which case the input values are the probabilities with which they play certain predetermined pure strategies.

input and output players are the *inputs* and *output* of the gadget. In every ϵ -well-supported Nash equilibrium of the gadget game, the output is equal to the result of an arithmetic operation on the inputs (up to small error). This arithmetic relation between the inputs and output is the *guarantee* of the gadget, and is achieved by choosing appropriate payoffs for the auxiliary and output players. Our reductions require gadgets with *linear* guarantees, which differ slightly from the *graphical* and *additive-graphical* gadgets used in previous works.

Definition 10 (Linear Gadgets). *A linear gadget is a polymatrix gadget game with payoffs in $[0, 1]$. Linear gadgets simulate linear arithmetic operations, i.e. their guarantee is a linear relation between the input and output values.*

Several gadgets can be combined into a single game, much like arithmetic gates are combined into a circuit to carry out involved calculations. We represent a combination of gadgets by a series of calculations on the inputs and outputs: For example, if the output player P of gadget G is set to be an input player of gadget G' whose output player is P' , then we write $p' = G'(p)$ (where in turn $p = G(\dots)$). Auxiliary players are never shared among gadgets. The following fact explains why the same player can be an input player of multiple gadgets, but can only be the output player of a single gadget (which determines her payoffs).

Fact 11 (Combining Gadgets). *For every game in which no player is the output player of more than one gadget, the guarantees of all gadgets hold simultaneously when the game is in ϵ -well-supported Nash equilibrium.*

Lemma 12 (Standard Linear Gadgets [8]). 

- For every rational $\zeta \in [0, 1]$, there exists a linear threshold gadget $G_{>\zeta}$ of size $O(1)$ with input p_1 , such that in an ϵ -well-supported Nash equilibrium the output is 1 if $p_1 > \zeta + \epsilon$ and 0 if $p_1 < \zeta - \epsilon$ (and otherwise is in $[0, 1]$).
- There exists a linear AND gadget G_{\wedge} of size $O(1)$ with inputs p_1, p_2 , such that in an ϵ -well-supported Nash equilibrium where $\epsilon < \frac{1}{4}$ the output is 1 if $p_1 = p_2 = 1$ and 0 if $(p_1 = 0) \vee (p_2 = 0)$ (and otherwise is in $[0, 1]$).
- For every rational $\zeta \in [0, 1]$, there exists a linear scaled-summation gadget $G_{+,*\zeta}$ of size $O(1)$ with inputs p_1, \dots, p_m , such that in an ϵ -well-supported Nash equilibrium the output is $\min\{\zeta(p_1 + \dots + p_m), 1\} \pm \epsilon$.

2.2 Construction and Correctness

We describe the construction of G_* that we shall use to prove Lemma 7. We show a construction for multiplying 2 inputs, and multiplying m inputs can be achieved by connecting $m - 1$ copies of G_* serially. Let P_1, P_2 be the input players representing values p_1, p_2 , and let P be the output player representing value p . Let $\tau = 3\epsilon$, and for simplicity assume that $1/\tau$ is integral. The construction first finds an encoding of every input in unary representation, up to precision of $\pm\tau$. This requires two sets $\{V_i^1\}$ and $\{V_i^2\}$ of $1/\tau$ auxiliary players each. The vectors $\mathbf{v}^1 = (v_1^1, \dots, v_{1/\tau}^1)$ and $\mathbf{v}^2 = (v_1^2, \dots, v_{1/\tau}^2)$ of values represented by $\{V_i^1\}$

³ There also exist standard, inherently nonlinear gadgets for multiplication.

and $\{V_i^2\}$ are the unary encodings of p_1 and p_2 respectively. The value of v_i^1 , the i 'th unary bit of p_1 , is set by applying the threshold gadget $G_{>\zeta}$ (Lemma 12) as follows: $v_i^1 = G_{>i\tau}(p_1)$. Similarly, $v_i^2 = G_{>i\tau}(p_2)$. The next component of G_* 's construction is performing unary multiplication among the two vectors $\mathbf{v}^1, \mathbf{v}^2$ using the AND gadget G_\wedge (Lemma 12). Another set of $1/\tau^2$ auxiliary players $\{U_{i,j}\}$ stores the result. Let U be a matrix of the values they represent, then $u_{i,j} = G_\wedge(v_i^1, v_j^2)$. The construction is complete by summing up and scaling U 's entries using the scaled-summation gadget $G_{+,*\zeta}$ (Lemma 12) as follows: $p = G_{+,*\tau^2}(u_{1,1}, u_{1,2}, \dots, u_{1/\tau,1/\tau})$. Note that the payoffs of all players are determined by the standard gadgets. We now show that the described construction establishes the guaranteed relation between inputs p_1, p_2 and output p of G_* .

Proof of Lemma 7: First we observe that G_* is a combination of linear gadgets and is thus itself linear. The size of G_* is $O(1/\tau^2)$, since this is the total size of the standard gadgets it combines ($2/\tau$ threshold gadgets $G_{>\zeta}$, $1/\tau^2$ AND gadgets G_\wedge , and 1 scaled-summation gadget $G_{+,*\zeta}$, all of size $O(1)$). Now assume G_* is in ϵ -well-supported Nash equilibrium where $\epsilon < 1/4$. We write the input values p_1, p_2 as integer multiples of τ plus a small error: Let $p_1 = i^*\tau + \delta_1$ and $p_2 = j^*\tau + \delta_2$, where $0 \leq i^*, j^* \leq 1/\tau$ and $0 \leq \delta_1, \delta_2 < \tau$. The following claim follows directly from the guarantee of the threshold gadget (Lemma 12). It states that while this gadget is *brittle* in the sense that for a small range of inputs it returns an arbitrary output, this cannot be the case for more than one unary bit of p_1 or p_2 . The proof of the claim utilizes the choice of $\tau = 3\epsilon$ (see full version). The rest of the proof of Lemma 7 is a straightforward corollary of the other gadget guarantees.

Claim 13 (Unary Encoding). \mathbf{v}^1 is of the form $(1, \dots, 1, ?, 0, \dots, 0)$, where $\|\mathbf{v}^1\| = i^* \pm 1$ and '?' denotes any value in $[0, 1]$. The same holds for \mathbf{v}^2 and j^* .

Example 14. Let $p_1 = 7\tau + \epsilon/4$ and $p_2 = 2\tau + (\tau - \epsilon/8)$. First G_* finds their unary encodings: $\mathbf{v}^1 = (1, 1, 1, 1, 1, 1, ?, 0, \dots, 0)$ and $\mathbf{v}^2 = (1, 1, ?, 0, \dots, 0)$. Then it performs unary multiplication and finds U (see below). Summing up and scaling U 's entries gives the output $p = 12\tau^2 + O(\epsilon)$, which is close to $p_1 p_2$ up to $O(\epsilon)$. □

$$U' = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & ? & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & ? & 0 \\ ? & ? & ? & ? & ? & ? & ? & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, U = \begin{pmatrix} U' & 0 \\ 0 & 0 \end{pmatrix}_{1/\tau \times 1/\tau}$$

3 Linearizing Multiplayer Games

In this section we show a direct reduction from k -player games to polymatrix games. Let G_k denote the input game to the reduction, and let G_m denote the corresponding output game. The first k players of G_m have the same pure strategies as the players of G_k . The reduction relies on the fact that, although G_k 's expected payoffs are nonlinear in its players' probabilities, they are linear in products of its players' probabilities. A key component of our reduction is a linear

multiplication gadget for computing these products, which exists according to Theorem 6. Let f be an increasing polynomial function as in Theorem 6.

Theorem 15 (A Linearizing Reduction). *For every $\epsilon_k < 1$, there exists a direct reduction from ϵ_k - k NASH to ϵ_m -LINEAR-NASH, where $\epsilon_m = \text{poly}(\epsilon_k/|G_k|)$. The reduction runs in polynomial time in $|G_k|$ and in $f(1/\epsilon_k)$.*

Lemma 16 (Recovering ϵ_k -Well-Supported Nash). *Let $(\mathbf{p}^1, \dots, \mathbf{p}^m)$ be an ϵ_m -well-supported Nash equilibrium of G_m . Then the first k mixed strategies $\mathbf{p}^1, \dots, \mathbf{p}^k$ form an ϵ_k -well-supported Nash equilibrium of G_k .*

The following lemma will be useful in designing the linearizing reduction.

Lemma 17 (Preserving Expected Payoffs). *If for every player $i \in [k]$, the expected payoff vectors $\mathbf{u}_{G_m}^i$ and $\mathbf{u}_{G_k}^i$ are entry-wise equal up to an additive factor of δ , and $(\mathbf{p}^1, \dots, \mathbf{p}^m)$ is an ϵ_m -well-supported Nash equilibrium of G_m , then $(\mathbf{p}^1, \dots, \mathbf{p}^k)$ is an ϵ_k -well-supported Nash of G_k where $\epsilon_k = 2\delta + \epsilon_m$.*

Proof. Let $j \in [n]$ be a pure strategy in the support of player i ($p_j^i > 0$). We know that j is an ϵ_m -best response in G_m . Assume for contradiction that j is not an ϵ_k -best response in G_k , i.e. there is a pure strategy $j' \in [n], j' \neq j$ such that $\mathbf{u}_{G_k}^i(j') > \mathbf{u}_{G_k}^i(j) + \epsilon_k$. So $\mathbf{u}_{G_m}^i(j') + \delta > \mathbf{u}_{G_m}^i(j) - \delta + \epsilon_k$. Since $\epsilon_k - 2\delta = \epsilon_m$, then $\mathbf{u}_{G_m}^i(j') > \mathbf{u}_{G_m}^i(j) + \epsilon_m$, contradiction. \square

3.1 The Linearizing Reduction and Correctness

Given an input pair (G_k, ϵ_k) , we find an output pair (G_m, ϵ_m) as follows. Let $\epsilon_0 < 1, c, d$ be the constant parameters of Theorem 6. Then $\epsilon_m = \min\{(\epsilon_k/3n^{k-1}dk)^{1/c}, \epsilon_0\}$. The players of G_m are:

- *Original players* - the first k players of G_m have the same pure strategies as G_k 's players. \mathbf{p}^i denotes the mixed strategy of original player i .
- *Mediator players* - for every $i \in [k]$, there is a set of n^{k-1} binary players that corresponds to the set of n^{k-1} pure strategy profiles of all original players except i . We denote by $Q_{\mathbf{s}^{-i}}$ the mediator player corresponding to pure strategy profile \mathbf{s}^{-i} and by $q_{\mathbf{s}^{-i}}$ the represented value.
- *Gadget players* - all auxiliary players belonging to kn^{k-1} copies of the linear multiplication gadget G_* .

Every mediator player is set to be the output player of a gadget G_* as follows: $q_{\mathbf{s}^{-i}} = G_*(p_{\mathbf{s}^{-i}[1]}^1, \dots, p_{\mathbf{s}^{-i}[k]}^k)$. Thus, $q_{\mathbf{s}^{-i}}$ will be approximately equal to the probability with which the original players play the pure strategy profile \mathbf{s}^{-i} . Let \mathbf{q}^i be the vector of values $\{q_{\mathbf{s}^{-i}}\}$, then it's approximately equal to $\tilde{\mathbf{p}}^{-i}$, the joint mixed strategy distribution of all original players except i .

To complete the description of G_m it remains to specify the non-zero payoff matrices of the original players (all other payoffs are determined by the gadgets). In G_k , the expected payoff vector of player i is $\mathbf{u}_{G_k}^i = M_{G_k}^i \tilde{\mathbf{p}}^{-i}$. In G_m , the payoff of original player i will be influenced only by the i 'th set of mediator players $\{Q_{\mathbf{s}^{-i}}\}$ who play \mathbf{q}^i . Instead of describing every payoff matrix $M^{i, Q_{\mathbf{s}^{-i}}}$

separately, we describe one large payoff matrix $M_{G_m}^i$ that contains all the others (or more precisely, all their nonzero columns) as submatrices. We want the expected payoffs in G_m to be as close as possible to those of G_k . Thus, we set $M_{G_m}^i = M_{G_k}^i$. This concludes the construction.

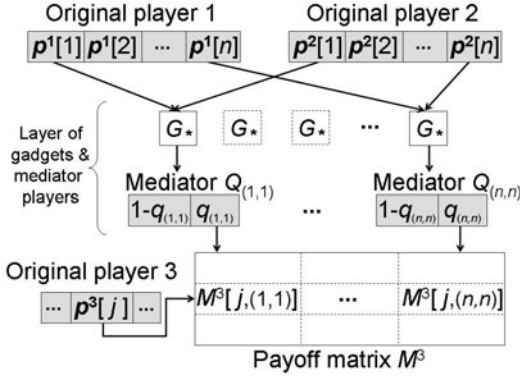


Fig. 1. Linearization of a 3-Player Game - Partial View of G_m

The arrows indicate how the probabilities of original players 1 and 2 influence the expected payoff of original player 3 via a layer of gadgets and mediator players.

Correctness. First note that the reduction runs in time polynomial in $|G_k| = \Theta(kn^k)$ and in $f(1/\epsilon_k)$: The running time depends on the size of the polymatrix game G_m , which is polynomial in the number of its players. There are k original players, kn^{k-1} mediator players and $kn^{k-1}O(|G_*|)$ auxiliary players. By Theorem 6, $|G_*| = O(k \cdot f(1/\epsilon_m))$. Since f is a polynomial function and $\epsilon_m = \text{poly}(\epsilon_k/|G_k|)$, the total number of players is indeed polynomial in $|G_k|$ and in $f(1/\epsilon_k)$. As described above, the expected payoff vector of original player i in G_m is $u_{G_m}^i = M_{G_m}^i \mathbf{q}^i$. The linear multiplication gadget G_* guarantees that vectors \mathbf{q}^i and $\tilde{\mathbf{p}}^{-i}$ are close to each other, and since all payoffs are in $[0, 1]$, the expected payoffs $u_{G_k}^i = M_{G_k}^i \tilde{\mathbf{p}}^{-i}$ are preserved $u_{G_m}^i = M_{G_m}^i \mathbf{q}^i$. The proof of Lemma 6 is then immediate by preservation of expected payoffs (Lemma 17).

4 Reducing Linear Games to Bimatrix Games

In this section we show how to replace the multiple players of a polymatrix game by two representative "super-players" of a bimatrix game. Let G_m denote the input game to the reduction, and let G_2 denote the corresponding output game.

Theorem 18 (Linear to Bimatrix). *For every $\epsilon_m < 1$, there exists a direct reduction from ϵ_m -LINEAR-NASH to ϵ_2 -2NASH, where $\epsilon_2 = \text{poly}(\epsilon_m/|G_m|)$. The reduction runs in polynomial time in $|G_m|$ and in $\log(1/\epsilon_m)$.*

Lemma 19 (Recovering Approximate Nash). *For every ϵ_2 -well-supported Nash (\mathbf{x}, \mathbf{y}) of G_2 , partitioning \mathbf{y} into subvectors of lengths n_i and normalizing gives an ϵ_m -well-supported Nash equilibrium $(\mathbf{y}^1/\|\mathbf{y}^1\|, \dots, \mathbf{y}^m/\|\mathbf{y}^m\|)$ of G_m .*

4.1 Imitation Games and Block ϵ -Uniform Games

The following definitions and lemmas will be useful in proving Theorem 18. An *imitation game* is a bimatrix game in which player 2's payoff matrix is the identity matrix. The following was proved in 20 for the case of exact Nash.

Lemma 20 (Imitation). *Let (\mathbf{x}, \mathbf{y}) be an ϵ_2 -well-supported Nash equilibrium of an imitation game G_2 where $\epsilon_2 \leq 1/N$. Then $\text{support}(\mathbf{y}) \subseteq \text{support}(\mathbf{x})$.*

A bimatrix game is *block ϵ -uniform* if player 1's payoff matrix A is as follows:

- *Block matrix:* A is composed of m^2 blocks $A^{i,i'}$ of size $n_i \times n_{i'}$ each;
- *Very negative diagonal:* The i 'th diagonal block $A^{i,i}$ is equal to $-\alpha E_{n_i}$, where $\alpha = 8m^2/\epsilon$ and E_{n_i} is the all-ones matrix of size $n_i \times n_i$;
- *[0, 1] entries:* All other entries of A are arbitrary values in $[0, 1]$.

For a similar construction see the generalized Matching Pennies game of 13. Let (\mathbf{x}, \mathbf{y}) be a mixed strategy profile of an ϵ -block-uniform game. We denote by $\mathbf{x}^1, \dots, \mathbf{x}^m$ and $\mathbf{y}^1, \dots, \mathbf{y}^m$ its separation to blocks of size n_1, \dots, n_m . We say block i belongs to the support of \mathbf{x} if there is some pure strategy in block i that belongs to it. The following lemma shows that in a block ϵ -uniform game, the weight of player 2 is ϵ -uniformly divided among all blocks in $\text{support}(\mathbf{x})$.

Lemma 21 (ϵ -Uniform Weights). *Let \mathbf{x}, \mathbf{y} be an ϵ_2 -well-supported Nash equilibrium of a block ϵ_2 -uniform game G_2 . If block $i \in [m]$ belongs to the support of \mathbf{x} , then for every $i' \in [m]$, $\|\mathbf{y}^i\| \leq \|\mathbf{y}^{i'}\| + (1 + \epsilon_2)/\alpha$.*

Proof. The expected payoff vector $u_{G_2}^1$ of player 1 is $A\mathbf{y}$. By construction of matrix A , the expected payoff vector for playing pure strategies in block i is $\sum_{i' \in [m]} A^{i,i'} \mathbf{y}^{i'}$. The dominant vector in this sum is $A^{i,i} \mathbf{y}^i$, whose entries are all $-\alpha \|\mathbf{y}^i\|$. The entries of every other vector $A^{i,i'} \mathbf{y}^{i'}$ in the sum are in the range $[0, \|\mathbf{y}^{i'}\|]$, and since \mathbf{y} is a distribution vector, the total contribution to the sum is at most $\sum_{i' \in [m]} \|\mathbf{y}^{i'}\| = 1$. Thus, the expected payoff for playing any pure strategy in block i is in the range $[-\alpha \|\mathbf{y}^i\|, -\alpha \|\mathbf{y}^i\| + 1]$. Assume for contradiction that $\|\mathbf{y}^i\| > \|\mathbf{y}^{i'}\| + (1 + \epsilon_2)/\alpha$. Then the expected payoff for playing a pure strategy in block i is at most $-\alpha(\|\mathbf{y}^i\| + (1 + \epsilon_2)/\alpha) + 1$, while the expected payoff for playing in block i' is at least $-\alpha(\|\mathbf{y}^{i'}\|)$. The difference is more than ϵ_2 , contradicting the assumption that i belongs to $\text{support}(\mathbf{x})$. \square

Corollary 22 (Imitation and Block ϵ -Uniform). *Let \mathbf{x}, \mathbf{y} be an ϵ_2 -well-supported Nash equilibrium of a block ϵ_2 -uniform imitation game G_2 , where $\epsilon_2 \leq 1/N$. Then for every two blocks $i, i' \in [m]$, $\|\mathbf{y}^i\| = \|\mathbf{y}^{i'}\| \pm (1 + \epsilon_2)/\alpha$.*

Proof. We show that if a game is both imitation and block ϵ -uniform, the weight of player 2 is divided ϵ -uniformly among all blocks in $[m]$. Since \mathbf{y} is a distribution vector, there exists a block $i \in [m]$ such that $\|\mathbf{y}^i\| \geq 1/m$. So i belongs to the support of \mathbf{y} , and by Lemma 20, i also belongs to the support of \mathbf{x} . By Lemma 21, $1/m \leq \|\mathbf{y}^i\| \leq \|\mathbf{y}^{i'}\| + (1 + \epsilon_2)/\alpha$ for every $i' \in [m]$. Since $(1 + \epsilon_2)/\alpha < 1/m$ we conclude that $0 < \|\mathbf{y}^{i'}\|$ for every i' . Thus by Lemma 20 all blocks are in $\text{support}(\mathbf{x})$ and get almost uniform weight. \square

4.2 The Reduction and Correctness

Given an input pair (G_m, ϵ_m) , we find an output pair (G_2, ϵ_2) , where G_2 has payoffs in the range $[-\alpha, 1]$. To complete the reduction, G_2 can then be normalized by adding α to all payoffs and scaling by $1/(\alpha + 1)$ (ϵ_2 is also scaled). Let $N = \sum_{i=1}^m n_i$ be the total number of pure strategies in G_m , and let $M^{i,i'}$ be the payoff matrix of player i for interacting with player i' . Set $\epsilon_2 = \epsilon_m/N$. The pure strategies of every player in G_2 are the set $[N]$. The payoffs are chosen such that G_2 is both an imitation game and a block ϵ_2 -uniform game:

$$A = \begin{pmatrix} -\alpha E_{n_1} & M^{1,2} & \dots & M^{1,m} \\ M^{2,1} & -\alpha E_{n_2} & & M^{2,m} \\ \vdots & & \ddots & \vdots \\ M^{m,1} & M^{m,2} & \dots & -\alpha E_{n_m} \end{pmatrix}_{N \times N}, B = \begin{pmatrix} I_{n_1} & 0 & \dots & 0 \\ 0 & I_{n_2} & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \dots & I_{n_m} \end{pmatrix}_{N \times N}$$

Correctness. The reduction runs in time polynomial in $|G_m| = \Theta(N^2)$ and in $\log(1/\epsilon_m)$: The running time depends on the size of the bimatrix game G_2 , whose payoff matrices are of size N^2 with entries of size $O(\log \alpha)$. It's left to prove Lemma 19 for the unnormalized game G_2 and $\epsilon_2 = \epsilon_m/N$; this immediately gives a proof for $\epsilon_2 = \epsilon_m/N(\alpha + 1)$ after normalizing the payoffs from $[-\alpha, 1]$ to $[0, 1]$.⁴ Since $\epsilon_m/N(\alpha + 1) = \text{poly}(\epsilon_m/N)$, Theorem 18 immediately follows. To prove Lemma 19, we define for every player i of G_m a mapping h_i , which maps the j 'th pure strategy of i to the j 'th pure strategy in block i of G_2 . When strategy profiles (\mathbf{x}, \mathbf{y}) and $(\mathbf{y}^1/\|\mathbf{y}^1\|, \dots, \mathbf{y}^m/\|\mathbf{y}^m\|)$ are played in G_2 and G_m respectively, then player 1's expected payoff for playing $h_i(j)$ in G_2 is closely related to player i 's expected payoff for playing j in G_m . In fact, the expected payoffs are the same up to shifting by $\alpha\|\mathbf{y}^i\|$ (the contribution from the diagonal of A), scaling by m (the number of blocks on which \mathbf{y} is uniformly distributed), and small additive errors (proof appears in the full version). As in Section 3, the fact that the expected payoffs are preserved, even up to shift and scale, is enough for one game's ϵ -well-supported Nash equilibrium to imply the other's. Thus, if (\mathbf{x}, \mathbf{y}) is an ϵ_2 -well-supported Nash equilibrium of G_2 , then $(\mathbf{y}^1/\|\mathbf{y}^1\|, \dots, \mathbf{y}^m/\|\mathbf{y}^m\|)$ is an ϵ_m -well-supported Nash equilibrium of G_m .

References

1. Althofer, I.: On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications* 240, 9–19 (1994)
2. Bosse, H., Byrka, J., Markakis, E.: New algorithms for approximate Nash equilibria in bimatrix games. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 17–29. Springer, Heidelberg (2007)
3. Bubelis, V.: On equilibria in finite games. *International Journal of Game Theory* 8(2), 65–79 (1979)
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)

⁴ Every $\epsilon_m/N(\alpha + 1)$ -well-supported Nash equilibrium of the normalized game is an ϵ_m/N -well-supported Nash equilibrium of the unnormalized game.

5. Chen, X., Deng, X.: Settling the complexity of 2-player Nash-equilibrium. In: FOCS (2006)
6. Chen, X., Deng, X., Teng, S.: Computing Nash equilibria: Approximation and smoothed complexity. In: FOCS (2006)
7. Daskalakis, C., Fabrikant, A., Papadimitriou, C.H.: The game world is flat: The complexity of Nash equilibria in succinct games. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4051, pp. 513–524. Springer, Heidelberg (2006)
8. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. SICOMP 39(1), 195–259 (2009)
9. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: Progress in approximate Nash equilibria. In: EC (2006)
10. Daskalakis, C., Papadimitriou, C.H.: On oblivious PTAS's for Nash equilibrium. In: STOC (2009)
11. Etessami, K., Yannakakis, M.: On the complexity of Nash equilibria and other fixed points. In: FOCS (2007)
12. Gale, D., Kuhn, H.W., Tucker, A.W.: On symmetric games. In: Kuhn, H.W., Tucker, A.W. (eds.) Contributions to the Theory of Games, Princeton, pp. 81–87 (1950)
13. Goldberg, P.W., Papadimitriou, C.H.: Reducibility among equilibrium problems. In: STOC (2006)
14. Govindan, S., Wilson, R.: A decomposition algorithm for n -player games. Economic Theory 42(1), 97–117 (2010)
15. Hazan, E., Krauthgamer, R.: How hard is it to approximate the best Nash equilibrium? In: SODA (2009)
16. Kontogiannis, S.C., Spirakis, P.G.: Well supported approximate equilibria in bimatrix games. Algorithmica 57(4), 653–667 (2010)
17. van der Laan, G., Talman, A.J.J.: On the computation of fixed points in the product space of unit simplices and an application to noncooperative n -person games. Mathematics of Operations Research 7(1), 1–13 (1982)
18. Lemke, C.E., Howson, J.T.: Equilibrium points of bimatrix games. SIAM Journal of Applied Mathematics 12, 413–423 (1964)
19. Lipton, R., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: EC (2003)
20. McLennan, A., Tourky, R.: Imitation games and computation. Games and Economic Behavior (2009)
21. Nash, J.F.: Non-cooperative games. Annals of Math. 54, 289–295 (1951)
22. Papadimitriou, C.H.: On the complexity of the parity argument and other inefficient proofs of existence. Journal of Computer and System Sciences 48(3), 498–532 (1994)
23. Roughgarden, T.: Computing equilibria: A computational complexity perspective. Economic Theory 42(1), 193–236 (2010)
24. Tsaknakis, H., Spirakis, P.G.: An optimization approach for approximate Nash equilibria. In: Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 42–56. Springer, Heidelberg (2007)
25. Tsaknakis, H., Spirakis, P.G.: A graph spectral approach for computing approximate Nash equilibria. In: ECCO (2009)

Responsive Lotteries

Uriel Feige^{1,*} and Moshe Tennenholtz²

¹ Department of Computer Science and Applied Mathematics,
Weizmann Institute, Rehovot, Israel

`uriel.feige@weizmann.ac.il`

² Microsoft R&D Center, Herzelia, Israel, and Faculty of Industrial Engineering and
Management, Technion, Haifa, Israel

`moshet@ie.technion.ac.il`

Abstract. Given a set of alternatives and a single player, we introduce the notion of a *responsive lottery*. These mechanisms receive as input from the player a reported utility function, specifying a value for each one of the alternatives, and use a lottery to produce as output a probability distribution over the alternatives. Thereafter, exactly one alternative wins (is given to the player) with the respective probability. Assuming that the player is not indifferent to which of the alternatives wins, a lottery rule is called *truthful dominant* if reporting his true utility function (up to affine transformations) is the unique report that maximizes the expected payoff for the player. We design truthful dominant responsive lotteries. We also discuss their relations with scoring rules and with VCG mechanisms.

1 Introduction

We consider a setting where there are n alternatives A_1, \dots, A_n and a single player. We assume that the player has a cardinal utility function over the alternatives, in the sense of Von-Neumann and Morgenstern. Namely, the player has a utility vector $U = (u_1, \dots, u_n)$, with utility value u_i associated with the respective alternative A_i , and this utility vector determines the preference of the player over different lotteries. Formally, given two lotteries, one that associates probabilities p_i with the respective alternative A_i , and the other associates probabilities q_i with the respective alternative A_i , the player prefers the former lottery if $\sum p_i u_i > \sum q_i u_i$, the latter lottery if $\sum p_i u_i < \sum q_i u_i$, and is indifferent over the choice of lotteries if $\sum p_i u_i = \sum q_i u_i$. Recall that Von-Neumann and Morgenstern show that if all that the player knows is his preferences over every conceivable pair of lotteries, and that these preferences are consistent in the sense that they satisfy a certain set of axioms (these axioms are natural, though there is a well known debate whether they actually reflect human behavior), then this in fact defines a utility function that is unique up to positive affine transformations (shift by a scalar and multiplication by a positive scalar – game theory

* Work done at Microsoft R&D Center, Herzelia, Israel.

literature often calls these linear transformations). We shall assume throughout that the player is not *indifferent* to the alternatives, namely, that there are at least two alternatives A_i and A_j with $u_i \neq u_j$. As utility functions are defined only up to affine transformations, we shall often represent utility functions in one of two canonical forms: either as *unit range*, meaning that $\min_i u_i = 0$ and $\max_i u_i = 1$, or as *unit sum*, meaning that $\min_i u_i = 0$ and $\sum_i u_i = 1$.

We introduce here a concept that we call a *responsive lottery*.

Definition 1. *Given a set of alternatives A_1, \dots, A_n and a single player, a responsive lottery is a mechanism that operates as follows:*

1. *The player provides a report $X = (x_1, \dots, x_n)$, where $X \in R^n$.*
2. *Using a function f from R^n to R^n , which is called the lottery rule, one computes a probability vector $f(X) = P = (p_1, \dots, p_n)$, with $p_i \geq 0$ and $\sum_i p_i = 1$.*
3. *A lottery is held and alternative A_i wins with probability p_i .*

The lottery is *responsive* in the sense that the corresponding probabilities are not given in advance, but rather determined in response to the report of the player. The notion of an alternative winning the lottery should be aligned with what the utility function of the player refers to. For example, if the alternatives are the choice of seat in a certain flight (say, a window seat, an aisle seat, or a middle seat) and the utility function refers to the value the player associates with sitting in such a seat, then following the lottery the player should be seated in a seat corresponding to the winning alternative.

Given a responsive lottery and a utility vector U for the player, we say that the report X of the player is *honest* if $X = U$. Note that since utility functions are defined only up to affine transformations, we assume here that both X and U are given in the same canonical form (say, unit-sum). We say that the report X of the player is *rational* if X is such that $f(X)U = \sum p_i u_i$ is maximized. Namely, the player chooses a report that maximizes his expected payoff.

Definition 2. *A lottery rule for a responsive lottery is truthful dominant if it has the property that for every utility function of the player, the honest report is rational, and every rational report is honest (or equivalently, given the first condition, the second condition is that the rational report is unique).*

The truthful dominance property can be seen to combine three properties.

1. *Rational invertibility.* For every report X there is at most one utility function U for which X is a rational report.
2. *Rational uniqueness.* For every utility function U , there is a unique rational report X .
3. *Incentive compatibility.* For every utility function U , the report $X = U$ is rational.

Observe that rational invertibility and rational uniqueness are properties of the range of the lottery rule. Given a lottery rule f that is rational invertible, obtaining a truthful dominant lottery out of it involves only appending in front of

it an appropriate permutation mapping π , that maps a report X to the report $Y = \pi(X)$ such that $Z = f(Y)$ maximize ZX . By rational invertibility, this now implies that given a utility function U , the unique rational report is $X = U$. This idea is similar to the revelation principle in mechanism design (see [5], for example). Note on the other hand that given a lottery rule that is incentive compatible, there does not seem to be a straightforward way to turn it into a truthful dominant rule. For example, the lottery rule that assigns $p_i = 1$ for the index i for which x_i is largest is incentive compatible, but there are many rational reports that give out no information beyond which is the preferred alternative for the player. Truthful dominance requires much more – that the rational report reveals the whole utility function.

A responsive lottery with a truthful dominant lottery rule may be viewed as a mechanism for elicitation of the utility function of the player. Recall that the work of Von-Neumann and Morgenstern already implies that utility functions can be inferred by observing preferences over lotteries. However, the procedure implicit in [7] involves a (potentially infinite) sequence of comparisons between pairs of lotteries (or a comparison among infinitely many lotteries, which is not feasible in practice). The mere fact that lottery comparisons are performed more than once is problematic for elicitation of utility functions. If the winning alternative is not actually given to the player after each lottery, the player might not have incentives to report the truth. And if the winning alternative is given to the player after each lottery (assuming that this can be practically done), then the issue of complementarities among the alternatives might distort the original utility function of the player. We circumvent these difficulties by having only a single lottery. The aspect of this lottery that allows the elicitation of the utility function (if the player is rational) is its responsive nature. In a sense, the player is choosing among infinitely many lotteries. The rational invertibility property implies that the choice of the player allows one to infer his utility function (assuming that the player is rational). The incentive compatibility property makes it easy for the rational player to choose one lottery out of the infinite set of lotteries.

We assume infinite precision in the values of the utility function, in the reports and in the probabilities assigned by the lottery rule to the alternatives. Namely, they are real numbers. Employing our lottery rules with finite precision will obviously introduce rounding errors. We ignore this issue in this paper.

1.1 Related Work

This manuscript refers only to utility functions as defined by Von-Neumann and Morgenstern [7]. It may be interesting to extend this work (if possible) to other notions of utility function (for the need for other notions, see for example [5] or [4]), but this is beyond the scope of the current work.

As far as we know, our notion of truthful dominant responsive lotteries is new. However, it is related to some other mechanisms for eliciting information from players.

Strictly proper scoring rules provide a mechanism for eliciting the belief of a player regarding the probabilities of future events. This is done by giving

monetary rewards that depend on the predictions of the player, and on the actual realization of the future events. See for example [6], [9], [11] and [3]. Our notion of truthful dominant lottery rules is related to the notion of strictly proper scoring rules for categorical variables. There is no immediate equivalence between these two concepts, but there are certain algebraic transformations between them. Moreover, there are geometric characterizations of scoring rules in a spirit similar to that of our geometric approach of Section 2. In particular, our spherical lottery rule is based on a high dimensional sphere, and so is the spherical scoring rule (though these are two different spheres).

VCG mechanisms are a method for eliciting the true value that a bidder has for items that are sold in an auction. The incentives are built into the monetary payments that the bidder makes if he wins the item.

Due to lack of space, most details on the relation between scoring rules, VCG mechanisms and responsive lotteries are deferred to the full version of this paper.

1.2 Our Results

We view the introduction of the concept of truthful dominant responsive lotteries as one of the contributions of this work. Our main results are as follows:

1. We present a geometric approach for designing truthful dominant responsive lotteries, and use it to design what we call the *spherical lottery rule*. This lottery rule is continuous – a small change in the reports results in a small change in the probabilities of the alternatives. See Section 2.
2. For three alternatives we present an algebraic approach for designing truthful dominant lottery rules. These rules are continuous. See Section 3.
3. We present methodologies for transforming any truthful dominant lottery rule over three alternatives to a truthful dominant lottery rule over $n > 3$ alternatives. The resulting lottery rule is not continuous. See Section 3.1.

Additional results discussed in the full version of this paper include:

1. We show a transformation from *bounded* proper scoring rules for n events to truthful dominant lottery rules for n alternatives. The resulting lottery rule is not continuous. The transformation does not apply if the scoring rule is unbounded (such as the *logarithmic score*).
2. We show how the VCG mechanism (which involves money and multiple agents) can be used to design truthful dominant lottery rules (that involves only one player and no money). Also here, the resulting lottery rules are not continuous.
3. We show a transformation from truthful dominant lottery rules for $n + 2$ alternatives to proper scoring rules for n events. By way of example, we use this transformation to derive a well known proper scoring rule, the *quadratic score*. We also show a transformation from truthful dominant lottery rules for $n + 1$ alternatives to proper scoring rules for n events. Combining this with item (5) implies a methodology for deriving strictly proper scoring rules from the VCG mechanism.

1.3 Some Remarks

In our lotteries exactly one alternative wins. Our lottery mechanisms do not assume that if no alternative wins then the player gets 0 utility. If we wish to encompass situations in which valid outcomes include the possibility that no alternative wins, or that more than one alternative wins, the lottery mechanism needs to add these possible outcomes as additional alternatives.

The way of providing incentives to the player is by the choice of the winning alternative. There is no transfer of money involved in our mechanisms. Money can be introduced into our mechanisms by specifying alternatives that involve receiving or paying money.

The incentives in a truthful dominant lottery rule only refer to reporting the exact true utility function. In our mechanisms, there will also be some correlation between how close a report is to the true utility function and the expected value of the report. However, we make no formal claims regarding the nature of this correlation, and do not exclude the possibility that among two different reports, the one “further away” from the true utility function (according to some metric to be chosen by the reader) results in higher expected payoffs.

Some of the lottery rules that we design are continuous – a small change in the reports results in a small change in the probabilities of the alternatives. We view continuity as a desirable property for lottery rules, if one wishes them to be used in practice. Discontinuity of the lottery rule might have negative psychological effects on players who are not sure about their utility functions. They might spend too much time deliberating among reports that are almost identical but that lead to very different probability vectors. We note that for all our lottery rules, even the discontinuous ones, the value of their expected payoff is continuous (even though the probability vector might not be continuous), provided that the reports are honest.

An important aspect of a lottery mechanism is its economic efficiency. Namely, we want the winning alternative to be the one that is actually preferred by the player. For more than two alternatives, there are no economically efficient lottery rules that are truthful dominant. However, we remark that there are truthful dominant lottery rules that achieve almost perfect economic efficiency, though we do not advocate using them (see Section 4).

Our lottery rules provide ex-ante incentives to reveal the true utility function. However, this does not exclude the possibility that the player will experience ex-post regret. This issue too will be discussed in Section 4.

1.4 Ordinal Utilities

Though our work is concerned with cardinal utilities, it may be instructive to consider first the case of ordinal utilities. In this case a responsive lottery is truthful dominant if the unique optimal report for the player is to report the alternatives in his order of preference (and specify ties, if there are any).

If $n = 1$, the problem is not interesting. The winning alternative is determined regardless of what the voter reports. If $n = 2$, the voter may report his preferred alternative, and the winning alternative is the reported alternative. If $n \geq 3$,

there is no deterministic lottery with the rational uniqueness property. There are $n!$ different ranking orders (in fact more, if one allows ties), and only n possible winners. For any deterministic mechanism with $n \geq 3$, there are different orders that result in the same winner. Even if with respect to both orders reporting the truth gives the best payoff to the player, there is no incentive to the player in distinguishing between these two orders in the report.

This motivates considering randomized mechanisms (lotteries). Given a report that ranks the alternatives from 1 to n , we may let the j th alternative win with probability $\frac{2(n-j)}{(n-1)(n-2)}$ (or any other probability distribution that decreases with rank). If the player possesses a complete order over the alternatives, then the dominant strategy for the player is to report his true ranking.

Note however what happens if the player views two of the alternatives as being equivalent (a tie). Then asking the player to report a total order (with no ties) forces the player not to be truthful. Hence we should allow the player to report a tie among some alternatives. It is natural in this case to redistribute the probability of winning equally among the tied alternatives. Note however that by now we lost both the rational uniqueness property and the rational invertibility property. In case of a tie in the ranking, reporting a tie is not the unique best strategy. Reporting an arbitrary order among the tied alternatives gives the same expected payoff to the player as reporting a tie. This illustrates part of the challenges in designing truthful dominant mechanisms.

In this work we shall be interested not only in learning ordinal utilities, but cardinal utilities. Hence we wish to learn more than just the ranking, but we also have greater control of the rewards. The player can be incentivized not only by the choice of order among winning alternatives (which alternatives have higher probability of winning than others), but also by the choice of the actual values of these probabilities.

2 A Geometric View of Truthful Dominant Lottery Rules

We present here a geometric view of truthful dominant lottery rules. We use it to design what we call the *spherical lottery rule* (which shares some common principles with the *spherical scoring rule*). More generally, the approach presented here leads to a geometric characterization of a wide class of truthful dominant lottery rules.

We shall use the following notation. The number of alternatives is n , the utility vector of the player is $U = (u_1, \dots, u_n)$, the report vector is $X = (x_1, \dots, x_n)$, and the probability vector is $P = (p_1, \dots, p_n)$, taken from an infinite set \mathcal{P} of *feasible* probability vectors (\mathcal{P} is the range of f for the lottery rule). Observe that all vectors $P \in \mathcal{P}$ are nonnegative and lie on the hyperplane $\sum p_i = 1$. Let $\bar{1} = \frac{1}{\sqrt{n}}(1, \dots, 1)$ denote the unit vector in the direction of the n -dimensional all 1 vector. Given an arbitrary vector Y , it can be decomposed into $Y = \alpha_Y \bar{1} + \beta_Y Y^\perp$, where Y^\perp is a unit vector orthogonal to $\bar{1}$, $\alpha_Y = \langle Y, \bar{1} \rangle$ and $\beta_Y = \langle Y, Y^\perp \rangle$. We assume w.l.o.g. that the sign of Y^\perp is chosen so that $\beta_Y \geq 0$. Observe that for

all $P \in \mathcal{P}$ we have that $\alpha_P = \frac{1}{\sqrt{n}}$. Given a utility function $U = \alpha_U \bar{1} + \beta_U U^\perp$ and a probability vector $P = \frac{1}{\sqrt{n}} \bar{1} + \beta_P P^\perp$ for the responsive lottery, the payoff to the player is $\langle U, P \rangle = \alpha_U / \sqrt{n} + \beta_U \beta_P \langle U^\perp, P^\perp \rangle$. The rational report X for the player is the one for which $P = f(X)$ maximizes $\langle U, P \rangle$, and hence the maximum expected payoff attainable by the player is $\max_{P \in \mathcal{P}} \{ \langle U, P \rangle \}$. Observe that the optimal choice of $P \in \mathcal{P}$ is the one maximizing $\beta_P \langle U^\perp, P^\perp \rangle$ (since β_U is positive). The optimal P is preserved under positive affine transformations to U , because these transformations only change α_U and β_U (without flipping its sign) but not U^\perp .

We now describe a methodology for deriving truthful dominant lottery rules. (Presumably this methodology characterizes all continuous truthful dominant rules. Proving this appears to be an exercise in formalities that does not add interesting insights, and hence will not be pursued here.) A compact convex body K will be called *nice* if (1) for every point z on its boundary ∂K there is a unique hyperplane H such that $H \cap K = z$, and (2) for every two points on ∂K , the line joining them lies entirely within K . For example, balls, ellipsoids and eggs are nice convex bodies, whereas polyhedrons are not. For nice convex bodies, for every vector v , there is a unique value $t(v)$ such that the closed halfspace $H(v) = \{x \mid \langle x, v \rangle \leq t\}$ (whose defining hyperplane is orthogonal to v) contains K and $\partial H \cap \partial K \neq \emptyset$. Moreover, ∂H and ∂K intersect in exactly one point.

Consider the $(n - 1)$ -dimensional subspace of R^n defined by the hyperplane $\sum p_i = 1$. Within its nonnegative orthant (satisfying $p_i \geq 0$ for every i) consider an arbitrary nice convex body K . Let \mathcal{P} (the set of feasible probability vectors for a responsive lottery) be precisely ∂K . Given a report X , consider the halfspace $H(X^\perp)$ as described above, and choose $P = f(X)$ to be the unique point $z \in \partial K$ intersecting $\partial H(X^\perp)$. This maximizes the projection of P on X^\perp , and hence maximizes $\langle P, X \rangle$. For this choice of lottery rule f , given a utility vector U , reporting $X = U$ maximizes the expected payoff. Moreover, for any report $X \neq U$, the probability vector $P = f(X)$ will be one that is strictly inferior to $f(U)$ in terms of the expected payoff.

It is natural to require (though not necessary) that the nice convex body K has geometric symmetries that reflect the intention that a-priori, all alternatives are treated symmetrically. In particular, in this case the center of mass of K will be at $\frac{1}{\sqrt{n}} \bar{1}$. Of all convex bodies, the most symmetric one is the ball, and its boundary is a sphere. Our *spherical lottery rule* uses a sphere centered at $\frac{1}{\sqrt{n}} \bar{1}$. To maximize the the variability in expected payoffs, this sphere has maximum possible radius. This radius is governed by the need to stay in the nonnegative orthant. A closest point P on the boundary of this orthant to the center of the sphere is $(0, 1/(n - 1), \dots, 1/(n - 1))$ for which $P^\perp = (-1/n, 1/n(n - 1), \dots, 1/n(n - 1))$. Hence the radius of the sphere is $1/\sqrt{n(n - 1)}$ (implying among other things that no entry in P is larger than $2/n$). Observe that using the spherical lottery rule, given a report $X = \alpha_X \bar{1} + \beta_X X^\perp$, the probability vector P is derived simply by projecting X^\perp on the sphere (along the line connecting X^\perp to the center of

the sphere). We can assume that $X^\perp \neq 0$, by the assumption that the player is not indifferent. Hence

$$f(X) = \frac{1}{\sqrt{n}}\bar{1} + \frac{1}{\sqrt{n(n-1)}}X^\perp = \left(\frac{1}{n}, \dots, \frac{1}{n}\right) + \frac{1}{\sqrt{n(n-1)}}X^\perp$$

One readily observes that $P (= f(X))$ is an affine transformation of the report X . Hence the spherical lottery rule can be viewed as a normalization of the utility vector U with $\alpha_U = 1/\sqrt{n}$ and $\beta_U = 1/\sqrt{n(n-1)}$, and following this normalization one simply takes $P = U$.

Theorem 1. *The spherical lottery rule described above is truthful dominant.*

The proof of Theorem 1 is implicit in the discussion preceding it. But let us sketch here yet another proof. Observe that for the spherical lottery rule, all vectors in \mathcal{P} have the same norm $1/\sqrt{n-1}$. Hence the inner product $\langle U, P \rangle$ is maximized by the vector in \mathcal{P} that minimizes the angle with U . This vector is precisely the projection of U on the sphere, and $f(X)$ is this projection if and only X is a positive affine transformation of U .

3 Three Alternatives

In this section we design truthful dominant responsive lottery mechanisms for three alternatives. We assume that the utility function of the player is normalized to be unit range $0 = u_1 \leq u_2 \leq u_3 = 1$, and so are his reports.

If all reports are identical, then we set $p_i = 1/3$ for every alternative. If the reports are not identical, let the reports (after normalization) be $0 = x_1 \leq x_2 \leq x_3 = 1$. For simplicity of notation, let $x = x_2$.

Theorem 2. *Any responsive lottery over three alternatives satisfying all the following conditions is truthful dominant.*

1. $p_i \geq 0$ for $i \in \{1, 2, 3\}$.
2. $\sum p_i = 1$.
3. $p_1 \leq p_2 \leq p_3$, with $p_1(x) = p_2(x)$ iff $x = 0$ and $p_2(x) = p_3(x)$ iff $x = 1$.
4. p_2 is strictly increasing in x and p_1 and p_3 are strictly decreasing in x .
5. The derivatives satisfy $x p'_2(x) + p'_3(x) = 0$ for every $0 \leq x \leq 1$.

Proof. Conditions 1 and 2 are satisfied by every responsive lottery. Condition 3 ensures that in the optimal reports, the alternatives are ranked in their true order of preference (satisfy the ordinal aspect of the utility function). Specifically, $u_1 \leq u_2 \leq u_3$, with $u_1 = u_2$ iff $x_1 = x_2$ and $u_2 = u_3$ iff $x_2 = x_3$. Note that we assume that the player is not indifferent, and hence $u_1 < u_3$. After normalization to unit range, we have $0 = u_1 \leq u_2 \leq u_3 = 1$. For simplicity of notation, let $u = u_2$. We need to prove that the optimal report x is $x = u$.

The payoff to the voter is $v = u_1 p_1(x) + u_2 p_2(x) + u_3 p_3(x) = u p_2(x) + p_3(x)$. The derivative of v with respect to x is $u p'_2(x) + p'_3(x)$ which equals 0 if $x = u$, by

Condition 5, and only if $x = u$, by Condition 4 that implies that the derivatives are nonzero. This is the unique extremum for v . Condition 4 implies that this is a maximum rather than a minimum. \square

We remark that replacing condition 5 in Theorem 2 by the weaker condition that $-p'_3(x)/p'_2(x)$ is strictly increasing in x , from 0 to 1, will result in a rational invertible mechanism, though not necessarily incentive compatible.

Theorem 2 allows for many truthful dominant mechanisms. It is natural to limit the possible choices by normalizing the rewards such that there is some report for which $p_1 = 0$. It is easy to see that in conjunction with Theorem 2 this amounts to postulating that when $x = 1$ we have $p_1 = 0$ and $p_2 = p_3 = 1/2$. By way of example, we present two mechanisms that satisfy Theorem 2 and this additional requirement. They are named after the largest degree in the polynomials that are involved.

- The 3-alternative quadratic lottery rule: $p_1 = \frac{1-2x+x^2}{6}$, $p_2 = \frac{1+2x}{6}$, $p_3 = \frac{4-x^2}{6}$.
- The 3-alternative cubic lottery rule: $p_1 = \frac{1-3x^2+2x^3}{8}$, $p_2 = \frac{1+3x^2}{8}$, $p_3 = \frac{6-2x^3}{8}$.

3.1 Extension to More Than Three Alternatives

Here we present two approaches for extending truthful dominant responsive lotteries over three alternatives to truthful dominant responsive lotteries with $n > 3$ alternatives.

The first of these approaches is as follows. Given the report of the player, pick uniformly at random three alternatives. If all three have the same reported utility, let each one of them win with probability $1/3$. Otherwise, normalize the part of the report of the player that refers to these three alternatives so that it becomes unit sum, and apply a 3-alternative truthful dominant responsive lottery on these three alternatives. It is not hard to see that this gives a truthful dominant responsive lottery.

In the full version of this paper we present another approach in more detail. When there are $n + 2$ alternatives A_0, \dots, A_{n+1} and the reports are $0 = x_0 \leq x_1 \leq \dots \leq x_n \leq x_{n+1} = 1$, it will give the $((n + 2)$ -alternative) quadratic lottery rule for which the probabilities are derived from the following expressions after dividing by $n^2 + 3n + 2$:

$$\begin{aligned}
 p_0 &= n - \sum_{i=1}^n 2x_i + \sum_{i=1}^n (x_i)^2; \\
 p_i &= n + 2x_i \text{ for } 1 \leq i \leq n; \\
 p_{n+1} &= 2n + 2 - \sum_{i=1}^n (x_i)^2.
 \end{aligned}$$

4 Convex Combinations

In general, given one truthful dominant mechanism, one can generate others by the method of taking *convex combinations*. We say that a mechanism is a convex combination of two mechanisms M_1 and M_2 if there is some probability

$0 < q < 1$ such that with probability q the mechanism employs M_1 and with probability $1 - q$ it employs M_2 . Equivalently, given the reports, the probability of a given alternative to win is the convex combination (with weights q and $1 - q$) of the respective probabilities in M_1 and M_2 . The following proposition is self evident.

Proposition 1. *A convex combination of a truthful dominant responsive lottery with an incentive compatible responsive lottery is truthful dominant.*

As an example for the use of Proposition 1, the following responsive lottery is truthful dominant. Choose the alternative with highest reported value as a winner with probability q , and with the remaining probability employ the spherical lottery rule. Observe that as q approaches 1, this voting rule converges to optimal economic efficiency, but at the cost of weakening the incentives for the player to distinguish in his report between the less desirable alternatives.

As another example, consider the issue of ex-post regret involved in responsive lotteries. Even though reporting the true utility function is optimal for the player ex-ante, the player may suffer ex-post regret after the lottery is held. For the 3-player quadratic lottery rule, given any report other than $(0, 1, 1)$ the least desirable alternative might win the lottery, and then the player may regret not having reported $(0, 1, 1)$ which would have avoided this possibility. To prevent this ex-post regret, one may take a convex combination of the quadratic rule with the uniform rule (each alternative equally likely to win), which ensures that regardless of the report of the player, the least desirable alternative has some probability of winning. This decrease in ex-post regret comes at the cost of economic efficiency (among other things).

The availability of several (infinitely many) mechanisms that are truthful dominant allows one to introduce some additional objective function, and select the mechanism that optimizes this additional property. For example, among all truthful dominant mechanisms one may want to select the mechanism minimizing the maximum probability with which the least desirable alternative wins. (This probability is $1/6$ for the 3-player quadratic lottery rule.) However, for this particular objective function, there is no truthful dominant mechanism that minimizes it. For every truthful dominant mechanism, the value of this objective function is strictly positive. But then it can be lowered by taking a convex combination with an incentive compatible mechanism. This relates to the fact that the notion of truthful dominant mechanisms is defined using strict inequalities, and its closure is the incentive compatible mechanisms.

5 Applications

The notion of truthful dominant responsive lotteries is a mathematical construct that may have practical applications. We believe that in choosing a truthful dominant lottery rule for a practical application, one would need to strike a careful balance among several considerations (such as ex-post regret, economic efficiency and the strength of the incentives, see Section 4), and this will be possible only if

the number of alternatives is fairly small (four alternatives appears to be a good number). Also, we believe that a carefully designed user interface may help the players understand the concept of responsive lotteries and the effect of the choice of report on the expected reward. For example, one may imagine an interface which includes sliding bar controls for each alternative, and a screen showing a pie chart for the relative probability of each alternative winning. As the player moves the sliding bars to indicate to which extent he values each alternative, the pie chart changes dynamically. Having such a user interface in mind is one of the reasons why we wish lottery rules to be continuous (and among our lottery rules for more than three alternative, only the spherical one is continuous).

Though the purpose of this manuscript is mainly to develop the mathematical theory of truthful dominant responsive lotteries, we briefly and informally discuss some potential applications. A formal treatment of these and other potential applications will hopefully be undertaken elsewhere. In all cases below we assume that the winning alternative can actually be given to the player after the lottery is held (or at least, that the player believes that this is what will happen).

Experimental Psychology. If one wishes to gain a quantitative understanding of preferences of people over a set of alternatives, one may in principle use truthful dominant responsive lotteries. For example, a psychophysical experiment may study the relative sensation of pleasure or pain associated with various temperatures, and the alternatives may be those of putting one's hand in containers of water of various (possibly unpleasant but not harmful) temperatures. One may not want to repeat such an experiment many times with the same subject, due to effects of adaptation, and responsive lotteries may serve as a way of eliciting more information in fewer experiments. As always in experimental settings, caution is needed in performing experiments and in interpreting the results (which at best indicate what were the true preferences of the subject under the conditions of the experiment).

Market Research. A company may use responsive lotteries to gain understanding of the preferences of its potential costumers. For example, an airline company may offer some passengers on a flight a responsive lottery over the choice of seat (say, a window seat, an aisle seat, or a middle seat) so as to get a sense of what the true preferences of costumers are.

Multiple-agent Mechanism Design. In many settings one is interested in designing a mechanism in which agents report their utilities, and then some global decision is taken so as to optimize some objective function that depends on the true utilities of the agents. The difficulty is often in incentivizing the agents to reveal their true utilities. Mechanisms based on statistical approaches often take a small sample of agents, ask them for their utility function, and use this output so as to reach a global decision that effects the agents not in the sample. In such a setting (and assuming no externalities), the agents in the sample have no incentive to not tell the truth, and hence are sometimes assumed to be truthful. See for example [2] for the use of a statistical approach in the design of combinatorial auctions. If this statistical approach is combined with truthful

dominant responsive lotteries then there is more justification in assuming that the agents in the sample are truthful. Let us provide a hypothetical example.

Suppose a company wants to reward each one of its employees with a \$100 gift certificate to some store chain. There are two store chains that are being considered (say one specializes in electronics, one in sports). However, (almost) all gift certificates should be to the same store chain, as then the company gets a big discount from the store chain. How can the company decide among the store chains? One option is to sample at random a small number of employees and offer each one of them a truthful dominant responsive lottery over four alternatives, where two of them are the gift certificates to the two chains, and the other two alternatives are \$50 and \$100 in cash (for calibration). Each employee in the sample actually does get the alternative that wins the respective lottery. All remaining employees get gift certificates to just one store chain, and this store chain is determined based on the information elicited by the responsive lotteries (and on the objective of the company, which may be for example to maximize welfare). Arguably, employees in the sample will actually reveal their true preferences. If the total number of employees is large, then with high probability this mechanism leads to almost optimal economic efficiency: the sample size may be chosen to be large enough to be representative, yet small enough to make the marginal inefficiencies small (inefficiencies resulting from giving sampled employees more expensive rewards, and from occasionally giving sampled employees less favorable rewards, due to the random nature of responsive lotteries).

Acknowledgements

We thank Ran Smorodinsky and Aviv Zohar for many helpful discussions.

References

1. Allen, F.: Discovering personal probabilities when utility functions are unknown. *Management Science* 33(4), 542–544 (1989)
2. Dobzinski, S., Nisan, N., Schapira, M.: Truthful randomized mechanisms for combinatorial auctions. In: *STOC 2006*, pp. 644–652 (2006)
3. Gneiting, T., Raftery, A.: Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102, 359–378 (2007)
4. Kahneman, D., Tversky, A.: Prospect Theory: An Analysis of Decision under Risk. *Econometrica* XLVII, 263–291 (1979)
5. Mas-Colell, A., Whinston, M., Green, J.: *Microeconomic Theory*. Oxford University Press, Oxford (1995)
6. McCarthy, J.: Measures of the value of information. *Proceedings of the National Academy of Sciences* 42, 654–655 (1956)
7. von Neumann, J., Morgenstern, O.: *Theory of games and economic behavior*, 3rd edn. Princeton University Press, Princeton (1954)
8. Nisan, N., Roughgarden, T.: Eva Tardos and Vijay Vazirani. *Algorithmic Game Theory*, Cambridge (2007)
9. Savage, L.J.: Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66, 783–801 (1971)

On the Existence of Optimal Taxes for Network Congestion Games with Heterogeneous Users^{*}

Dimitris Fotakis¹, George Karakostas², and Stavros G. Kolliopoulos³

¹ School of Electrical and Computer Engineering,
National Technical University of Athens, 157 80 Athens, Greece
`fotakis@cs.ntua.gr`

² Department of Computing and Software, and
School of Computational Engineering and Science,
McMaster University, Hamilton, ON, Canada.
`karakos@mcmaster.ca`

³ Department of Informatics and Telecommunications,
National and Kapodistrian University of Athens, Athens 157 84, Greece
`www.di.uoa.gr/~sgk`

Abstract. We consider network congestion games in which a finite number of non-cooperative users select paths. The aim is to mitigate the inefficiency caused by the selfish users by introducing taxes on the network edges. A tax vector is *strongly (weakly)-optimal* if all (at least one of) the equilibria in the resulting game minimize(s) the total latency. The issue of designing optimal tax vectors for selfish routing games has been studied extensively in the literature. We study for the first time taxation for networks with atomic users which have unsplittable traffic demands and are *heterogeneous*, i.e., have different sensitivities to taxes. On the positive side, we show the existence of weakly-optimal taxes for single-source network games. On the negative side, we show that the cases of homogeneous and heterogeneous users differ sharply as far as the existence of strongly-optimal taxes is concerned: there are parallel-link games with linear latencies and heterogeneous users that do not admit strongly-optimal taxes.

1 Introduction

We consider atomic network congestion games with unsplittable traffic demands, where a finite number of non-cooperative users select each a path from a specified source to a sink in an underlying network. The users experience a load-dependent latency on their chosen paths. Being selfish, they want to choose a minimum-latency path. The solution concept we study is that of a *pure Nash equilibrium*, where no user has an incentive to unilaterally switch to a different path. It is well-known that this type of game always has at least one pure Nash equilibrium [13].

^{*} Research partially supported by an NTUA Basic Research Grant (PEBE 2009) and by an NSERC Discovery grant.

The users induce a *social cost* to the system, which in this work we define as the total latency. Selfish behavior leads typically to suboptimal social cost at equilibrium. A long series of papers has studied the inefficiency of Nash equilibria for congestion games as quantified by the price of anarchy. See the surveys [10,8] for an introduction to the very rich literature on the topic.

In order to offset the inefficiency of uncoordinated users, a common approach is to introduce fixed *taxes* (or *tolls*) on the edges of the network. The users will experience the taxes as part of their individual disutility, in addition to their latency. The aim is to design an *optimal* tax vector steering the selfish users to an equilibrium with desirable characteristics; in our case the desired target is minimum total latency.

Related Work. In the non-atomic setting, where there is an infinite number of users and each user controls an infinitesimal amount of traffic demand, the problem of designing optimal tax vectors has been studied extensively. A classic result going all the way back to Pigou [12] states that marginal cost taxes induce the optimal traffic pattern for homogeneous users [2]. A significant volume of recent work on optimal taxes for non-atomic congestion games considers the more intriguing and realistic case of *heterogeneous* users, which may have different valuations of time (latency) in terms of money (taxes). Yang and Huang [17] established the existence of optimal taxes for non-atomic asymmetric network congestion games¹ with heterogeneous users. Subsequently, their result was re-discovered by Fleischer, Jain, and Mahdian [5], and Karakostas and Kolliopoulos [9]. Previously the the single-source special case had been investigated by Cole, Dodis, and Roughgarden [4]. The existence of optimal taxes for non-atomic congestion games with heterogeneous users follows from Linear Programming duality, and thus an optimal tax vector can be computed efficiently by solving a linear program.

For non-atomic games, under mild assumptions on the latency functions the edge flow at equilibrium is unique. Hence the taxes of [2,4,5,9,17] induce the optimal solution as the unique edge flow of the equilibria of the game with taxes. On the other hand, atomic congestion games, even with splittable traffic, may admit many different Nash equilibria, possibly with different edge flows. Therefore, when considering atomic games, one has to distinguish between *weakly-optimal* tax vectors, for which at least one Nash equilibrium of the game with taxes minimizes the total latency, and *strongly-optimal* tax vectors, for which all Nash equilibria of the game with taxes minimize the total latency.

For atomic congestion games with splittable traffic and heterogeneous players, Swamy [14] proved that weakly-optimal tax vectors exist and can be computed efficiently by solving a convex program. As for atomic congestion games with unsplittable traffic, the existence and efficient computation of optimal taxes has been studied only in the restricted setting of homogeneous users. Caragiannis, Kaklamanis, and Kanellopoulos [3] considered atomic games with linear latency functions and homogeneous users, and investigated how much taxes can improve

¹ A network congestion game is symmetric if all users share the same source and sink and, in the case of atomic games, have the same traffic demand.

the price of anarchy. On the negative side, they established that if the users either do not share the same source and sink or have different traffic demands, then strongly-optimal taxes may not exist. In particular, Caragiannis et al. presented a non-symmetric game for which any tax vector induces a Nash equilibrium of total latency at least $6/5$ times the optimum, and a parallel-link game with user-specific traffic demands for which any tax vector induces an equilibrium of total latency at least $9/8$ times the optimum. On the positive side, they presented an efficient construction of strongly-optimal taxes for parallel-link games with linear latencies and unit-demand users. Subsequently, Fotakis and Spirakis [7] proved that weakly-optimal taxes exist and can be computed efficiently for atomic symmetric network congestion games, and that such taxes are strongly-optimal if the network is series-parallel.

Contribution. Despite the considerable interest in optimal taxes for non-atomic games with heterogeneous users and for atomic games with homogeneous users, it is unknown whether weakly- or strongly-optimal taxes exist for *atomic* network games with *heterogeneous* users. The case of heterogeneous users is substantially different, and more complicated, than that of homogeneous users, since the game with taxes is a congestion game with player-specific additive constants [11].

In this work, we study for the first time the existence of optimal taxes for atomic network games with heterogeneous users, and present two complementary and essentially best-possible results. On the positive side, we prove the existence of weakly-optimal taxes in single-source network congestion games with heterogeneous users (cf. Section 3). To establish this result, we follow the proof technique of [9], and show that any acyclic traffic pattern is induced as a Nash equilibrium of the game with the taxes calculated as in [9, Theorem 1]. Our result is significantly stronger than any previously known positive result on weakly-optimal taxes for atomic congestion games. In particular, our result generalizes previous results of [3,7] not only in the direction of considering heterogeneous users, but also in the direction of considering non-symmetric games on single-source multiple-sink networks.

On the negative side, we show that users' heterogeneity precludes the existence of strongly-optimal taxes even on the simplest topology of parallel-link networks. More specifically, we present a parallel-link game with linear latency functions and heterogeneous users for which any tax vector induces an equilibrium with total latency at least $28/27$ times the optimum. Hence, we establish a dichotomy between the general case of heterogeneous users and the special case of homogeneous ones, as far as the existence of strongly-optimal taxes is concerned.

To the best of our knowledge, this is the first time in congestion games that a dichotomy is established (i) between the cases of homogeneous and heterogeneous users with respect to the existence of optimal taxes, and (ii) between the cases of non-atomic and atomic users on parallel links with respect to the efficiency of a price-of-anarchy-reducing mechanism. For the latter, we note that the worst-case price of anarchy for atomic games on parallel links is the same as the worst-case price of anarchy for non-atomic congestion games (see e.g. [15,6]), and

that the two classes of congestion games have similar behaviour with respect to their worst-case price of anarchy under some common price-of-anarchy-reducing mechanisms, such as Stackelberg strategies (see e.g. the bounds in [14,6] on the efficiency of Stackelberg strategy LLF for non-atomic and atomic parallel-link games) and taxes for homogeneous users.

2 Preliminaries

We consider a network congestion game $\mathcal{G}(l)$ defined on a directed graph $G = (V, E)$ with a nondecreasing latency function $l_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ on each edge $e \in E$. A set N of *users* is given, each with an amount of traffic (flow) to be routed from an origin node to a destination node of G . The users are *non-atomic* if each has infinitesimal demand and *atomic* otherwise. The game is single-source (resp. single-sink) if all users share the same origin (resp. destination) node, and *symmetric* if all users share the same origin-destination pair and have the same traffic demand.

Each user α has a positive *tax-sensitivity* factor $a(\alpha) > 0$. We will assume that the tax-sensitivity factors for all users come from a finite set of possible positive values. We call the users *heterogeneous* if there are at least two distinct sensitivity values and *homogeneous* otherwise. Unless we declare them explicitly to be heterogeneous, the users are assumed to be homogeneous. We can bunch together into a single *user class* all the users with the same origin-destination pair and with the same tax-sensitivity factor; let k be the number of different such classes. We denote by $d_i, \mathcal{P}_i, a(i)$ the total traffic demand of class i , the paths that can be used by class i , and the tax-sensitivity of class i , for all $i = 1, \dots, k$ respectively. Thus each user in class i selects a path in \mathcal{P}_i and routes her traffic through it. We set $\mathcal{P} \doteq \cup_{i=1, \dots, k} \mathcal{P}_i$ the union of paths used by all classes. In the following, we assume that the game is single-source and the users are atomic and have unit demands, unless it is stated otherwise.

A configuration f is a tuple $f = (f^j)_{j \in N}$ consisting of a path f^j from the corresponding origin node to the corresponding destination node for each user j . Given a configuration f , we let f_P denote the total traffic routed through any path $P \in \mathcal{P}$, and let $f_e = \sum_{e \in P} f_P$ denote the total traffic routed through any edge $e \in E$. Given a configuration f , we refer to the traffic vector $(f_e)_{e \in E}$ as the (edge-)flow induced by f . We note that different configurations may induce the same edge-flow. We say that a flow f is *feasible* (with respect to an atomic network congestion game $\mathcal{G}(l)$) if there is a configuration f of $\mathcal{G}(l)$ which routes traffic f_e through any edge e . We slightly abuse the notation by letting the same symbol denote both a configuration and the feasible flow induced by it. A configuration (or the corresponding flow) f is acyclic if for any cycle C in the underlying network G , there is an edge $e \in C$ with $f_e = 0$.

The *latency function* $l_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ assigned to each edge e gives the latency experienced by any user on e due to the congestion caused by the traffic routed through e . We assume that the functions l_e are nondecreasing, and that $l_e(f_e) > 0$ when $f_e > 0$, i.e., the function l_e is *positive*.

For any configuration f and path $P \in \mathcal{P}$, the latency of P is $l_P(f) = \sum_{e \in P} l_e(f_e)$. The *individual cost* of a user j in a configuration f is $c^j(f) = \sum_{e \in f^j} l_e(f_e)$, i.e., the latency on her path in f . A configuration f is a *pure Nash equilibrium* of $\mathcal{G}(l)$ if no user can improve her individual cost by unilaterally deviating from f . Formally, for a tuple $x = (x_1, \dots, x_n)$, let $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ and $(x_{-i}, x'_i) = (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$. Configuration f is a pure Nash equilibrium if $c^j(f) \leq c^j(f_{-j}, P)$ for any user j in any class i and any path $P \in \mathcal{P}_i$.

A flow f satisfies the *Wardrop principle* [16] if for each class i , the latency on all paths in \mathcal{P}_i used by f is no greater than the latency on any other path in \mathcal{P}_i . A non-atomic (atomic) *Wardrop equilibrium* is a (feasible) flow f that satisfies the Wardrop principle. We distinguish between atomic and non-atomic Wardrop equilibria, depending on whether the users are atomic or not. An atomic Wardrop equilibrium is also a pure Nash equilibrium, while the converse may not be true.

If every edge is assigned a *tax* (also called *toll*) $\beta_e \geq 0$, the resulting game is denoted as $\mathcal{G}(l + \beta)$. Given a configuration f in $\mathcal{G}(l + \beta)$, the individual cost of a user j included in a class i is: $c^j_\beta(f) = \sum_{e \in f^j} l_e(f_e) + a(i) \sum_{e \in f^j} \beta_e$.

Let \hat{f} be a configuration that minimizes the total latency $\sum_e f_e l_e(f_e)$ over all configurations of $\mathcal{G}(l)$. Although in certain cases (e.g., when the functions $f_e l_e(f_e)$ are convex) the flow \hat{f} can be computed efficiently, for more general latency functions it may be intractable to compute \hat{f} . We will assume that \hat{f} is given to us off-line and that it induces a finite latency on every edge. A tax vector β *weakly induces* a feasible (non-atomic) flow f if f is a pure Nash (non-atomic Wardrop) equilibrium of $\mathcal{G}(l + \beta)$. A tax vector β is called *weakly-optimal* if it weakly induces a pure Nash equilibrium f whose total latency $\sum_{e \in E} f_e l_e(f_e)$ is equal to the optimal total latency $\sum_e \hat{f}_e l_e(\hat{f}_e)$. A tax vector β is called *strongly-optimal*, if every pure Nash equilibrium it induces in $\mathcal{G}(l + \beta)$ has total latency equal to the optimal total latency $\sum_e \hat{f}_e l_e(\hat{f}_e)$.

Let $F(x) = (F_1(x), F_2(x), \dots, F_n(x))$ be a vector-valued function from the n -dimensional space \mathbb{R}^n into itself. Then the *nonlinear complementarity problem* of mathematical programming is to find a vector x that satisfies the following system:

$$x^T F(x) = 0, \quad x \geq 0, \quad F(x) \geq 0.$$

3 Existence of Weakly-Optimal Taxes

In this section we consider networks with a single-source s and heterogeneous users. Each user class i consists of a single user who wishes to route d_i units of traffic through a single $s - t_i$ path. We show that if $d_i = 1$ (or more generally, if d_i are arbitrary and the optimal configuration is acyclic), there exists a vector of weakly-optimal taxes. In particular, we establish the existence of a tax vector that weakly induces any acyclic flow \hat{f} as an atomic Wardrop equilibrium. Since single-source network congestion games with unit-demand users admit an acyclic optimal flow \hat{f} , this implies the existence of weakly optimal taxes for such games.

The proof follows closely [9], where the existence of weakly-optimal taxes is shown for the non-atomic case, and here we give a sketch with the new elements added for our case. In [9], it is shown that, if we add to the network artificial capacity constraints, $f_e \leq \hat{f}_e, \forall e \in E$, there is a tax-vector β^* that induces as a non-atomic Wardrop equilibrium a flow f^* that satisfies demands d_i and respects the capacities. In particular, [9] shows that the following nonlinear complementarity problem always has a solution (details omitted). Moreover, if \hat{f} is given offline, this solution can be computed in polynomial time.

$$\begin{aligned}
 f_P(T_P(f) - u_i) &= 0 & \forall i, \forall P \in \mathcal{P}_i & \quad (\text{BIG CP}) \\
 T_P(f) &\geq u_i & \forall i, \forall P \in \mathcal{P}_i \\
 u_i(\sum_{P \in \mathcal{P}_i} f_P - d_i) &= 0 & \forall i \\
 \sum_{P \in \mathcal{P}_i} f_P &\geq d_i & \forall i \\
 \beta_e(f_e - \hat{f}_e) &= 0 & \forall e \in E \\
 f_e &\leq \hat{f}_e & \forall e \in E \\
 f_P, \beta_e, u_i &\geq 0 & \forall P, e, i
 \end{aligned}$$

Here the function $T_P(f)$ is set to $l_P(f)/a(i) + \sum_{e \in P} \beta_e^*, \forall P \in \mathcal{P}_i, \forall i$.

Lemma 1. *Let \hat{f} be an acyclic feasible flow for demands d_i , and let (f^*, β^*, u^*) be any solution of (BIG CP). Then $\sum_{P \in \mathcal{P}_i} f_P^* = d_i, \forall i$ and $f_e^* = \hat{f}_e, \forall e \in E$.*

Proof. The proof of the first part is essentially the same as the proof by contradiction of Proposition 4.1 in [1] and is omitted.

Vector f^* is a non-atomic flow, that satisfies the following set of constraints:

$$\sum_{P \in \mathcal{P}_i} f_P = d_i \quad \forall i \in \{1, \dots, k\} \tag{1}$$

$$f_e = \sum_{P \in \mathcal{P}: e \in P} f_P \quad \forall e \in E \tag{2}$$

$$f_e \leq \hat{f}_e \quad \forall e \in E \tag{3}$$

$$f_P \geq 0 \quad \forall P \in \mathcal{P} \tag{4}$$

Consider the network which consists only of the edges e of G with $\hat{f}_e > 0$. Augment this network by adding a super-sink t and an edge (t_i, t) from each of the old sinks to t . Call $G_{\hat{f}}$ the resulting network. Extend f^* to an s - t flow in $G_{\hat{f}}$ by setting $f_{(t_i, t)}^* = d_i$. Let (S, T) be any cut that separates s from t in $G_{\hat{f}}$. Since \hat{f} is acyclic, it must be that $\sum_{e \in \delta(S)} \hat{f}_e = \sum_{i=1}^k d_i$. Because of (1), it must be that $\sum_{e \in \delta(S)} f_e^* \geq \sum_{i=1}^k d_i = \sum_{e \in \delta(S)} \hat{f}_e$. By the capacity constraints (3), we conclude that $\sum_{e \in \delta(S)} f_e^* = \sum_{e \in \delta(S)} \hat{f}_e$, and in particular, that $f_e^* = \hat{f}_e$

for all edges e that cross the cut. The only edges of G on which f^* might send positive flow are the edges of $G_{\hat{f}}$. Any such edge e belongs to at least one s - t cut in $G_{\hat{f}}$. By applying the previous argument to such a cut, it follows that $f_e^* = \hat{f}_e, \forall e \in E$. □

Then [9] establishes that we can use β^* as a tax vector to weakly induce f^* as a non-atomic Wardrop equilibrium in the original network *without* the capacity constraints. This follows from the fact that we can use β^* as a tax vector to ensure that (f^*, u^*) is also a solution to the following complementarity problem:

$$\begin{aligned}
 (T_P(f) - u_i)f_P &= 0 & \forall i, \forall P \in \mathcal{P}_i & \tag{CP} \\
 T_P(f) - u_i &\geq 0 & \forall i, \forall P \in \mathcal{P}_i & \\
 u_i \left(\sum_{P \in \mathcal{P}_i} f_P - d_i \right) &= 0 & \forall i & \\
 \sum_{P \in \mathcal{P}_i} f_P &\geq d_i & \forall i & \\
 f, u &\geq 0 & &
 \end{aligned}$$

If the path latency functions are continuous and positive, Aashtiani and Magnanti [1] show that the Wardrop equilibria of the game $\mathcal{G}(l + \beta)$ can be described as the solutions to (CP). T_P above denotes the cost of a user that uses path P , f_P is the flow through path P , and $u = (u_1, \dots, u_k)$ is the vector of shortest travel times for the commodities. The first two equations model Wardrop’s principle by requiring that for any origin-destination pair i , the travel cost for all paths in \mathcal{P}_i with nonzero flow is the same and equal to u_i . The remaining equations ensure that the demands are met and that the variables are nonnegative.

The fact that for all $e \in E, f_e^* = \hat{f}_e$ proves that the tax vector β^* we compute weakly induces as an equilibrium the *atomic* solution \hat{f} as well. We have thus shown the following theorem, which is the main result of this section.

Theorem 1. *Let all atomic heterogeneous users share the same source, and let \hat{f} be any acyclic feasible flow. If for every edge $e \in E, l_e(\cdot)$ is a nondecreasing positive function, then there is a tax vector $\beta \in \mathbb{R}_+^{|E|}$ such that, there is an atomic Wardrop traffic equilibrium \bar{f} for the game $\mathcal{G}(l + \beta)$, where $\bar{f}_e = \hat{f}_e, \forall e \in E$. Given \hat{f}, β can be computed in polynomial time.*

If the latency functions are strictly increasing, the uniqueness results from [1] yield that \hat{f} is the only Wardrop atomic equilibrium induced by the tax vector of the theorem.

Single-source network congestion games with unit-demand users and nondecreasing latency functions admit an acyclic optimal flow \hat{f} . Moreover, if for all $e \in E, x l_e(x)$ are convex, such an optimal flow can be computed in polynomial time by a min-cost flow computation. Therefore, we obtain the following corollary of Theorem 1:

Corollary 1. *Let $\mathcal{G}(l)$ be an atomic network congestion with nondecreasing latency functions and heterogeneous users, where all users share the same source and have the same traffic demand. Then $\mathcal{G}(l)$ admits a weakly-optimal tax vector β . Furthermore, if for all edges e , $xl_e(x)$ is convex, β can be computed in polynomial time.*

Theorem 1 states that computing the weakly-optimal tax vector β for an acyclic optimal flow \hat{f} is not substantially harder than computing \hat{f} : if \hat{f} can be computed in polynomial time, β can also be computed in polynomial time. An interesting question is whether computing the tax vector β of Theorem 1 is substantially *easier* than computing the corresponding acyclic optimal flow \hat{f} . The following theorem practically excludes this possibility. In particular, we show that given the weakly-optimal tax vector β of Theorem 1, we can decide in polynomial time whether the optimal total latency is bounded from above by a given number. So the problem of computing the weakly-optimal tax vector β is at least as difficult as the problem of determining the optimal total latency.

Theorem 2. *For atomic games with user-specific demands, if the optimal flow \hat{f} is not given, it is NP-hard to compute the taxes whose existence is established by Theorem 1. This holds even for parallel-link games with homogeneous users.*

Proof. We employ a Turing reduction from PARTITION. We consider an instance of the decision version of PARTITION, i.e., a set of integers $\{a_1, \dots, a_n\}$ whose total sum is $2B$ for some $B > 0$. For every integer a_i , we create a user with demand a_i and tax-sensitivity 1. Every partition of the users into two sets, one with total sum $B - T$ and the other with total sum $B + T$ for some $T \geq 0$, induces in the network with two parallel links and latency function $l(x) = x$, a corresponding acyclic routing of the users whose total cost is

$$(B - T)^2 + (B + T)^2 = 2B^2 + 2T^2.$$

This quantity is minimized for $T = 0$, i.e., when the PARTITION instance is a YES-instance. Assume now that you can compute in polynomial time the taxes of Theorem 1. Because the latency functions are strictly increasing, the Wardrop equilibrium is unique in terms of edge flows 1. Moreover, it is well-known that the equilibrium solution f can be computed in polynomial time by solving a convex quadratic mathematical program 2. By Theorem 1, on each of the two parallel edges, the value of f will be equal to value of the optimal unsplitable solution. Checking these values, we can determine whether the PARTITION instance is a YES-instance. □

Unfortunately, it is known that the taxes of Theorem 1 are not in general strongly-optimal. Note that for homogeneous users, our taxes are cost-balancing in the sense of Fotakis and Spirakis 7. They give an example of a symmetric network congestion game, with homogeneous users, where the cost-balancing taxes induce an a pure Nash equilibrium of total latency 1.13 times the optimum. In the full version of the paper we give another such example where the cost-balancing taxes induce a pure Nash equilibrium of total latency $(1.2 - \epsilon)$ times the optimum.

4 Inexistence of Strongly-Optimal Taxes

We proceed to show that that atomic congestion games with heterogeneous users may not admit strongly-optimal taxes, even for parallel-link games with linear latencies and unit-demand users.

Theorem 3. *There exists a parallel-link game with linear latencies and heterogeneous unit-demand users, for which any tax vector induces an equilibrium with total latency at least $28/27$ times the optimal total latency.*

Proof. We consider a game $\mathcal{G}(l)$ on 3 parallel links with latency functions $l_1(x) = 7$, $l_2(x) = 2x$, and $l_3(x) = x + 1$. There are 6 unit-demand users, 2 users with tax-sensitivity 1 and 4 users with tax-sensitivity $1/2$. The unique optimal flow assigns a single user to link 1, 2 users to link 2, and 3 users to link 3, and achieves a total latency of 27. Any other feasible flow has total latency at least 28. In the following, we show that any weakly-optimal tax vector β induces an equilibrium of total latency at least 28, and thus this game does not admit strongly-optimal taxes. The proof proceeds by considering different cases depending on the 5 optimal allocations of heterogeneous users.

Case I: We consider an optimal flow that assigns a user with tax-sensitivity 1 to link 1, the other user with tax-sensitivity 1 and a user with tax-sensitivity $1/2$ to link 2, and 3 users with tax-sensitivity $1/2$ to link 3 (we denote such a configuration as $\langle (1), (1, 1/2), (1/2, 1/2, 1/2) \rangle$). Let $\beta = (\beta_1, \beta_2, \beta_3)$ be any (weakly-optimal) tax vector that induces the particular configuration as an equilibrium of $\mathcal{G}(l + \beta)$. No user has an incentive to deviate from its assigned link; writing down the corresponding inequalities, we obtain that β must satisfy the following:

$$1 + \beta_1 \leq \beta_2 \leq 3 + \beta_1 \tag{5}$$

$$\beta_2 - 1 \leq \beta_3 \leq 4 + \beta_2 \tag{6}$$

$$2 + \beta_1 \leq \beta_3 \leq 6 + \beta_1 \tag{7}$$

If β is strongly-optimal, configuration $\langle (1, 1), (1/2, 1/2), (1/2, 1/2) \rangle$ is not an equilibrium of $\mathcal{G}(l + \beta)$. Therefore at least one user in that configuration has an incentive to deviate, and β must satisfy *at least one* of the following:

$$\beta_2 < 1 + \beta_1 \tag{8} \qquad \qquad \qquad \beta_3 < \beta_2 \tag{11}$$

$$\beta_3 < 3 + \beta_1 \tag{9} \qquad \qquad \qquad 8 + \beta_1 < \beta_3 \tag{12}$$

$$6 + \beta_1 < \beta_2 \tag{10} \qquad \qquad \qquad 6 + \beta_2 < \beta_3 \tag{13}$$

We observe that (8) contradicts (5), (10) contradicts (5), (12) contradicts (7), and (13) contradicts (6). Hence, if β is strongly optimal, either $\beta_3 < 3 + \beta_1$ or $\beta_3 < \beta_2$ (ie. β_3 must be “small”).

Moreover, if β is strongly-optimal, $\langle (\cdot), (1, 1), (1/2, 1/2, 1/2, 1/2) \rangle$ is not an equilibrium of $\mathcal{G}(l + \beta)$, and β must satisfy at least one of the following:

$$3 + \beta_1 < \beta_2 \tag{14} \qquad 4 + \beta_1 < \beta_3 \tag{16}$$

$$\beta_3 < \beta_2 - 2 \tag{15} \qquad 2 + \beta_2 < \beta_3 \tag{17}$$

We observe that (14) contradicts (5) and (15) contradicts (6). Hence, if β is strongly optimal, either $\beta_3 > 4 + \beta_1$ or $\beta_3 > 2 + \beta_2$ (ie. β_3 must be “large”).

If $\beta_3 < 3 + \beta_1$, neither $\beta_3 > 4 + \beta_1$ nor $\beta_3 > 2 + \beta_2$ is possible (note that $3 + \beta_1 > \beta_3 > 2 + \beta_2$, which contradicts (5)). If $\beta_3 < \beta_2$, neither $\beta_3 > 2 + \beta_2$ nor $\beta_3 > 4 + \beta_1$ is possible (note that $\beta_2 > \beta_3 > 4 + \beta_1$, which contradicts (5)). Therefore, any tax vector that induces optimal configuration $\langle (1), (1, 1/2), (1/2, 1/2, 1/2) \rangle$ as an equilibrium of $\mathcal{G}(l + \beta)$ also induces either $\langle (1, 1), (1/2, 1/2), (1/2, 1/2) \rangle$ or $\langle (\cdot), (1, 1), (1/2, 1/2, 1/2, 1/2) \rangle$ (both of total latency 28) as an equilibrium.

Case II: We consider optimal configuration $\langle (1), (1/2, 1/2), (1, 1/2, 1/2) \rangle$. Working as in Case I, we obtain that any tax vector $\beta = (\beta_1, \beta_2, \beta_3)$ that induces this configuration as an equilibrium of $\mathcal{G}(l + \beta)$ must satisfy the following:

$$1 + \beta_1 \leq \beta_2 \leq 5 + \beta_1 \tag{18}$$

$$\beta_2 - 2 \leq \beta_3 \leq 2 + \beta_2 \tag{19}$$

$$2 + \beta_1 \leq \beta_3 \leq 3 + \beta_1 \tag{20}$$

In fact, the right-hand side of (18) follows from $\beta_2 - 2 \leq \beta_3 \leq 3 + \beta_1$.

Considering configuration $\langle (1, 1), (1/2, 1/2), (1/2, 1/2) \rangle$ and working as in Case I, we obtain that if β is strongly-optimal, either $\beta_3 < 3 + \beta_1$ or $\beta_3 < \beta_2$. Considering configuration $\langle (\cdot), (1, 1), (1/2, 1/2, 1/2, 1/2) \rangle$, we obtain that if β is strongly-optimal, either $\beta_2 > 3 + \beta_1$ (note that (14) does not contradict (18)), or $\beta_3 > 4 + \beta_1$, or $\beta_3 > 2 + \beta_2$.

Working as in Case I, we show that if β is strongly-optimal, it must satisfy both $\beta_2 > 3 + \beta_1$ and $\beta_3 < \beta_2$ (since $\beta_2 > 3 + \beta_1$, $\beta_3 < 3 + \beta_1$ implies $\beta_3 < \beta_2$, so β_3 must be smaller than β_2 in any case), in addition to (18), (19), (20).

Moreover, if β is strongly-optimal, configuration $\langle (1, 1), (1/2), (1/2, 1/2, 1/2) \rangle$ is not an equilibrium of $\mathcal{G}(l + \beta)$, and β must satisfy at least one of the following:

$$\beta_2 < 3 + \beta_1 \tag{21} \qquad 6 + \beta_3 < \beta_2 \tag{24}$$

$$\beta_3 < 2 + \beta_1 \tag{22} \qquad 6 + \beta_1 < \beta_3 \tag{25}$$

$$10 + \beta_1 < \beta_2 \tag{23} \qquad \beta_2 < \beta_3 \tag{26}$$

We observe that (21) contradicts $\beta_2 > 3 + \beta_1$, (22) contradicts (20), (23) contradicts (18), (25) contradicts (20), and (26) contradicts $\beta_3 < \beta_2$. Furthermore, (18) and (19) imply that $\beta_2 \leq 5 + \beta_1 \leq 3 + \beta_3$, which contradicts (24). Hence, any tax vector that induces optimal configuration $\langle (1), (1/2, 1/2), (1, 1/2, 1/2) \rangle$ as an equilibrium also induces either configuration $\langle (1, 1), (1/2, 1/2), (1/2, 1/2) \rangle$, or $\langle (\cdot), (1, 1), (1/2, 1/2, 1/2, 1/2) \rangle$, or $\langle (1, 1), (1/2), (1/2, 1/2, 1/2) \rangle$ (all of total latency 28) as an equilibrium.

Case III: We consider optimal configuration $\langle (1/2), (1, 1), (1/2, 1/2, 1/2) \rangle$. Any tax vector $\beta = (\beta_1, \beta_2, \beta_3)$ that induces this configuration as an equilibrium of $\mathcal{G}(l + \beta)$ must satisfy the following:

$$2 + \beta_1 \leq \beta_2 \leq 3 + \beta_1 \quad (27)$$

$$\beta_2 - 1 \leq \beta_3 \leq 4 + \beta_2 \quad (28)$$

$$4 + \beta_1 \leq \beta_3 \leq 6 + \beta_1 \quad (29)$$

Therefore, any tax vector that induces $\langle (1/2), (1, 1), (1/2, 1/2, 1/2) \rangle$ as an equilibrium of $\mathcal{G}(l + \beta)$ must satisfy (5), (6), and (7), and by Case I, is not strongly-optimal.

Cases IV and V: Optimal configurations $\langle (1/2), (1, 1/2), (1, 1/2, 1/2) \rangle$ and $\langle (1/2), (1/2, 1/2), (1, 1, 1/2) \rangle$ are not induced as an equilibrium of $\mathcal{G}(l + \beta)$ by any tax vectors. In particular, applying the inequalities for possible deviations between link 1 and link 3, we obtain that any tax vector β that induces any of the configurations above as an equilibrium must satisfy $4 + \beta_1 \leq \beta_3 \leq 3 + \beta_1$.

Thus we have considered all optimal allocations of heterogeneous users and all weakly-optimal tax vectors β , and have shown that any of them induces a configuration of total latency at least 28 as an equilibrium of $\mathcal{G}(l + \beta)$. \square

Remark 1. For the atomic game with homogeneous users corresponding to the parallel-link game in the proof of Theorem 3, the tax vector $(0, 3 - \delta, 3 - \delta)$, for a sufficiently small $\delta > 0$, is a strongly-optimal tax vector (a slightly different strongly-optimal tax vector is given by [3, Theorem 1]). For the corresponding non-atomic game with heterogeneous users, the tax vector $(0, 3, 3)$ is a strongly-optimal one.

5 Open Problems

It is known that for homogeneous users with unit-demands on multicommodity networks there exist no strongly-optimal taxes [3]. Series-parallel networks is the largest class for which such taxes have been shown so far to exist [7]. In this work, we established that when the users are heterogeneous, there are no strongly-optimal taxes even on the very specialized topology of parallel links. The challenging open problem stated in [3] remains for future work: determine the largest class of network congestion games for which strongly-optimal taxes exist. The candidate class is that of symmetric network games [3], i.e., when users are homogeneous, have identical demands, and share the same source and destination on a general-topology network.

Acknowledgement. G. Karakostas and S. Kolliopoulos thank Ioannis Caragiannis for introducing them to the problem and for valuable discussions.

References

1. Aashtiani, H.Z., Magnanti, T.L.: Equilibria on a congested transportation network. *SIAM Journal of Algebraic and Discrete Methods* 2, 213–226 (1981)
2. Beckmann, M., McGuire, C.B., Winsten, C.B.: *Studies in the Economics of Transportation*. Yale University Press, New Haven (1956)
3. Caragiannis, I., Kaklamani, C., Kanellopoulos, P.: Taxes for linear atomic congestion games. In: *Proceedings of the 14th Annual European Symposium on Algorithms*, pp. 184–195 (2006)
4. Cole, R., Dodis, Y., Roughgarden, T.: Pricing network edges for heterogeneous selfish users. In: *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pp. 521–530 (2003)
5. Fleischer, L., Jain, K., Mahdian, M.: Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In: *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 277–285 (2004)
6. Fotakis, D.: Stackelberg strategies for atomic congestion games. *Theory of Computing Systems* 47(1), 218–249 (2010)
7. Fotakis, D., Spirakis, P.G.: Cost-balancing tolls for atomic network congestion games. *Internet Mathematics* 5(4), 343–364 (2008)
8. 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)
9. Karakostas, G., Kolliopoulos, S.G.: Edge pricing of multicommodity networks for heterogeneous selfish users. In: *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 268–276 (2004)
10. Kontogiannis, S.C., Spirakis, P.G.: Atomic selfish routing in networks: a survey. In: *Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828*, pp. 989–1002. Springer, Heidelberg (2005)
11. Mavronicolas, M., Milchtaich, I., Monien, B., Tiemann, K.: Congestion games with player-specific constants. In: *Kučera, L., Kučera, A. (eds.) MFCS 2007. LNCS, vol. 4708*, pp. 633–644. Springer, Heidelberg (2007)
12. Pigou, A.C.: *The Economics of Welfare*. Macmillan and Co., London (1920)
13. Rosenthal, R.W.: A class of games possessing pure Nash strategy equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
14. Swamy, C.: The effectiveness of Stackelberg strategies and tolls for network congestion games. In: *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1133–1142 (2007)
15. Roughgarden, T.: The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences* 67(2), 341–364 (2003)
16. Wardrop, J.G.: Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers, Part II* 1, 325–378 (1952)
17. Yang, H., Huang, H.-J.: The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Research B* 38, 1–15 (2004)

Computing Stable Outcomes in Hedonic Games

Martin Gairing and Rahul Savani

Department of Computer Science
University of Liverpool
{m.gairing,rahul.savani}@liverpool.ac.uk

Abstract. We study the computational complexity of finding stable outcomes in symmetric additively-separable hedonic games. These coalition formation games are specified by an undirected edge-weighted graph: nodes are players, an outcome of the game is a partition of the nodes into coalitions, and the utility of a node is the sum of incident edge weights in the same coalition. We consider several natural stability requirements defined in the economics literature. For all of them the existence of a stable outcome is guaranteed by a potential function argument, so local improvements will converge to a stable outcome and all these problems are in PLS. The different stability requirements correspond to different local search neighbourhoods. For different neighbourhood structures, our findings comprise positive results in the form of polynomial-time algorithms for finding stable outcomes, and negative (PLS-completeness) results.

1 Introduction

Hedonic games were introduced in the economics literature as a model of coalition formation where each player cares only about those within the same coalition [9]. Such games can be used to model a variety of settings ranging from multi-agent coordination to group formation in social networks. This paper studies the computational complexity of finding stable outcomes in hedonic games. We consider the stability requirements introduced in [5], which includes a detailed discussion of real-life situations in which hedonic models are reasonable.

The literature has focused almost exclusively on the issue of the *existence* of stable outcomes. When computational complexity has been addressed, it has been in the context of deciding whether a stable outcome exists. This has been done under different utility functions and stability requirements [5, 3, 18]. An outcome is called *Nash-stable* if no player prefers to be in a different coalition. This is the most stringent stability requirement we consider: here a deviation depends only on the preferences of the deviating player. Less stringent stability requirements are achieved by restricting feasible deviations: a coalition may try to hold on to an attractive player or block the entry of an unattractive player.

We consider the case of hedonic games with *additively-separable* utilities, as they allow a succinct representation which is suitable for studying computational complexity. In this representation, a game is specified by an *edge-weighted graph*. In general this graph is directed, which allows non-symmetric preferences, but

then a stable outcome might not exist [5, 3]. In the sequel, a hedonic game is an *undirected* edge-weighted graph, so that preferences are *symmetric*. Every node is a player, and an outcome is a partition of the players into coalitions. For a given outcome, the utility of a player is the sum of the edges weights of the incident edges to nodes in the same coalition. For a symmetric additively-separable hedonic game, a Nash-stable outcome always exists by a simple potential function argument: the potential function is the total happiness of an outcome, i.e., the sum of players' utilities. Nash-deviations improve the potential. We define the problem NASHSTABLE as that of computing a Nash-stable outcome for an additively-separable hedonic game.

We also consider a less stringent stability requirement, called *individual stability*. Here the set of all feasible deviations for a given outcome is a subset of Nash deviations: a player can deviate to another coalition only if everyone in this coalition is happy to have her. We also consider an even less stringent stability requirement, called *contractual individual stability*. Here the set of all feasible deviations for a given outcome is a subset of individually-stable deviations: (in addition to the above requirement) a player can deviate only if everyone in the coalition she leaves is happy for her to leave. These stability requirements were introduced in [5]. The same potential function argument shows that individually-stable outcomes and contractually-individually-stable outcomes exist for symmetric additively-separable hedonic games, and indeed every Nash-stable outcome is also individually-stable and contractually-individually-stable. In each case, local improvements will find a stable outcome, and all the problems we consider are in the complexity class PLS (polynomial local search) [12]. Local search dynamics are desirable because they are *distributed*. Are they also efficient for hedonic games? If not, can we find efficient dynamics or centralized algorithms for finding stable outcomes?

Symmetric additively-separable hedonic games are closely related to party affiliation games, which are also specified by an undirected edge-weighted graph. In a party affiliation game each player must choose between one of two "parties"; each player's happiness is the sum of her edges to nodes in the same party; in a stable outcome no player would prefer to be in the other party. The problem PARTYAFFILIATION is to find a stable outcome in such a game. If such an instance has only negative edges then it is equivalent to the problem LOCAL-MAXCUT, which is to find a stable outcome of a local max-cut game. In party affiliation games there are at most two coalitions, while in hedonic games any number of coalitions is allowed. Thus, whereas PARTYAFFILIATION for instances with only negative edges is PLS-complete [16], NASHSTABLE is trivial in this case, as the outcome where all players are in singleton coalitions is Nash-stable. Both problems are trivial when all edges are non-negative, in which case the grand coalition of all players is Nash-stable. Thus, interesting hedonic games contain both positive and negative edges.

Our Contribution. In this paper, we examine the complexity of computing stable outcomes in symmetric additively-separable hedonic games. We observe that NASHSTABLE, i.e., the problem of computing a Nash-stable outcome, is

PLS-complete. Here, we give a simple reduction from PARTYAFFILIATION, which was shown to be PLS-complete in [16]. Our reduction relies on a method to ensure that all stable outcomes use exactly two coalitions (where in general there can be as many coalitions as players). In contrast, the problem CIS of finding a contractually-individually-stable outcome can be solved in polynomial time.

Moreover, we study IS, i.e., the problem of finding an individually-stable outcome. We show that if the outcome is restricted to contain at most two coalitions, an individually-stable outcome can be found in polynomial time. This suggests that a reduction showing PLS-hardness for IS cannot be as simple as for NASH-STABLE: one would need to construct hedonic games that allow three or more coalitions. In order to prove a hardness result, we increase the size of the neighbourhood, defining the search problem ISWITHSWAPS, which is similar to IS, but in addition to one-player deviations, two players can switch coalitions.

We define a restricted version of PARTYAFFILIATION, called ONEENEMY-PARTYAFFILIATION, in which each player dislikes at most one other player. Our main result is that ONEENEMYPARTYAFFILIATION is PLS-complete. This reduction is from CIRCUITFLIP and is rather involved. We reduce ONEENEMY-PARTYAFFILIATION to ISWITHSWAPS, which shows it is PLS-complete; we leave the complexity of IS open.

Related Work. Hedonic coalition formation games were first considered by [9]. [11] later surveyed coalition structures in game theory and economics. Based on [9], [5] formulated different stability concepts in the context of hedonic games, which are the basic definitions we use here. The general focus in the game theory community has been on characterizing the conditions for which stable outcomes exist. [6] showed that additively-separable and symmetric preferences guarantee the existence of a Nash-stable partition. They also showed that under certain different conditions on the preferences, the set of Nash-stable partitions can be empty but the set of individually-stable partitions is always non-empty.

[7] surveys algorithmic problems related to stable partitions. [3] showed that for hedonic games represented by an *individually rational list of coalitions*, the complexity of checking whether core-stable, Nash-stable or individual-stable partitions exist is NP-complete, and that every hedonic game has a contractually-individually-stable solution. Recently, [18] showed that for additively-separable hedonic games checking whether a core-stable, strict-core-stable, Nash-stable or individually-stable partition exists is NP-hard. [10] characterize the complexity of problems related to coalitional stability for hedonic games represented by hedonic nets, a succinct, rule-based representation based on marginal contribution nets.

The definition of party affiliation games we use appears in [2]. Recent work on local max cut and party affiliation games has focused on approximation [4, 8]; see also [15]. For surveys on the computational complexity of local search and the complexity class PLS, see [1, 14]. Our main PLS-completeness result (Theorem 2) uses ideas from [19] which in turn builds on [16].

2 Preliminaries

A symmetric additively-separable hedonic game is an undirected edge-weighted graph $G = (V, E, w)$. Every node $i \in V$ is a player. An outcome is a partition p of V into coalitions. Denote by $p(i)$ the coalition to which $i \in V$ belongs under p , and by $E(p(i))$ the set of edges $\{\{i, j\} \in E \mid j \in p(i)\}$. The utility of $i \in V$ under p is the sum of edges to others in the same coalition, $\sum_{e \in E(p(i))} w(e)$. We consider different levels of restrictions for player deviations; see [5].

Definition 1. Consider an outcome p of a game $G = (V, E, w)$. The outcome is Nash-stable if and only if there exists no player i and coalition $c \neq p(i)$, possibly empty, such that

$$\sum_{e \in E(p(i))} w(e) < \sum_{\{\{i, j\} \in E \mid j \in c\}} w(\{i, j\}) . \tag{1}$$

The outcome is individually-stable if and only if there exists no player i and coalition $c \neq p(i)$, possibly empty, such that (1) holds and

$$w(\{i, j\}) > 0 , \quad \forall j \in c . \tag{2}$$

The outcome is contractually-individually-stable if and only if there exists no player i and coalition $c \neq p(i)$, possibly empty, such that (1) and (2) hold and $w(\{i, j\}) < 0$ for all $j \in p(i)$.

The search problems NASHSTABLE, IS, and CIS are to find a stable outcome of a hedonic game for the corresponding definition of stability. For Nash-stability, a player is allowed to deviate based only on her own utility, irrespective of others. Individual stability allows any player to block an unattractive individual from entering her coalition, i.e., a single negative edge to a coalition prevents a player switching to that coalition (although a stable outcome may contain negative edges). Contractual individual stability also allows any player to prevent an attractive player leaving her coalition, i.e., a single positive edge prevents a player leaving a coalition. Recall that all games we consider contain both positive and negative edges, else the problem of finding a stable outcome is easy.

3 Computational Complexity of Finding Stable Outcomes

We start with the least restrictive condition under which player deviations are allowed, i.e., Nash deviations. Here a player is allowed to change her coalition whenever this improves her utility. By a very simple reduction from PARTYAFFILIATION we observe the following:

Observation 1. NASHSTABLE is PLS-complete.

Proof. Consider an instance of PARTYAFFILIATION which is represented as an edge weighted graph $G = (V, E, w)$. We augment G by introducing two new

players, called *supermodels*. Every player $i \in V$ has an edge of weight $W > \sum_{e \in E} |w_e|$ to each of the supermodels. The two supermodels are connected by an edge of weight $-M$, where $M > |V| \cdot W$. By the choice of M the two supermodels will be in different coalitions in any Nash-stable outcome of the resulting hedonic game. Moreover, by the choice of W , each player will be in a coalition with one of the supermodels. The fact that edges to supermodels have all the same weight directly implies a one-to-one correspondence between the Nash-stable outcomes in the hedonic game and in the party affiliation game. \square

Now that we have PLS-completeness under the least restrictive deviation condition, it is natural to ask about stable outcomes under more restrictive conditions. We proceed with the most restrictive version that we consider.

Proposition 1. *CIS can be solved in $\mathcal{O}(|E|)$ time. Moreover, local improvements converge in at most $2|V|$ steps.*

It is easy to construct stable CIS partitions. The reason for this comes from the very restrictive conditions under which deviations are allowed. We now study deviation conditions which are less restrictive than in CIS but more restrictive than Nash deviations. Recall that in an individually-stable outcome a player is always allowed to leave a coalition but only allowed to enter if no player in the new coalition is connected to her by a negative edge. It is an interesting open problem whether IS is PLS-hard. The following result implies that for a PLS-hardness reduction we need to use at least three coalitions (unless $\text{PLS} \subseteq \text{P}$), unlike the reduction for NASHSTABLE (Observation [1](#)). Let 2-IS be the problem of computing an individually-stable outcome when at most two coalitions are allowed.

Proposition 2. *2-IS can be solved in polynomial time.*

Proof. We assume that there is at least one negative edge. Otherwise, the grand coalition is Nash-stable. The algorithm goes as follows:

Start with any bipartition. Move nodes with incident negative edges so that they have a negative edge to the other coalition. In each of the two coalitions, contract all nodes with negative incident edges into a single node and call the contracted nodes s and t . For any other node the new edge weights to s and t are the sum of the original edge weights. Now (ignoring all edges between s and t) compute a min cut between s and t via a max flow algorithm and assign the nodes accordingly.

After the first stage, all nodes that we are about to contract have a negative edge to the other coalition. So they are not allowed to join the other coalition. This property is preserved by contraction. Afterwards, the flow algorithm operates only on positive edges and computes a global minimum cut between s and t . Thus, the cut also maximizes the total happiness of all non-contracted nodes, so none of these nodes has an incentive to switch coalitions. All performed steps of the algorithm can be done in polynomial time. \square

What makes the problem easy in the case of two coalitions? The reason is simple: negative edges block deviations. This leads to an interesting question. What happens when we allow players to *swap* coalitions? Certainly, this increases the PLS-neighbourhood, and (in general) reduces the number of stable outcomes. We define an extended neighbourhood that includes *swaps*.

Definition 2. *In a swap two players swap coalitions. A swap is improving if at least one of the players becomes strictly better off and neither gets worse off.*


The new neighbourhood is comprised of single-player (IS type) deviations *and* swaps. Observe that in a solution where no player can improve by a single-player deviation, only swaps of two players connected by a negative edge can give rise to local improvements. With this larger neighbourhood we prove the following, which is the main result.

Theorem 2. *ISWITHSWAPS is PLS-complete.*

We prove this in two parts. We use the fact that in ONEENEMYPARTYAFFILIATION every node is incident to at most one negative edge to reduce this problem to ISWITHSWAPS by replacing negative edges with a simple local gadget (Lemma 1). Then our main result is that ONEENEMYPARTYAFFILIATION is PLS-complete (Theorem 3).

Lemma 1. *ONEENEMYPARTYAFFILIATION can be reduced in polynomial time to ISWITHSWAPS.*

Proof. We start with a party affiliation game where every player dislikes at most one other player. We add supermodels to enforce only two coalitions. We replace a negative edge (a, b) of weight $-w$ with the following gadget.



Here M is sufficiently large so that a and a' (as well as b and b') have to be in different coalitions and thus can only swap coalitions. Thus, if $a = b$, then both a and b receive a payoff of $-w$ from the original edge and 0 from the gadget.

On the other hand, if $a \neq b$, then both a and b receive a payoff of 0 from the original edge and w from the gadget. So we shifted the payoffs of a and b by w . Observe that the payoff of a' and b' is always w . So they will never block a swap. Thus, we didn't change the PLS neighbourhood of a and b . □

In order to complete the proof of Theorem 2, we show that ONEENEMYPARTYAFFILIATION is PLS-complete. Our proof is by reduction from the well known PLS-complete CIRCUITFLIP problem (cf. [16]).

Definition 3. *An instance of CIRCUITFLIP is a boolean circuit with n inputs and n outputs. A feasible solution is an assignment to the inputs and the value of a solution is the output treated as a binary number. The neighbourhood of an assignment consists of all assignments obtained by flipping exactly one input bit. The objective is to maximise the value.*

Theorem 3. ONEENEMYPARTYAFFILIATION is PLS-complete.

Proof. We reduce from CIRCUITFLIP. Let C be an instance of CIRCUITFLIP with inputs V_1, \dots, V_n , outputs C_1, \dots, C_n , and gates G_1, \dots, G_N . We make the following simplifying assumptions about C : (i) The gates are topologically ordered so that if the output of G_i is an input to G_j then $i > j$. (ii) All gates are NOR gates with fan-in 2. (iii) G_1, \dots, G_n is the output and G_{n+1}, \dots, G_{2n} is the (bitwise) negated output of C with G_1 and G_{n+1} being the most significant bits. (iv) G_{2n+1}, \dots, G_{3n} outputs a (canonical) better neighbouring solution if V_1, \dots, V_n is not locally optimal.

We use two complete copies of C . One of them represents the current solution while the other ones represents the next (better) solution. Each copy gives rise to a graph. We will start by describing our construction for one of the two copies and later show how they interact. Given C construct a graph G_C as follows:

We have nodes v_1, \dots, v_n representing the inputs of C , and nodes g_i representing the output of the gates of C . We will also use g_i to refer to the whole gate. For $i \in [n]$, denote by $w_i := g_{2n+i}$ the nodes representing the better neighbouring solution. Recall that g_1, \dots, g_n represent the output of C while g_{n+1}, \dots, g_{2n} correspond to the negated output.

In our party affiliation game we use 0 and 1 to denote the two coalitions. We slightly abuse notation by using $u = \kappa$ for $\kappa \in \{0, 1\}$ to denote that node u is in coalition κ . In the construction, we assume the existence of nodes with a fixed coalition. This can be achieved as in the proof of Observation 1 with the help of supermodels. We use 0 and 1 to refer to those constant nodes. In the graphical representation (cf. Figure 1), we represent those constants by square nodes.

We follow the exposition of [16] and [19] and use types to introduce our construction. Nodes may be part of multiple types. In general types are ordered w.r.t. decreasing edge weights. So earlier types are more important. Different types will serve different purposes.

Type 1: Check Gates. For each gate g_i we have a three-part component as depicted in Figure 1(a). The inputs of g_i , denoted $I_1(g_i)$ and $I_2(g_i)$, are either inputs of the circuit or outputs of some gate with larger index. The main purpose of this component is to check if g_i is correct, i.e., $g_i = \neg(I_1(g_i) \wedge I_2(g_i))$, and to set $z_i = 1$ if g_i is incorrect. The $\alpha, \beta, \gamma, \delta$ and λ nodes are local nodes for the gate. A gate can be in two operational modes, called *gate push regimes*. Type 7 will determine in which of the following push regimes a gate is.

Definition 4 (Gate push regimes). In the RESET GATE regime $\alpha_{i,1}, \alpha_{i,2}, \gamma_{i,1}$ and $\gamma_{i,2}$ get a bias towards 1 while $\lambda_{i,1}, \lambda_{i,2}, \beta_{i,1}, \beta_{i,2}, \beta_{i,3}, \delta_{i,1}, \delta_{i,2}$ and $\gamma_{i,3}$ get a bias towards 0. In the FIX GATE regime we have opposite biases.

Type 2: Propagate Flags. In order to propagate incorrect values for the z variables we interconnect them as in Figure 1(b) by using the topological order on the gates. Observe that for any locally optimal solution $z_i = 1$ enforces $z_j = 1$ for all $j < i$. The component is also used to (help to) fix the gates in order and to RESET them in the opposite order. Node z_{N+1} is for technical convenience.

Type 1 and 2 components are the same for both copies. In the following we describe how the copies interact. We denote the two copies of C by C^0 and C^1 and also use superscripts to distinguish between them for nodes of type 1 and 2.

Type 3: Set/Reset Circuits. The component of type 3 interconnects the z -flags from the two circuits C^0, C^1 . This component is depicted in Figure 1(c) and has multiple purposes. First, it ensures that in a local optimum d^0 and d^1 are not both 1. Second, at the appropriate time, it triggers to reset the circuit with smaller output. And third, it locks d^0 or d^1 to 1 and resets them back to 0 at the appropriate times.

The z and y nodes can also be in two different operational modes called COMPUTE regime and RESET regime which is determined by Type 6.

Definition 5 (Circuit push regimes). *Let $\kappa \in \{0, 1\}$. In the COMPUTE regime for z^κ all z_i^κ get a bias to 0 for all $0 \leq i \leq N + 1$ and y^κ gets a bias to 1. In the RESET regime for z^κ we give opposite biases.*

Type 4: Check Outputs. This component compares the current output of the two circuits and gives incentive to set one of the nodes d^0 or d^1 to 1 for which the output of the corresponding circuit is smaller. For all $i \in [n]$, we have edges $(d^0, g_{n+i}^0), (d^0, g_i^1), (d^1, g_{n+i}^1), (d^1, g_i^0)$ and $(0, g_{n+i}^0), (1, g_i^1), (0, g_{n+i}^1), (1, g_i^0)$ of weight 2^{2n+1-i} . To break symmetry we have edges $(0, d^0), (1, d^1)$ of weight 2^n .

Type 5: Feedback Better Solution. This component is depicted in Figure 1(d). It is used to feedback the improving solution of one circuit to the input of the other circuit. Its operation is explained in Lemma 3.

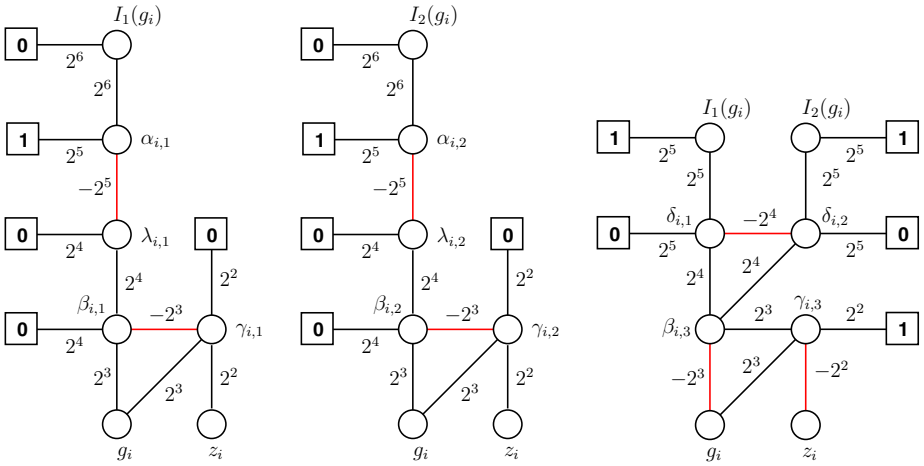
For the remaining types we use the following lemma and definition which are analogous to those in [19, 13].

Lemma 2. *For any polynomial-time computable function $f : \{0, 1\}^k \mapsto \{0, 1\}^m$ one can construct a graph $G_f(V_f, E_f, w)$ having the following properties: (i) there exist $s_1, \dots, s_k, t_1, \dots, t_m \in V_f$ with no negative incident edge, (ii) each node in V_f is only incident to at most one negative edge, (iii) $f(s) = t$ in any Nash-stable solution of the party affiliation game defined by G_f .*


Definition 6. *For a polynomial-time computable function $f : \{0, 1\}^k \mapsto \{0, 1\}^m$ we say that G_f as constructed in Lemma 2 is a graph that looks at $s_1, \dots, s_k \in V_f$ and biases $t_1, \dots, t_m \in V_f$ according to the function f .*

In the final three types we look at and bias nodes from the lower types already defined. For the final types we do not give explicit edge weights. In order that the “looking” has no side-effects on the operation of the lower types, we scale edge weights in these types such that any edge weight of lower type is larger than the sum of the edge weights of all higher types. More precisely, for $j \in \{5, 6, 7\}$, the weight of the smallest edge of type j is larger than the sum of weights of all edges of types $(j + 1), \dots, 8$.


In the following, denote by $C(v)$ the value of circuit C of the CIRCUITFLIP instance on input $v = (v_i)_{i \in [n]}$ and $w(v)$ the better neighbouring solution. Both are functions as in Definition 6.




(a) Type 1. Extra factor: 2^{2n+5i}



(b) Type 2. Extra factor: 2^{2n}



(c) Type 3. Extra factor: 2^{2n}



(d) Type 5. No extra factor

Fig. 1. Components of type 1,2,3, and 5. Edge weights have to be multiplied by the factors given above.

Type 6: Change Push Regimes for z . The component of type 6 looks at $v^0, v^1, d^0, d^1, \eta^0$ and η^1 (type 5) and biases z_i^0, z_i^1, y^0 and y^1 as follows. z^0 is put in the COMPUTE regime if at least one of the following 3 conditions is fulfilled: (i) $C(v^0) \geq C(v^1)$, (ii) $w(v^1) = v^0$, or (iii) $w(v^1) \neq \eta^1 \wedge d^0 = 1$. Else z^0 is put into the RESET regime. Likewise z^1 is put in the COMPUTE regime if at least one of the following three conditions is fulfilled: (i) $C(v^0) < C(v^1)$, (ii) $w(v^0) = v^1$, or (iii) $w(v^0) \neq \eta^0 \wedge d^1 = 1$. Else z^1 is put into the RESET regime. Note that conditions (i) and (ii) are important for normal computation, while (iii) is needed to overcome bad starting configurations.

Type 7: Change Push Regimes for Gates. For each $i \in [N]$ and $\kappa \in \{0, 1\}$, if $z_{i+1}^\kappa = 0$ we put the local variable of g_i^κ in the FIX GATE regime and in the RESET GATE regime otherwise.

Type 8: Fix Incorrect Gate. For each $i \in [N]$ and $\kappa \in \{0, 1\}$, the components of type 8 give a tiny offset to g_i^κ for computing correctly. For each gate g_i^κ we look at $\alpha_{i,1}^\kappa, \alpha_{i,2}^\kappa$ and bias g_i^κ to $\neg(\alpha_{i,1}^\kappa \wedge \alpha_{i,2}^\kappa)$.

This completes our construction. We proceed by showing properties of Nash-stable outcomes. Each of the following six lemmas should be read with the implicit clause: “In every Nash-stable outcome.”

Lemma 3. *Let $\kappa \in \{0, 1\}$, then the following holds for all $i \in [n]$:*

- (a) *If $d^\kappa = 0$ then w_i^κ is indifferent w.r.t. edges of type 5.*
- (b) *If $d^\kappa = 1$ then $\eta_i^\kappa = w_i^\kappa$.*

Lemma 4. *If g_i^κ is incorrect then $z_i^\kappa = 1$. If $z_i^\kappa = 1$ then $z_j^\kappa = 1$ for all $0 \leq j \leq i$ and $y^\kappa = 0$.*

Lemma 5. *If $z_{i+1}^\kappa = 1$ then the inputs $I_1(g_i^\kappa)$ and $I_2(g_i^\kappa)$ are indifferent with respect to the type 1 edges of gate g_i^κ .*

Lemma 6. *Suppose $z_{i+1}^\kappa = 0$ and $z_i^\kappa = 1$ for some index $1 \leq i \leq N$.*

- (a) *If g_i^κ is correct then $\gamma_{i,1}^\kappa = \gamma_{i,2}^\kappa = 0$ and $\gamma_{i,3}^\kappa = 1$.*
- (b) *If g_i^κ is not correct then g_i^κ is indifferent w.r.t. edges of type 1 but w.r.t. the edges only in type 8 deviating would improve her happiness.*

Lemma 7. *If $d^\kappa = 1$ and $d^{\bar{\kappa}} = 0$ then for all $1 \leq i \leq 2n$, node g_i^κ is indifferent w.r.t. edges in type 4.*

Lemma 8. *Suppose $d^\kappa = 1$ and $d^{\bar{\kappa}} = 0$.*

- (a) *If z^κ is in the COMPUTE regime then $z_i^\kappa = 0$ for all $0 \leq i \leq N + 1$ and $y^\kappa = 1$.*
- (b) *If z^κ is in the RESET regime then $z_i^\kappa = 1$ for all $0 \leq i \leq N + 1$ and $y^\kappa = 0$.*

We now continue with the proof of Theorem 3. Suppose we are in a Nash-stable outcome of the party affiliation game. For our proof we assume $C(v^0) \geq$

$C(v^1)$. We will point out the small differences of the other case afterwards. Since $C(v^0) \geq C(v^1)$, z^0 is in the COMPUTE regime, i.e., all z_i^0 are biased to 0 and y^0 is biased to 1 (by type 6). Thus, $z_{N+1}^0 = 0$.

The remainder of the proof splits depending on the coalition of z_1^0 and z_1^1 . By Lemma 4 we know that $z_1^k = 0$ implies that all gates in C^k are correct.

$z_1^0 = 1$: By Lemma 4 we have $z_0^0 = 1$ and $y^0 = 0$. If $d^0 = d^1 = 0$ then d^0 is better off changing to 1 (by inspection of type 3 edges). If $d^0 = 1$ then Lemma 8(a) implies $z_1^0 = 0$, a contradiction. If $d^1 = 1$ and z^1 is in the RESET regime then by Lemma 8(b) and Lemma 5, v^1 is indifferent w.r.t. type 1 edges. Thus $v^1 = \eta^0$. But then either condition (ii) or (iii) for putting z^1 in the COMPUTE regime (cf. type 6) are fulfilled. So z^1 has to be in the COMPUTE regime. Lemma 8(a) then implies $z_1^1 = 0$. But then the neighbourhood of d^1 in type 3 is dominated by 0, a contradiction to $d^1 = 1$.

$z_1^0 = 0$ and $z_1^1 = 1$: By Lemma 4 we have $z_0^1 = 1$ and $y^1 = 0$. Since $C(v^0) \geq C(v^1)$ we know that z^0 is in the COMPUTE regime. So $z_1^0 = 0$ enforces $z_0^0 = 0$ and $y^0 = 0$. By inspection of type 3 edges we have $d^0 = 0$ and thus $d^1 = 1$. First assume that z^1 is in the RESET regime, then $z_i^1 = 1$ for all $0 \leq i \leq N + 1$ and Lemma 5 says that the inputs of all gates g_i^1 are indifferent w.r.t. type 1 edges. In particular this holds for $v^1 = (v_i^1)_{i \in [n]}$, so $v^1 = \eta^0$. By Lemma 3(b), $\eta^0 = w^0$. Since $z_1^0 = 0$, C^0 is computing correctly and thus $w^0 = w(v^0)$. Combining this we get $v_1 = w(v^0)$ which contradicts our assumption that z^1 is in the RESET regime. Thus z^1 is in the COMPUTE regime. Since $d^1 = 1$ we can apply Lemma 8(a) to conclude $z_1^1 = 0$, a contradiction.

$z_1^0 = 0$ and $z_1^1 = 0$: By Lemma 4 we have $z_0^0 = z_0^1 = 0$ and $y^0 = y^1 = 1$. Moreover we know that both circuits are computing correctly. If $d^0 = 1$ then $d^1 = 0$ and d^0 is indifferent w.r.t. type 3 edges. Since both circuits are computing correctly and $C(v^0) \geq C(v^1)$, the type 4 edges enforce $d^0 = 0$. But then d^1 is indifferent w.r.t. type 3 edges and the type 4 edges enforce $d^1 = 1$. So, $d^0 = 0$ and $d^1 = 1$. If z^1 is in the RESET regime then Lemma 8(b) gives $z_1^1 = 1$, a contradiction. Thus, z^1 is in the COMPUTE regime. Since $d^1 = 1$ we can apply Lemma 3(b). This and the fact that C^0 is computing correctly implies $\eta^0 = w(v^0)$. So z^1 can only be in the COMPUTE regime if $v^1 = w(v^0)$. Since $C(v^0) \geq C(v^1)$ this implies that $v^0 = v^1$ is a local optimum for the circuit C .

This finishes the proof in case $C(v^0) \geq C(v^1)$. The case $C(v^0) < C(v^1)$ is completely symmetric except here the conclusion $v^0 = v^1$ in the very last sentence contradicts $C(v^0) < C(v^0)$. So this case can't happen in a local optimum.

Note that throughout the construction we made sure that no node is incident to more than one negative edge. This completes the proof of Theorem 3. \square

The instance produced by this reduction has the property that no node is indifferent between the two coalitions. This might be useful for other reductions.

Corollary 1. ONEENEMYPARTYAFFILIATION is PLS-complete even if restricted to instances where no player is ever indifferent between the two coalitions.

Theorem 3 and Lemma 1 together establish that ISWITHSWAPS is PLS-complete (Theorem 2). Throughout the proof we used only two coalitions. Since 2-IS can be solved in polynomial time (Proposition 2), a PLS-hardness result for IS would require more than two coalitions (unless $PLS \subseteq P$). We leave the complexity of IS as an interesting open problem.

References

- [1] Aarts, E.H.L., Lenstra, J.K.: *Local Search in Combinatorial Optimization*. Wiley Interscience, Hoboken (1997)
- [2] Balcan, M.-F., Blum, A., Mansour, Y.: Improved equilibria via public service advertising. In: *SODA*, pp. 728–737 (2009)
- [3] Ballester, C.: NP-completeness in hedonic games. *Games and Economic Behavior* 49(1), 1–30 (2004)
- [4] Bhalgat, A., Chakraborty, T., Khanna, S.: Approximating pure Nash equilibrium in cut, party affiliation and satisfiability games. In: *ACM Conference in Electronic Commerce (EC)*, pp. 132–146 (2010)
- [5] Bogomolnaia, A., Jackson, M.O.: The stability of hedonic coalition structures. *Games and Economic Behavior* 38(2), 201–230 (2002)
- [6] Burani, N., Zwicker, W.S.: Coalition formation games with separable preferences. *Mathematical Social Sciences* 45(1), 27–52 (2003)
- [7] Ceclároová, K.: Stable partition problem. In: *Encyclopedia of Algorithms*. Springer, Heidelberg (2008)
- [8] Christodoulou, G., Mirrokni, V.S., Sidiropoulos, A.: Convergence and approximation in potential games. In: *STACS*, pp. 349–360 (2006)
- [9] Dreze, J.H., Greenberg, J.: Hedonic coalitions: Optimality and stability. *Econometrica* 48(4), 987–1003 (1980)
- [10] Elkind, E., Wooldridge, M.: Hedonic coalition nets. In: *Int. Conference on Autonomous Agents and Multiagent Systems*, pp. 417–424 (2009)
- [11] Greenberg, J.: Coalition structures. In: Aumann, R.J., Hart, S. (eds.) *Handbook of Game Theory with Economic Applications II*. Elsevier, Amsterdam (1994)
- [12] Johnson, D.S., Papadimitriou, C.H., Yannakakis, M.: How easy is local search? *Journal of Computer and System Sciences* 37, 79–100 (1988)
- [13] Monien, B., Tscheuschner, T.: On the power of nodes of degree four in the local max-cut problem. In: *7th International Conference on Algorithms and Complexity, CIAC* (2010)
- [14] Monien, B., Dumrauf, D., Tscheuschner, T.: Local search: Simple, successful, but sometimes sluggish. In: *37th International Colloquium on Automata, Languages and Programming, ICALP* (2010)
- [15] Orlin, J.B., Punnen, A.P., Schulz, A.S.: Approximate local search in combinatorial optimization. *SIAM Journal of Computing* 33(5), 1201–1214 (2004)
- [16] Schäffer, A.A., Yannakakis, M.: Simple Local Search Problems that are Hard to Solve. *SIAM Journal of Computing* 20(1), 56–87 (1991)
- [17] Sipser, M.: *Introduction to the Theory of Computation*. Thomson (2006)
- [18] Sung, S.C., Dimitrov, D.: Computational complexity in additive hedonic games. *European Journal of Operational Research* 203(3), 635–639 (2010)
- [19] Tscheuschner, T.: The local max-cut problem is PLS-complete even on graphs with maximum degree five (2010), <http://arxiv.org/abs/1004.5329>

A Perfect Price Discrimination Market Model with Production, and a (Rational) Convex Program for It*

Gagan Goel and Vijay Vazirani

Georgia Institute of Technology, Atlanta
{gagang,vazirani}@cc.gatech.edu

Abstract. Recent results showed PPAD-completeness of the problem of computing an equilibrium for Fisher’s market model under additively separable, piecewise-linear, concave utilities. We show that introducing perfect price discrimination in this model renders its equilibrium polynomial time computable. Moreover, its set of equilibria are captured by a convex program that generalizes the classical Eisenberg-Gale program, and always admits a rational solution.

We also introduce production into our model; our goal is to carve out as big a piece of the general production model as possible while still maintaining the property that a single (rational) convex program captures its equilibria, i.e., the convex program must optimize individually for each buyer and each firm.

1 Introduction

The search for efficient algorithms for computing market equilibria started with much interest within theoretical computer science about a decade ago. The goal was not only academic, i.e., providing an algorithmic ratification of Adam Smith’s “invisible hand of the market,” but was also motivated by potential applications to the plethora of new and highly lucrative markets that have emerged on the Internet.

This study started with the simple case of linear utility functions, for which polynomial time algorithms were obtained [7, 11], and gradually moved on to more general and realistic utility functions. However, the latter program had limited success (most notably, an efficient algorithm for approximating equilibria for the Fisher model under Leontief utilities [6, 21]), and was recently dealt a serious blow, with results showing that the problem of computing an equilibrium under even additively separable, piecewise-linear, concave utilities (plc utilities) is PPAD-complete for both Arrow-Debreu and Fisher market models [4, 5, 19]. Assuming $P \neq \text{PPAD}$, this effectively rules out the existence of efficient algorithms for almost all general and interesting classes of “traditional” market models.

* Research supported by NSF Grants CCF-0728640 and CCF-0914732, ONR Grant N000140910755, and a Google Research Grant.

On the other hand, markets in the West, based on Adam Smith’s free market principle, seem to do a good job of finding prices that maintain parity between supply and demand¹. This has prompted the question (see [18]) of whether we have failed to capture some essential elements of real markets in our models, and what is the “right” model which is not only realistic and admits equilibria but is also amenable to efficient computation of equilibria.

In this context, we show that Fisher markets with plc utilities can be rendered computationally efficient by introducing perfect price discrimination². Additionally, we introduce firms into our model which act as sellers of goods, suppliers of labor and producers of goods, or any combination of these activities. These firms have initial endowments of goods and labor, and their goal is to maximize profits by optimally producing and selling goods. Traditionally in economics, and in the model studied by Arrow and Debreu [2], production satisfies non-increasing returns to scale. This is the case in our model as well, though it is imposed in a “piecewise-linear manner”. As a consequence of linearity, for given prices of goods, the optimal operation of a firm is captured by an LP.

We show that equilibrium production and allocation for this market model is captured via a single convex program, a generalization of the classic Eisenberg-Gale convex program. The optimal dual of this program yields equilibrium prices. Two interesting theorems are the following. First, each buyer gets a utility maximizing bundle of goods under the rules of price discrimination we have assumed. Second, for each firm, the operation specified by the optimal solution to this convex program is also the optimal solution for its LP, with the equilibrium prices substituted into it. The idea behind the latter theorem goes back to [12]; however, our model of production is considerably more general than that in [12].

The notion of a *rational convex program* was introduced recently in [16], i.e., a nonlinear convex program that always has a rational solution of polynomial bit size, if all its parameters are rational numbers. Starting with the celebrated Eisenberg-Gale program, several convex programs arising in mathematical economics and game theory are now known to be rational, see details in [17].

We prove that the program capturing our market model is also rational. In particular, this implies that the ellipsoid method will yield the exact equilibrium in polynomial time [10]. For the special case that firms act only as sellers of goods, in the full paper, we will give a combinatorial polynomial time algorithm as well. We will also generalize this model and assume that buyers have utility for money, given by a piecewise-linear, concave function for each buyer. Now, at equilibrium, a buyer may choose to not spend all of her money. We show how to extend our polynomial time algorithm to this case as well. The solution still turns out to be rational; however, we do not know of a convex program

¹ For example, in the West, it is hard to see a sight that was commonplace in the Soviet Union, with massive surpluses of some goods and empty shelves of others.

² The model described in this paper was obtained in the process of attempting the open problem, posed in [16], of obtaining a combinatorial algorithm for solving the extension of game **ADNB** to plc utilities; linear utilities were assumed in **ADNB**. Combinatorial insights obtained in the process led to the model.

that captures this enhanced model. By exploiting the combinatorial structure discovered in obtaining our algorithm, we will give a characterization of the entire set of equilibria of this market. This will reveal the range of equilibrium prices of each good and the range of profit that the middleman can accrue from each buyer.

Using our convex program for the model with production, in the full paper, we obtain surprisingly simple proofs of both welfare theorems – simpler even than those for a normal Arrow-Debreu market model. We now give the key idea behind our convex program. First, consider the simpler version of our model in which firms are simply sellers of goods. Our convex program for this model is the “natural” generalization of the Eisenberg-Gale program from linear to plc utilities; however, it does not capture equilibria for Fisher markets with these utility functions. The latter statement follows directly from the observation that equilibria in the latter market model can be disconnected, however, the optimal solutions of a convex program form a convex set. Intuitively, introducing price discrimination in this market model renders the set of equilibria convex. The situation is somewhat analogous to that of Nash equilibrium. The set of optimal strategies of the pure equilibria of a bimatrix game can be disconnected; however, by introducing mixed strategies, it is rendered a convex set and hence suitable for study with the Kakutani fixed point theorem.

Regarding extending the model to include production, our main goal was to carve out as big a piece of the general production model as possible while maintaining the property stated above, i.e., that a single convex program optimizes individually for each firm, and at the same time, all constraints of the convex program are linear (this leads to a proof of its rationality). We leave the open problem of extending our combinatorial algorithm to the model with production.

2 The Market Model

2.1 Perfect Price Discrimination

Most businesses today charge different prices from different consumers for essentially the same goods or services in order to maximize their revenues. This practice is called *price discrimination*. It is not only widespread but also essential for survival of certain businesses, e.g., in the airline industry. Price discrimination has been extensively studied in economics from many different angles; see [20, 15, 14, 9, 8, 1, 3, 13] for just a small sampling of papers on this topic.

A monopolistic situation in which the business separates the market into individual consumers and charges each one *prices that they are willing and able to pay* is called *perfect price discrimination*. Of course, the business needs to have complete information about each consumer’s preferences. An interesting feature of our model is that in it, it is the consumers who decide at what *rate* they want utility.

Our market model consists of buyers, firms and a middleman. As stated in the Introduction, buyers have initial endowments of money with which they wish to buy goods and maximize the utility accrued, and firms act as sellers of goods,

suppliers of labor, producers of goods or any combination of these activities. The firms have initial endowments of goods and labor, and their goal is to produce and sell goods in a way that maximizes their profits at current prices of goods.

The middleman buys goods from the firms, which charge the middleman in accordance with the prices set and the amounts bought. As stated above, buyers decide at what rate they want utility, and the middleman sells to them goods accordingly, with the only condition that he never sells any part of any good at a loss – the fact that the middleman knows the buyer’s utility function enables him to do this. We show below that under these circumstances, for any given prices of goods, there is a unique optimal rate for each buyer. This is also the rate that ensures that the marginal utility accrued by the buyer per unit of money spent is equal to the marginal cost of the goods she is receiving, as desired under perfect price discrimination.

Thus, in our model, the elasticity among consumers leads to profit for the middleman. If all buyers had linear utility functions for goods, then the middleman will make no profit. In our model we assume that the buyers have *plc* utilities.

Prices for goods and labor are said to be *equilibrium prices* if with optimal operation of each firm at these prices and optimal rates for each buyer, the market clears, i.e., all the goods get sold to buyers and all of their money gets spent.

The set of all goods and types of labor in the system is denoted by G , $|G| = n_G$. Each good is assumed to be divisible, as is each unit of labor. Let B denote the set of buyers, $|B| = n_B$, and F denote the set of firms, $|F| = n_F$. Assume that the buyers are numbered from 1 to n_B and are indexed by i , goods are numbered from 1 to n_G and are indexed by j , and the firms are numbered from 1 to n_F and are indexed by f .

An application to online display advertising marketplaces. Currently, there are companies that sell ad slots on web sites to advertisers. In keeping with our model, we will view such a company as the middleman, the owners of web sites as sellers and the advertisers as buyers. We will view ad slots on different web sites as different items, which need to be priced. An advertiser’s utility for a particular ad slot is determined by the probability that her ad will get clicked if it is shown on that slot; her total utility is additive over all the slots she is allocated. Advertisers typically pay at fixed rate to the middleman for the expected number of clicks they get, i.e., they are paying at fixed rate for every unit of utility they get. Using knowledge of the utility function of buyers, the middleman is able to price discriminate. Clearly, this setup is captured by our model.

2.2 The Buyers and Their Utility Functions

Let $m_i \in \mathbf{Q}^+$ dollars denote the initial amount of money possessed by buyer $i \in B$. For each buyer i and good j we are specified a function $f_j^i : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ which gives the utility that i derives as a function of the amount of good j that she receives.

Each function f_j^i is a non-negative, non-decreasing, piecewise-linear, concave function. The overall utility of buyer i , $u_i(\mathbf{x})$ for a bundle $\mathbf{x} = (x_1, \dots, x_g)$ of goods, is additively separable over the goods, i.e., $u_i(\mathbf{x}) = \sum_{j \in G} f_j^i(x_j)$.

We will call each piece of f_j^i a *segment*. Number the segments of each function in order of decreasing slope. Let s_{ijk} , $k = 1, 2, \dots$ denote the k th segment of f_j^i . Let l_{ijk} denote the amount of good j represented by segment s_{ijk} and u_{ijk} denote the rate at which buyer i accrues utility per unit of good j received, when she is getting an allocation corresponding to segment s_{ijk} . Clearly, the maximum utility she can receive corresponding to segment s_{ijk} is $u_{ijk} \cdot l_{ijk}$. We will assume that u_{ijk} and l_{ijk} are rational numbers. Let S_{ij} denote the set of segments of function f_j^i and let S_i denote the set of all segments of buyer i , i.e., $S_i = \cup_{j=1}^g S_{ij}$.

2.3 The Middleman and Determining Buyers' Rates

Assume that the prices of goods are set at $\mathbf{p} = (p_1, \dots, p_g)$. Define the *bang-per-buck* of segment s_{ijk} to be u_{ijk}/p_j and denote it by $\text{bpb}(s_{ijk})$; clearly, this is the amount of utility accrued by i per dollar spent for an allocation corresponding to segment s_{ijk} .

Suppose buyer i fixes her rate at r_i which is the amount of utility she wants per dollar. Then, for an allocation corresponding to segment s_{ijk} , the middleman is effectively charging the buyer $\frac{u_{ijk}}{r_i}$ dollars per unit of j . In particular, if $\text{bpb}(s_{ijk}) < r_i$, then the middleman will be allocating this segment at a loss, i.e., at a price smaller than p_j dollars per unit of j . Moreover, the larger $\text{bpb}(s_{ijk})/r_i$ is, the higher is the profit the middleman can make from allocations corresponding to this segment. Therefore, once i announces her rate, the middleman removes from consideration all segments $s \in S_i$ such that $\text{bpb}(s) < r_i$, and allocates to i goods corresponding to segments that gives him the highest profit, until i exhausts her money.

Now, given how the middleman responds to prices and rates, what rate maximizes the utility of a buyer i ? We will define a rate r_i^* , which we will call the *optimal* rate of buyer i , as a function of prices \mathbf{p} , and show that this rate maximizes buyer i 's utility. The overall objective is to find prices for goods such that if the buyers report their optimal rates, the market clears under above transactions, i.e., there is no surplus or deficiency of any good. This is our notion of *equilibrium* for the market.

Rate r_i^* is obtained as follows. Sort all segments in S_i by decreasing bang-per-buck and start with a sufficiently large number α . Consider all segments $s \in S_i$ such that $\text{bpb}(s) \geq \alpha$, and add up their total utilities. We will denote this by $t(\alpha)$, i.e., $t(\alpha) = \sum_{s \in S_i: \text{bpb}(s) \geq \alpha} \text{utility}(s)$. Now the cost of buying goods corresponding to all these segments at rate α is $t(\alpha)/\alpha$. When α is very large, this will be less than m_i . Observe that as α is decreased, this number increases monotonically. Now r_i^* is the largest value of α such that this number is $\geq m_i$.

Formally,

$$r_i^*(\mathbf{p}) = \arg \max_{\alpha} \left\{ \frac{t(\alpha)}{\alpha} \geq m_i \right\}.$$

We will denote $r_i^*(\mathbf{p})$ by simply r_i^* when its meaning is clear from the context. The following lemma is straightforward.

Lemma 1. *Rate r_i equals r_i^* if and only if when buyer i picks r_i as her rate, each segment s such that $\text{bpb}(s) > r_i$ is fully allocated to her, and corresponding to segments s such that $\text{bpb}(s) = r_i$, i is allocated just the right amount of goods so that her total utility adds up to $r_i \cdot m_i$.*

Lemma 2. *For any prices \mathbf{p} , rate r_i^* maximizes the utility for buyer i .*

Proof. If the rate is fixed at $\alpha < r_i^*$, then $\alpha \cdot m_i < r_i^* \cdot m_i$ and therefore i will be allocated smaller utility. Next consider fixing the rate at $\beta > r_i^*$. Let s be the smallest bang-per-buck of an allocated segment at rate r_i^* . If $\text{bpb}(s) = r_i^*$, then at rate β she will accrue smaller utility. Otherwise, for $\text{bpb}(s) \geq \beta \geq r_i^*$, i will still be allocated $r_i^* \cdot m_i$ utility and for $\beta > \text{bpb}(s)$, she will accrue strictly smaller utility. This proves the lemma.

2.4 The Firms and Their Capabilities

Our model allows firms to have a rich set of capabilities, in particular, allowing them to model non-increasing returns to scale as was assumed in the Arrow-Debreu model. For sake of clarity, we first present the model assuming constant returns to scale. In this model, each firm $f \in F$ has variables y_{jf} corresponding to each good $j \in G$ which represent the amount of this good that it sells or buys in the market; y_{jf} is positive if the firm sells good j , negative if it buys it, and zero otherwise. The objective of the firm is to maximize its profit, which at prices $\mathbf{p} = (p_1, \dots, p_g)$ will be $\sum_{j \in G} p_j \cdot y_{jf}$. Let c_{jf} denote this firm's initial endowment of good j . In our model there is no need to partition the goods into raw materials and manufactured goods or to differentiate between goods and labor.

In order to formally state the various production processes of this firm, we will use auxiliary variables which are local to this firm. We will denote these by z_{lf} , i.e., they are indexed by l . The constraints imposed on production are all assumed to be linear and are indexed by m . Thus the set of constraints for firm f are:

$$\forall m : \sum_{j \in G} a_{jf}^m \cdot y_{jf} + \sum_l b_{lf}^m \cdot z_{lf} \leq d_f^m,$$

where a_{jf}^m , b_{lf}^m and d_f^m are constants determined by the production processes of firm f . In particular, some of the d_f^m 's may be the initial endowments, i.e., c_{jf} 's. The variables y_{jf} are unconstrained; however, z_{lf} 's are constrained to be

nonnegative. Clearly, the optimal operation of a firm can be stated as a linear program.

Next, we give some illustrative examples. First, consider a firm that does not produce anything but only acts as a seller. It sells its initial endowment of goods in the market at the going prices. Clearly, its LP only needs constraints of the form $y_{jf} \leq c_{jf}$.

Second, consider a firm that has an initial endowment, c_{1f} of good 1, and is able to produce goods 8 and 9. However, for this, it will need to buy goods 2 and 3. The production also requires good 1. If the amount of good 1 needed for production is less than c_{1f} , firm f sells the excess in the market, and if it is greater than c_{1f} , firm f will need to buy additional amounts of good 1 from the market. Assume that good 8 is produced using goods 1 and 2, and that good 9 is produced using goods 1 and 3. However, goods 8 and 9 are produced via qualitatively different processes. To produce a unit of good 8, the firm uses up α units of good 1 and β units of good 2. On the other hand, good 9 can be produced using either good 1 or good 3, with a unit of good 1 producing γ units of good 9 and a unit of good 3 producing δ units of good 9. These production constraints are captured by the following linear constraints, using auxiliary variables z_{1f} and z_{2f} : $y_{1f} + z_{1f} + z_{2f} \leq c_{1f}$, $y_{8f} \leq \alpha \cdot z_{1f}$, $y_{8f} \leq \beta \cdot y_{2f}$, and $y_{9f} \leq \gamma z_{2f} + \delta \cdot y_{3f}$. The objective of this firm is to maximize $p_1 \cdot y_{1f} + p_2 \cdot y_{2f} + p_3 \cdot y_{3f} + p_8 \cdot y_{8f} + p_9 \cdot y_{9f}$.

Next, we introduce non-increasing returns to scale in our model, though in a “piecewise-linear” manner. Thus the production of good j by firm f is partitioned into schedules, as a function of the amount of j produced. The schedules are indexed by r . Let y_{jfr} denote the amount of good j produced in the r th schedule and let ρ_{jf} denote the total number of schedules for producing good j in firm f . For each schedule, possibly other than the last one, there is a bound on the amount of good that can be produced in that schedule, i.e., a constraint of the form $y_{jfr} \leq \alpha$, for some constant α . Each raw material and labor required is non-decreasing as a function of the schedule, so that the earlier schedules produce goods at higher profits. The enhanced constraints now required are:

$$\forall m : \sum_{j \in G} a_{jf}^m \cdot y_{jf} + \sum_{j \in G, r \leq \rho_{jf}} e_{jfr}^m \cdot y_{jfr} + \sum_l b_{lf}^m \cdot z_{lf} \leq d_f^m,$$

where e_{jfr}^m 's are constants. Since the overall goal of the firm is to maximize profit, it will produce good j up to capacity in earlier schedules before starting production in the next schedule.

3 A Rational Convex Program for the Fixed Supply Case

In this section we give a convex program whose optimal primal and dual variables capture the equilibrium prices \mathbf{p} and rates \mathbf{r} for the price discrimination market model when the goods are given in a fixed supply. Then we will show the existence of equilibrium prices and rates that can be represented using rational numbers.

Let x_{ijk} denote the amount of good j that is allocated to buyer i from the k^{th} segment s_{ijk} of S_{ij} . Consider the following convex program:

$$\begin{aligned}
 &\text{maximize} && \sum_{i \in B} m_i \log(u_i) && (1) \\
 &\text{subject to} && \forall i \in B : u_i = \sum_{j \in G} \sum_{k \in S_{ij}} u_{ijk} x_{ijk} \\
 &&& \forall j \in G : \sum_{i \in B} \sum_{k \in S_{ij}} x_{ijk} \leq 1 \\
 &&& \forall i \in B, \forall j \in G, \forall k \in S_{ij} : x_{ijk} \leq l_{ijk} \\
 &&& \forall i \in B, \forall j \in G, \forall k \in S_{ij} : x_{ijk} \geq 0
 \end{aligned}$$

Here the first constraint is ensuring that u_i is the total utility of buyer i , the second constraint is saying that the total good sold should not be more than one, and the third constraint is saying that the amount allocated from each segment should not exceed its size.

Let p_j be the dual variable corresponding to good j in the second set of constraints above. We will prove the following theorem in this section:

Theorem 1. *Prices \mathbf{p} are equilibrium prices if and only if they form an optimal dual solution to convex program [A](#).*

We will make two mild **assumptions**:

1. For every good j , $\sum_{i,k: u_{ijk} > 0} l_{ijk} > 1$; that is, the supply of every good is limited w.r.t the total demand of the buyers if there were no prices.
2. Each buyer i desires some good; that is, $u_{ijk} > 0$ for some segment s_{ijk} of every buyer i .

Note that, because of the 2nd assumption, in the optimal solution of the above convex program, $u_i > 0$ for every buyer i . Also the optimal solution satisfies the following property: $x_{ijk} > 0 \Rightarrow x_{ijt} = l_{ijt} \ \forall t < k$. This is so because if the property is not true, we can transfer some quantity from the segment s_{ijk} to a segment s_{ijt} (for some $t < k$) and get a strictly higher objective function value. Thus the final allocation obtained is a *valid* allocation.

The KKT conditions for the above convex program are:

- (1) $\forall j \in G : p_j \geq 0$,
- (2) $\forall i \in B, \forall j \in G, \forall k \in S_{ij} : q_{ijk} \geq 0$,
- (3) $\forall j \in G : p_j > 0 \Rightarrow \sum_{i \in B} \sum_{k \in S_{ij}} x_{ijk} = 1$.
- (4) $\forall i \in B, \forall j \in G, \forall k \in S_{ij} : q_{ijk} > 0 \Rightarrow x_{ijk} = l_{ijk}$,
- (5) $\forall i \in B, \forall j \in G \forall k \in S_{ij} : p_j + q_{ijk} \geq \frac{m_i \cdot u_{ijk}}{u_i}$.
- (6) $\forall i \in B, \forall j \in G \forall k \in S_{ij} : x_{ijk} > 0 \Rightarrow p_j + q_{ijk} = \frac{m_i \cdot u_{ijk}}{u_i}$.

We will call p_j to be the *price* of good j , and q_{ijk} to be the *price differential*, which is unique for each buyer i , good j , and segment $k \in S_{ij}$. Also define the rate of a buyer i (r_i) to be $\frac{u_i}{m_i}$. Note that from equation (6), for any segment s_{ijk} for which x_{ijk} is positive: $r_i = \frac{u_i}{m_i} = \frac{u_{ijk}}{p_j + q_{ijk}} \leq \frac{u_{ijk}}{p_j}$. Thus if a segment s is allocated, fully or partially, to buyer i , its bang-per-buck value is at least r_i .

Lemma 3. *Corresponding to prices \mathbf{p} given by the KKT conditions, the rate $r_i = \frac{u_i}{m_i}$ is optimal for each buyer i .*

Proof. Suppose that for some segment s_{ijk} , $\frac{u_{ijk}}{p_j} > r_i = \frac{u_i}{m_i}$, then from the 5th KKT condition, we get that q_{ijk} is strictly positive, which from the 4th KKT condition implies that $x_{ijk} = l_{ijk}$. Thus, from lemma 1 and lemma 2, we get that the rate $r_i = \frac{u_i}{m_i}$ is optimal for each buyer i .

Lemma 4. *Under assumption 1, for every good j , p_j is strictly positive and j is exactly sold, i.e., $\sum_{i,k} x_{ijk} = 1$.*

Proof. Suppose that the price of some good j is zero. Then from the 5th condition above, $q_{ijk} > 0$ for every segment s_{ijk} for which $u_{ijk} > 0$. Thus along with 4th condition, this will imply that $\sum_{i,k} x_{ijk} > 1$ which violates the constraint $\sum_{i,k} x_{ijk} \leq 1$ in the convex program. Thus the price of every good is strictly positive. Using 3rd condition, this implies that every good is completely sold, i.e., $\sum_{i,k} x_{ijk} = 1$.

From the above observations, finding an optimal solution of the above convex program (1) is equivalent to finding a price vector \mathbf{p} , a rate vector \mathbf{r} , and allocations of the goods to the buyers (vector \mathbf{x}) that satisfy the following equilibrium conditions.

1. For prices \mathbf{p} , the rate $r_i = \frac{u_i}{m_i}$ is optimal for each buyer i . Moreover, since $u_i > 0$, buyer i spends his money completely.
2. No portion of a segment s is sold to a buyer i , if $bpb(s) < r_i$.
3. All goods are sold out completely.

Thus, if prices \mathbf{p} are optimal dual variables of the above convex program then these prices are also equilibrium prices.

Now, suppose we are given equilibrium prices \mathbf{p} ; we will show that these prices are also optimal dual variables. Given equilibrium prices, there exists optimal rates \mathbf{r} of buyers and allocation of goods \mathbf{x} to the buyers, so that a segment s is allocated to buyer i only if $bpb(s) \geq r_i$. Moreover, if $bpb(s) > r_i$, then the segment s is fully allocated to buyer i . We will show the existence of variables q_{ijk} 's such that the KKT conditions are satisfied: if $\frac{u_{ijk}}{p_j} \geq \frac{u_i}{m_i}$, set q_{ijk} so that $\frac{u_{ijk}}{p_j + q_{ijk}} = \frac{u_i}{m_i}$, else set $q_{ijk} = 0$. It is not difficult to see that these q_{ijk} 's along with \mathbf{p} , \mathbf{r} , and \mathbf{x} satisfies the KKT conditions. Thus prices \mathbf{p} are optimal dual variables. This finishes the proof of theorem 1.

Theorem 2. *If all the utilities u_{ijk} 's are rational, then there exists prices which are rational and the rates are rational. Moreover, they can be written using polynomially many bits in the length of the instance.*

4 Introducing Production into the Convex Program

In this section we will introduce production into the convex program. Instead of assuming that there is a fixed supply of goods, we will assume that the goods

are produced by firms as discussed in section 1 and 2. We will show that a single convex program can still be used to simultaneously optimize for each firm's objective function. For the ease of exposition, we will first assume constant returns to scale. All the results can easily be extended for non-increasing returns to scale. Suppose given prices \mathbf{c} , firm f is optimizing the following linear program:

$$\begin{aligned} &\text{maximize} && \sum_{j \in G} c_j \cdot y_{jf} && (2) \\ &\text{subject to} && \forall m : \sum_{j \in G} a_{jf}^m \cdot y_{jf} + \sum_l b_{lf}^m \cdot z_{lf} \leq d_f^m \end{aligned}$$

Let x_{ijk} denote the amount of good j which is allocated to buyer i corresponding to the k^{th} segment s_{ijk} of S_{ij} . Consider the following convex program:

$$\begin{aligned} &\text{maximize} && \sum_{i \in B} m_i \log(u_i) && (3) \\ &\text{subject to} && \forall i \in B : u_i = \sum_{j \in G} \sum_{k \in S_{ij}} u_{ijk} x_{ijk} \\ &&& \forall j \in G : \sum_{i \in B} \sum_{k \in S_{ij}} x_{ijk} \leq \sum_{f \in F} y_{jf} \\ &&& \forall f \in F, \forall m : \sum_{j \in G} a_{jf}^m \cdot y_{jf} + \sum_l b_{lf}^m \cdot z_{lf} \leq d_f^m \\ &&& \forall i \in B, \forall j \in G, \forall k \in S_{ij} : x_{ijk} \leq l_{ijk} \\ &&& \forall i \in B, \forall j \in G, \forall k \in S_{ij} : x_{ijk} \geq 0 \\ &&& \forall f \in F, \forall l : z_{lf} \geq 0 \end{aligned}$$

Here the first constraint is ensuring that u_i is the total utility of buyer i , the second constraint is ensuring that the total amount of any good sold to the buyers should not be more than what is produced by the firms, the third constraint is capturing the production constraints of the firms, and the fourth constraint is saying that the amount allocated in each segment should not exceed its size.

Our main theorem is the following:

Theorem 3. *Prices \mathbf{p} are equilibrium prices if and only if they form an optimal dual solution to convex program (3). Moreover, equilibrium production is captured by an optimal solution to primal variables y_{jk} 's.*

To prove the above theorem, we will again consider KKT equations. We will show that the KKT equations have two different components, one which corresponds to optimization of the buyers as was shown in previous section and other which

corresponds to optimization of the firms. Following are the KKT conditions of convex program (3).

- (1) $\forall j \in G : p_j \geq 0$,
- (2) $\forall f \in F, \forall m : \alpha_f^m \geq 0$,
- (3) $\forall i \in B, \forall j \in G, \forall k \in S_{ij} : q_{ijk} \geq 0$,
- (4) $\forall j \in G : p_j > 0 \Rightarrow \sum_{i \in B} \sum_{k \in S_{ij}} x_{ijk} = \sum_{f \in F} y_{jf}$.
- (5) $\forall i \in B, \forall j \in G, \forall k \in S_{ij} : q_{ijk} > 0 \Rightarrow x_{ijk} = l_{ijk}$,
- (6) $\forall f \in F, \forall m : \alpha_f^m > 0 \Rightarrow \sum_{j \in G} a_{jf}^m \cdot y_{jf} + \sum_l b_{lf}^m \cdot z_{lf} = d_f^m$,
- (7) $\forall i \in B, \forall j \in G \forall k \in S_{ij} : p_j + q_{ijk} \geq \frac{m_i \cdot u_{ijk}}{u_i}$.
- (8) $\forall i \in B, \forall j \in G \forall k \in S_{ij} : x_{ijk} > 0 \Rightarrow p_j + q_{ijk} = \frac{m_i \cdot u_{ijk}}{u_i}$.
- (9) $\forall f \in F, \forall l : \sum_m b_{lf}^m \cdot \alpha_f^m \geq 0$
- (10) $\forall f \in F, \forall l : z_{lf} > 0 \Rightarrow \sum_m b_{lf}^m \cdot \alpha_f^m = 0$
- (11) $\forall j \in G, \forall f \in F : \sum_m a_{jf}^m \cdot \alpha_f^m = p_j$

As earlier, we will call p_j to be the *price* of good j , and q_{ijk} to be the *price differential* which is unique for each buyer i , good j , and segment $k \in S_{ij}$.

Lemma 5. *Convex program (3) optimizes each firm's profit at equilibrium.*

Since the supply of goods is not fixed, we won't make an assumption similar to the first assumption in the previous section. Instead, we will work with a slightly different, but standard, definition of equilibrium that if a good is not sold completely, then its price must be zero. Following lemma easily follows from the 4th KKT condition.

Lemma 6. *For every good j , either price p_j is zero or j is exactly sold, i.e., $\sum_{i,k} x_{ijk} = 1$. For prices \mathbf{p} given by the KKT conditions, the rate $r_i = \frac{u_i}{m_i}$ is optimal for each buyer i .*

Thus we have shown that optimal solution of the convex program satisfies the following equilibrium conditions:

1. For prices \mathbf{p} , the rate $r_i = \frac{u_i}{m_i}$ is optimal for each buyer i . Moreover, since $u_i > 0$, buyer i spends his money completely.
2. No portion of a segment s is sold to a buyer i , if $bpb(s) < r_i$.
3. If the price of some good is strictly positive, it is sold out completely.
4. Each firm's production optimizes its profit.

The proof of the other direction, that any equilibrium solution is also a solution to the convex program, is similar to that in Section 3. This completes the proof of Theorem 3. The following theorems will be proved in the full paper.

Theorem 4 (First Welfare). *The utilities accrued by buyers at equilibrium prices \mathbf{p} and rates \mathbf{r} are Pareto efficient.*

Theorem 5 (Second Welfare). *For any Pareto efficient utilities \mathbf{u}^* , there exists a choice of money vector of buyers under which equilibrium utilities are \mathbf{u}^* .*

References

- [1] Anderson, S.P., de Palma, A.: Spatial price discrimination with heterogeneous products. *The Review of Economic Studies* 55(4), 573–592 (1988)
- [2] Arrow, K., Debreu, G.: Existence of an equilibrium for a competitive economy. *Econometrica* 22, 265–290 (1954)
- [3] Bhaskar, V., To, T.: Is perfect price discrimination really efficient? An analysis of free entry. *The RAND Journal of Economics* 35(4), 762–776 (2004)
- [4] Chen, X., Dai, D., Du, Y., Teng, S.-H.: Settling the complexity of arrow-debreu equilibria in markets with additively separable utilities. In: *FOCS* (2009)
- [5] Chen, X., Teng, S.-H.: Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. *Journal of the ACM* 56(3) (2009)
- [6] Codenotti, B., Varadarajan, K.: Efficient computation of equilibrium prices for markets with Leontief utilities. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) *ICALP 2004*. LNCS, vol. 3142, pp. 371–382. Springer, Heidelberg (2004)
- [7] Devanur, N., Papadimitriou, C.H., Saberi, A., Vazirani, V.V.: Market equilibrium via a primal-dual-type algorithm. *JACM* 55(5) (2008)
- [8] Edlin, A., Epelbaum, M., Heller, W.: Surplus maximization and price discrimination in general equilibrium: Part I. Technical Report 9405, Centro de Investigacion Economica, ITAM (1994)
- [9] Heller, M., Edlin, W.P., Epelbaum, A.S.: Is perfect price discrimination really efficient? Welfare and existence in general equilibrium. *Econometrica* 66(4), 897–922 (1998)
- [10] Grotschel, M., Lovasz, L., Schrijver, A.: *Geometric Algorithms and Combinatorial Optimization*. Springer, Heidelberg (1988)
- [11] Jain, K.: A polynomial time algorithm for computing the Arrow-Debreu market equilibrium for linear utilities. In: *FOCS* (2004)
- [12] Jain, K., Vazirani, V.V., Ye, Y.: Market equilibrium for homothetic, quasi-concave utilities and economies of scale in production. In: *SODA* (2005)
- [13] Liu, Q., Serfes, K.: Imperfect price discrimination, market structure, and efficiency. *Canadian Journal of Economics* 38(4), 1191–1203 (2005)
- [14] Varian, H.: Differential pricing and efficiency. *First Monday* 1(2) (1996)
- [15] Varian, H.R.: Price discrimination and social welfare. *American Economic Review* 75(4), 870–875 (1985)
- [16] Vazirani, V.V.: Nash bargaining via flexible budget markets (submitted 2009)
- [17] Vazirani, V.V.: 2-player Nash and nonsymmetric bargaining via market equilibria. In: *SAGT* (2010)
- [18] Vazirani, V.V.: Spending constraint utilities, with applications to the Adwords market. *Mathematics of Operations Research* 35(2) (2010)
- [19] Vazirani, V.V., Yannakakis, M.: Market equilibria under separable, piecewise-linear, concave utilities. In: *Proceedings of The First Symposium on Innovations in Computer Science* (2010)
- [20] Norman, G., MacLeod, W.B., Thisse, J.-F.: Price discrimination and equilibrium in monopolistic competition. *International Journal of Industrial Organization* 6(4), 429–446 (1988)
- [21] Ye, Y.: Exchange market equilibria with Leontief’s utility: Freedom of pricing leads to rationality. *Theoretical Computer Science* 378, 134–142 (2007)

The Computational Complexity of Trembling Hand Perfection and Other Equilibrium Refinements

Kristoffer Arnsfelt Hansen¹, Peter Bro Miltersen^{1,*},
and Troels Bjerre Sørensen^{2,**}

¹ Department of Computer Science
Aarhus University, Åbogade 34
8200 Århus N, Denmark

{arnsfelt,bromille}@cs.au.dk

² Department of Computer Science
University of Warwick, Gibbet Hill Rd
Coventry CV4 7AL, United Kingdom
trold@dcs.warwick.ac.uk

Abstract. The king of refinements of Nash equilibrium is trembling hand perfection. We show that it is **NP**-hard and **SQRT-SUM**-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form with integer payoffs is trembling hand perfect. Analogous results are shown for a number of other solution concepts, including proper equilibrium, (the strategy part of) sequential equilibrium, quasi-perfect equilibrium and CURB.

The proofs all use a reduction from the problem of comparing the minmax value of a three-player game in strategic form to a given rational number. This problem was previously shown to be **NP**-hard by Borgs *et al.*, while a **SQRT-SUM** hardness result is given in this paper. The latter proof yields bounds on the algebraic degree of the minmax value of a three-player game that may be of independent interest.

1 Introduction

Celebrated recent results [10,6,13] concern the computational hardness of *finding* a Nash equilibrium of a given finite game in strategic form, i.e., a game given by a finite payoff matrix for each of the players. In contrast, the problem of *deciding* whether a *given* strategy profile of a game in strategic form is a Nash equilibrium is trivial to solve efficiently. This latter fact can be regarded as an important feature of Nash equilibrium as a scientific concept: It is feasible to verify or falsify that a particular pure strategy profile we observe “in nature” is

* Research supported by Center for Algorithmic Game Theory, funded by the Carlsberg Foundation.

** Research supported by EPSRC award EP/G064679/1 and by the Centre for Discrete Mathematics and its Applications (DIMAP).

in equilibrium. The main message of the present paper is that this feature is not shared by standard *refinements* of Nash equilibrium.

Arguably [9], the most important refinement of Nash equilibrium for games in strategic form is Selten's [25] notion of *trembling hand perfection*. The set of trembling hand perfect equilibria is in general a subset of the Nash equilibria of a game and many "unreasonable" Nash equilibria are not trembling hand perfect, thus justifying the notion. However, we prove that this added degree of rationality of the solution concept comes at a cost. We prove: *It is NP-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form is trembling hand perfect.* In particular, unless $\mathbf{P}=\mathbf{NP}$, there is no polynomial time algorithm for deciding if a given equilibrium of a given three-player game in strategic form is trembling hand perfect. In contrast to the above hardness result, one may efficiently determine if a given equilibrium of a *two*-player game is trembling hand perfect. Indeed, for the two-player case, an equilibrium is trembling hand perfect if and only if it is undominated [9] and this can be checked by linear programming in polynomial time.

The hardness result is extended to a number of other refinements, including properness [24], sequential equilibrium¹ [21] and quasi-perfect equilibrium [8] of extensive form games, and the discrete solution concept CURB (Closed Under Rational Behavior) [1], where the proof yields **coNP**-hardness. In all cases, the hardness result is shown for games with three players. As is the case with trembling hand perfection, CURB sets of two-player game can be verified and found in polynomial time, using linear programming techniques [3]. In contrast, we do not know if the two-player case is easy for properness and quasi-perfection, and leave this as an open problem.

After establishing the **NP**-hardness result, we next ask if the problem of deciding whether an equilibrium is trembling hand perfect (or satisfies any of the other refinements notions we consider) is even *in NP*. An **NP**-membership result would be somewhat beneficial for the status of an equilibrium concept as a useful scientific concept, as it would mean that we can at least, with some ingenuity, verify that a situation is in equilibrium, even if we can not in general falsify this efficiently. For deciding trembling hand perfection, it seems that an obvious nondeterministic algorithm would be to guess and verify a *lexicographic belief structure* and appeal to the characterizations of Blume *et al* [4] and Govindan and Klumpp [15] of trembling hand perfection in terms of these. However, it is not clear if the real numbers involved in such a belief structure can be represented as polynomial length strings over a finite alphabet in a way that yields to efficient verification. To argue that it is fact *not* possible to do so using current knowledge, we apply the notion of SQRT-SUM hardness introduced by Etessami and Yannakakis [12]. In particular, we show that deciding trembling hand perfection (and all the other refinements considered) is SQRT-SUM hard and therefore not in **NP** unless SQRT-SUM is in **NP**. Hence, devising a compact representation of belief structures witnessing trembling hand perfection would solve a long standing open problem of numerical analysis.

¹ To be precise, the "strategy part" of a sequential equilibrium.

The hardness proofs all use a reduction from the problem of comparing the minmax value of a game in a strategic form to a given rational number. This problem was previously shown to be **NP**-hard by Borgs *et al.*, [5], while the **SQRT-SUM** hardness result is given in this paper. The latter proof yields bounds on the algebraic degree of the minmax value of a three-player game that may be of independent interest.

1.1 Related Work

As mentioned, there is a lot of work on hardness of finding equilibria when the game is given as input, or checking whether equilibria with certain properties exists when the game is given as input. In contrast, we are not aware of much previous work on the complexity of determining whether a *given* equilibrium satisfies a refined stability notion. An exception is Etessami and Lochbihler [11] who show that it is **NP**-hard to determine if a given strategy in a symmetric game in strategic form is an evolutionarily stable strategy.

2 NP-hardness of Trembling Hand Perfect and Proper Equilibrium

We recall the definitions of Selten [25]. For a motivation and discussion of the solution concept, we refer to the excellent monograph of van Damme [9].

Definition 1 (ϵ -perfect equilibrium). *A strategy profile σ is an ϵ -perfect equilibrium iff it assigns strictly positive probability to all pure strategies, and only pure strategies that are best replies get probability more than ϵ .*

Definition 2 (Trembling hand perfect equilibrium). *A strategy profile σ is a trembling hand perfect equilibrium iff is the limit point of a sequence of ϵ -perfect equilibria with $\epsilon \rightarrow 0+$.*

Theorem 1. *It is **NP**-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form is trembling hand perfect.*

Proof. Our proof is a reduction from the problem of approximately computing minmax values of 3-player games with 0-1 payoffs. The minmax value of a 3-player game is the smallest number v so that player 2 and player 3 can guarantee, using *uncorrelated* mixed strategies, that player 1 does not get an expected payoff larger than v . The problem of approximately computing this value was recently shown to be **NP**-hard by Borgs *et al* [5]. In particular, it follows from Borgs *et al.* that the following *promise problem* **MINMAX** is **NP**-hard:

MINMAX:

1. **YES**-instances: Pairs (G, r) for which the minmax value for Player 1 in the 3-player game G is strictly smaller than the rational number r .
2. **NO**-instances: Pairs (G, r) for which the minmax value for Player 1 in G is strictly greater than r .

In fact, by multiplying the payoffs of the game with the denominator of r , we can without loss of generality assume that r is an integer. We now reduce MINMAX to deciding trembling hand perfection.

Let G be a three-player game in strategic form and let r be an integer. We define G' be the game where the strategy space of each player is as in G , except that it is extended by a single pure strategy, \perp . The payoffs of G' are defined as follow. The payoff to Players 2 and 3 are 0 for all strategy combinations. The payoff to Player 1 is r for all strategy combinations where at least one player plays \perp . For those strategy combinations where no player plays \perp , the payoff to player 1 is the same as it would have been in the game G . Obviously, $\mu = (\perp, \perp, \perp)$ is a Nash equilibrium of G' .

We claim that if the minmax value for Player 1 in G is strictly smaller than r , then μ is a trembling hand perfect equilibrium of G' . Indeed, let (τ_2, τ_3) be a minmax strategy profile of Players 2 and 3 in G . Let τ be any profile of G' where Players 2 and 3 play (τ_2, τ_3) . Also, let u be the strategy profile of G' where each player mixes all pure strategies uniformly. Now define

$$\sigma_k = \left(1 - \frac{1}{k} - \frac{1}{k^2}\right)\mu + \frac{1}{k}\tau + \frac{1}{k^2}u$$

We have that σ_k is a fully mixed strategy profile of G' converging to μ as $k \rightarrow \infty$. Also, for sufficiently large k , the strategies of μ are best replies to σ_k . This follows from the fact that Players 2 and 3 are indifferent about the outcome and the fact that Player 1 gets payoff r by playing \perp while he gets a payoff strictly smaller than r for large values of k by playing any other strategy. We conclude, using Theorem 2.2.5 in van Damme [9], that μ is trembling hand perfect, as desired.

On the other hand, we claim that if the minmax value for Player 1 in G is strictly greater than r , then μ is a *not* a trembling hand perfect equilibrium of G' . Indeed, let $(\sigma_{k,1}, \sigma_{k,2}, \sigma_{k,3})_k$ be any sequence of fully mixed strategy profiles converging to (\perp, \perp, \perp) . Since $\sigma_{k,2}$ and $\sigma_{k,3}$ do not put all their probability mass on \perp , Player 1 has a reply to $(\sigma_{k,2}, \sigma_{k,3})$ with an expected payoff strictly greater than r . Therefore, \perp is not a best reply of Player 1 to $(\sigma_{k,2}, \sigma_{k,3})$ and we conclude that (\perp, \perp, \perp) is not trembling hand perfect.

That is, we have reduced the promise problem MINMAX to deciding trembling hand perfection and are done. □

We now refine the proof so that it applies to *proper equilibrium*. Proper equilibrium was introduced by Myerson [24] as a further refinement of trembling hand perfect equilibrium. For a motivation and discussion of the solution concept, we refer to the excellent monograph of van Damme [9] or the survey of Hillas and Kohlberg [18].

Definition 3 (ϵ -proper equilibrium). *A strategy profile σ is an ϵ -proper equilibrium iff it assigns strictly positive probability to all pure strategies, and the following condition holds: Given two pure strategies, p_i and p_j , of the same player. If p_i is a worse reply against σ than p_j , then σ must assign a probability to p_i that is at most ϵ times the probability it assign to p_j .*

Definition 4 (Proper equilibrium). A strategy profile σ is a proper equilibrium iff it is the limit point of a sequence of ϵ -proper equilibria with $\epsilon \rightarrow 0+$.

Theorem 2. It is NP-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form is proper.

Proof. We only need to make minor changes to the proof of NP-hardness of trembling hand perfection to get the same result for proper equilibria. Construct the game in the same way, with a new strategy \perp for each player. Define the strategy $\tau_{1,k}$ for Player 1 to be a permutation of $(1 - \sum_i k^{-i}, k^{-1}, k^{-2}, \dots, k^{-n+1})$, such that worse replies against (τ_2, τ_3) get more negative powers of k . In case two pure strategies are equal against (τ_2, τ_3) , compare against the uniform mix u of Players 2 and 3, again with the worse reply getting the more negative powers of k . This can be achieved by sorting the strategies of Player 1 lexicographically on payoff against (τ_2, τ_3) and the uniform strategy u , and then assigning powers in decreasing order to the lower indices. Define $\tau_k = (\tau_{1,k}, \tau_2, \tau_3)$, $\nu_k = (\tau_{1,k}, u, u)$, and $\mu = (\perp, \perp, \perp)$. Now define

$$\sigma_k = \left(1 - \frac{1}{k} - \frac{1}{k^2}\right)\mu + \frac{1}{k}\tau_k + \frac{1}{k^2}\nu_k$$

σ_k is fully mixed of all finite k . Furthermore, if the minmax value for Player 1 in G is less than r , then for any sufficiently large k , better replies of Player 1 gets k' times higher probability than worse replies, thus satisfying the condition for being a $\frac{1}{k'}$ -proper equilibrium, with $k' = k/(1 - \sum_i k^{-i})$. Since σ_k tends towards μ as $k \rightarrow \infty$, we therefore have that μ is a proper equilibrium. If the minmax value for Player 1 in G is greater than r , μ is not even trembling hand perfect, and therefore not proper either. μ is therefore proper if and only if the minmax value for Player 1 in G is less than r . \square

3 NP-hardness of Refinements of Nash Equilibria for Extensive form Games

An extensive form game is given by a finite tree with payoffs for each player at the leaves, information sets partitioning nodes of the tree and with some of the nodes having predefined moves of chance. An information set is a collection of nodes of the same player, where the player cannot distinguish between them. This can be used to model information hidden from the player, both as actively hidden information in a game over time, and as a way of modelling simultaneous moves. A player is said to have *perfect recall* if for each of the player's information sets, all nodes in the set share the same sequence of actions and information sets of the player on the path from the root to the nodes. A game is said to be of perfect recall, if all players have perfect recall. This is a standard assumption to make, and one that the game produced by our reduction will satisfy.

Actions of a player are denoted by labels on edges of the tree. A *behavior strategy* assigns probabilities to actions such that it forms probability distributions over the actions for each of the information sets. A Nash equilibrium in

behavior strategies is a profile of behavior strategies so that no player wants to deviate, given that other players play according to the profile. As is the case of games in strategic form, it is straightforward to verify in polynomial time that a given profile is a Nash equilibrium. For details, see e.g., Koller, Megiddo and von Stengel [20] or any textbook on game theory.

The most important refinement of Nash equilibrium for game in sequential form is the notion of *sequential equilibrium* due to Kreps and Wilson [21] is based on the notion of *beliefs*. Formally, a belief of a player is a probability distribution on each of his information sets. Intuitively, the belief should indicate the subjective probability of the player of being in each of the nodes in the information set, given that he has arrived at this information set. An *assessment* (ρ, μ) is a strategy profile ρ , and a *belief profile* μ : a belief for each of the players. A sequential equilibrium is an assessment which is (1) *consistent* and (2) *a sequential best reply against itself*, the former notion capturing that the beliefs are sensible given the strategies, and the latter notion capturing that the strategies are sensible given the beliefs. We define these two notions formally next.

We first define consistency for *fully mixed* strategy profiles, i.e., ones where every action in every information set has a strictly positive probability of being taken. For such a strategy profile, the *induced belief profile* is the unique one consistent with the strategy profile: The strategies being played out against each other induces a probability distribution on possible plays; the induced belief assigns to information set u the conditional probability distribution on u derived from this probability distribution. This is well-defined as at most one node in u may be reached during each particular play (due to the perfect recall property) and u has a non-zero probability of being reached (as the strategies are fully mixed). The contribution of Kreps and Wilson is a generalization of this consistency notion to strategy profiles where some of the information sets may be reached with probability 0: For this general case, we say that an assessment is consistent if it is the limit point of a sequence of consistent assessments with fully mixed strategy profiles.

We next define what it means to be a sequential best reply against itself. For each player, a strategy profile will assign an expected value to each node of the tree, which is the expectation over the leaves given than play starts at that node and follows the given probabilities of play. Given a belief as well, we can assign an expected value to each action, being the expected value of the node reached by taking the action for each node, weighted by the probability given by the belief. An assessment is said to be a sequential best reply against itself, if all players only assign positive probability to actions with maximal expected payoff, given the strategy profile and belief.

Note that a sequential equilibrium is an assessment, i.e., a behavior strategy profile and a belief profile. Our NP-hardness result applies when the input is the “strategy-part” only.

Theorem 3. *Given a pure strategy profile of an extensive form three-player game, it is NP-hard to decide if it is part of an assessment that is a sequential equilibrium.*

Proof. The reduction is again similar to that of trembling hand perfection. Given a game G in strategic form, construct an extensive form game G' where the players choose an action of G in turn, but without revealing the choice to the other players. Player 2 chooses first, then Player 3, and finally Player 1. Each Player now has a single information set, and the game is strategically equivalent to G . Now give each player a new action \perp . If this action is chosen, the game ends immediately without having the remaining players choose an action. If \perp is chosen by either player, the payoff to Player 1 is r , otherwise it is simply the payoff from G .




Fig. 1. The extensive form game G'

We now argue that $\mu = (\perp, \perp, \perp)$ is part of a sequential equilibrium iff the minmax value of Player 1 in G is less than r .

Define τ to be some strategy profile where Players 2 and 3 play minmax against Player 1, and let u be the strategy profile with all players playing the uniform distribution. As in the previous proofs, let

$$\sigma_k = \left(1 - \frac{1}{k} - \frac{1}{k^2}\right)\mu + \frac{1}{k}\tau + \frac{1}{k^2}u$$

σ_k is fully mixed of all finite k . Furthermore, if the minmax value for Player 1 in G is less than r , then for any sufficiently large k , \perp will be the unique best reply of Player 1 against σ_k . This also means that the expected value for Player 1 of choosing \perp given the induced belief of σ_k will be strictly higher than for all other actions, and this will also hold for the limit.

On the other hand, if the minmax value is greater than r , no strategy of Players 2 and 3 will make \perp be the best reply of Player 1. Therefore, no belief (consistent with a strategy of Players 2 and 3) will give a maximal expected payoff to Player 1 playing \perp .

μ is therefore part of a sequential equilibrium if and only if the minmax value for Player 1 in G is less than r . □

Theorem 3 begs the following question: Can one check in polynomial time if an entire assessment (a strategy profile and a belief profile) given as input is a

sequential equilibrium? Kohlberg and Reny [19] present a finite-step algorithm performing this task, but as they state it, their algorithm is exponential. It is not clear to us if this problem is in **P** or if it is **NP**-hard and we consider this an interesting open problem. It is interesting to note that this is in some contrast with the situation for strategic form games: Perfect and proper equilibrium of strategic form games can also be “backed up” by belief structures [4,15] and if a rational-valued belief structure is given as part of the input, it is straightforward to verify the equilibrium condition (however, as we argue in Section 5, a given perfect strategy profile may require belief structures with no polynomial-size representation - in particular, using algebraic numbers of very high degree may be necessary).

A refinement of (the strategy part of) sequential equilibrium is *quasi-perfect* equilibrium [8]. Despite the fact that quasi-perfect equilibrium is a lesser known refinement than sequential equilibrium, it has been argued strongly by Mertens [22] (see also [18]) that quasi-perfect equilibrium is the “right” equilibrium notion of extensive form games. We omit the technically involved definition of quasi-perfection, but note that it is straightforward to check that the reduction in the proof of Theorem 3 maps “yes”-instance to equilibria that are not only sequential but also quasi-perfect. Since quasi-perfect equilibrium refines sequential equilibrium, we also have that the reduction maps “no”-instances to equilibria that are not quasi-perfect. Therefore we have the following corollary.

Corollary 1. *It is NP-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in extensive form is quasi-perfect.*

4 coNP-hardness of CURB Sets

A set valued solution concept is Strategy Sets Closed Under Rational Behavior (CURB) [1].

Definition 5 (CURB set). *In an m -player game, a family of sets of pure strategies, S_1, S_2, \dots, S_m with S_i being a subset of the strategy set of player i , is closed under rational behavior (CURB) iff for all pure strategies x of Player i so that x is a best reply to some product distribution on $S_1 \times S_2 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_m$, we have that $x \in S_i$.*

CURB sets are guaranteed always to exist, as the set of all pure strategies is trivially CURB, as there are no pure strategies outside the set. The CURB condition is usually paired with a minimality condition, so as not to get unnecessarily large solutions. This minimality condition would be the obvious place to look for **coNP**-hardness, but we show here that simply checking the CURB condition is **coNP**-hard. This also implies that it is not obvious that minimality should even be contained in **coNP**.

Theorem 4. *It is coNP-hard to check whether a set of n pure strategies of each player is CURB in an $(n + 1) \times n \times n$ strategic form game with integer payoffs.*

Proof. We again reduce from MINMAX, so let G be a three-player game in strategic form and let r be an integer. We define G' be the game where the strategy space of each player is as in G , except that Player 1 gets an additional strategy, \perp . The payoffs of G' are defined as follow. The payoff to Players 2 and 3 are 0 for all strategy combinations. The payoff to Player 1 is r , if he plays \perp , and otherwise the payoff to player 1 is the same as it would have been in the game G .

Now, the minmax value of G is less than r iff the set of all pure strategies except \perp is CURB in G' . Indeed, if the minmax value of G is less than r , then Player 1's best reply to the optimal treat of Players 2 and 3 is \perp in G' . The set of all pure strategies except \perp is therefore *not* CURB. If the minmax value of G is greater than r , then \perp is never a best reply in G , and the set of all other strategies is CURB. \square

5 Sqrt-Sum-hardness

SQRT-SUM is the following decision problem [16][14]: Given positive integers a_1, a_2, \dots, a_n, k , decide whether $\sum_{i=1}^n \sqrt{a_i} < k$.

Though it is not unlikely that this problem is in \mathbf{P} , we do not even know if it is in \mathbf{NP} at the moment. A decision problem is called SQRT-SUM-hard if SQRT-SUM reduces to it by a polynomial time many-one reduction. Etesami and Yannakakis [12] pioneered the use of SQRT-SUM-hardness to argue that certain problems are hard “given current state of the art”. It is important to notice that unlike \mathbf{NP} -hardness, SQRT-SUM-hardness should *not* be used as an indication that a problem is *actually* hard, only as an indication that we do not know if it is easy. In this section we show SQRT-SUM-hardness of the minmax value of a 3-player game and thus by the previously described reductions give evidence that it is not possible to decide the refined solution concepts in \mathbf{NP} “given current state of the art”.

Lemma 1. *For every pair of probability distributions x and y on $\{1, \dots, n\}$ there exists another probability distribution z such that $x_i y_i \leq z_i^2$ for all i .*

Proof. If $x_i y_i = 0$ for all i we may pick z arbitrarily. Otherwise, define $w_i = \sqrt{x_i y_i}$ for all i . By the Cauchy-Schwarz inequality we have

$$\sum_{i=1}^n w_i = \sum_{i=1}^n \sqrt{x_i} \sqrt{y_i} \leq \sqrt{\sum_{i=1}^n x_i} \sqrt{\sum_{i=1}^n y_i} = 1 .$$

We may thus obtain the required z by letting $z_j = w_j / (\sum_{i=1}^n w_i)$. \square

Given positive numbers a_1, \dots, a_n define the payoff to player 1 in an $n \times n \times n$ game $G(a_1, \dots, a_n)$ by letting $u_1(i, j, k) = -1/a_i$ if $i = j = k$ and $u_1(i, j, k) = 0$ otherwise.

Proposition 1. *The minmax value for player 1 in the game $G(a_1, \dots, a_n)$ is $-1/(\sum_{i=1}^n \sqrt{a_i})^2$.*

Proof. If player 2 and player 3 play strategies p and q , player 1 may obtain payoff $\max_i -p_i q_i / a_i$. Let v be the minmax value for player 1. For optimal strategies for player 2 and 3 we may assume by Lemma [1](#) that $p = q$, and furthermore we must then have that $v = -p_i^2 / a_i$ for all i , and thus $p_i = \sqrt{-v} \sqrt{a_i}$ for all i . Summing over i gives

$$1 = \sum_{i=1}^n p_i = \sum_{i=1}^n \sqrt{-v} \sqrt{a_i} = \sqrt{-v} \sum_{i=1}^n \sqrt{a_i} .$$

Squaring and rearranging gives

$$v = -\frac{1}{\left(\sum_{i=1}^n \sqrt{a_i}\right)^2} ,$$

as stated. □

Theorem 5. *Deciding whether the minmax value for player 1 in a $n \times n \times n$ game is less than a given rational k is SQRT-SUM hard.*

Proof. Deciding whether $\sum_{i=1}^n \sqrt{a_i} < k$ reduces to decide for the minmax value v for player 1 in the game $G(a_1, \dots, a_n)$ whether $v < -\frac{1}{k^2}$ by Proposition [1](#). □

Corollary 2. *It is SQRT-SUM hard to determine whether a given pure equilibrium in a 3-player game in strategic form with integer payoffs is trembling-hand perfect or proper and whether a given pure equilibrium in a 3-player game in extensive form with integer payoffs is quasi-perfect or the strategy part of a sequential equilibrium. It is also SQRT-SUM hard to test whether a given set of pure strategies is not CURB. In particular, neither of these problems are in NP unless SQRT-SUM is in NP.*

Finally, we show that our reduction can also be used to give lower bounds on the algebraic degree of the minmax value of a 3-player game. Such a result is interesting for computational reasons: They indicate that if we want to compute the *exact* minmax value of a 3-player game and want to represent the exact irrational but algebraic answer in, say, a standard representation such as Thom encoding [7](#), exponential space is needed even to represent the output.

For providing the lower bound of the algebraic degree of the minmax value we use basic results from the theory of field extensions.

Proposition 2. *The algebraic degree of the minmax value for player 1 in a $n \times n \times n$ game can be 2^{n-1} .*

Proof. Let a_1, \dots, a_n be arbitrary relatively prime positive integers, and let v be the minmax value of the game $G(a_1, \dots, a_n)$. We shall calculate the degree $[\mathbf{Q}(v) : \mathbf{Q}]$ of the field extension $\mathbf{Q}(v)$ of \mathbf{Q} . It is well known that for relatively prime positive integers a_1, \dots, a_n we have $[\mathbf{Q}(\sqrt{a_1}, \dots, \sqrt{a_n}) : \mathbf{Q}] = 2^n$ (e.g. [23](#), Example 11.5). Furthermore, we have $\mathbf{Q}(\sqrt{a_1} + \dots + \sqrt{a_n}) = \mathbf{Q}(\sqrt{a_1}, \dots, \sqrt{a_n})$. By Proposition [1](#) we have that $-1/\sqrt{v} = \sum_{i=1}^n \sqrt{a_i}$, and thus $[\mathbf{Q}(\sqrt{v}) : \mathbf{Q}] = 2^n$. Finally using $[\mathbf{Q}(\sqrt{v}) : \mathbf{Q}] = [\mathbf{Q}(\sqrt{v}) : \mathbf{Q}(v)][\mathbf{Q}(v) : \mathbf{Q}] \leq 2[\mathbf{Q}(v) : \mathbf{Q}]$ the result follows. □

One can give an almost matching upper bound using the general tool of quantifier elimination for the first order theory of the reals.

Proposition 3. *The algebraic degree of the minmax value for player 1 in a $n \times n \times n$ game is $2^{O(n)}$.*

Proof. We may describe the minmax value by a first order formula $P(v)$ with free variable v , as $P(v) := A(v) \wedge B(v)$, where

$$A(v) := (\exists p, q \in \mathbf{R}^n) \bigwedge_{i=1}^n \left(\sum_{j=1}^n \sum_{k=1}^n u_1(i, j, k) p_j q_k \leq v \right) \wedge C(p, q) ,$$

$$B(v) := (\forall p, q \in \mathbf{R}^n) \bigvee_{i=1}^n \left(\sum_{j=1}^n \sum_{k=1}^n u_1(i, j, k) p_j q_k \geq v \right) \wedge C(p, q) ,$$

and $C(p, q) := (\bigwedge_{i=1}^n p_i \geq 0) \wedge (\sum_{i=1}^n p_i = 1) \wedge (\bigwedge_{i=1}^n q_i \geq 0) \wedge (\sum_{i=1}^n q_i = 1)$.

We note that the degree of each polynomial in the formula is at most 2. Thus applying the quantifier elimination procedure of Basu, Pollack and Roy [2] to each of the formulas $A(v)$ and $B(v)$ yields equivalent quantifier free formulas $A'(v)$ and $B'(v)$ wherein each polynomial is of degree $2^{O(n)}$. It follows that $A'(v) \wedge B'(v)$ is quantifier free formula equivalent to $P(v)$ wherein each polynomial are univariate polynomials in v of degree $2^{O(n)}$. Now, since the actual minmax value v is an isolated solution to this formula, it must satisfy one of the polynomial equations involving a nonconstant polynomial with equality. We can thus conclude it must be a root of a polynomial of degree $2^{O(n)}$. \square

Remark 1. The above bound is especially relevant for the special case of $k \times n \times n$ games, where k is considered a constant [17]. For this case one may find the minmax value by considering all $k \times k \times k$ subgames and the minmax value of those. This also means that for fixed k one can in polynomial time compute the Thom encoding of the minmax value of a given $k \times n \times n$ game, employing general algorithms for the first-order theory of the reals [2].

References

1. Basu, K., Weibull, J.W.: Strategy subsets closed under rational behavior. *Economics Letters* 2(36), 141–146 (1991)
2. Basu, S., Pollack, R., Roy, M.F.: On the combinatorial and algebraic complexity of quantifier elimination. *Journal of the ACM* 43(6), 1002–1045 (1996)
3. Benisch, M., Davis, G.B., Sandholm, T.: Algorithms for closed under rational behavior (CURB) sets. *Journal of Artificial Intelligence Research* (Forthcoming 2010)
4. Blume, L., Brandenburger, A., Dekel, E.: Lexicographic probabilities and equilibrium refinements. *Econometrica* 59, 81–98 (1991)
5. Borgs, C., Chayes, J.T., Immorlica, N., Kalai, A.T., Mirrokni, V.S., Papadimitriou, C.H.: The myth of the folk theorem. In: *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pp. 365–372. ACM, New York (2008)

6. Chen, X., Deng, X.: Settling the complexity of two-player Nash equilibrium. In: Proceedings of 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), pp. 261–272 (2006)
7. Coste, M., Roy, M.: Thom’s lemma, the coding of real algebraic numbers and the computation of the topology of semi-algebraic sets. *Journal of Symbolic Computation* 5(1-2), 121–129 (1988)
8. van Damme, E.: A relation between perfect equilibria in extensive form games and proper equilibria in normal form games. *International Journal of Game Theory* 13, 1–13 (1984)
9. van Damme, E.: *Stability and Perfection of Nash Equilibria*, 2nd edn. Springer, Heidelberg (1991)
10. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. In: Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC 2006), pp. 71–78 (2006)
11. Etessami, K., Lochbihler, A.: The computational complexity of evolutionarily stable strategies. *International Journal of Game Theory* 31(1), 93–113 (2008)
12. Etessami, K., Yannakakis, M.: Recursive markov decision processes and recursive stochastic games. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 891–903. Springer, Heidelberg (2005)
13. Etessami, K., Yannakakis, M.: On the complexity of Nash equilibria and other fixed points (extended abstract). In: Proc. 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), pp. 113–123 (2007)
14. Garey, M.R., Graham, R.L., Johnson, D.S.: Some NP-complete geometric problems. In: Proceedings of the 8th Annual ACM Symposium on Theory of Computing, STOC 1976, Hershey, PA, May 3-5, pp. 10–22. ACM Press, New York (1976)
15. Govindan, S., Klumpp, T.: Perfect equilibrium and lexicographic beliefs. *International Journal of Game Theory* 31(2), 229–243 (2003)
16. Graham, R.L.: 10 - Problems and Solutions. In: P73: Euclidian Minimum Spanning Trees. *Bulletin of the EATCS*, vol. 24, pp. 205–206. EATCS (October 1984)
17. Hansen, K.A., Hansen, T.D., Miltersen, P.B., Sørensen, T.B.: Approximability and parameterized complexity of minmax values. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 684–695. Springer, Heidelberg (2008)
18. Hillas, J., Kohlberg, E.: Foundations of strategic equilibria. In: Aumann, R.J., Hart, S. (eds.) *Handbook of Game Theory*, ch. 42, vol. 3, pp. 1597–1663. Elsevier Science, Amsterdam (2002)
19. Kohlberg, E., Reny, P.J.: Independence on relative probability spaces and consistent assessments in game trees. *Journal of Economic Theory* 75(2), 280–313 (1997)
20. Koller, D., Megiddo, N., von Stengel, B.: Efficient computation of equilibria for extensive form games. *Games and Economic Behavior* 14, 247–259 (1996)
21. Kreps, D.M., Wilson, R.: Sequential equilibria. *Econometrica* 50(4), 863–894 (1982)
22. Mertens, J.F.: Two examples of strategic equilibrium. *Games and Economic Behavior* 8, 378–388 (1995)
23. Morandi, P.: *Field and Galois Theory*. Graduate Texts in Mathematics, vol. 167. Springer, Heidelberg (1996)
24. Myerson, R.B.: Refinements of the Nash equilibrium concept. *International Journal of Game Theory* 15, 133–154 (1978)
25. Selten, R.: A reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4, 25–55 (1975)

Complexity of Safe Strategic Voting

Noam Hazon¹ and Edith Elkind²

¹ Department of Computer Science, Bar Ilan University, Israel

² School of Physical and Mathematical Sciences,
Nanyang Technological University, Singapore

Abstract. We investigate the computational aspects of *safe manipulation*, a new model of coalitional manipulation that was recently put forward by Slinko and White [10]. In this model, a potential manipulator v announces how he intends to vote, and some of the other voters whose preferences coincide with those of v may follow suit. Depending on the number of followers, the outcome could be better or worse for v than the outcome of truthful voting. A manipulative vote is called *safe* if for some number of followers it improves the outcome from v 's perspective, and can never lead to a worse outcome. In this paper, we study the complexity of finding a safe manipulative vote for a number of common voting rules, including Plurality, Borda, k -approval, and Bucklin, providing algorithms and hardness results for both weighted and unweighted voters. We also propose two ways to extend the notion of safe manipulation to the setting where the followers' preferences may differ from those of the leader, and study the computational properties of the resulting extensions.

1 Introduction

Computational aspects of voting, and, in particular, voting manipulation, is an active topic of current research. While the complexity of the manipulation problem for a single voter is quite well understood (specifically, this problem is known to be efficiently solvable for most common voting rules with the notable exception of STV [12]), the more recent work has mostly focused on coalitional manipulation, i.e., manipulation by multiple, possibly weighted voters. In contrast to the single-voter case, coalitional manipulation tends to be hard. Indeed, it has been shown to be NP-hard for weighted voters even when the number of candidates is bounded by a small constant [3]. For unweighted voters, nailing the complexity of coalitional manipulation proved to be more challenging. However, Faliszewski et al. [4] have recently established that this problem is hard for most variants of Copeland, and Zuckerman *et al* [12] showed that it is easy for Veto and Plurality with Runoff. Further, a very recent paper [11] makes substantial progress in this direction, showing, for example, that unweighted coalitional manipulation is hard for Maximin and Ranked Pairs, but easy for Bucklin (see Section 2 for the definitions of these rules).

All of these papers (as well as the classic work of Bartholdi et al. [1]) assume that the set of manipulators is given exogenously, and the manipulators are not endowed with preferences over the entire set of candidates; rather, they simply want to get a particular candidate elected, and select their votes based on the non-manipulators' preferences that are publicly known. That is, this model abstracts away the question of how

the manipulating coalition forms. However, to develop a better understanding of coalitional manipulation, it is desirable to have a plausible model of the coalition formation process. In such a model the manipulators would start out by having the same type of preferences as sincere voters, and then some agents—those who are not satisfied with the current outcome and are willing to submit an insincere ballot—would get together and decide to coordinate their efforts.

However, it is quite difficult to formalize this intuition so as to obtain a realistic model of how the manipulating coalition forms. In particular, it is not clear how the voters who are interested in manipulation should identify each other, and then reach an agreement as to which candidate to promote. Indeed, the latter decision seems to call for a voting procedure, and therefore is itself vulnerable to strategic behavior. Further, even assuming that suitable coalition formation and decision-making procedures exist, their practical implementation may be hindered by the absence of reliable two-way communication among the manipulators.

In a recent paper [10], Slinko and White put forward a model that provides a partial answer to these questions. They consider a setting where a single voter v announces his manipulative vote L (the truthful preferences of all agents are, as usual, common knowledge) to his set of associates F , i.e., the voters whose true preferences coincide with those of v . As a result, some of the voters in F switch to voting L , while others (as well as all voters not in F) vote truthfully. This can happen if, e.g., v 's instructions are broadcast via an unreliable channel, i.e., some of the voters in F simply do not receive the announcement, or if some voters in F consider it unethical to vote non-truthfully. Such a situation is not unusual in politics, where a public figure may announce her decision to vote in a particular manner, and may be followed by a subset of like-minded voters. That is, in this model, the manipulating coalition always consists of voters with identical preferences (and thus the problem of which candidate to promote is trivially resolved), and, moreover, the manipulators always vote in the same way. Further, it relies on minimal communication, i.e., a single broadcast message. However, due to lack of two-way communication, v does not know how many voters will support him in his decision to vote L . Thus, he faces a dilemma: it might be the case that if x voters from F follow him, then the outcome improves, while if some $y \neq x$ voters from F switch to voting L , the outcome becomes even less desirable to v than the current alternative (we provide an example in Section 2). If v is conservatively-minded, in such situations he would choose not to manipulate at all. In other words, he would view L as a successful manipulation only if (1) there exists a subset $U \subseteq F$ such that if the voters in U switch to voting L , the outcome improves; (2) for any $W \subseteq F$, if the voters in W switch to voting L the outcome does not get worse. Paper [10] calls any manipulation that satisfies (1) and (2) *safe*. The main result of [10] is a generalization of the Gibbard–Satterthwaite theorem [69] to safe manipulation: the authors prove that any onto, non-dictatorial voting rule with at least 3 alternatives is safely manipulable, i.e., there exists a profile in which at least one voter has a safe manipulation. However, paper [10] does not explore the computational complexity of the related problems.

In the first part of this paper, we focus on algorithmic complexity of safe manipulation, as defined in [10]. We first formalize the relevant computational questions and discuss some basic relationships between them. We then study the complexity of these

questions for several classic voting rules, such as Plurality, Veto, k -approval, Bucklin, and Borda, for both weighted and unweighted voters. For instance, we show that finding a safe manipulation is easy for k -approval and for Bucklin, even if the voters are weighted. In contrast, for Borda, finding a safe manipulation—or even checking that a given vote is safe—turns out to be hard for weighted voters even if the number of candidates is bounded by a small constant.

We then explore whether it is possible to extend the model of safe manipulation to settings where the manipulator may be joined by voters whose preferences differ from his own. Indeed, in real life a voter may follow advice to vote in a certain way if it comes from a person whose preferences are similar (rather than identical) to hers, or simply because she thinks that voting in this manner can be beneficial to her. For instance, in politics, a popular personality may influence many different voters at once by announcing his decision to vote in a particular manner. We propose two ways of formalizing this idea, which differ in their approach to defining the set of a voter’s potential followers, and provide initial results on the complexity of safe manipulation in these models.

In our first extension, a manipulator v may be followed by all voters who rank the same candidates above the current winner as v does. That is, in this model a voter u may follow v if any change of outcome that is beneficial to v is also beneficial to u . We show that some of the positive algorithmic results for the standard model also hold in this more general setting. In our second model, a voter u may follow a manipulator v that proposes to vote L , if, roughly, there are circumstances when voting L is beneficial to u . This model tends to be computationally more challenging: we show that finding a safe strategic vote in this setting is hard even for very simple voting rules.

We conclude the paper by summarizing our results and proposing several directions for future research. Due to space constraints, most of the proofs are omitted.

2 Preliminaries and Notation

An *election* is given by a set of *candidates* (or, *alternatives*) $C = \{c_1, \dots, c_m\}$ and a set of *voters* $V = \{1, \dots, n\}$. Each voter i is represented by his *preference* R_i , which is a total order over C ; we will also refer to total orders over C as *votes*. For readability, we will sometimes denote the order R_i by \succ_i . The vector $\mathcal{R} = (R_1, \dots, R_n)$ is called a *preference profile*. We say that two voters i and j are of the same *type* if $R_i = R_j$; we write $V_i = \{j \mid R_j = R_i\}$. A *voting rule* \mathcal{F} is a mapping from the set of all preference profiles to the set of candidates; if $\mathcal{F}(\mathcal{R}) = c$, we say that c *wins* under \mathcal{F} in \mathcal{R} . A voting rule is said to be *anonymous* if $\mathcal{F}(\mathcal{R}) = \mathcal{F}(\mathcal{R}')$, where \mathcal{R}' is a preference profile obtained by permuting the entries of \mathcal{R} . To simplify the presentation, in this paper we consider anonymous voting rules only. In addition, we restrict ourselves to voting rules that are polynomial-time computable. During the election, each voter i submits a vote L_i ; the outcome of the election is then given by $\mathcal{F}(L_1, \dots, L_n)$. We say that a voter i is *truthful* if $L_i = R_i$. For any $U \subseteq V$ and a vote L , we denote by $\mathcal{R}_{-U}(L)$ the profile obtained from \mathcal{R} by replacing R_i with L for all $i \in U$.

Voting rules. We will now define the voting rules considered in this paper. All of these rules assign scores to all candidates; the winner is then selected among the candidates

with the highest score using a *tie-breaking rule*, i.e., a mapping $T : 2^C \rightarrow C$ that satisfies $T(S) \in S$. Unless specified otherwise, we assume that the tie-breaking rule is *lexicographic*, i.e., given a set of tied alternatives, it selects one that is maximal with respect to a fixed ordering \succ .

Given a vector $\alpha = (\alpha_1, \dots, \alpha_m)$ with $\alpha_1 \geq \dots \geq \alpha_m$, the *score* $s_\alpha(c)$ of a candidate $c \in C$ under a *positional scoring rule* F_α is given by $\sum_{i \in V} \alpha_{j(i,c)}$, where $j(i, c)$ is the position in which voter i ranks candidate c . Many classic voting rules can be represented using this framework. Indeed, *Plurality* is the scoring rule with $\alpha = (1, 0, \dots, 0)$, *Veto* (also known as *Anti-plurality*) is the scoring rule with $\alpha = (1, \dots, 1, 0)$, and *Borda* is the scoring rule with $\alpha = (m-1, m-2, \dots, 1, 0)$. Further, *k-approval* is the scoring rule with α given by $\alpha_1 = \dots = \alpha_k = 1, \alpha_{k+1} = \dots = \alpha_m = 0$; we will also refer to $(m-k)$ -approval as *k-veto*.

Bucklin rule can be viewed as an adaptive version of *k-approval*. We say that $k, 1 \leq k \leq m$, is the *Bucklin winning round* if for any $j < k$ no candidate is ranked in top j positions by at least $\lceil n/2 \rceil$ voters, and there exists some candidate that is ranked in top k positions by at least $\lceil n/2 \rceil$ voters. We say that the candidate c 's *score in round* j is his j -approval score, and his *Bucklin score* $s_B(c)$ is his k -approval score, where k is the Bucklin winning round. The *Bucklin winner* is the candidate with the highest Bucklin score. Observe that the Bucklin score of the Bucklin winner is at least $\lceil n/2 \rceil$.

Weighted voters. Our model can be extended to the situation where not all voters are equally important by assigning an integer *weight* w_i to each voter i . To compute the winner on a profile (R_1, \dots, R_n) under a voting rule \mathcal{F} given voters' weights $\mathbf{w} = (w_1, \dots, w_n)$, we apply \mathcal{F} on a modified profile which for each $i = 1, \dots, n$ contains w_i copies of R_i . As an input to our problems we usually get a *voting domain*, i.e., a tuple $S = \langle C, V, \mathbf{w}, \mathcal{R} \rangle$, together with a specific voting rule. When $\mathbf{w} = (1, \dots, 1)$, we say that the voters are *unweighted*. For each $U \subseteq V$, let $|U|$ be the number of voters in U and let $w(U)$ be the total weight of the voters in U .

Safe manipulation. We will now formally define the notion of safe manipulation. For the purposes of our presentation, we can simplify the definitions in [10] considerably.

As before, we assume that the voters' true preferences are given by a preference profile $\mathcal{R} = (R_1, \dots, R_n)$.

Definition 1. We say that a vote L is an *incentive to vote strategically*, or a *strategic vote* for i at \mathcal{R} under \mathcal{F} , if $L \neq R_i$ and for some $U \subseteq V_i$ we have $\mathcal{F}(\mathcal{R}_{-U}(L)) \succ_i \mathcal{F}(\mathcal{R})$. Further, we say that L is a *safe strategic vote* for a voter i at \mathcal{R} under \mathcal{F} if L is a strategic vote at \mathcal{R} , and for any $U \subseteq V_i$ either $\mathcal{F}(\mathcal{R}_{-U}(L)) \succ_i \mathcal{F}(\mathcal{R})$ or $\mathcal{F}(\mathcal{R}_{-U}(L)) = \mathcal{F}(\mathcal{R})$.

To build intuition for the notions defined above, consider the following example.

Example 1. Suppose $C = \{a, b, c, d\}$, $V = \{1, 2, 3, 4\}$, the first three voters have preference $b \succ a \succ c \succ d$, and the last voter has preference $c \succ d \succ a \succ b$. Suppose also that the voting rule is 2-approval. Under truthful voting, a and b get 3 points, and c and d get 1 point each. Since ties are broken lexicographically, a wins. Now, if voter 1 changes his vote to $L = b \succ c \succ a \succ d$, b gets 3 points, a gets 2 points, and c gets 2 points, so b wins. As $b \succ_1 a$, L is a strategic vote for 1. However, it is not a safe

strategic vote: if players in $V_1 = \{1, 2, 3\}$ all switch to voting L , then c gets 4 points, while b still gets 3 points, so in this case c wins and $a \succ_1 c$.

3 Computational Problems: First Observations

The definition of safe strategic voting gives rise to two natural algorithmic questions. In the definitions below, \mathcal{F} is a given voting rule and the voters are assumed to be unweighted.

- $\text{ISSAFE}(\mathcal{F})$: Given a voting domain, a voter i and a linear order L , is L a safe strategic vote for i under \mathcal{F} ?
- $\text{EXISTSAFE}(\mathcal{F})$: Given a voting domain and a voter i , can voter i make a safe strategic vote under \mathcal{F} ?

The variants of these problems for weighted voters will be denoted, respectively, by $\text{wISSAFE}(\mathcal{F})$ and $\text{wEXISTSAFE}(\mathcal{F})$. Note that, in general, it is not clear if an efficient algorithm for $(\text{w})\text{EXISTSAFE}(\mathcal{F})$ can be used to solve $(\text{w})\text{ISSAFE}(\mathcal{F})$, or vice versa. However, if the number of candidates is constant, $(\text{w})\text{EXISTSAFE}(\mathcal{F})$ reduces to $(\text{w})\text{ISSAFE}(\mathcal{F})$. We state the following two results (the easy proofs are omitted) for weighted voters; clearly, they also apply to unweighted voters.

Proposition 1. *Consider any voting rule \mathcal{F} . For any constant k , if $|C| \leq k$, then a polynomial-time algorithm for $\text{wISSAFE}(\mathcal{F})$ can be used to solve $\text{wEXISTSAFE}(\mathcal{F})$ in polynomial time.*

A similar reduction exists when each voter only has polynomially many “essentially different” votes.

Proposition 2. *Consider any scoring rule \mathcal{F}_α that satisfies either (i) $\alpha_j = 0$ for all $j > k$ or (ii) $\alpha_j = 1$ for all $j \leq m - k$, where k is a given constant. For any such rule, a polynomial-time algorithm for $\text{wISSAFE}(\mathcal{F}_\alpha)$ can be used to solve $\text{wEXISTSAFE}(\mathcal{F}_\alpha)$ in polynomial time.*

Observe that the class of rules considered in Proposition 2 includes Plurality and Veto, as well as k -approval and k -veto when k is bounded by a constant.

Further, for unweighted voters it is easy to check if a given manipulation is safe.

Proposition 3. *The problem $\text{ISSAFE}(\mathcal{F})$ is in P for any (anonymous) voting rule \mathcal{F} .*

Together with Propositions 1 and 2, Proposition 3 implies that $\text{EXISTSAFE}(\mathcal{F})$ is in P for Plurality, Veto, k -veto and k -approval for constant k , as well as for any voting rule with a constant number of candidates.

Note that when voters are weighted, the conclusion of Proposition 3 no longer holds. Indeed, in this case the number of subsets of V_i that have different weights (and thus may have a different effect on the outcome) may be exponential in n . However, it is not hard to show that the problem remains easy when all weights are small (polynomially bounded).

4 Plurality, Veto, and k -approval

We will now show that the easiness results for k -approval and k -veto extend to arbitrary $k \leq m$ and weighted voters (note that the distinction between k -veto and $(m - k)$ -approval only matters for constant k).

Theorem 1. *For k -approval, the problems WISSAFE and WEXISTSAFE are in P.*

Proof. Fix a voter $v \in V$. To simplify notation, we renumber the candidates so that v 's preference order is given by $c_1 \succ_v \dots \succ_v c_m$. Denote v 's truthful vote by R . Recall that V_v denotes the set of voters who have the same preferences as v . Suppose that under truthful voting the winner is c_j . For $i = 1, \dots, m$, let $s_i(\mathcal{R}')$ denote the k -approval score of c_i given a profile \mathcal{R}' , and set $s_i = s_i(\mathcal{R})$.

We start by proving a useful characterization of safe strategic votes for k -approval.

Lemma 1. *A vote L is a safe strategic vote for v if and only if the winner in $\mathcal{R}_{-V_v}(L)$ is a candidate c_i with $i < j$.*

Proof. Suppose that L is a safe strategic vote for v . Then there exists an $i < j$ and a $U \subseteq V_v$ such that the winner in $\mathcal{R}_{-U}(L)$ is c_i . It must be the case that each switch from R to L increases c_i 's score or decreases c_j 's score: otherwise c_i cannot beat c_j after the voters in U change their vote from R to L . Therefore, if c_i beats c_j when the preference profile is $\mathcal{R}_{-U}(L)$, it continues to beat c_j after the remaining voters in V_v switch, i.e., when the preference profile is $\mathcal{R}_{-V_v}(L)$. Hence, the winner in $\mathcal{R}_{-V_v}(L)$ is not c_j ; since L is safe, this means that the winner in $\mathcal{R}_{-V_v}(L)$ is c_ℓ for some $\ell < j$.

For the opposite direction, suppose that the winner in $\mathcal{R}_{-V_v}(L)$ is c_i for some $i < j$. Note that if two candidates gain points when some subset of voters switches from R to L , they both gain the same number of points; the same holds if both of them lose points.

Now, if $j > k$, a switch from R to L does not lower the score of c_j , so it must increase the score of c_i for it to win in $\mathcal{R}_{-V_v}(L)$. Further, if a switch from R to L grants points to some $c_\ell \neq c_i$, then either $s_\ell < s_i$ or $s_\ell = s_i$ and the tie-breaking rule favors c_i over c_ℓ : otherwise, c_i would not be the winner in $\mathcal{R}_{-V_v}(L)$.

Similarly, if $j \leq k$, a switch from R to L does not increase the score of c_i , so it must lower the score of c_j . Further, if some $c_\ell \neq c_i$ does not lose points from a switch from R to L , then either $s_\ell < s_i$ or $s_\ell = s_i$ and the tie-breaking rule favors c_i over c_ℓ : otherwise, c_i would not be the winner in $\mathcal{R}_{-V_v}(L)$.

Now, consider any $U \subseteq V_v$. If $s_j(\mathcal{R}_{-U}(L)) > s_i(\mathcal{R}_{-U}(L))$, then c_j is the winner. If $s_i(\mathcal{R}_{-U}(L)) > s_j(\mathcal{R}_{-U}(L))$, then c_i is the winner. Finally, suppose $s_i(\mathcal{R}_{-U}(L)) = s_j(\mathcal{R}_{-U}(L))$. By the argument above, no other candidate can have a higher score. So, suppose that $s_\ell(\mathcal{R}_{-U}(L)) = s_i(\mathcal{R}_{-U}(L))$, and the tie-breaking rule favors c_ℓ over c_i and c_j . Then this would imply that c_ℓ wins in \mathcal{R} or $\mathcal{R}_{-V_v}(L)$ (depending on whether a switch from R to L causes c_ℓ to lose points), a contradiction. Thus, in this case, too, either c_i or c_j wins. \square

Lemma 1 immediately implies an algorithm for WISSAFE: we simply need to check that the input vote satisfies the conditions of the lemma. We now show how to use it to construct an algorithm for WEXISTSAFE. We need to consider two cases.

$\mathbf{j} > \mathbf{k}$: In this case, the voters in V_v already do not approve of c_j and approve of all $c_i, i \leq k$. Thus, no matter how they vote, they cannot ensure that some $c_i, i \leq k$, gets more points than c_j . Hence, the only way they can change the outcome is by approving of some candidate $c_i, k < i < j$. Further, they can only succeed if there exists an $i = k + 1, \dots, j - 1$ such that either $s_i + w(V_v) > s_j$ or $s_i + w(V_v) = s_j$ and the tie-breaking rule favors c_i over c_j . If such an i exists, v has an incentive to manipulate by swapping c_1 and c_i in his vote. Furthermore, it is easy to see that any such manipulation is safe, as it only affects the scores of c_1 and c_i .

$\mathbf{j} \leq \mathbf{k}$: In this case, the voters in V_v already approve of all candidates they prefer to c_j , and therefore they cannot increase the scores of the first $j - 1$ candidates. Thus, their only option is to try to lower the scores of c_j as well as those of all other candidates whose score currently matches or exceeds the best score among s_1, \dots, s_{j-1} .

Set $C_g = \{c_1, \dots, c_{j-1}\}$, $C_b = \{c_j, \dots, c_m\}$. Let C_0 be the set of all candidates in C_g whose k -approval score is maximal, and let s_{\max} be the k -approval score of the candidates in C_0 . For any $c_\ell \in C_b$, let s'_ℓ denote the number of points that c_ℓ gets from all voters in $V \setminus V_v$; we have $s'_\ell = s_\ell$ for $k < \ell \leq m$ and $s'_\ell = s_\ell - w(V_v)$ for $\ell = j, \dots, k$. Now, it is easy to see that v has a safe manipulation if and only if the following conditions hold:

- For all $c_\ell \in C_b$ either $s'_\ell < s_{\max}$, or $s'_\ell = s_{\max}$ and there exists a candidate $c \in C_0$ such that the tie-breaking rule favors c over c_ℓ ;
- There exist a set $C_{\text{safe}} \subseteq C_b$, $|C_{\text{safe}}| = k - j + 1$, such that for all $c_\ell \in C_{\text{safe}}$ either $s'_\ell + w(V_v) < s_{\max}$ or $s'_\ell + w(V_v) = s_{\max}$ and there exists a candidate $c \in C_0$ such that the tie-breaking rule favors c over c_ℓ .

Note that these conditions can be easily checked in polynomial time by computing s_ℓ and s'_ℓ for all $\ell = 1, \dots, m$.

Indeed, if such a set C_{safe} exists, voter v can place the candidates in C_{safe} in positions j, \dots, k in his vote; denote the resulting vote by L . Clearly, if all voters in V_v vote according to L , they succeed to elect some $c \in C_0$. Thus, by Lemma 1, L is safe. Conversely, if a set C_{safe} with these properties does not exist, then for any vote $L \neq R$ the winner in $\mathcal{R}_{-V_v}(L)$ is a candidate in C_b , and thus by Lemma 1 L is not safe. \square

We remark that Theorem 1 crucially relies on the fact that we break ties based on a fixed priority ordering over the candidates. Indeed, it can be shown that there exists a (non-lexicographic) tie-breaking rule such that finding a safe vote with respect to k -approval combined with this tie-breaking rule is computationally hard (assuming k is viewed as a part of the input). As the focus of this paper is on lexicographic tie-breaking, we omit the formal statement and the proof of this fact.

In contrast, we can show that any scoring rule with 3 candidates is easy to manipulate safely, even if the voters are weighted and arbitrary tie-breaking rules are allowed.

Theorem 2. $\text{WISSAFE}(\mathcal{F})$ is in P for any voting rule \mathcal{F} obtained by combining a positional scoring rule with at most three candidates with an arbitrary tie-breaking rule.

5 Bucklin and Borda

Bucklin rule is quite similar to k -approval, so we can use the ideas in the proof of Theorem 10 to design a polynomial-time algorithm for finding a safe manipulation with respect to Bucklin. However, the proof becomes significantly more complicated.

Theorem 3. *For the Bucklin rule, WEXISTSAFE is in P.*

Interestingly, despite the intuition that WISSAFE should be easier than WEXISTSAFE, it turns out that WISSAFE for Bucklin is coNP-hard.

Theorem 4. *For the Bucklin rule, WISSAFE is coNP-hard, even for a constant number of candidates.*

For Borda, unlike k -approval and Bucklin, both of our problems are hard when the voters are weighted.

Theorem 5. *For the Borda rule, WISSAFE and WEXISTSAFE are coNP-hard. The hardness result holds even if there are only 5 candidates.*

6 Extensions of the Safe Strategic Voting Model

So far, we followed the model of [10] and assumed that the only voters who may change their votes are the ones whose preferences exactly coincide with those of the manipulator. Clearly, in real life this assumption does not always hold. Indeed, a voter may follow a suggestion to vote in a certain way as long as it comes from someone he trusts (e.g., a well-respected public figure), and this does not require that this person's preferences are completely identical to those of the voter. For example, if both the original manipulator v and his would-be follower u rank the current winner last, it is easy to see that following v 's recommendation that leads to displacing the current winner is in u 's best interests.

In this section, we will consider two approaches to extending the notion of safe strategic voting to scenarios where not all manipulators have identical preferences. In both cases, we define the set of potential followers for each voter (in our second model, this set may depend on the vote suggested), and define a vote L to be safe if, whenever a subset of potential followers votes L , the outcome of the election does not get worse (and sometimes gets better) from the manipulator's perspective. However, our two models differ in the criteria they use to identify a voter's potential followers.

Preference-Based Extension. Our first model identifies the followers of a given voter based on the similarities in voters' preferences.

Fix a preference profile \mathcal{R} and a voting rule \mathcal{F} , and let c be the winner under truthful voting. For any $v \in V$, let $I(v, c)$ denote the set of candidates that v ranks strictly above c . We say that two voters u and v are *similar* if $I(u, c) = I(v, c)$. A *similar set* S_v of a voter v for a given preference profile \mathcal{R} and a voting rule \mathcal{F} is given by $S_v = \{u \mid I(u, c) = I(v, c)\}$. (The set S_v depends on \mathcal{R} and \mathcal{F} ; however, for readability we omit \mathcal{R} and \mathcal{F} from the notation).

Note that if u and v are similar, they rank c in the same position. Further, a change of outcome from c to another alternative is positive from u 's perspective if and only if it is positive from v 's perspective. Thus, intuitively, any manipulation that is profitable for u is also profitable for v . Observe also that similarity is an equivalence relation, and the sets S_v are the corresponding equivalence classes. In particular, this implies that for any $u, v \in V$ either $S_u = S_v$ or $S_u \cap S_v = \emptyset$.

We can now adapt Definition 1 to our setting by replacing V_v with S_v .

Definition 2. A vote L is a strategic vote in the preference-based extension for v at \mathcal{R} under \mathcal{F} if for some $U \subseteq S_v$ we have $\mathcal{F}(\mathcal{R}_{-U}(L)) \succ_v \mathcal{F}(\mathcal{R})$. Further, we say that L is a safe strategic vote in the preference-based extension for a voter v at \mathcal{R} under \mathcal{F} if L is a strategic vote at \mathcal{R} under \mathcal{F} , and for any $U \subseteq S_v$ either $\mathcal{F}(\mathcal{R}_{-U}(L)) \succ_v \mathcal{F}(\mathcal{R})$ or $\mathcal{F}(\mathcal{R}_{-U}(L)) = \mathcal{F}(\mathcal{R})$.

Observe that if L is a (safe) strategic vote for v at \mathcal{R} under \mathcal{F} , then it is also a (safe) strategic vote for any $u \in S_v$. Indeed, $u \in S_v$ implies $S_u = S_v$ and for any $a \in C$ we have $a \succ_u \mathcal{F}(\mathcal{R})$ if and only if $a \succ_v \mathcal{F}(\mathcal{R})$. Note also that we do not require $L \neq R_v$: indeed, in the preference-based extension $L = R_v$ may be a non-trivial manipulation, as it may induce voters in $S_v \setminus \{v\}$ to switch their preferences to R_v . That is, a voter may manipulate the election simply by asking other voters with similar preferences to vote like he does. Finally, it is easy to see that for any voter v , the set S_v of similar voters is easy to compute.

The two computational problems considered throughout this paper, i.e., the safety of a given manipulation and the existence of a safe manipulation remain relevant for the preference-based model. We will refer to these problems in this setting as ISSAFE^{pr} and EXISTSAFE^{pr} , respectively, and use prefix w to denote their weighted variants. The problems $(w)\text{ISSAFE}^{pr}$ and $(w)\text{EXISTSAFE}^{pr}$ appear to be somewhat harder than their counterparts in the original model. Indeed, while voters in S_v have similar preferences, their truthful votes may be substantially different, so it now matters *which* of the voters in S_v decide to follow the manipulator (rather than just *how many* of them, as in the original model). In particular, it is not clear if $\text{ISSAFE}^{pr}(\mathcal{F})$ is polynomial-time solvable for any voting rule \mathcal{F} . However, it turns out that both of our problems are easy for k -approval, even with weighted voters.

Theorem 6. For k -approval, the problems $w\text{ISSAFE}^{pr}$ and $w\text{EXISTSAFE}^{pr}$ are in P .

In the preference-based model, a voter v follows a recommendation to vote in a particular way if it comes from a voter whose preferences are similar to those of v . However, this approach does not describe settings where a voter follows a recommendation not so much because he trusts the recommender, but for pragmatic purposes, i.e., because the proposed manipulation advances her own goals. Clearly, this may happen even if the overall preferences of the original manipulator and the follower are substantially different. We will now propose a model that aims to capture this type of scenarios.

Goal-Based Extension. If the potential follower's preferences are different from those of the manipulator, his decision to join the manipulating coalition is likely to depend on the specific manipulation that is being proposed. Thus, in this subsection we will define the set of potential followers F in a way that depends both on the original manipulator's

identity i and his proposed vote L , i.e., we have $F = F_i(L)$. Note, however, that it is not immediately obvious how to decide whether a voter j can benefit from following i 's suggestion to vote L , and thus should be included in the set $F_i(L)$. Indeed, the benefit to j depends on which other voters are in the set $F_i(L)$, which indicates that the definition of the set $F_i(L)$ has to be self-referential.

In more detail, for a given voting rule \mathcal{F} , an election (C, V) with a preference profile \mathcal{R} , a voter $i \in V$ and a vote L , we say that a voter j is *pivotal for a set* $U \subseteq V$ with respect to (i, L) if $j \notin U$, $R_j \neq L$ and $\mathcal{F}(\mathcal{R}_{-(U \cup \{j\})}(L)) \succ_j \mathcal{F}(\mathcal{R}_{-U}(L))$. That is, a voter j is pivotal for a set U if when the voters in U vote according to L , it is profitable for j to join them. Now, it might appear natural to define the follower set for (i, L) as the set that consists of i and all voters $j \in V$ that are pivotal with respect to (i, L) for some set $U \subseteq V$. However, this definition is too broad: a voter is included as long as it is pivotal for some subset $U \subseteq V$, even if the voters in U cannot possibly benefit from voting L . To exclude such scenarios, we need to require that U itself is also drawn from the follower set. Formally, we say that $F_i(L)$ is a *follower set* for (i, L) if it is a maximal set F that satisfies the following condition:

$$\forall j \in F [(j = i) \vee (\exists U \subseteq F \text{ s. t. } j \text{ is pivotal for } U \text{ with respect to } (i, L))] \quad (*)$$

Observe that this means that $F_i(L)$ is a fixed point of a mapping from 2^V to 2^V , i.e., this definition is indeed self-referential. To see that the follower set is uniquely defined for any $i \in V$ and any vote L , note that the union of any two sets that satisfy condition [\(28\)](#) also satisfies [\(28\)](#); note also that we always have $i \in F_i(L)$.

We can now define what it means for L to be a *strategic vote in the goal-based extension* and a *safe strategic vote in the goal-based extension* by replacing the condition $U \subseteq S_i$ in Definition [2](#) with $U \subseteq F_i(L)$. We will denote the computational problems of checking whether a given vote is a safe strategic vote for a given voter in the goal-based extension and whether a given voter has a safe strategic vote in the goal-based extension by ISSAFE^{gl} and EXISTSAFE^{gl} , respectively, and use the prefix w to refer to weighted versions of these problems.

Two remarks are in order. First, it may be the case that even though i benefits from proposing to vote L , he is never pivotal with respect to (i, L) (this can happen, e.g., if i 's weight is much smaller than that of the other voters). Thus, we need to explicitly include i in the set $F_i(L)$, to avoid the paradoxical situation where $i \notin F_i(L)$. Second, our definition of a safe vote only guarantees safety to the original manipulator, but not to her followers. In contrast, in the preference-based extension, any vote that is safe for the original manipulator is also safe for all similar voters.

The definition of a safe strategic vote in the goal-based extension captures a number of situations not accounted for by the definition of a safe strategic vote in the preference-based extension. However, computationally it is considerably harder to deal with than the preference-based extension. Indeed, it is not obvious how to compute the set $F_i(L)$, as its definition is non-algorithmic in nature. While one can consider all subsets of V and check whether they satisfy condition [\(28\)](#), this approach is obviously inefficient. We can avoid full enumeration if we have access to a procedure $\mathcal{A}(i, L, j, W)$ that for each pair (i, L) , each voter $j \in V$ and each set $W \subseteq V$ can check if $j = i$ or there is a set $U \subseteq W$ such that j is pivotal for U with respect to (i, L) . Indeed, if this is the case, we

can compute $F_i(L)$ as follows. We start with $W = V$, run $\mathcal{A}(i, L, j, W)$ for all $j \in W$, and let W' to be the set of all voters for which $\mathcal{A}(i, L, j, W)$ outputs “yes”. We then set $W = W'$, and iterate this step until $W = W'$. In the end, we set $F_i(L) = W$. The correctness of this procedure can be proven by induction on the number of iterations and follows from the fact that if a set W contains no subset U that is pivotal for j , then no smaller set $W' \subset W$ can contain such a subset. Moreover, since each iteration reduces the size of W , the process converges after at most n iterations. Thus, this algorithm runs in polynomial time if the procedure $\mathcal{A}(i, L, j, W)$ is efficiently implementable. We will now show that this is indeed the case for Plurality (with unweighted voters).

Theorem 7. *Given an election (C, V) with a preference profile \mathcal{R} and unweighted voters, a manipulator i , and a vote L , we can compute the set $F_i(L)$ with respect to Plurality in time polynomial in the input size.*

We can use Theorem 7 to show that under Plurality one can determine in polynomial time whether a given vote L is safe for a player i , as well as find a safe strategic vote for i if one exists, as long as the voters are unweighted.

Theorem 8. *ISSAFE^{gl} and EXISTS SAFE^{gl} are polynomial-time solvable for Plurality.*

For weighted voters, computing the follower set is hard even for Plurality. While this result does not directly imply that WISSAFE^{gl} and WEXISTS SAFE^{gl} are also hard for Plurality, it indicates that these problems are unlikely to be easily solvable.

Theorem 9. *Given an instance $(C, V, \mathbf{w}, \mathcal{R})$ of Plurality elections, voters $i, j \in V$ and a vote L , it is NP-hard to decide whether $j \in F_i(L)$.*

Just a little further afield, checking whether a given vote is safe with respect to 3-approval is computationally hard even for unweighted voters. This is in contrast with the standard model and the preference-based extension, where safely manipulating k -approval is easy for arbitrary k .

Theorem 10. *ISSAFE^{gl} is coNP-hard for 3-approval.*

Thus, while the preference-based extension appears to be similar to the original model of [10] from the computational perspective, the goal-based extension is considerably more difficult to work with.

7 Conclusions

In this paper, we started the investigation of algorithmic complexity of safe manipulation, as defined by Slinko and White [10]. We showed that finding a safe manipulation is easy for k -approval for an arbitrary value of k and for Bucklin, even with weighted voters. Somewhat surprisingly, checking whether a given manipulation is safe appears to be a more difficult problem, at least for weighted voters: while this problem is polynomial-time solvable for k -approval, it is coNP-hard for Bucklin. For the Borda rule, both checking whether a given manipulation is safe and identifying a safe manipulation is hard when the voters are weighted.

We also proposed two ways of extending the notion of safe manipulation to heterogeneous groups of manipulators, and initiated the study of computational complexity of related questions. Our first extension of the model of [10] is very simple and natural, and seems to behave similarly to the original model from the algorithmic perspective. However, arguably, it does not capture some of the scenarios that may occur in practice. Our second model is considerably richer, but many of the associated computational problems become intractable.

A natural open question is determining the complexity of finding a safe strategic vote for voting rules not considered in this paper, such as Copeland, Ranked Pairs, or Maximin. Moreover, for some of the voting rules we have investigated, the picture given by this paper is incomplete. In particular, it would be interesting to understand the computational complexity of finding a safe manipulation for Borda (and, more generally, for all scoring rules) for unweighted voters. The problem for Borda is particularly intriguing as this is perhaps the only widely studied voting rule for which the complexity of unweighted coalitional manipulation in the standard model is not known.

Other exciting research directions include formalizing and investigating the problem of selecting the best safe manipulation (is it the one that succeeds more often, or one that achieves better results when it succeeds?), and extending our analysis to other types of tie-breaking rules, such as, e.g., randomized tie-breaking rules. However, the latter question may require modifying the notion of a safe manipulation, as the outcome of a strategic vote becomes a probability distribution over the alternatives.

References

1. Bartholdi III, J.J., Tovey, C.A., Trick, M.: The computational difficulty of manipulating an election. *Social Choice and Welfare* 6, 227–241 (1989)
2. Bartholdi III, J.J., Orlin, J.B.: Single transferable vote resists strategic voting. *Social Choice and Welfare* 8(4), 341–354 (1991)
3. Conitzer, V., Sandholm, T., Lang, J.: When are elections with few candidates hard to manipulate? *J. ACM* 54, 1–33
4. Faliszewski, P., Hemaspaandra, E., Schnoor, H.: Copeland voting: ties matter. In: *AAMAS 2008* (2008)
5. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, New York (1979)
6. Gibbard, A.F.: Manipulation of voting schemes: a general result. *Econometrica* 41, 597–601 (1973)
7. Moulin, H.: Choice functions over a finite set: a summary. *Social Choice and Welfare* 2, 147–160 (1985)
8. Papadimitriou, C.: *Computational Complexity*. Addison-Wesley, Reading (1994)
9. Satterthwaite, M.A.: Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10, 187–217 (1975)
10. Slinko, A., White, S.: Non-dictatorial social choice rules are safely manipulable. In: *COMSOC 2008* (2008)
11. Xia, L., Zuckerman, M., Procaccia, A.D., Conitzer, V., Rosenschein, J.S.: Complexity of unweighted coalitional manipulation under some common voting rules. In: *IJCAI 2009* (2009)
12. Zuckerman, M., Procaccia, A.D., Rosenschein, J.S.: Algorithms for the coalitional manipulation problem. *Artificial Intelligence* 173(2), 392–412 (2009)

Bottleneck Congestion Games with Logarithmic Price of Anarchy^{*}

Rajgopal Kannan and Costas Busch

Department of Computer Science, Louisiana State University, Baton Rouge, LA 70803, USA
rkannan@csc.lsu.edu, busch@csc.lsu.edu,
www.csc.lsu.edu/~rkannan, www.csc.lsu.edu/~busch

Abstract. We study *bottleneck congestion games* where the social cost is determined by the worst congestion on any resource. In the literature, bottleneck games assume player utility costs determined by the worst congested resource in their strategy. However, the Nash equilibria of such games are inefficient since the price of anarchy can be very high and proportional to the number of resources. In order to obtain smaller price of anarchy we introduce *exponential bottleneck games*, where the utility costs of the players are exponential functions of their congestions. In particular, the delay function for any resource r is \mathcal{M}^{C_r} , where C_r denotes the number of players that use r , and \mathcal{M} is an integer constant. We find that exponential bottleneck games are very efficient and give the following bound on the price of anarchy: $O(\log |R|)$, where R is the set of resources. This price of anarchy is tight, since we demonstrate a game with price of anarchy $\Omega(\log |R|)$. We obtain our tight bounds by using two novel proof techniques: *transformation*, which we use to convert arbitrary games to simpler games, and *expansion*, which we use to bound the price of anarchy in a simpler game.

1 Introduction

We consider non-cooperative congestion games with n players, where each player has a *pure strategy profile* from which it selfishly selects a strategy that minimizes the player's utility cost function (such games are also known as *atomic* or *unsplittable-flow* games). We focus on *bottleneck congestion games* where the objective for the social outcome is to minimize C , the maximum congestion on any resource. Typically, the congestion on a resource is a non-decreasing function on the number of players that use the resource; here, we consider the congestion to be simply the number of players that use the resource.

Bottleneck congestion games have been studied in the literature [14, 3] in the context of routing games, where each player's utility cost is the worst resource congestion on its strategy. In particular, player i has utility cost function $C_i = \max_{r \in S_i} C_r$, where S_i is the strategy of the player and C_r denotes the congestion of resource r . Note that $C = \max_i C_i$. In [1] the authors observe that bottleneck games are important in networks for various practical reasons. In networks, each resource corresponds to a network link, each

^{*} This work was supported in part by NSF grants #CNS-1018273, #IIS-0905478 and a BBN subcontract from #CNS-0940805.

player corresponds to a packet, and a strategy represents a path for the packet. In wireless networks, the maximum congested link is related to the lifetime of the network since the nodes adjacent to high congestion links transmit large number of packets which results to higher energy utilization. High congestion links also result to congestion hot-spots which may slow-down the network throughput. Hot spots also increase the vulnerability of the network to malicious attacks which aim to to increase the congestion of links in the hope to bring down the network. Thus, minimizing the maximum congested edge results to hot-spot avoidance and more load-balanced and secure networks.

In networks, bottleneck games are also important from a theoretical point of view since the maximum resource congestion is immediately related to the optimal packet scheduling. In a seminal result, Leighton *et al.* [13] showed that there exist packet scheduling algorithms that can deliver the packets along their chosen paths in time very close to $C + D$, where D is the maximum chosen path length. When $C \gg D$, the congestion becomes the dominant factor in the packet scheduling performance. Thus, smaller C immediately implies faster packet delivery time.

A natural problem that arises concerns the effect of the players' selfishness on the welfare of the whole system measured with the *social cost* C . We examine the consequence of the selfish behavior in pure *Nash equilibria* which are stable states of the game in which no player can unilaterally improve her situation. We quantify the effect of selfishness with the *price of anarchy* (PoA) [12][18], which expresses how much larger is the worst social cost in a Nash equilibrium compared to the social cost in the optimal coordinated solution. The price of anarchy provides a measure for estimating how closely do Nash equilibria of bottleneck routing games approximate the optimal C^* of the respective coordinated optimization problem.

Ideally, the price of anarchy should be small. However, the current literature results have only provided weak bounds for bottleneck games. In [1] it is shown that if the resource congestion delay function is bounded by some polynomial with degree k (with respect to the packets that use the resource) then $PoA = O(|R|^k)$, where R is the set of links (resources) in the graph. In [4] the authors consider bottleneck routing games for the case $k = 1$ and they show that $PoA = O(L + \log |V|)$, where L is the maximum path length (maximum number of resources) in the players' strategies and V is the set of nodes in the network. This bound is asymptotically tight since there are game instances with $PoA = \Omega(L)$. Note that L can be proportional to $|R|$ (when $L = |R| = O(|V|)$), and thus the price of anarchy can be large.

1.1 Contributions

The lower bound in [4] suggests that in order to obtain better price of anarchy in bottleneck games (where the social cost is the bottleneck congestion C), we need to consider alternative player utility cost functions. Towards this goal, we introduce *exponential bottleneck games* where the social cost is the bottleneck C and the player cost functions are exponential expressions of the congestions along the resources. In particular, the player utility cost function for player i is: $C'_i = \sum_{r \in S_i} \mathcal{M}^{C_r}$, for some integer constant $\mathcal{M} \geq 2$. Note that the new utility cost is a sum of exponential terms on the congestion of the resources in the chosen strategy (instead of the max that we described

earlier). For the bottleneck social cost C we prove that the price of anarchy of exponential games is:

$$PoA = O(\log |R|).$$

We show that this bound is tight by providing an instance of an exponential bottleneck game with $PoA = \Omega(\log |R|)$. Our price of anarchy bound is a significant improvement over the price of anarchy from the regular utility cost functions described above.

Exponential games are interesting variations of bottleneck games not only because they can provide good price of anarchy but also because they represent real-life problems. It is well shown that in wireless communication networks the power used by individual nodes to transmit messages along an edge with guaranteed rate is exponentially proportional to the flow of the link. Thus, exponential game equilibria represent power game equilibria in wireless networks, where small price of anarchy translates to small power utilization by the nodes. Exponential cost functions on resource congestion have been used before in a different context for online routing optimization problems [2][Chapter 13]. However, here we use the exponential functions for the first time in the context of congestion games.

In our analysis, we obtain the price of anarchy upper bound by developing two new techniques: *transformation* and *expansion*. Consider a game G with a Nash equilibrium S and congestion C . We identify two kinds of players in S : type-A players which use only one resource in their strategies, and type-B players which use two or more resources. In our first technique, transformation, we show how to convert G to a simpler game \tilde{G} , having a Nash equilibrium \tilde{S} with congestion \tilde{C} , such that $\tilde{C} = O(C)$, and all players in \tilde{S} with congestion above a threshold T are of type-A; that is, we transform type-B players to type-A players. Having type-A players is easier to bound the price of anarchy. Then, we develop a second technique, expansion, that is used to give an upper bound on the price of anarchy of game \tilde{G} , which implies an upper bound on the price of anarchy of the original game G .

1.2 Related Work

Congestion games were introduced and studied in [17][19]. In [19], Rosenthal proves that congestion games have always pure Nash equilibria. Koutsoupias and Papadimitriou [12] introduced the notion of price of anarchy in the specific *parallel link networks* model in which they provide the bound $PoA = 3/2$. Roughgarden and Tardos [22] provided the first result for splittable flows in general networks in which they showed that $PoA \leq 4/3$ for a player cost which reflects to the sum of congestions of the resources of a path. Pure equilibria with atomic flow have been studied in [4][5][14][24] (our work fits into this category), and with splittable flow in [20][21][22][23]. Mixed equilibria with atomic flow have been studied in [7][9][11][12][15][16][18], and with splittable flow in [6][8].

Most of the work in the literature uses a cost metric measured as the sum of congestions of all the resources of the player's path [5][21][22][23][24]. Our work differs from these approaches since we adopt the exponential metric for player cost. The vast majority of the work on congestion games has been performed for parallel link networks, with only a few exceptions on network topologies [4][5][6][20], which we consider here.

In [4], the authors consider bottleneck routing games in networks with player cost C_i and social cost C . They prove that the price of stability is 1. They show that the price of

anarchy is bounded by $O(L + \log |V|)$, where L is the maximum allowed path length, and V is the set of nodes. They also prove that $\kappa \leq PoA \leq c(\kappa^2 + \log^2 |V|)$, where κ is the size of the largest resource-simple cycle in the graph and c is a constant. That work was extended in [3] to the $C + D$ routing problem. Bottleneck routing games have also been studied in [1], where the authors consider the maximum congestion metric in general networks with splittable and atomic flow (but without considering path lengths). They prove the existence and non-uniqueness of equilibria in both the splittable and atomic flow models. They show that finding the best Nash equilibrium that minimizes the social cost is a NP-hard problem. Further, they show that the price of anarchy may be unbounded for specific resource congestion functions. In [10], the authors prove the existence of strong Nash equilibria (which concern coalitions of players) for games with the lexicographic improvement property; such games include Bottleneck congestion games and our exponential games.

2 Definitions

A *congestion game* is a strategic game $G = (\Pi_G, R, \mathbb{S}, (d_r)_{r \in R}, (pc_\pi)_{\pi \in \Pi_G})$ where:

- $\Pi_G = \{\pi_1, \dots, \pi_n\}$ is a non-empty and finite set of players.
- $R = \{r_1, \dots, r_z\}$ is a non-empty and finite set of resources.
- $\mathbb{S} = \mathbb{S}_{\pi_1} \times \mathbb{S}_{\pi_2} \times \dots \times \mathbb{S}_{\pi_n}$, where \mathbb{S}_{π_i} is a strategy set for player π_i , such that $\mathbb{S}_{\pi_i} \subseteq \text{powerset}(R)$; namely, each strategy $S_{\pi_i} \in \mathbb{S}_{\pi_i}$ is pure, and it is a collection of resources. A *game state* (or *pure strategy profile*) is any $S \in \mathbb{S}$. We consider *finite games* which have finite \mathbb{S} (finite number of states).
- In any game state S , each resource $r \in R$ has a *delay cost* denoted $d_r(S)$.
- In any game state S , each player $\pi \in \Pi_G$ has a *player cost* $pc_\pi(S) = \sum_{r \in S_\pi} d_r(S)$.

Consider a game G with a state $S = (S_{\pi_1}, \dots, S_{\pi_n})$. The (*congestion*) of a resource r is defined as $C_r(S) = |\{\pi_i : r \in S_{\pi_i}\}|$, which is the number of players that use r in state S . The (*bottleneck*) *congestion of a set of resources* $Q \subseteq R$ is defined as $C_Q(S) = \max_{r \in Q} C_r(S)$, which is the maximum congestion over all resources in Q . The (*bottleneck*) *congestion of state* S is denoted $C(S) = C_R(S)$, which is the maximum congestion over all resources in R . The *length of state* S is defined to be $L(S) = \max_i |S_{\pi_i}|$, namely, the maximum number of resources used in any player. When the context is clear, we will drop the dependence on S . We examine exponential congestion games:

- *Exponential games*: The delay cost function for any resource r is $d_r = \mathcal{M}^{C_r}$, for some integer constant $\mathcal{M} \geq 2$.

For any state S , we use the standard notation $S = (S_{\pi_i}, S_{-\pi_i})$ to emphasize the dependence on player π_i . Player π_i is *locally optimal* (or *stable*) in state S if $pc_{\pi_i}(S) \leq pc_{\pi_i}((S'_{\pi_i}, S_{-\pi_i}))$ for all strategies $S'_{\pi_i} \in \mathbb{S}_{\pi_i}$. A greedy move by a player π_i is any change of its strategy from S'_{π_i} to S_{π_i} which improves the player's cost, that is, $pc_{\pi_i}((S_{\pi_i}, S_{-\pi_i})) < pc_{\pi_i}((S'_{\pi_i}, S_{-\pi_i}))$. *Best response dynamics* are sequences of greedy moves by players. A state S is in a *Nash Equilibrium* if every player is locally

optimal. Nash Equilibria quantify the notion of a stable selfish outcome. In the games that we study there could exist multiple Nash Equilibria.

For any game G and state S , we will consider a *social cost* (or *global cost*) which is simply the bottleneck congestion $C(S)$. A state S^* is called *optimal* if it has minimum attainable social cost: for any other state S , $C(S^*) \leq C(S)$. We will denote $C^* = C(S^*)$. We quantify the quality of the states which are Nash Equilibria with the *price of anarchy* (*PoA*) (sometimes referred to as the coordination ratio). Let \mathcal{P} denote the set of distinct Nash Equilibria. Then the price of anarchy of game G is:

$$PoA(G) = \sup_{S \in \mathcal{P}} \frac{C(S)}{C^*},$$

We continue with some more special definitions that we use in the proofs. Consider a game G with a socially optimal state $S^* = (S_{\pi_1}^*, \dots, S_{\pi_n}^*)$, and let $S = (S_{\pi_1}, \dots, S_{\pi_n})$ denote the equilibrium state. We consider two special kinds of players with respect to states S and S^* :

- *Type-A players*: any player π_i with $|S_{\pi_i}| = 1$.
- *Type-B players*: any player π_i with $|S_{\pi_i}| \geq 2$.

For any resource $r \in R$, we will let Π_r and Π_r^* denote the set of players with r in their equilibrium and socially optimal strategies respectively, i.e $\Pi_r = \{\pi_i \in \Pi_G | r \in S_{\pi_i}\}$ and $\Pi_r^* = \{\pi_i \in \Pi_G | r \in S_{\pi_i}^*\}$.

Let $G = (\Pi_G, R, \mathbb{S}, d, (pc_\pi)_{\pi \in \Pi_G})$ and $\tilde{G} = (\Pi_{\tilde{G}}, \tilde{R}, \tilde{\mathbb{S}}, \tilde{d}, (\tilde{p}c_\pi)_{\pi \in \Pi_{\tilde{G}}})$ be two games. We say that G *dominates* \tilde{G} if the following conditions hold between them for the highest cost Nash equilibrium and optimal states and : $|\tilde{R}| \leq |R|$, $d = \tilde{d}$, $\tilde{C} = C$, $\tilde{C}^* \leq C^* \mathcal{M}^2$, where C, C^* and \tilde{C}, \tilde{C}^* represent the bottleneck congestions in the highest cost Nash equilibrium and optimal states of G and \tilde{G} , respectively.

Corollary 1. $PoA(G) \leq \mathcal{M}^2 \cdot PoA(\tilde{G})$ for an arbitrary game G and dominated game \tilde{G} .

Let $S, S^*, \tilde{S}, \tilde{S}^*$ denote the equilibrium and socially optimal states of G and \tilde{G} . Henceforth for notational convenience and when there is no ambiguity, for player specific states, we will drop the π subscript and use S_i, \tilde{S}_i, \dots for $S_{\pi_i}, \tilde{S}_{\pi_i} \dots$. In the next section, we will describe how an arbitrary game G in Nash equilibrium state S can be transformed into a dominated game \tilde{G} containing type *A* players of arbitrary cost and type *B* players restricted to costs below a given threshold in Nash equilibrium state \tilde{S} .

3 Type-B to Type-A Game Transformation

We first state our main result in this section.

Theorem 1. *Every game G with arbitrary type-B players can be transformed into a dominated game \tilde{G} where all players with costs $\geq \psi = \mathcal{M}^{(\mathcal{M}+2)C^*}$ are exclusively type-A.*

In order to prove Theorem 1, we first require the following results on base \mathcal{M} arithmetic and resource partitioning.

Lemma 1. *Let a, i_1, i_2, \dots, i_a be non-negative integers such that $\sum_{j=1}^a i_j \mathcal{M}^{a-j} \geq \mathcal{M}^a$. Then $\exists k, l_k$ with $k \leq a, l_k \leq i_k$ such that $\sum_{j=1}^{k-1} i_j \mathcal{M}^{a-j} + l_k \mathcal{M}^{a-k} = \mathcal{M}^a$.*

We will use this lemma in the following context: For any integer a , given a collection of $i_j \geq 0$ resources of cost \mathcal{M}^{a-j} , for $1 \leq j \leq a$ and total cost at least \mathcal{M}^a , we can always select a subset of resources whose cost will add up to exactly \mathcal{M}^a .

Let $\pi_i \in \Pi_G$ be an arbitrary type B -player using k resources r_1, r_2, \dots, r_k in its equilibrium strategy S_i that are *distinct* from the m resources r_1^*, \dots, r_m^* in its socially optimal strategy S_i^* , with $C_{r_j}(S), C_{r_j^*}(S)$ denoting the congestion on these resources in equilibrium state S

Define procedure **PMS – Partition**(π_i) as follows:

Procedure 1 *Partition S_i and S_i^* into t pairs $(L_1, L_1^*), (L_2, L_2^*), \dots, (L_t, L_t^*)$ where 1) the L_j 's form a disjoint resource partition of S_i , i.e $L_k \cap L_l = \emptyset$ with $\bigcup_{j=1}^t L_j = S_i$, 2) the L_j^* 's are disjoint subsets of S_i^* , i.e $L_k^* \cap L_l^* = \emptyset$ with $\bigcup_{j=1}^t L_j^* \subseteq S_i^*$, and 3)*

$$\sum_{r^* \in L_j^*} \mathcal{M}^{C_{r^*}+1} \geq \sum_{r \in L_j} \mathcal{M}^{C_r}, \quad 1 \leq j \leq t \tag{1}$$

Lemma 2. *There exists an implementation of **PMS – Partition**(π_i) in which either $|L_j| = 1$ or $|L_j^*| = 1$ or both, $j = 1, 2, \dots, t$. In the case $|L_j| > 1$ and $|L_j^*| = 1$ with $L_j^* = \{r_p^*\}$ for some $r_p^* \in S_i^*$, we must have $C_{r_p^*} \geq \max\{C_{r_q} | r_q \in L_j\}$.*

Due to space considerations, we omit the proof above.

Procedure **PMS – Partition**() forms the basic step in our transformation of G to \tilde{G} . We start with a restricted version of game G labeled \tilde{G} and iteratively transform it by using procedure **PMS – Partition**() to convert type- B players of cost $\psi \geq \mathcal{M}^{(\mathcal{M}+2)C^*}$ into type- A players, one at a time in decreasing order of player costs until all type- B players remaining either fall below the threshold cost function ψ or no type- B players exist. We add and delete players/resources from \tilde{G} iteratively and have a working set of players. However \tilde{G} will always remain in equilibrium state \tilde{S} at every step of the transformation process. Note that when a new player π_k is created, we assign two strategy sets to π_k : an ‘equilibrium’ strategy \tilde{S}_k and a socially optimal strategy \tilde{S}_k^* . Thus $\tilde{S} = \tilde{S} \cup \tilde{S}_k$ and $\tilde{S}^* = \tilde{S}^* \cup \tilde{S}_k^*$.

We are now ready to prove our main result.

Proof of Theorem 1: We assume that our input is the restricted version \tilde{G} of game G with exactly two strategies per player: $\tilde{S}_\pi = S_\pi$ and $\tilde{S}_\pi^* = S_\pi^*$. We also assume that each type- B player π has distinct resources in its equilibrium and optimal strategies i.e $\tilde{S}_\pi \cap \tilde{S}_\pi^* = \emptyset$. If not already true, this can be achieved by creating $|\tilde{S}_\pi \cap \tilde{S}_\pi^*|$ new type- A players with identical and one type- B player with disjoint equilibrium and optimal strategies for each original player π .

Let $K = \max_{\pi_j \in \Pi_{\tilde{G}}} \text{of type-}B \{pc_{\pi_j}(\tilde{S})\}$, where $K \geq \psi$. This implies that all players with cost $> K$ in \tilde{G} are type- A players. To begin the transformation,

choose any type- B player π_i with cost K . Let $\tilde{S}_i = (r_1, r_2, \dots, r_k)$ and $\tilde{S}_i^* = (r_1^*, \dots, r_m^*)$ denote the *distinct* resources in π_i 's equilibrium and optimal strategies and let $(L_1, L_1^*), \dots, (L_t, L_t^*)$ denote the output of PMS-Partition(π_i).

Consider the case when the number of partitions $t > 1$. First, delete the strategies of π_i from \tilde{G} i.e $\tilde{S} = \tilde{S} - \tilde{S}_i$ and $\tilde{S}^* = \tilde{S}^* - \tilde{S}_i^*$. The equilibrium congestion on \tilde{S}_i and optimal congestion on \tilde{S}_i^* now decrease by one in \tilde{G} . For each partition member L_q , with $|L_q| = 1, 1 \leq q \leq t$, we now add a new type- A player π_{L_q} to \tilde{G} with two strategy sets: equilibrium strategy $\tilde{S}_{\pi_{L_q}} = \{r_a\}$ where $L_q = \{r_a\}$, for some $a, 1 \leq a \leq k$ and optimal strategy $\tilde{S}_{\pi_{L_q}}^* = L_q^*$, where $L_q^* \subseteq \tilde{S}_i^*$. By claim [2](#) $\mathcal{M} \cdot \sum_{r \in L_q^*} \mathcal{M}^{C_r^*} \geq \mathcal{M}^{C_{r_a}}$ and thus π_{L_q} is in equilibrium in state $\tilde{S}_{\pi_{L_q}}$. Similarly, for each partition member $L_q \subseteq \tilde{S}_i$ with $|L_q| > 1, 1 \leq q \leq t$, we add a new type- B player π_{L_q} to \tilde{G} with equilibrium strategy $\tilde{S}_{\pi_{L_q}} = L_q$ and socially optimal strategy $\tilde{S}_{\pi_{L_q}}^* = \{r_a\}$ where $L_q^* = \{r_a\}$. Note that since $t > 1$, the cost $pc_{\pi_{L_q}}(\tilde{S}) < K$ for all $1 \leq q \leq t$ and thus π_i has been transformed into new type- A and type- B players of lower cost. The new type- B players will be considered for further transformation during subsequent iterations. Also note that congestion in the equilibrium and socially optimal states on resources in \tilde{G} is unchanged after adding these new players.

Now consider the case $t = 1$ where there is only one partition pair (L_1, L_1^*) with $L_1 = \tilde{S}_{\pi_i}$ and $L_1^* = \{r_l^*\} = \tilde{S}_{\pi_i}^*$. Let $|L_1| = \alpha$, where $1 \leq \alpha \leq |R|$. We can replace π_i with α new type- A players each containing a distinct resource in L_1 in their equilibrium strategies and r_l^* as their optimal strategy. However this might increase the optimal congestion \tilde{C}^* of \tilde{G} to as much as $C^* + |R|$, thereby violating the domination of \tilde{G} . Thus we need to find a larger optimal strategy set from among existing resources and assign them to these players without increasing the optimal congestion beyond $C^* \mathcal{M}^2$. Finding such a set forms the core of our transformation algorithm.

We first note that $C_{r_l^*} \geq \lceil \log_{\mathcal{M}} K - 1 \rceil$ and every resource r with $C_r > C_{r_l^*}$ must be occupied only by type- A players in equilibrium. Let $\Pi_{r_l^*}$ and $\Pi_{r_l^*}^*$ denote the players in \tilde{G} with r_l^* in their equilibrium and optimal strategy sets respectively, where $|\Pi_{r_l^*}| = C_{r_l^*}$ and $|\Pi_{r_l^*}^*| \leq C^*$. Let $\Pi'_{r_l^*} = \Pi_{r_l^*} - \Pi_{r_l^*}^*$ denote the set of players in $\Pi_{r_l^*}$ for whom r_l^* is **not** also in their optimal strategy, where $|\Pi'_{r_l^*}| \geq C_{r_l^*} - C^* + 1 > (\mathcal{M} + 1)C^*$. We will obtain a larger optimal strategy set for π_i by considering the optimal strategy sets of some of these $\Pi'_{r_l^*}$ players. Let $\Pi'_{r_l^*} = \{\pi_{i_1}, \pi_{i_2}, \dots, \pi_{i_r}\}$, where $r > (\mathcal{M} + 1)C^*$. We will construct a resource chain that yields a larger optimal strategy set for π_i from one out of these r players. Without loss of generality, we describe the chain for player $\pi(i_1)$.

Let PMS-chain(r_l^*, π_{i_1}) = $r_{j_1} \rightarrow r_{j_2} \rightarrow \dots \rightarrow r_{j_q}$ denote a chain of q *distinct* resources starting from $r_{j_1} = r_l^*$ and satisfying the following constraints:

1. r_{j_k} and $r_{j_{k+1}}, 1 \leq k \leq q - 1$, are in partition pairs (L_a, L_a^*) , with $|L_a| > 1, |L_a^*| = 1$ in the PMS-Partition of some type- B player $\hat{\pi}_{j_k} \in \Pi_{r_{j_k}}$, where $\Pi_{r_{j_k}}$ denotes the set of players for whom r_{j_k} is in the equilibrium strategy. Specifically, in PMS-Partition($\hat{\pi}_{j_k}$), we must have partition pairs $L_a^* = \{r_{j_{k+1}}\}$ and $r_{j_k} \in L_a$, where $|L_a| > 1$ and $C_{r_{j_k}} = C_{r_l^*}, 1 \leq k \leq q - 1$.

2. $\Pi_{r_{j_q}}$ has $< (\mathcal{M}^2 - 1)C^*$ **marked** players and an **unmarked** player $\hat{\pi}_{j_q}$ which is either a type-*A* player or PMS–Partition($\hat{\pi}_{j_q}$) contains r_{j_q} as a singleton resource (We will define marked players below). Specifically, in PMS–Partition($\hat{\pi}_{j_q}$), we must have partition pairs (L_b, L_b^*) where $L_b = \{r_{j_q}\}$, $C_{r_{j_q}} \geq C_{r_l^*}$ and $|L_b^*| \geq 1$ with $r_l^* \notin L_b^*$. Note that if π_{i_1} on resource $r_{j_1} = r_l^*$ itself satisfies condition (b) then the chain terminates with $q = 1$ and $\hat{\pi}_{j_q} = \pi_{i_1}$. Essentially the chain continues as long as condition 2 is not met for players on r_{j_k} .

First, note that since $|L_a| > 1, |L_a^*| = 1$ in the partition pairs above, we must have $C_{r_{j_{k+1}}} \geq C_{r_{j_k}} \geq C_{r_l^*}$, $1 \leq k \leq q - 1$, by lemma 2. Second, note that the chain consists only of $q - 1$ type-*B* players from r_{j_1} up to $r_{j_{q-1}}$. No type-*B* player in \tilde{G} can have a resource with congestion $\geq C_{r_l^*} + 1$ in its equilibrium strategy. Such a player π would have cost $pc_\pi(\tilde{S}) > \mathcal{M}^{C_{r_l^*} + 1} \geq K$, which is a contradiction. Thus the first $q - 1$ resources on the chain must have congestion exactly $C_{r_l^*}$. For the chain to terminate at some resource r_{j_q} , we either need a type-*A* player whose equilibrium strategy is r_{j_q} and socially optimal strategy does not contain r_{j_q} , or a type-*B* player whose equilibrium strategy contains r_{j_q} but has a partition pair in which r_{j_q} is a singleton set. If the terminating player is a type-*A* player, then it is possible that the congestion on r_{j_q} is $C_{r_{j_q}} > C_{r_l^*}$, otherwise if a type-*B* player then $C_{r_{j_q}} = C_{r_l^*}$. In both cases, by lemma 2, we must have

$$\mathcal{M} \cdot \sum_{r \in L_b^*} \mathcal{M}^{C_r} \geq \mathcal{M}^{C_{r_{j_q}}} \geq \mathcal{M}^{C_{r_l^*}} \geq \mathcal{M}^{C_{r_p}}, \tag{2}$$

where $C_{r_p} = \max\{C_{r_q} | r_q \in L_1\}$ and the last inequality arises from the fact that $|L_1| > 1, |L_1^*| = 1$ and lemma 2.

L_b^* is distinct from r_l^* as specified in condition 2 above. The optimal strategy for player π_i is now updated to $\tilde{S}_{\pi_i} = \{r_l^*, L_b^*\}$. We claim that this is a better optimal strategy set for π_i than the previous singleton set r_l^* because any PMS–partition of π_i will yield at least 2 partitions: one matching L_b^* with at least one resource from L_1 (by Eq. 2) and another matching r_l^* with the rest of the resources.

The player set $\Pi_{\tilde{G}}$ is further modified in the following manner. Let $\hat{\pi}_d \in \Pi_{r_{j_q}}$ denote the player whose PMS–partition resulted in L_b^* . Remove r_{j_q} and L_b^* from its equilibrium and socially optimal strategy sets, i.e

$$\tilde{S}_{\hat{\pi}_d} = \tilde{S}_{\hat{\pi}_d} - \{r_{j_q}\} \quad \tilde{S}_{\hat{\pi}_d}^* = \tilde{S}_{\hat{\pi}_d}^* - \tilde{L}_b^* \tag{3}$$

If player $\hat{\pi}_d$ was a type-*A* player, it will disappear from \tilde{G} , otherwise it will remain as a lower-cost player in equilibrium in \tilde{G} after this.

Finally, we create a new **marked** type-*A* player $\tilde{\pi}_d$ and add it to $\Pi_{\tilde{G}}$ along with the following strategies:

$$\tilde{S}_{\tilde{\pi}_d} = \tilde{S}_{\tilde{\pi}_d}^* = \{r_{j_q}\}, \tag{4}$$

Note that the equilibrium congestion on r_{j_q} is now the same as before and thus all players π with $r_{j_q} \in \tilde{S}_\pi^*$ remain in equilibrium in \tilde{S} . Also note that the optimal congestion on resources in L_b^* does not increase on being added to $\tilde{S}_{\pi_i}^*$ due to their

simultaneous removal from $\tilde{S}_{\pi_d}^*$. The optimal congestion on non-terminal resources in the chain, r_{j_k} , $1 \leq k \leq q - 1$ also does not change. The number of marked players on terminal resource r_{j_q} increases by one simultaneously with its optimal congestion. Since a resource can be marked at most $(\mathcal{M}^2 - 1)C^*$ times, the final optimal congestion on any resource and hence \tilde{G} is bounded by $\tilde{C}^* \leq \mathcal{M}^2 C^*$. The number of resources in \tilde{G} does not increase over G in any step and hence \tilde{G} is dominated by G .

The algorithm now continues onto the next iteration where we repeat the transformation process as before, starting with the highest cost type- B player remaining in \tilde{G} . Eventually the algorithm terminates when the cost of all type- B players in \tilde{G} is bounded by threshold ψ . □

Due to space considerations, we leave the proof of the following lemma which completes the proof above, for the expanded version of the paper.

Lemma 3. *For any resource r_l^* and player π_{i_j} , there always exists a PMS-chain(r_l^*, π_{i_j}) which terminates at a resource r_{j_q} with $< (\mathcal{M}^2 - 1)C^*$ marked players.*

4 Price of Anarchy

4.1 Price of Anarchy for Type-A Players

Let $\tau = (\mathcal{M} + 2)C^*$ be a *threshold* congestion value. Consider a game with optimal solution $S^* = (S_{\pi_1}^*, \dots, S_{\pi_n}^*)$ and congestion C^* . Consider also a Nash equilibrium state $S_E = (S_{\pi_1}, \dots, S_{\pi_n})$ which has the highest congestion C among all Nash equilibria states, and all players on resources r with $C_r \geq \tau$ are of type-A. We will give a price of anarchy result by bounding the ratio C/C^* .

We define the *expansion tree* \mathcal{T} for state S , which will help us to obtain the price of anarchy bound. We first define a set of nodes V which will be used in the construction of tree \mathcal{T} . Each resource $r \in R$ corresponds to C_r^* distinct nodes $V_r = \{x_1^r, \dots, x_{C_r^*}^r\}$, one node for each player π_i whose strategy $S_{\pi_i}^*$ contains r (there are C_r^* such players); thus, each $x_j^r \in V_r$ has a respective *owner* player. Note that for any two distinct resources $r_i, r_j \in R$, where $r_i \neq r_j$, the respective node sets are distinct, that is, $V_{r_i} \cap V_{r_j} = \emptyset$. The set V consists of the union of all the sets V_r , that is, $V = \bigcup_{r \in R} V_r$. For convenience, for any node $x \in V$ we denote by r_x the respective resource of x .

Let Π_r denote the set of players that use the resource r in S_E . Let Π_r^l be the set of players that use r in S_E and their optimal strategy in S^* does not contain r . Note that $|\Pi_r^l| \geq |\Pi_r| - C_r^*$. We divide Π_r^l into C_r^* (almost) equal size disjoint sets $\Pi_{x_1^r}^l, \dots, \Pi_{x_{C_r^*}^r}^l$, where each set consists of at least $\lfloor |\Pi_r^l|/C_r^* \rfloor \geq \lfloor |\Pi_r|/C_r^* - 1 \rfloor = \lfloor C_r/C_r^* - 1 \rfloor$ players. We say that the players in $\Pi_{x_j^r}^l$ are *assigned* to node $x_j^r \in V_r$.

The *expansion of a node* $x \in V$ is a tree of two levels rooted at x such that for every player π assigned to node x , that is, $\pi \in \Pi_x^l$, we add as children to x all the nodes $y \in V$ where the optimal strategy of π in S^* contains r_y and π is the owner of y . We build the expansion tree \mathcal{T} starting with an arbitrary root node $\rho \in V$ with $C_{r_\rho} = C$, and we recursively expand all nodes x with $C_{r_x} \geq \tau$. The final tree \mathcal{T} is obtained by removing any link that points to the root ρ (note that there could be at most one such

link). Thus, all leaves y of the tree have the property $C_{r_y} < \tau$. We define a *proper subtree* of \mathcal{T} to be a subtree where each node is either fully expanded or not expanded at all (and does not contain a link to the root).

Lemma 4. *For any node $x \in \mathcal{T}$ with $C_{r_x} \geq \tau$, it holds that $\sum_{y \in A} \mathcal{M}^{C_{r_y}} \geq \mathcal{M}^{C_{r_x}}$, where A are the children of x in \mathcal{T} .*

Proof. Let π_ρ be the owner player of root ρ . Let Π'_x be the set of players that are assigned to x . Let $\Pi''_x = \Pi'_x \setminus \{\pi_\rho\}$. By construction of the expansion tree, $|\Pi'_x| \geq \lfloor C_{r_x}/C^* - 1 \rfloor > \tau/C^* - 2 = \mathcal{M}$; hence, $|\Pi'_x| \geq \mathcal{M} + 1$, and $|\Pi''_x| \geq \mathcal{M}$. For every player $\pi \in \Pi''_x$, let A_π denote the respective set of nodes in A that are owned by π . Note that $A_{\pi_i} \cap A_{\pi_j} = \emptyset$ for any $\pi_i, \pi_j \in \Pi''_x$, and $\bigcup_{\pi \in \Pi''_x} A_\pi \subseteq A$. Since any player $\pi \in \Pi''_x$ is locally optimal in Nash equilibrium S_E and also of type-A, it holds that $\sum_{y \in A_\pi} \mathcal{M}^{C_{r_y}+1} \geq \mathcal{M}^{C_{r_x}}$. Therefore:

$$\sum_{y \in A} \mathcal{M}^{C_{r_y}+1} \geq \sum_{\pi \in \Pi''_x} \sum_{y \in A_\pi} \mathcal{M}^{C_{r_y}+1} \geq |\Pi''_x| \cdot \mathcal{M}^{C_{r_x}} \geq \mathcal{M} \cdot \mathcal{M}^{C_{r_x}} \geq \mathcal{M}^{C_{r_x}+1}.$$

The result follows by factoring out \mathcal{M} in the above expressions.

We can then prove by induction the following lemma, the proof of which we omit for space considerations.

Lemma 5. *For any proper subtree of \mathcal{T} with leaves H , it holds: $\sum_{y \in H} \mathcal{M}^{C_{r_y}} \geq \mathcal{M}^C$.*

Theorem 2 (Price of Anarchy for Type-A players). *The price of anarchy is $PoA = O(\log |R|)$.*

Proof. Let H be the leaves of \mathcal{T} . Since each resource in R corresponds to at most C^* nodes in \mathcal{T} , we have that $|H| \leq C^*|R|$. For any leaf $x \in H$ it holds $C_{r_x} < \tau$. From Lemma 5, $\sum_{y \in H} \mathcal{M}^{C_{r_y}} \geq \mathcal{M}^C$. For every $y \in H$ it holds that $C_{r_y} < \tau$. Thus, $|H| \cdot \mathcal{M}^\tau > \sum_{y \in H} \mathcal{M}^{C_{r_y}}$. Consequently, $|H| \cdot \mathcal{M}^\tau > \mathcal{M}^C$. Equivalently $|H| > \mathcal{M}^{C-\tau}$. Thus, $\log_{\mathcal{M}} |H| \geq C - (\mathcal{M}+2)C^*$. Hence, $\log_{\mathcal{M}}(C^*|R|) \geq C - (\mathcal{M}+2)C^*$. Therefore,

$$PoA = \frac{C}{C^*} \leq 3 + \mathcal{M} + \log_{\mathcal{M}}(|R|).$$


4.2 Price of Anarchy for Arbitrary Games

We can consider now an arbitrary game. Note that if all player costs are less than $\psi = \mathcal{M}^{(\mathcal{M}+2)C^*}$ then it then $PoA = O(1)$ since the base \mathcal{M} is a constant. Therefore, by Theorem 1, we only need to consider games with type-A players with costs above ψ , which implies that the congestion on any resource of these players is at least the threshold value τ (which is necessary for Theorem 2). Finally, by combining Theorem 1, Theorem 2, and Corollary 1 we obtain the main result for price of anarchy:

Theorem 3 (Price of Anarchy for Arbitrary Exponential Games). *The price of anarchy is $PoA = O(\log |R|)$.*

5 Lower Bound

We show that the upper bound of $O(\log |R|)$ in the price of anarchy is tight by demonstrating a congestion game with a lower bound on the price of anarchy of $\Omega(\log |R|)$. We construct a game instance represented as a graph in the figure below, such that each edge in the graph corresponds to a resource, and each player π_i has two strategies available: either the path from u to v through the direct edge $e = (u, v)$, or an alternative path $p_i = (u, x_i, \dots, y_i, v)$ (note that different player paths are edge-disjoint). Each path has length $|p_i| = \mathcal{M}^{n-1}$ edges and the number of players is $n = \log_{\mathcal{M}} |R| - \log_{\mathcal{M}} \log_{\mathcal{M}} |R|$.



Let S be the state depicted on the left part of the figure, where each player chooses the first strategy, and let S^* be the state on the right part of the figure, where each player chooses the alternative path. We have that $C(S^*) = 1$, which is the smallest congestion possible. Thus, S^* represents a socially optimal solution. For state S we have that $C(S) = n$, since all players use edge (u, v) . Note that S is a Nash Equilibrium, since each player π_i has cost $pc_{\pi_i}(S) = \mathcal{M}^n = |R| / \log_{\mathcal{M}} |R|$, and the cost of switching to path p_i would be $\mathcal{M}^1 \cdot |p_i| = \mathcal{M}^n = |R| / \log_{\mathcal{M}} |R|$, which is the same as the cost of using edge (u, v) . Consequently, a lower bound on the price of anarchy is $C(S)/C(S^*) = n/1 = \log_{\mathcal{M}} |R| - \log_{\mathcal{M}} \log_{\mathcal{M}} |R| = \Omega(\log |R|)$.

References

1. Banner, R., Orda, A.: Bottleneck routing games in communication networks. *IEEE Journal on Selected Areas in Communications* 25(6), 1173–1179 (2007); Also appears in *INFOCOM 2006*
2. Borodin, A., El-Yaniv, R.: *Online computation and competitive analysis*. Cambridge University Press, New York (1998)
3. Busch, C., Kannan, R., Vasilakos, A.V.: Quality of routing congestion games in wireless sensor networks. In: *Proc. 4th International Wireless Internet Conference (WICON)*, Maui, Hawaii (November 2008)
4. Busch, C., Magdon-Ismael, M.: Atomic routing games on maximum congestion. *Theoretical Computer Science* 410(36), 3337–3347 (2009)
5. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 67–73. ACM, Baltimore (May 2005)
6. Correa, J.R., Schulz, A.S., Moses, N.E.S.: Computational complexity, fairness and the price of anarchy of the maximum latency problem. In: *Bienstock, D., Nemhauser, G.L. (eds.) IPCO 2004*. LNCS, vol. 3064, pp. 59–73. Springer, Heidelberg (2004)

7. Czumaj, V.: Tight bounds for worst-case equilibria. In: ACM Transactions on Algorithms (TALG), vol. 3. ACM, New York (2007)
8. Fotakis, D., Kontogiannis, S.C., Spirakis, P.G.: Selfish unsplittable flows. *Theoretical Computer Science* 348(2-3), 226–239 (2005)
9. Gairing, M., Lüking, T., Mavronicolas, M., Monien, B.: Computing Nash equilibria for scheduling on restricted parallel links. In: Proceedings of the 36th Annual ACM Symposium on the Theory of Computing (STOC), Chicago, Illinois, USA, pp. 613–622 (June 2004)
10. Harks, T., Klimm, M., Möhring, R.H.: Strong Nash equilibria in games with the lexicographical improvement property. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 463–470. Springer, Heidelberg (2009)
11. Koutsoupias, E., Mavronicolas, M., Spirakis, P.G.: Approximate equilibria and ball fusion. *Theory Comput. Syst.* 36(6), 683–693 (2003)
12. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
13. Leighton, F.T., Maggs, B.M., Rao, S.B.: Packet routing and job-scheduling in $O(\textit{congestion} + \textit{dilation})$ steps. *Combinatorica* 14, 167–186 (1994)
14. Libman, L., Orda, A.: Atomic resource sharing in noncooperative networks. *Telecommunication Systems* 17(4), 385–409 (2001)
15. Lüking, T., Mavronicolas, M., Monien, B., Rode, M.: A new model for selfish routing. *Theoretical Computer Science* 406(3), 187–206 (2008)
16. Mavronicolas, S.: The price of selfish routing. *Algorithmica* 48 (2007)
17. Monderer, D., Shapely, L.S.: Potential games. *Games and Economic Behavior* 14, 124–143 (1996)
18. Papadimitriou, C.: Algorithms, games and the Internet. In: ACM (ed.) Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC), Hersonissos, Crete, Greece, pp. 749–753 (July 2001)
19. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
20. Roughgarden, T.: The maximum latency of selfish routing. In: Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), New Orleans, Louisiana (USA), pp. 980–981 (January 2004)
21. Roughgarden, T.: Selfish routing with atomic players. In: Proc. 16th Symp. on Discrete Algorithms (SODA), pp. 1184–1185. ACM/SIAM (2005)
22. Roughgarden, T.: Éva Tardos: How bad is selfish routing. *Journal of the ACM* 49(2), 236–259 (2002)
23. Roughgarden, T.: Éva Tardos: Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior* 47(2), 389–403 (2004)
24. Suri, S., Toth, C.D., Zhou, Y.: Selfish load balancing and atomic congestion games. *Algorithmica* 47(1), 79–96 (2007)

Single-Parameter Combinatorial Auctions with Partially Public Valuations^{*}

Gagan Goel¹, Chinmay Karande², and Lei Wang²

¹ Google Inc.

`gagangoel@google.com`

² Georgia Institute of Technology, Atlanta

`{ckarande,lwang}@cc.gatech.edu`

Abstract. We consider the problem of designing truthful auctions, when the bidders' valuations have a public and a private component. In particular, we consider combinatorial auctions where the valuation of an agent i for a set S of items can be expressed as $v_i f(S)$, where v_i is a private single parameter of the agent, and the function f is publicly known. Our motivation behind studying this problem is two-fold: (a) Such valuation functions arise naturally in the case of ad-slots in broadcast media such as Television and Radio. For an ad shown in a set S of ad-slots, $f(S)$ is, say, the number of *unique* viewers reached by the ad, and v_i is the valuation per-unique-viewer. (b) From a theoretical point of view, this factorization of the valuation function simplifies the bidding language, and renders the combinatorial auction more amenable to better approximation factors. We present a general technique, based on maximal-in-range mechanisms, that converts any α -approximation non-truthful algorithm ($\alpha \leq 1$) for this problem into $\Omega(\frac{\alpha}{\log n})$ and $\Omega(\alpha)$ -approximate truthful mechanisms which run in polynomial time and quasi-polynomial time, respectively.

1 Introduction

A central problem in computational mechanism design is that of combinatorial auctions, in which an auctioneer wants to sell a heterogeneous set of items \mathcal{J} to interested agents. Each agent i has a valuation function $v_i(\cdot)$ which describes her valuation $v_i(S)$ for every set $S \subseteq \mathcal{J}$ of items. In its most general form, the entire valuation function is assumed to be private information which may not be revealed truthfully by the agents. Maximizing the social welfare in a combinatorial auction with an incentive-compatible mechanism is an important open problem. However, recent results [5, 4] have established polynomial lower bounds on the approximation ratio of maximal-in-range mechanisms - which account for a majority of positive results in mechanism design - even when all the valuations are assumed to be submodular. On the other hand, in the non-game-theoretic case, if all the agents' valuations are public knowledge and hence

^{*} Research supported by NSF grants CCF-0728640 and CCF-0914732.

truthfully known, then we can maximize the social welfare to much better factors [6, 7, 17], under varying degree of restrictions on the valuations. In this paper, we introduce a model that lies in between these two extremes.

We wish to explore the setting when some inherent property of the items induces a common and publicly known *partial* information about the valuation function of the buyers. For instance, in position auctions in sponsored search, the agents' valuation for a position consists of a private value-per-click as well as a public click-through rate, that is known to the auctioneer. Another situation where such private/public factorization of valuations arises is advertisements in broadcast media such as Television and Radio. Suppose we are selling TV ad-slots on a television network. There are m ad-slots and n advertisers interested in them. Let us define a function $f : 2^{[m]} \rightarrow \mathbb{Z}_+$, such that for any set S of ad-slots $f(S)$ is the number of *unique* viewers who will see the ad¹ if the ad is shown on each slot in S . If an advertiser i is willing to pay v_i dollars per unique viewer reached by her ad, then her total valuation of the set S is $v_i f(S)$.

With this background, we define the following class of problems which we call *single-parameter combinatorial auctions with partially public valuations*: We are given a set \mathcal{J} of m items and a global *public* valuation function² $f : 2^{\mathcal{J}} \rightarrow \mathbb{R}$. The function f can either be specified explicitly or via an oracle which takes a set S as input and returns $f(S)$. In addition, we have n agents each of whom has a *private* multiplier v_i such that the item set S provides $v_i f(S)$ amount of utility to agent i . The goal is to design a truthful mechanism which maximizes $\sum_i v_i f(S_i)$, where $S_1 \cdots S_n$ is a partition of \mathcal{J} .

One can think of this model as combinatorial auctions with *simplified bidding language*. The agents only need to specify one parameter v_i as their bid. Moreover, our problem has deeper theoretical connections to the area of single-parameter mechanism design in general. For single-parameter domains such as ours, it is known that *monotone* allocation rules characterize the set of all truthful mechanisms. An allocation rule or algorithm is said to be monotone if the allocation parameter of an agent ($f(S_i)$ in our case) is non-decreasing in his reported bid v_i . Unfortunately, often it is the case that good approximation algorithms known for a given class of valuation functions are not monotonic. It is an important and well-known open question in algorithmic mechanism design to resolve whether the design of monotone algorithms is fundamentally harder than the non-monotone ones. In other words, it is not known if, for single-parameter problems, we can always convert any α -approximation algorithm into a truthful mechanism with the same factor. We believe that our problem is a suitable

¹ For a single ad-slot j , the function $f(\{j\})$ is nothing but the television rating for that slot as computed by rating agencies such as Nielsen. In fact, their data collection through set-top boxes results in a TV slot-viewer bipartite graph on the sample population, from which $f(S)$ can be estimated for any set S of ad slots.

² We do not make any explicit assumptions such as non-negativity or free disposal about the function f . We provide a method to convert any non-truthful black-box algorithm into a truthful mechanism. This black-box algorithm may make some implicit assumptions about f .

candidate to attack this question as it gives a lot of flexibility in defining the complexity of function f . From this discussion, it follows that the only lower bound known for the approximation factor of a truthful mechanism in our setting is the hardness of approximation of the underlying optimization problem.

Our Results and techniques. We give a general technique which accepts any (possibly non-truthful) α -approximation algorithm for our problem as a black-box and uses it to construct a truthful mechanism with an approximation factor of $\Omega\left(\frac{\alpha}{\log n}\right)$. We also give a truthful mechanism with factor $\Omega(\alpha)$ which runs in time $O\left(n^{\log \log n} \cdot \text{poly}(m)\right)$. Both these results are corollaries obtained by setting parameters appropriately in Theorem [1](#) to achieve desired trade-off between the approximation factor and the running time. Our results can also be interpreted as converting non-monotone algorithms into monotone ones for the above model.

Our mechanisms are *maximal-in-range*, *i.e.*, they fix a range \mathcal{R} of allocations and compute the allocation $\mathbf{S} \in \mathcal{R}$ that maximizes the social welfare. The technical core of our work lies in careful construction of this range.

While the black-box algorithm may be randomized, our mechanism does not introduce any further randomization. Depending upon whether the black-box algorithm is deterministic or randomized, our mechanism is deterministically truthful or universally truthful respectively (See [8](#) for definitions). The approximation factor of our mechanism is deterministic (or with high probability or in expectation) if the black-box algorithm also provides the approximation guarantees deterministically (or with high probability or in expectation).

Note that we don't need to worry about how the public valuation function f is specified. This is plausible since the function is accessed only from within the black-box algorithm. Hence, our mechanism can be applied to any model of specification - whether it is specified explicitly or through a value or demand oracle - using the corresponding approximation algorithm from that model.

Submodular valuations arise naturally in practice from economies of scale or the law of diminishing returns. Hence, we make a special note of our results when the public valuation is submodular. Using the algorithm of [17](#) as black-box, our results imply a $\Omega(1/\log n)$ and $\Omega(1)$ approximation factors in polynomial time and quasi-polynomial time, respectively. We would like to note that the standard greedy algorithm for submodular welfare maximization is not monotone (See [8](#) for a simple example) and hence, not truthful. Similarly, the optimal approximation algorithm of [17](#) is also not known to be non-monotone. The best known truthful mechanism for combinatorial auctions with entirely private submodular valuations [6](#) has $\Omega(1/\sqrt{m})$ approximation factor.

Future Directions. As shown in [5, 4](#), it seems that designing a truthful mechanism with good approximation factor for maximizing social welfare is a difficult problem. In light of this, our work suggests an important research direction to pursue in combinatorial auctions- to divide the valuation function into a part

which is common among all the agents and can be estimated by the auctioneer and a part which is unique and private to individual agents.

Also, it would be interesting to see if for submodular public functions (or even more specifically, for coverage functions), which have concrete motivation in TV ad auctions, one can design a constant factor polynomial time truthful mechanism.

Related Work. When agents have a general multi-parameter valuation function, the best known truthful approximation of social welfare in the value oracle model is $\Omega(\sqrt{\log m}/m)$ [10]. Under subadditive valuation functions, [6] gave $\Omega(1/\sqrt{m})$ -approximate truthful mechanism. It is known that no maximal-in-range mechanism making polynomially many calls to the value oracle can have an approximation factor better than $\Omega(1/m^{1/6})$ [5] even for the case of submodular valuation functions. A similar $\Omega(1/\sqrt{m})$ hardness result for maximal-in-range algorithms based on $\text{NP} \not\subseteq \text{P/poly}$ appears in [4]. See [3] for a comprehensive survey of the results, and [16, 4] for other more recent work. Previous work on the single-parameter case of combinatorial auctions have primarily focused on the *single-minded* bidders. In this setting, any bidder i is only interested in single set S_i and has a valuation v_i for it. Lehmann et al. [12] gave a truthful mechanism which achieves an essentially best-possible approximation factor of $\Omega(1/\sqrt{m})$. For other results in single-minded combinatorial auction, see [14, 1]. When the desired set is publicly known and only the valuation is private, [2] gave a general technique which converts any α -approximation algorithm into a truthful mechanism with factor $\alpha/\log(v_{max})$. This result is very much in spirit to our work, however the model and the techniques used in the two papers are very different. Similarly, [11] present a general framework which uses a gap-verifying linear program as black-box to construct mechanisms that are truthful in expectation.

For the non-truthful optimization, we note that our problem is hard up to a constant factor (see [13]) even when all the agents have private value equal to 1 and with common valuation function being submodular. For designing monotone algorithms from non-monotone algorithms in the Bayesian setting, see [9]. We also note that TV ad auctions are in use by Google Inc. (see [15]), although currently they treat the valuations for a set of ad-slots as additive with budget constraints, which yields a multi-parameter auction.

Organization: The full version of this paper [8] provides preliminary section containing a brief introduction to mechanism design with a few concepts relevant to our work, such as different notions of truthfulness and maximal-in-range mechanisms. In Section 2 in which we state some basic properties and assumptions about single parameter combinatorial auctions. Section 3 introduces our vector-fitting technique and presents our main result, a vector-fitting mechanism formalized by Theorem 1. Due to space constraints, we omit the proofs of Observation 3 and 4, as well as Lemma 1 here. These proofs follow largely from the definitions and can be found in the full version of this paper [8].

2 Notations and Basic Properties

By boldface \mathbf{v} , we will denote a vector of private multipliers of the agents, where v_i is the multiplier of agent i . For a constant $\beta \geq 0$, let $\beta\mathbf{v} = (\beta v_1, \beta v_2, \dots, \beta v_n)$. By boldface \mathbf{S} , we will denote the vector of allocations, where S_i is the set of items allocated to agent i . We will overload the function symbol f to express the social welfare as: $f(\mathbf{v}, \mathbf{S}) = \sum_i v_i f(S_i)$. An allocation \mathbf{S} is *optimal* for a multiplier vector \mathbf{v} if it maximizes $f(\mathbf{v}, \mathbf{S})$.

We begin by observing two simple properties of our problem and its solutions: *symmetry* and *scale-freeness*. Our problem and its solutions are symmetric, *i.e.*, invariant under relabeling of agents in the following sense: Let \mathbf{v} be any multiplier vector, \mathbf{S} be any allocation and π be any permutation of $[n]$. Let \mathbf{u} and \mathbf{T} be such that $u_i = v_{\pi(i)}$ and $T_i = S_{\pi(i)}$. Then clearly, $f(\mathbf{v}, \mathbf{S}) = f(\mathbf{u}, \mathbf{T})$. The problem and its solutions are also invariant under scaling, since we have $f(\beta\mathbf{v}, \mathbf{S}) = \beta \cdot f(\mathbf{v}, \mathbf{S})$.

The above properties lead us to:

Observation 1. *Without loss of generality, every multiplier vector \mathbf{v} has non-increasing entries $v_1 \geq v_2 \geq \dots \geq v_n$ such that $\sum_i v_i = 1$.*

Given a multiplier vector \mathbf{v} , let $A(\mathbf{v})$ be the optimal allocation for \mathbf{v} and $\text{OPT}(\mathbf{v}) = f(\mathbf{v}, A(\mathbf{v}))$. Moreover, if $f(\mathbf{v}, \mathbf{S}) \geq \alpha \cdot \text{OPT}(\mathbf{v})$ for some $\alpha \leq 1$ then the allocation \mathbf{S} is said to be α -optimal or α -approximate for \mathbf{v} .

We note a simple property of $A(\mathbf{v})$: Let \mathbf{v} be a multiplier vector with $v_1 \geq v_2 \geq \dots \geq v_n$. Let \mathbf{S} be any allocation. If \mathbf{T} is a permutation of \mathbf{S} such that $f(T_1) \geq f(T_2) \geq \dots \geq f(T_n)$, then $f(\mathbf{v}, \mathbf{T}) \geq f(\mathbf{v}, \mathbf{S})$. In particular, if $\mathbf{S} = A(\mathbf{v})$ then $f(S_1) \geq f(S_2) \geq \dots \geq f(S_n)$.

Finally, we assume the existence of a poly-time black-box algorithm that computes an α -approximate allocation $B(\mathbf{v})$ for the multiplier vector \mathbf{v} . We express the performance guarantees of our truthful mechanisms in terms of α and other parameters of the problem. Although the output allocation \mathbf{S} of such an algorithm may not obey $f(S_1) \geq f(S_2) \geq \dots \geq f(S_n)$, it is easy to construct a non-decreasing permutation of \mathbf{S} which only improves the objective function value, as discussed above.

Observation 2. *Without loss of generality, any allocation \mathbf{S} output by the black-box algorithm obeys $f(S_1) \geq f(S_2) \geq \dots \geq f(S_n)$.*

Henceforth, we enforce assumptions from Observation [1](#) and [2](#).

Definition 1 (u dominates w). *We say that a multiplier vector \mathbf{u} dominates \mathbf{w} if there exists an index i such that for $k < i$, $u_k \geq w_k$ and for $k \geq i$, $u_k \leq w_k$.*

Lemma 1. *If \mathbf{u} dominates \mathbf{w} , then $f(\mathbf{u}, \mathbf{S}) \geq f(\mathbf{w}, \mathbf{S})$ for any allocation \mathbf{S} satisfying $f(S_1) \geq f(S_2) \geq \dots \geq f(S_n)$.*

Proof. Refer to the full version of this paper [8](#).

Staircase Representation: Suppose we represent a multiplier vector \mathbf{v} as a histogram, which consists of n vertical bars corresponding to v_1, \dots, v_n , in that

order from left to right. Since multiplier vectors have non-increasing components, such a histogram looks like a staircase descending from left to right (Refer to Figure 1 for an example). We will refer to it as the *staircase representation* of \mathbf{v} and use it mainly as a visual tool.




Fig. 1. The staircase representation of $\mathbf{v} = (v_1, \dots, v_n)$

3 A Vector-Fitting Mechanism

Consider the following candidate approach to single parameter combinatorial auctions with partially public valuations: Fix a set \mathcal{U} of some multiplier vectors. Using the black-box algorithm, compute an α -approximate allocation $B(\mathbf{v})$ for each vector $\mathbf{v} \in \mathcal{U}$ and populate the range $\mathcal{R} = \{ B(\mathbf{v}) : \mathbf{v} \in \mathcal{U} \}$. Run the maximal-in-range mechanism which given a multiplier vector \mathbf{v} , chooses the allocation $\mathbf{S} \in \mathcal{R}$ that maximizes $f(\mathbf{v}, \mathbf{S})$.

Let’s consider the merits and demerits of this mechanism. If the input multiplier vector happens to be in \mathcal{U} , then the mechanism will indeed return an output allocation that is at least α -approximate. But we have no guarantees otherwise. If \mathcal{U} consisted of all possible vectors, we would have an α -approximate truthful mechanism that could be computationally infeasible due to the size of \mathcal{U} . We handle this trade-off with *vector-fitting*. The intuition behind vector-fitting is as follows: If two multiplier vectors \mathbf{u} and \mathbf{v} are ‘very similar’ to each other, then $B(\mathbf{u})$ and $B(\mathbf{v})$ should be ‘similar’ as well. In particular, $B(\mathbf{u})$ should be a reasonably good allocation for \mathbf{v} and vice versa.

Our mechanism will be the same as the candidate mechanism outlined above, except that we will construct the set of vectors \mathcal{U} very carefully. For any input vector of multipliers \mathbf{v} , we will guarantee that a reasonably similar vector \mathbf{v}' can be found in \mathcal{U} , and hence an allocation \mathbf{S}' is in the range \mathcal{R} with provably large objective value $f(\mathbf{v}, \mathbf{S}')$. We will prove the following theorem:

Theorem 1. *There exists a truthful mechanism for maximizing welfare in a single parameter combinatorial auction with partially public valuations that runs in time $O((\log_a n)^{\log_b n} \cdot \text{poly}(m, n))$ and produces an allocation with total welfare at least $\frac{3\alpha}{4ab} \cdot \text{OPT}(\mathbf{v})$ - where α is the approximation factor of the black-box optimization algorithm and $a, b > 1$ are parameters of the mechanism.*

Setting $a = b = 2$ we get: (Henceforth, all logarithms are to base 2)

Corollary 1. *There exists a $\frac{3\alpha}{16}$ -factor truthful mechanism running time $O(n^{\log \log n} \cdot \text{poly}(m))$, i.e. quasi-polynomial time.*

Similarly, setting $a = 2$ and $b = \log n$ we get:

Corollary 2. *There exists a truthful mechanism with factor $\Omega\left(\frac{\alpha}{\log n}\right)$ and polynomial running time.*

When the public valuation f is submodular, we have $\alpha = (1 - \frac{1}{e})$ and the above corollaries yield factors $\Omega(1)$ and $\Omega\left(\frac{1}{\log n}\right)$ respectively.

3.1 Constructing the Range \mathcal{R}

Overview: Recall the staircase representation of a multiplier vector \mathbf{v} , such as in Figure [1](#). Depending upon the entries of \mathbf{v} , the steps of the staircase may have varying heights. We can construct a discretization of the space of all multiplier vectors by restricting the values the height of any step can take. That is, we populate the initial set \mathcal{U} with all vectors whose components take values of the form b^{-k} for some constant $b > 1$ and for all $k \geq 0$. Now given any input vector \mathbf{v} , we can find a vector $\mathbf{u} \in \mathcal{U}$ such that u_i is at most a multiplicative factor b away from v_i . Thus, \mathbf{u} can serve as a vector ‘similar’ to \mathbf{v} . We need more complex machinery to ensure that the size of \mathcal{U} does not blow up, and that the vectors in \mathcal{U} still have unit norm.

Let $a, b > 1$ be suitably chosen parameters of the mechanism. Let $Q = \{ b^{-k} : 0 \leq k < \log_b n \}$ be a set of values discretizing the interval $(\frac{1}{n}, 1]$ and q be the minimum element of Q . For a multiplier $v_i \geq q$, we define $\lfloor v_i \rfloor$ to be the largest element of Q that is no greater than v_i . For a multiplier vector \mathbf{v} we define the floor of \mathbf{v} , $\lfloor \mathbf{v} \rfloor$ as follows:

Definition 2 (Floor $\lfloor \mathbf{v} \rfloor$). *The floor $\lfloor \mathbf{v} \rfloor$ of a multiplier vector \mathbf{v} is the vector \mathbf{u} constructed by Algorithm [7](#).*

In short, to find the ‘floor’ of a multiplier vector, we successively round down the ‘large’ components into elements of Q , until we need to set all the remaining components equal due the monotonicity and unit norm requirement or only ‘small’ components are remaining. When represented as a staircase (Refer Figure [1](#)), all the steps of $\lfloor \mathbf{v} \rfloor$ except the last one must have height that belongs to Q .

Observation 3. The floor of a vector \mathbf{v} is a valid multiplier vector itself, i.e. it has non-increasing components and unit l_1 norm. Moreover, \mathbf{v} dominates $\lfloor \mathbf{v} \rfloor$.

Proof. Refer to the full version of this paper [8](#).

Intuitively, the floor of a vector is (in a sense formalized by Lemma [2](#)) ‘similar’ to the vector, and the similarity is parametrized by b .

Lemma 2. *For any multiplier vector \mathbf{v} and allocation \mathbf{S} , $f(\lfloor \mathbf{v} \rfloor, \mathbf{S}) \geq \frac{3}{4b} \cdot f(\mathbf{v}, \mathbf{S})$.*

Algorithm 1. ConstructFloor

```

for  $i = 1$  to  $n$  do
     $r \leftarrow \frac{(1 - \sum_{k=1}^{i-1} u_k)}{(n - i + 1)}$ ;

    /*  $r$  is the minimum
       permissible value of
        $u_i$  due to
       monotonicity. */

    if  $v_i \geq q$  and  $\lfloor v_i \rfloor > r$ 
    then
         $u_i \leftarrow \lfloor v_i \rfloor$ ;
    else
        for  $j = i$  to  $n$  do
             $u_j \leftarrow r$ ;
        break
    
```

Algorithm 2. ConstructCore

```

 $i_1 \leftarrow 1; j_1 \leftarrow 1;$ 
while  $i_1 \leq n$  do
     $r \leftarrow (1 - \sum_{i=1}^{j_1-1} u_i) / (n - j_1 + 1)$ ;

    if  $v_{i_1} > r$  then
        Find the largest index  $i_2$ 
        such that  $v_{i_1} = v_{i_2}$ ;

        Find largest integer  $k$  such
        that  $\lceil a^k \rceil \leq (i_2 - j_1 + 1)$ ;
        for  $i = j_1$  to  $j_1 + \lceil a^k \rceil - 1$ 
        do
             $u_i \leftarrow v_{i_1}$ ;
             $i_1 \leftarrow i_2 + 1$ ;
             $j_1 \leftarrow j_1 + \lceil a^k \rceil$ ;
        else
            for  $i = j_1$  to  $n$  do
                 $u_i \leftarrow r$ ;
            break
    
```

Proof. Refer to Appendix [A](#).

We will construct our preliminary set of vectors \mathcal{U}' as

$$\mathcal{U}' = \{ \mathbf{u} : \mathbf{u} = \lfloor \mathbf{v} \rfloor \text{ for some multiplier vector } \mathbf{v} \}$$

It turns out that \mathcal{U}' is too large for our purposes. Hence we construct a subset $\mathcal{U} \subseteq \mathcal{U}'$, which is small enough. Referring back to the staircase representation of a multiplier vector (Figure 2), we constructed \mathcal{U}' by discretizing the ‘height’ of each step - by fitting the vectors vertically. Since rounding down the components of \mathbf{v} might lead to many components of $\mathbf{u} = \lfloor \mathbf{v} \rfloor$ having the same value, \mathbf{u} also looks like a staircase, perhaps with ‘wider’ steps. Each step of \mathbf{u} may have any integral width - at most n .

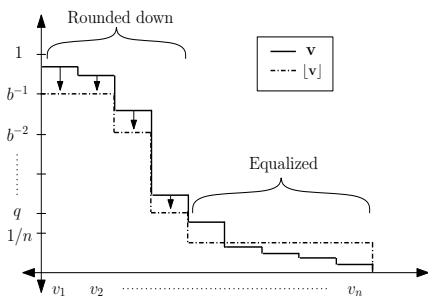


Fig. 2. Vertical fitting of \mathbf{v}




Fig. 3. Horizontal fitting of \mathbf{v}

We construct \mathcal{U} from \mathcal{U}' by further restricting how wide a step can be - by horizontal fitting (See Figure 3). We allow each step (except the last) to be of width $\lceil a^k \rceil$ for some integer $k \geq 0$ - where $a > 1$ is a suitably chosen parameter of the mechanism. To this end, we need to slightly formalize the staircase representation of a multiplier vector, which till now we only used as a visual aid. By a *step* of the staircase of \mathbf{v} , we will mean a maximal interval $[i_1, \dots, i_2] \subseteq [1, \dots, n]$ such that $v_{i_1} = v_{i_2}$. All the indices $i_1 \leq i \leq i_2$ will be said to belong to the step, whereas i_1 and i_2 are the first and last indices of the step. The *height* of the step is given by v_{i_1} and the *width* by $i_2 - i_1 + 1$.

Remark: Notice that just as a multiplier vector can be specified by the n -tuple (v_1, \dots, v_n) , it can also be identified by specifying the height and width of each step of its staircase representation. In fact, specifying all but the last step of a staircase fixes the last step due to the unit norm requirement.

For a multiplier vector \mathbf{v} , we define the core $\overleftarrow{\mathbf{v}}$ of \mathbf{v} as:

Definition 3 (Core $\overleftarrow{\mathbf{v}}$). *The core $\overleftarrow{\mathbf{v}}$ of a multiplier vector \mathbf{v} is the vector \mathbf{u} constructed by Algorithm 2.*

Operation of Algorithm 2: Each iteration of the **while** loop processes one step of \mathbf{v} and \mathbf{u} . i_1 and j_1 hold the first index of the current step of \mathbf{v} and \mathbf{u} respectively. r is the minimum height of the current step of \mathbf{u} by monotonicity. If $r \geq v_{i_1}$, then the requirement for unit l_1 norm forces us to introduce the last step of the staircase of \mathbf{u} . Otherwise, $[i_1, \dots, i_2]$ is the current step of \mathbf{v} and we set the width of the current step of \mathbf{u} to be $\lceil a^k \rceil$.

Observation 4. The core of a vector \mathbf{v} is a multiplier vector itself, *i.e.* it has non-increasing components and unit l_1 norm. Moreover, \mathbf{v} dominates $\overleftarrow{\mathbf{v}}$.

Proof. Refer to the full version of this paper [8].

Lemma 3. *For any multiplier vector \mathbf{v} and allocation \mathbf{S} , $f(\overleftarrow{\mathbf{v}}, \mathbf{S}) \geq f(\mathbf{v}, \mathbf{S})/a$.*

Proof. Refer to Appendix B.

We now define our set of vectors \mathcal{U} as follows: $\mathcal{U} = \{ \overleftarrow{\mathbf{v}} : \mathbf{v} \in \mathcal{U}' \}$. We populate the range \mathcal{R} of allocations as $\mathcal{R} = \{ \mathbf{B}(\mathbf{v}) : \mathbf{v} \in \mathcal{U} \}$ where $\mathbf{B}(\mathbf{v})$ is the α -approximate allocation returned by the black box algorithm.

3.2 Proof of Theorem 1

We run the following maximal-in-range mechanism: Given an input multiplier vector \mathbf{v} we return the allocation $\mathbf{T} \in \mathcal{R}$ that maximizes $f(\mathbf{v}, \mathbf{T})$. We need to prove that $f(\mathbf{v}, \mathbf{T}) \geq \frac{3\alpha}{4ab} \cdot \text{OPT}(\mathbf{v})$

Let $\mathbf{S} = \mathbf{A}(\mathbf{v})$ be the optimal allocation for \mathbf{v} and $\overleftarrow{\lfloor \mathbf{v} \rfloor}$ be the core of the floor of \mathbf{v} . Combining Lemmas 2 and 3, we conclude that $f(\overleftarrow{\lfloor \mathbf{v} \rfloor}, \mathbf{S}) \geq \frac{3}{4ab} \cdot \text{OPT}(\mathbf{v})$. Since $\overleftarrow{\lfloor \mathbf{v} \rfloor} \in \mathcal{U}$, there exists an allocation $\mathbf{X} \in \mathcal{R}$ such that

$$f(\overleftarrow{\lfloor \mathbf{v} \rfloor}, \mathbf{X}) \geq \alpha \cdot \text{OPT}(\overleftarrow{\lfloor \mathbf{v} \rfloor}) \geq \alpha \cdot f(\overleftarrow{\lfloor \mathbf{v} \rfloor}, \mathbf{S}) \geq \frac{3\alpha}{4ab} \cdot \text{OPT}(\mathbf{v}) \quad (1)$$

Since \mathbf{v} dominates $\lfloor \mathbf{v} \rfloor$ which in turn dominates $\overleftarrow{\lfloor \mathbf{v} \rfloor}$ (Refer to Observation 3 and 4), application of Lemma 1 yields:

$$f(\mathbf{v}, \mathbf{X}) \geq f(\lfloor \mathbf{v} \rfloor, \mathbf{X}) \geq f(\overleftarrow{\lfloor \mathbf{v} \rfloor}, \mathbf{X}) \quad (2)$$

Using equations (1) and (2),

$$f(\mathbf{v}, \mathbf{T}) \geq f(\mathbf{v}, \mathbf{X}) \geq f(\overleftarrow{\lfloor \mathbf{v} \rfloor}, \mathbf{X}) \geq \frac{3\alpha}{4ab} \cdot \text{OPT}(\mathbf{v})$$

The running time of the mechanism is established by Lemma 4, which finishes the proof of Theorem 1.

Lemma 4. $|\mathcal{R}| = O((\log_a n)^{\log_b n})$

Proof. $|\mathcal{R}|$ is bounded by $|\mathcal{U}|$. \mathcal{U} consists of only those vectors which are cores of floors of some multiplier vectors. We have seen that each step of the staircase of $\mathbf{v} \in \mathcal{U}$ except the last must be of width $w = \lceil a^k \rceil$ for some integer k . Moreover, there can be only $|Q| = O(\log_b n)$ such steps and at most one of each height. We have also remarked that specifying all but the last step of a staircase fixes it. Therefore there can be at most $O((\log_a n)^{\log_b n})$ distinct staircases in \mathcal{U} .

References

- [1] Archer, A., Papadimitriou, C., Talwar, K., Tardos, E.: An approximate truthful mechanism for combinatorial auctions with single parameter agents. In: SODA 2003: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 205–214 (2003)
- [2] Babaioff, M., Lavi, R., Pavlov, E.: Single-value combinatorial auctions and algorithmic implementation in undominated strategies. *J. ACM* 56(1), 1–32 (2009)
- [3] Blumrosen, L., Nisan, N.: *Algorithmic Game Theory*, ch. 11. Cambridge University Press, Cambridge (2007)
- [4] Buchfuhrer, D., Dughmi, S., Fu, H., Kleinberg, R., Mossel, E., Papadimitriou, C., Schapira, M., Singer, Y., Umans, C.: Inapproximability for vcg-based combinatorial auctions. In: SODA 2010: Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 518–536 (2010)
- [5] Dobzinski, S., Nisan, N.: Limitations of vcg-based mechanisms. In: STOC 2007: Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing, pp. 338–344 (2007)
- [6] Dobzinski, S., Nisan, N., Schapira, M.: Approximation algorithms for combinatorial auctions with complement-free bidders. In: STOC 2005: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, pp. 610–618 (2005)
- [7] Dobzinski, S., Schapira, M.: An improved approximation algorithm for combinatorial auctions with submodular bidders. In: SODA 2006: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm, pp. 1064–1073 (2006)
- [8] Goel, G., Karande, C., Wang, L.: Single-parameter combinatorial auctions with partially public valuations. CoRR abs/1007.3539 (2010)

- [9] Hartline, J.D., Lucier, B.: Bayesian algorithmic mechanism design. In: STOC 2010: Proceedings of the 42nd ACM Symposium on Theory of Computing, pp. 301–310 (2010)
- [10] Holzman, R., Kfir-dahav, N., Monderer, D., Tennenholtz, M.: Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior* 47(1), 104–123 (2001)
- [11] Lavi, R., Swamy, C.: Truthful and near-optimal mechanism design via linear programming. In: FOCS 2005: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, pp. 595–604 (2005)
- [12] Lehmann, D., O’callaghan, L.I., Shoham, Y.: Truth revelation in approximately efficient combinatorial auctions. *J. ACM* 49(5), 577–602 (2002)
- [13] Mirrokni, V., Schapira, M., Vondrak, J.: Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In: EC 2008: Proceedings of the 9th ACM Conference on Electronic Commerce, pp. 70–77 (2008)
- [14] Muálem, A., Nisan, N.: Truthful approximation mechanisms for restricted combinatorial auctions: extended abstract. *Games and Economic Behavior* 64(2), 612–631 (2008)
- [15] Nisan, N., Bayer, J., Chandra, D., Franji, T., Gardner, R., Matias, Y., Rhodes, N., Seltzer, M., Tom, D., Varian, H., Zigmond, D.: Google’s auction for tv ads. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikolettseas, S., Thomas, W. (eds.) ICALP 2009. LNCS, vol. 5556, pp. 309–327. Springer, Heidelberg (2009)
- [16] Papadimitriou, C., Schapira, M., Singer, Y.: On the hardness of being truthful. In: FOCS 2008: Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science, pp. 250–259 (2008)
- [17] Vondrak, J.: Optimal approximation for the submodular welfare problem in the value oracle model. In: STOC 2008: Proceedings of the 40th Annual ACM Symposium on Theory of Computing, pp. 67–74 (2008)

A Proof of Lemma 2

Define $\mathbf{u} = \lfloor \mathbf{v} \rfloor$. Let p be the highest index such that v_p is rounded down by the procedure that constructs \mathbf{u} , *i.e.* $u_p = \lfloor v_p \rfloor$ and $u_p > r = u_{p+1}$. Since, $\sum_{i=1}^p u_i \leq \sum_{i=1}^p v_i$, it is clear that $p < n$. Now for $i \leq p$, we have $u_i = \lfloor v_i \rfloor \geq v_i/b$. Consider two cases about v_{p+1} :

Case 1 - $v_{p+1} \geq q$: In this case, $u_{p+1} = r \geq \lfloor v_{p+1} \rfloor \geq v_{p+1}/b$. For $i \geq p+1$, we have $v_i \leq v_{p+1}$ and $u_i = u_{p+1}$ implying $u_i \geq v_i/b$. Therefore,

$$f(\mathbf{u}, \mathbf{S}) = \sum_{i=1}^n u_i f(S_i) \geq \frac{1}{b} \sum_{i=1}^n v_i f(S_i) = \frac{1}{b} \cdot f(\mathbf{v}, \mathbf{S})$$

Case 2 - $v_{p+1} < q$: Let $h = \sum_{i=1}^p v_i$ and $H = \left(\sum_{i=1}^p v_i f(S_i) \right) / f(\mathbf{v}, \mathbf{S})$. From the monotonicity of \mathbf{S} , we conclude that

$$H \cdot f(\mathbf{v}, \mathbf{S}) = \sum_{i=1}^p v_i f(S_i) \geq h \cdot f(\mathbf{v}, \mathbf{S})$$

and hence $H \geq h$.

Since $u_i \leq v_i$ for all $i \leq p$, and both \mathbf{u} and \mathbf{v} must have unit l_1 norm, we have $\sum_{i>p} u_i \geq \sum_{i>p} v_i = (1 - h)$. Hence, $u_i \geq \frac{1-h}{n}$ for $i > p$. By definition, $v_i < q \leq \frac{b}{n}$ for $i > p$. Together, these imply $u_i \geq (1 - h)v_i/b$. Finally, using $H \geq h$, we conclude

$$\sum_{i>p} u_i f(S_i) \geq \frac{1-h}{b} \left(\sum_{i>p} v_i f(S_i) \right) \geq \frac{1-H}{b} [(1-H)f(\mathbf{v}, \mathbf{S})]$$

Combining these pieces together, we get:

$$\begin{aligned} f(\lfloor \mathbf{v} \rfloor, \mathbf{S}) &= \sum_{i=1}^p u_i f(S_i) + \sum_{i>p} u_i f(S_i) \\ &\geq \frac{1}{b} \sum_{i=1}^p v_i f(S_i) + \frac{(1-H)^2}{b} \cdot f(\mathbf{v}, \mathbf{S}) \\ &= \frac{H + (1-H)^2}{b} \cdot f(\mathbf{v}, \mathbf{S}) \geq \frac{3}{4b} \cdot f(\mathbf{v}, \mathbf{S}) \end{aligned}$$

B Proof of Lemma 3

Suppose the staircase of \mathbf{v} has s_1 steps and that of $\mathbf{u} = \overleftarrow{\mathbf{v}}$ has s_2 steps. Then the following four properties follow directly from the algorithm:

1. $s_2 \leq s_1$
2. For $1 \leq i < s_2$, the i 'th step of \mathbf{v} is at most a times as wide as the i 'th step of \mathbf{u} and both have the same height.
3. For $1 \leq i \leq s_2$, let i_1 and j_1 be the first indices of the i 'th steps of \mathbf{v} and \mathbf{u} respectively. Then $i_1 \geq j_1$.
4. If $[j, \dots, n]$ is the last step of \mathbf{u} then $u_i \geq v_i$ for $i \geq j$.

To prove the lemma, we will compare the the contributions of corresponding steps of the staircases of \mathbf{v} and \mathbf{u} to the objective functions.

For $i < s_2$, let $[i_1, \dots, i_2]$ be the i 'th step of \mathbf{v} , $[j_1, \dots, j_2]$ be the i 'th step of \mathbf{u} and $h = v_{i_1} = u_{j_1}$ be their common height. We have

$$\sum_{k=j_1}^{j_2} u_k f(S_k) = h \sum_{k=j_1}^{j_2} f(S_k) \geq h \sum_{k=i_1}^{i_1+j_2-j_1} f(S_k)$$

by the third property. The monotonicity of \mathbf{S} and the second property then imply

$$\sum_{k=j_1}^{j_2} u_k f(S_k) \geq \frac{1}{a} \sum_{k=i_1}^{i_2} v_k f(S_k)$$

So the i 'th step of \mathbf{v} contributes at most a times value to $f(\mathbf{v}, \mathbf{S})$ as the i 'th step of \mathbf{u} contributes to $f(\mathbf{u}, \mathbf{S})$, where $i < s_2$.

Finally by the fourth property, the step s_2 of \mathbf{u} contributes more to $f(\mathbf{u}, \mathbf{S})$ than the corresponding contribution of steps s_2, \dots, s_1 of \mathbf{v} to $f(\mathbf{v}, \mathbf{S})$ combined. The result therefore follows.

On the Efficiency of Markets with Two-Sided Proportional Allocation Mechanisms

Volodymyr Kuleshov and Adrian Vetta*

Department of Mathematics and Statistics,
and School of Computer Science, McGill University
volodymyr.kuleshov@mail.mcgill.ca, vetta@math.mcgill.ca

Abstract. We analyze the performance of resource allocation mechanisms for markets in which there is competition amongst both consumers and suppliers (namely, two-sided markets). Specifically, we examine a natural generalization of both Kelly’s proportional allocation mechanism for demand-competitive markets [9] and Johari and Tsitsiklis’ proportional allocation mechanism for supply-competitive markets [7].

We first consider the case of a market for one divisible resource. Assuming that marginal costs are convex, we derive a tight bound on the price of anarchy of about 0.5887. This worst case bound is achieved when the demand-side of the market is highly competitive and the supply-side consists of a duopoly. As more firms enter the market, the price of anarchy improves to 0.64. In contrast, on the demand side, the price of anarchy improves when the number of consumers decreases, reaching a maximum of 0.7321 in a monopsony setting. When the marginal cost functions are concave, the above bound smoothly degrades to zero as the marginal costs tend to constants. For monomial cost functions of the form $C(x) = cx^{1+\frac{1}{d}}$, we show that the price of anarchy is $\Omega(\frac{1}{d^2})$.

We complement these guarantees by identifying a large class of two-sided single-parameter market-clearing mechanisms among which the proportional allocation mechanism uniquely achieves the optimal price of anarchy. We also prove that our worst case bounds extend to general multi-resource markets, and in particular to bandwidth markets over arbitrary networks.

1 Introduction

How to produce and allocate scarce resources is the most fundamental question in economics [1]. The standard tool for guiding production and allocation is a pricing mechanism. However, different mechanisms will have different performance attributes: no two mechanisms are equal. Of particular interest to computer scientists is the fact that there will typically be an inherent trade-off between the economic efficiency of a mechanism (measured in terms of social welfare) and its

* The authors were supported in part by NSERC grant 28833.

¹ In fact, economics is often defined as “the study of scarcity”.

computational efficiency (both time and communication complexity). Socially optimal allocations can be achieved using pricing mechanisms based on classical VCG results, but implementing such mechanisms generally induces excessively high informational and computational costs [12]. In this paper, we study this tradeoff from the opposite viewpoint: we examine the level of social welfare that can be achieved by mechanisms performing *minimal* amounts of computation. In particular, we restrict our attention to so-called *scalar-parametrized* pricing mechanisms. Each participant submits only a single scalar bid that is used to set a unique market-clearing price for each good. Evidently, such mechanisms are computationally trivial to handle; more surprisingly, they can produce high welfare.

The chief practical motivation for considering scalar-parametrized mechanisms (both in our work and in the existing literature) is the problem of bandwidth sharing. Namely, how should we allocate capacity amongst users that want to transmit data over a network link? The use of market mechanisms for this task has been studied in Asynchronous Transfer Mode (ATM) networks [15] and the Internet [14]. The Internet is made up of smaller interconnected networks that buy capacities from each other, and the market mechanisms we consider are closely inspired by the structure of the Internet. Specifically, we are restricting our attention to mechanisms that are scalable to very large networks. This requirement for scalability forces us consider only simple mechanisms, such as those that set a unique market clearing price. The computational requirements of more complex systems, e.g. mechanisms that perform price discrimination, become impractical on large networks [1].

We remark that unique price mechanisms are also intuitively “fair”, as every participant is treated equally. This fairness is appealing from a social and political perspective, and indeed these systems are used in many real-world settings, such as electricity markets [16].

1.1 Background and Previous Work

A basic method for resource allocation is the *proportional allocation mechanism* of Kelly [9]. In the context of networks, it operates as follows: each potential consumer submits a bid b_q ; bandwidth is then allocated to the consumers in proportion to their bids. This simple idea has also been studied within economics by Shapley and Shubik [13] as a model for understanding pricing in market economies. In a groundbreaking result, Johari and Tsitsiklis [5] showed that the welfare loss incurred by this mechanism is at most 25% of optimal.

Observe that Kelly’s is a scalar-parametrized mechanism for a one-sided market: every participant is a consumer. Johari and Tsitsiklis [7] also examined one-sided markets with supply-side competition only. There, under a corresponding single-parameter mechanism, the welfare loss tends to zero as the level of competition increases. We remark that we cannot simply analyze supply-side competition by trying to model suppliers as demand-side consumers [3].

Of course, competition in markets typically occurs on *both* sides. Consequently, understanding the efficiency of *two-sided market*² mechanisms is an important problem. In this work, we analyze the price of anarchy in a mechanism for a two-sided market in which consumers and producers compete simultaneously to determine the production and allocation of goods. This mechanism was first proposed by Neumayer [10] and is the natural generalization of both the demand-side model of Kelly [9] and the supply-side model of Johari and Tsitsiklis [7].

In order to examine how the generalized proportional allocation mechanism performs in a two-sided market, it is important to note that there are three primary causes of welfare loss. First, the underlying allocation problem may be computationally hard. In other words, in some settings (such as combinatorial auctions, for example), it may be hard to compute the optimal allocation even when the players' utilities are known. Secondly, even if the allocation problem is computationally simple, the mechanism itself may still be insufficiently sophisticated to solve it. Thirdly, the mechanism may be susceptible to gaming; namely, the mechanism may incentivize selfish agents to behave in a manner that produces a poor overall outcome. As we will see in Section 4, the first two causes do not arise here: as long as the users do not behave strategically, the proportional allocation mechanism can quickly find optimal allocations in two-sided markets. Thus, we are concerned only with the third factor: how adversely is the proportional allocation mechanism affected by gaming agents? That is, the mechanism may be capable of producing an optimal solution, but how will the agents' selfish behaviour affect social welfare at the resultant equilibria?

In this paper we prove that the proportional allocation mechanism does perform well in two-sided markets. Specifically, under quite general assumptions, the mechanism admits a constant factor price of anarchy guarantee. Moreover, there exists a large family of mechanisms among which the proportional allocation mechanism uniquely achieves the best possible price of anarchy guarantee. We state our exact results in Section 3, after we have described the model and our assumptions.

2 The Model

2.1 The Two-Sided Proportional Allocation Mechanism

We now formally present the two-sided proportional allocation mechanism due to Neumayer [10]. There are Q consumers and R suppliers in the market. Each consumer q has a valuation function $V_q(d_q)$, where d_q is the amount of the resource allocated to consumer q , and each supplier r has a cost function $C_r(s_r)$, where s_r is the amount produced by supplier r . Consumers and suppliers respectively

² It should be noted that “two-sided market” often has a different meaning in the economics literature than the one we use here. There it refers to a specific class of markets where externalities occur between groups on the two sides of the market.

input bids b_q and b_r to the mechanism. Doing so, consumers are implicitly selecting b_q -parametrized demand functions of the form $D(b_q, p) = \frac{b_q}{p}$, and suppliers are selecting b_r -parametrized supply functions of the form $S(b_r, p) = 1 - \frac{b_r}{p}$. We can also interpret a high consumer bid as an indicator of high willingness to pay for the product, and a low supplier bid as an indicator of a high willingness to supply (alternatively, a high bid indicates a high cost supplier). The actual choice of constant used for the supply functions does not affect our results, and so we choose it to be 1.

Observe that the parametrized demand functions are identical to the ones in the demand-side mechanism of Kelly [9], and the supply functions are identical to the ones in the supply-side mechanism of Johari and Tsitsiklis [7]. The peculiar form of the supply functions comes from the interesting fact that for most scalar-parametrized mechanisms, in order to have a non-zero welfare ratio, the supply functions have to be bounded from above. In other words, suppliers' strategies must necessarily be constrained in order to obtain high welfare; see the full version of the paper for the precise statement of this fact. This rules out, for instance, Cournot-style mechanisms where suppliers directly submit the quantities they wish to produce.

More detailed justifications for this choice of model can be found in [10], as well as in [9] and [7]. Further justification for the mechanism will be provided by our results. Specifically, the proportional allocation mechanism generally produces high welfare allocations and, in addition, it is the optimal mechanism amongst a class of single-parameter mechanisms for two-sided markets.

Given the bids, the mechanism sets a price $p(\mathbf{b})$ that clears the market; i.e. that satisfies the supply equals demand equation: $\sum_{q=1}^Q \frac{b_q}{p} = \sum_{r=1}^R (1 - \frac{b_r}{p})$. The price therefore gets set to $p(\mathbf{b}) = \frac{\sum_q b_q + \sum_r b_r}{R}$. Consumer q then receives d_q units of the resource, and pays pd_q , while supplier r produces s_r units and receives a payment of ps_r . In the game induced by this mechanism, the payoff (or utility) to consumer q placing a bid b_q is defined to be

$$\Pi_q(b_q) = \begin{cases} V_q \left(\frac{b_q}{\sum_{q \in Q} b_q + \sum_r b_{r \in R}} R \right) - b_q & \text{if } b_q > 0 \\ V_q(0) & \text{if } b_q = 0 \end{cases}$$

and the payoff to supplier r placing a bid b_r is defined as

$$\Pi_r(b_r) = \begin{cases} \frac{\sum_{q \in Q} b_q + \sum_{r \in R} b_r}{R} - b_r - C_r \left(1 - \frac{b_r}{\sum_{q \in Q} b_q + \sum_{r \in R} b_r} R \right) & \text{if } b_r > 0 \\ \frac{\sum_{q \neq r} b_q + \sum_{r \in R} b_r}{R} - C_r(1) & \text{if } b_r = 0 \end{cases}$$

2.2 The Welfare Ratio

Given a vector of bids \mathbf{b} , the *social welfare* at the resulting mechanism allocation is defined to be

$$\mathcal{W}(\mathbf{b}) = \sum_{q=1}^Q V_q(d_q(\mathbf{b})) - \sum_{r=1}^R C_r(s_r(\mathbf{b}))$$

If the agents do not strategically anticipate the effects of their actions on the price, that is if they act as “price-takers”, we show in Section 4 that the mechanism maximizes social welfare. However, since the price is a function of their bid, each agent is a “price-maker”. If agents attempt to exploit this market power, then a welfare loss may occur at a Nash equilibrium. Consequently we are interested in maximizing (over all equilibria) the *welfare ratio*, more commonly known as the *price of anarchy*, $\frac{\mathcal{W}^{\text{NE}}}{\mathcal{W}^{\text{OPT}}}$. Equivalently, we wish to minimize the *welfare loss*, $1 - \frac{\mathcal{W}^{\text{NE}}}{\mathcal{W}^{\text{OPT}}}$.

2.3 Assumptions

We make the following assumption on the valuation and cost functions.

Assumption 1. *For each consumer q , the valuation function $V_q(d_q) : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is strictly increasing and concave. For each supplier r , the cost function $C_r(s_r) : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is strictly increasing and convex.*

Assumption 1 corresponds to decreasing marginal valuations and increasing marginal costs. The assumption is standard in the literature. It certainly may not hold in every market 3, but without it there will be a natural incentive for the number of agents to decline on both sides of the market. In this paper, we will also assume that our functions are differentiable over their entire domain; this property is assumed primarily for clarity and is not essential.

Assumption 1, however, is not sufficient to ensure a large welfare ratio. In fact, the welfare ratio depends upon the curvature of the *marginal cost functions*. Specifically, if the marginal cost functions are convex, then we show in Section 4 that the welfare ratio is at least 0.58. Concave marginal cost functions also exhibit constant welfare ratios, provided the corresponding total cost function is sufficiently non-linear. However, in the limit as the total cost functions become linear, the welfare ratio degrades to zero (see Section 5 for more details).

Our main result thus concerns convex marginal cost functions. Formally, for most of the paper, we assume that

Assumption 2. *For each supplier r , the marginal cost function $C'_r(s_r)$ is convex. Furthermore, we assume that $C_r(0) = C'_r(0) = 0$.*

Convex marginal cost functions are extremely common in both the theoretical and the practical literature on industrial theory [17], so this assumption is not

³ For example, in markets exhibiting economies of scale.

particularly restrictive. In Assumption 2 we also set $C'_r(0) = 0$, but as we show in the full version of the paper, constant welfare ratios still arise whenever $C'_r(0)$ is bounded below one (it cannot be higher than one or the firm is uncompetitive).

We also remark that Assumption 2 was used in Johari, Mannor and Tsitsiklis 4 in their analysis of the demand-side proportional allocation mechanism with elastic supply. Most of the results of Johari and Tsitsiklis 6 and Tobias and Harks 2 on demand-side Cournot competition with elastic supply also hold under the assumption of convex marginal costs.

3 Our Results

Our first results are concerned with the performance of the mechanism when the users act as price-takers. Under Assumption 1, we prove that:

Theorem 1. *A unique competitive equilibrium exists for the two-sided proportional allocation mechanism. The social welfare attained at the competitive equilibrium is optimal.*

This property was exhibited by Kelly’s original proportional allocation mechanism, and has been a feature of all subsequent generalizations by Johari and Tsitsiklis. It is very appealing from a practical point of view, as in actual networks, users are likely to have little information about each other, making it difficult to manipulate the system.

In many other settings however, users will be incentivized to act strategically. In that case, we need to use the stronger solution concept of a Nash equilibrium to analyze the resulting game. Our second result establishes the existence and uniqueness of such equilibria under Assumption 1.

Theorem 2. *The two-sided proportional allocation mechanism has a unique Nash equilibrium for $R \geq 2$.*

Our main result measures the loss of welfare at that unique Nash equilibrium under Assumption 2.

Theorem 3. *The worst case welfare ratio for the mechanism involving $R \geq 2$ suppliers equals*

$$\frac{s^2((R - 1)^2 + 4(R - 1)s + 2s^2)}{(R - 1)(R - 1 + 2s)}$$

where s is the unique positive root of the quartic polynomial $\gamma(s) = 16s^4 + (R - 1)s^2(49s - 24) + 10(R - 1)^2s(3s - 2) + (R - 1)^3(5s - 4)$. Furthermore, this bound is tight.

It follows that the mechanism admits a constant bound on the price of anarchy. Moreover, Theorem 3 allows us to measure the effects of market competition on social welfare. The following two corollaries are concerned with that relationship.

Corollary 1. *The worst possible price of anarchy is achieved when the supply side is a duopoly ($R = 2$). It evaluates numerically to about 0.588727.*

Corollary 2. *When the supply side is fully competitive ($R \rightarrow \infty$) the price of anarchy equals precisely 0.64.*

Consequently, as supply-side competition increases, the welfare ratio improves. In contrast, the welfare ratio *decreases* as demand-side competition increases. Although this fact may seem surprising at first, it turns out to have a simple intuitive explanation. The optimal demand-side allocation consists in giving the entire production to the user which derives from it the highest utility. When more consumers are present in the market, they selfishly request more of the resource for themselves, leaving less for the most needy user and reducing the overall social welfare.

The best welfare ratios thus arise when there is only one consumer ($Q = 1$), that is, in the case of a *monopsony*. In the two-sided proportional allocation mechanism, the best possible price of anarchy over all possible values of Q and R is given by the next corollary.

Corollary 3. *In a market in which a monopsonist faces a fully competitive supply side, the price of anarchy equals $\sqrt{3} - 1$, which is about 0.7321.*

Recall that in the one-sided proportional allocation mechanism for suppliers facing a fixed demand, the welfare loss tends to zero when the supply side is fully competitive [7]. In contrast, Corollary 3 implies that in two-sided markets, that result no longer holds and that full efficiency cannot be achieved.

So far, our results assumed the convexity of marginal costs. Dropping that assumption, we find that the welfare ratio equals zero when the providers' total cost functions are linear. However, the price of anarchy remains bounded for a class of concave marginal cost functions, and degrades smoothly to zero as the total costs become linear.

Corollary 4. *The welfare ratio for cost functions $C_r(s_r) = c_r s_r^{1+\frac{1}{d}}$ where $c_r > 0$ and $d \geq 1$ is $\Omega(\frac{1}{d^2})$.*

Like its one-sided versions, the two-sided mechanism can be generalized to multi-resource markets. An important multi-resource setting is that of bandwidth shared on a network of links. The same guarantees as in the single-resource setting hold for the network version of our market, as well as for more general multi-resource markets (see the full version of the paper for more details).

Theorem 4. *The welfare ratio in networks equals that of the single-resource model.*

Theorem 5. *The welfare guarantees hold for more general multi-resource markets.*

Finally, we show that the proportional allocation mechanism is optimal in the following way:

Theorem 6. *In two-sided markets, the proportional allocation mechanism provides the best welfare ratio amongst a class of single-parameter market-clearing mechanisms.*

Our proof techniques are inspired by the approaches and techniques developed to analyze single-sided markets by Johari [3], Johari and Tsitsiklis ([5], [8] and [7]), Johari, Mannor and Tsitsiklis [4], Tobias and Harks [2], and Roughgarden [11]. Due to space limitations, most of our results will be deferred to the full version of the paper. Here, we will focus upon the proof of Theorem 3.

4 Optimization in Eight Steps

The proof of the main result, Theorem 3, is presented below in eight steps. We formulate the efficiency loss problem as an optimization program in Step III. To be able to formulate this we first need to understand the structure of optimal solutions and of equilibria under this mechanism. This we do in Steps I and II, where we give necessary and sufficient conditions for optimal solutions and for equilibrium. This leads us to an optimization problem that initially appears slightly formidable, so we then attempt to simplify it. In Steps IV and V, we show how to simplify the demand constraints in the program, and in Steps VI and VII, we simplify the supply constraints. This produces an optimization program in a form more amenable to quantitative analysis; we perform this analysis in Step VIII.

Step I: Optimality Conditions. The best possible allocation is the solution to the system:

$$\begin{aligned}
 \text{(OPT)} \quad & \max \sum_{q=1}^Q V_q(d_q^{\text{OPT}}) - \sum_{r=1}^R C_r(s_r^{\text{OPT}}) \\
 & \text{s.t.} \quad \sum_{q=1}^Q d_q^{\text{OPT}} = \sum_{r=1}^R s_r^{\text{OPT}} \\
 & \quad 0 \leq s_r^{\text{OPT}} \leq 1 \\
 & \quad d_q^{\text{OPT}} \geq 0
 \end{aligned}$$

Since the constraints are linear, there exists an optimal solution at which the Karush-Kuhn-Tucker (KKT) conditions hold. As the objective function is concave, the following first order conditions are both necessary and sufficient:

$$\begin{aligned}
 C'_r(s_r^{\text{OPT}}) &\leq \lambda && \text{if } 0 < s_r^{\text{OPT}} \leq 1 \\
 C'_r(s_r^{\text{OPT}}) &\geq \lambda && \text{if } 0 \leq s_r^{\text{OPT}} < 1 \\
 V'_q(d_q^{\text{OPT}}) &\leq \lambda && \text{if } d_q^{\text{OPT}} = 0 \\
 V'_q(d_q^{\text{OPT}}) &= \lambda && \text{if } d_q^{\text{OPT}} > 0
 \end{aligned}$$

We have used λ to denote the dual variable corresponding to the equality constraint.

Step II: Equilibria Conditions. Here we describe necessary and sufficient conditions for a set of bids \mathbf{b} to form Nash equilibrium.

First, observe that there must be at least two suppliers, that is $R \geq 2$. If not, then we have a monopolist k whose payoff is strictly increasing in b_k . Specifically,

$$\Pi_k(b_k, b_{-k}) = \sum_q b_q - C_k \left(1 - \frac{b_k}{b_k + \sum_q b_q}\right) = \sum_q b_q - C_k \left(\frac{\sum_q b_q}{b_k + \sum_q b_q}\right)$$

Next, we show that if \mathbf{b} is a Nash equilibrium, then at least two bids must be positive. Suppose for a contradiction that we have a supplier k and $\sum_{r \neq k} b_r = \sum_q b_q = 0$. Then $\Pi_k(0) = -C_k(1)$, and $\Pi_k(b_k) = -\frac{R-1}{R} b_k$ when $b_k > 0$. For the second expression, we used the fact that $C_k(x) = 0$ for any $x \leq 0$. Observe that if $b_k = 0$ then the firm can profitably deviate by increasing b_k infinitesimally; on the other hand, if $b_k > 0$ then the firm should infinitesimally decrease b_k . Thus, there is no equilibrium in which either all bids are zero, or a single supplier is the only agent to make a positive bid. Thus there must be at least two positive bids at equilibrium.

Since at least two bids are positive, the payoffs Π_k are differentiable and concave, and the following conditions are necessary and sufficient for the existence of a Nash equilibrium. For the suppliers,

$$C'_r(s_r) \left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \geq p \quad \text{if } 0 < b_r \leq p \quad C'_r(s_r) \left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \leq p \quad \text{if } 0 \leq b_r < p$$

For the consumers, $V'_q(0) \leq p$ and $V'_q(d_q^{\text{NE}}) \left(1 - \frac{d_q}{R}\right) = p$ if $d_q^{\text{NE}} > 0$.

Step III: An optimization problem. We can now formulate the welfare ratio as an optimization problem.

$$\min \frac{\sum_{q=1}^Q V_q(d_q^{\text{NE}}) - \sum_{r=1}^R C_r(s_r^{\text{NE}})}{\sum_{q=1}^Q V_q(d_q^{\text{OPT}}) - \sum_{r=1}^R C_r(s_r^{\text{OPT}})} \tag{1}$$

$$\text{s.t. } V'_q(d_q^{\text{NE}}) \left(1 - \frac{d_q^{\text{NE}}}{R}\right) \geq p \quad \forall q \text{ s.t. } d_q^{\text{NE}} > 0 \tag{2}$$

$$V'_q(d_q^{\text{NE}}) \left(1 - \frac{d_q^{\text{NE}}}{R}\right) \leq p \quad \forall q \tag{3}$$

$$C'_r(s_r^{\text{NE}}) \left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \leq p \quad \forall r \text{ s.t. } 0 < s_r^{\text{NE}} \leq 1 \tag{4}$$

$$C'_r(s_r^{\text{NE}}) \left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \geq p \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{NE}} < 1 \tag{5}$$

$$\sum_{q=1}^Q d_q^{\text{NE}} = \sum_{r=1}^R s_r^{\text{NE}} \tag{6}$$

$$C'_r(s_r^{\text{OPT}}) \leq \lambda \quad \forall r \text{ s.t. } 0 < s_r^{\text{OPT}} \leq 1 \tag{7}$$

$$C'_r(s_r^{\text{OPT}}) \geq \lambda \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{OPT}} < 1 \tag{8}$$

$$V'_q(d_q^{\text{OPT}}) \leq \lambda \quad \forall q \text{ s.t. } d_q^{\text{OPT}} = 0 \tag{9}$$

$$V'_q(d_q^{\text{OPT}}) = \lambda \quad \forall q \text{ s.t. } d_q^{\text{OPT}} > 0 \tag{10}$$

$$\sum_{q=1}^Q d_q^{\text{OPT}} = \sum_{r=1}^R s_r^{\text{OPT}} \tag{11}$$

$$d_q^{\text{OPT}}, d_q^{\text{NE}} \geq 0 \quad \forall q \tag{12}$$

$$0 \leq s_r^{\text{NE}}, s_r^{\text{OPT}} \leq 1 \quad \forall q, r \tag{13}$$

$$p, \lambda \geq 0 \tag{14}$$

Given the cost and valuation functions, the constraints (2)-(6) are necessary and sufficient conditions for a Nash equilibrium by Step II, and constraints (7)-(11) are the optimality conditions from Step I. We now want to find the worst-case cost and valuation functions for the mechanism.

Step IV: Linear Valuation Functions.

To evaluate this intimidating looking program we attempt to simplify it. First, efficiency loss is worst when each consumer has a linear valuation function. This is simple to show using a standard trick (see, for example, [5]). Thus, we restrict ourselves to linear functions of the form $V_q(d_q) = \alpha_q d_q$. Without loss of generality, we may assume that $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_Q$ and that $\max_q \alpha_q = 1$ after we normalize the functions by $1/\max_q \alpha_q$. Observe that this implies that $d_1^{\text{OPT}} = \sum_r s_r^{\text{OPT}}$ and $d_q^{\text{OPT}} = 0$ for $q > 1$. As a result the objective function becomes $(d_1^{\text{NE}} + \sum_{q=2}^Q \alpha_q d_q^{\text{NE}} - \sum_{r=1}^R C_r(s_r^{\text{NE}})) / (\sum_{r=1}^R s_r^{\text{OPT}} - \sum_{r=1}^R C_r(s_r^{\text{OPT}}))$, and the optimality constraints become $C'_r(s_r^{\text{OPT}}) \leq 1, \forall r \text{ s.t. } 0 < s_r^{\text{OPT}} \leq 1$ and $C'_r(s_r^{\text{OPT}}) \geq 1, \forall r \text{ s.t. } 0 \leq s_r^{\text{OPT}} < 1$. With linear valuations, the new optimality constraints ensure s_r^{OPT} is optimal by setting the marginal cost of each supplier to the marginal valuation, $\alpha_1 = 1$, of the first consumer.

Step V: Eliminating the Demand Constraints.

In this step, we describe how to eliminate the demand constraints from the program. First we show that we can transform constraint (14) into $0 \leq p < 1$. Since $\alpha_q \leq 1, \forall q$, we see that constraint (2) implies that $p \leq 1$. Furthermore, if $p = 1$, then (2) can never be satisfied, and so we must have $d_q^{\text{NE}} = 0, \forall q$. The supply equals demand constraint (6) then gives $s_r^{\text{NE}} = 0, \forall r$. This gives a contradiction as the resulting allocation is not a Nash equilibrium: any supplier can increase its profits by providing a bid slightly smaller than p (remember that $C'_r(0) = 0$ by Assumption 2). Thus $p < 1$. This, in turn, implies that $d_1^{\text{NE}} > 0$. To see this, note that if $d_1^{\text{NE}} = 0$ then (3) cannot be satisfied for $q = 1$. Consequently, constraints (2) and (3)

must hold with equality for $q = 1$. In fact, without loss of generality, constraints (2) and (3) hold with equality for $q > 1$. If constraint (2) does not hold with equality, we can reduce α_q , and this does not increase the value of the objective function. If $d_q^{\text{NE}} = 0$ and constraint (3) does not hold with equality, we can set $\alpha_q = p$ and the objective function will be unaffected. So, $\alpha_q = \frac{p}{1-d_q^{\text{NE}}/R}$ for all q . Substituting into the objective function:

$$\min \frac{d_1^{\text{NE}} + p \sum_{q=2}^Q \frac{d_q^{\text{NE}}}{1-d_q^{\text{NE}}/R} - \sum_{r=1}^R C_r(s_r^{\text{NE}})}{\sum_{r=1}^R s_r^{\text{OPT}} - \sum_{r=1}^R C_r(s_r^{\text{OPT}})} \tag{15}$$

$$\text{s.t.} \quad \left(1 - \frac{d_1^{\text{NE}}}{R}\right) = p \tag{16}$$

$$C'_r(s_r^{\text{NE}}) \left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \leq p \quad \forall r \text{ s.t. } 0 < s_r^{\text{NE}} \leq 1 \tag{17}$$

$$C'_r(s_r^{\text{NE}}) \left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \geq p \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{NE}} < 1 \tag{18}$$

$$\sum_{q=1}^Q d_q^{\text{NE}} = \sum_{r=1}^R s_r^{\text{NE}} \tag{19}$$

$$C'_r(s_r^{\text{OPT}}) \leq 1 \quad \forall r \text{ s.t. } 0 < s_r^{\text{OPT}} \leq 1 \tag{20}$$

$$C'_r(s_r^{\text{OPT}}) \geq 1 \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{OPT}} < 1 \tag{21}$$

$$d_q^{\text{NE}} \geq 0 \quad \forall q \geq 2 \tag{22}$$

$$d_1^{\text{NE}} > 0 \tag{23}$$

$$0 \leq s_r^{\text{NE}}, s_r^{\text{OPT}} \leq 1 \quad \forall r \tag{24}$$

$$0 \leq p < 1 \tag{25}$$

Now, observe that the objective function is convex and symmetric in the variables d_2, \dots, d_Q , when all the other variables are held fixed. Convexity holds because our function is a sum of functions $\frac{d_q^{\text{NE}}}{1-d_q^{\text{NE}}/R}$, $q = 2, \dots, Q$, that are convex on the range $[0, R]$; note that $d_q^{\text{NE}} \leq R$ by (6), (12) and (13). Therefore, for any given fixed assignment to the other variables, we must have $d_2 = \dots = d_Q := x$. Otherwise, we could reshuffle the variable labels and obtain a second minimum, which is impossible by the convexity of the objective function. So, after replacing every d_q by x , constraint (19) becomes $x = (\sum_{r=1}^R s_r^{\text{NE}} - d_1^{\text{NE}}) / (Q - 1)$. After inserting constraint (16) and the new constraint (19), the numerator of the objective function (15) becomes

$$(1 - p)R + p(Q - 1) \frac{x}{1 - x/R} - \sum_{r=1}^R C_r(s_r^{\text{NE}})$$

$$\begin{aligned}
 &= (1 - p)R + p(Q - 1) \frac{\left(\sum_{r=1}^R s_r^{\text{NE}} - d_1^{\text{NE}}\right) / (Q - 1)}{1 - \frac{1}{R} \left(\sum_{r=1}^R s_r^{\text{NE}} - d_1^{\text{NE}}\right) / (Q - 1)} - \sum_{r=1}^R C_r(s_r^{\text{NE}}) \\
 &= (1 - p)R + p \frac{\sum_{r=1}^R s_r^{\text{NE}} - (1 - p)R}{1 - \left(\sum_{r=1}^R s_r^{\text{NE}} - d_1^{\text{NE}}\right) / R(Q - 1)} - \sum_{r=1}^R C_r(s_r^{\text{NE}})
 \end{aligned}$$

Finally, observe that if we increase Q by one, the objective function (11) cannot increase, since we can set $d_{Q+1} = 0$ and at least keep the same objective function value as before. Therefore, without loss of generality, we can take the limit as $Q \rightarrow \infty$. Note that this only changes the objective function, as all the constraints that contained Q have been inserted into the function and can be eliminated. After these changes, the optimization problem becomes

$$\min \frac{(1 - p)^2 R + p \sum_{r=1}^R s_r^{\text{NE}} - \sum_{r=1}^R C_r(s_r^{\text{NE}})}{\sum_{r=1}^R s_r^{\text{OPT}} - \sum_{r=1}^R C_r(s_r^{\text{OPT}})} \tag{26}$$

$$\text{s.t. } C'_r(s_r^{\text{NE}}) \left(1 + \frac{s_r^{\text{NE}}}{R - 1}\right) \leq p \quad \forall r \text{ s.t. } 0 < s_r^{\text{NE}} \leq 1 \tag{27}$$

$$C'_r(s_r^{\text{NE}}) \left(1 + \frac{s_r^{\text{NE}}}{R - 1}\right) \geq p \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{NE}} < 1 \tag{28}$$

$$C'_r(s_r^{\text{OPT}}) \leq 1 \quad \forall r \text{ s.t. } 0 < s_r^{\text{OPT}} \leq 1 \tag{29}$$

$$C'_r(s_r^{\text{OPT}}) \geq 1 \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{OPT}} < 1 \tag{30}$$

$$0 \leq s_r^{\text{NE}}, s_r^{\text{OPT}} \leq 1 \quad \forall r \tag{31}$$

$$0 \leq p < 1 \tag{32}$$

Hence, we have achieved our goal and completely eliminated the demand side of the optimization problem. Specifically, all the demand constraints have been replaced with an expression that is a function of the supply-side allocation. Now we must find the worst such allocation.

Step VI: Linear Marginal Cost Functions. The next step is to show that, in searching for a worst case allocation, we can restrict our attention to linear marginal cost functions of the form $C'_r(s_r) = \beta_r s_r$ where $\beta_r > 0$. In this section, we briefly sketch the proof of this fact and defer the full treatment to the full version of the paper. Our proof technique is based on the work of Johari, Mannor and Tsitsiklis on demand-side markets with elastic supply [4], [6].

The proof consists in exhibiting, for any family of cost functions $C_r(s_r), r \in R$, two new families $\hat{C}_r(\cdot)$ and $\bar{C}_r(\cdot)$ with the property that the C_r have a better performance ratio than the \bar{C}_r which, in turn, have a better performance ratio than the \hat{C}_r . Furthermore, the \hat{C}_r will be a family with linear marginal costs, as desired. The cost functions are defined as

$$\bar{C}'_r(s_r) = \begin{cases} C'_r(s_r) & \text{if } s_r < s_r^{\text{NE}} \\ \frac{C'_r(s_r^{\text{NE}})}{s_r^{\text{NE}}} s_r & \text{if } s_r \geq s_r^{\text{NE}} \end{cases} \quad \text{and}$$

$$\hat{C}'_r(s_r) = \frac{C'_r(s_r^{\text{NE}})}{s_r^{\text{NE}}} s_r$$

where s_r^{NE} is the Nash equilibrium allocation to supplier r when the cost functions are $C_r(s_r)$. Observe that the s_r^{NE} still satisfy the Nash equilibrium conditions (27) and (28) for both \bar{C}_r and \hat{C}_r . Thus $\bar{s}_r^{\text{NE}} = \hat{s}_r^{\text{NE}} = s_r^{\text{NE}}$. The heart of the proof consists in showing that the optimal welfare can only improve when going from one family to the next.

Step VII: Eliminating the Supply Constraints. Assuming linear marginal cost functions, the optimization problem (26)-(32) becomes

$$\min \frac{(1-p)^2 R + p \sum_{r=1}^R s_r^{\text{NE}} - \frac{1}{2} \sum_{r=1}^R \beta_r (s_r^{\text{NE}})^2}{\sum_{r=1}^R s_r^{\text{OPT}} - \frac{1}{2} \sum_{r=1}^R \beta_r (s_r^{\text{OPT}})^2} \tag{33}$$

$$\text{s.t. } \beta_r s_r^{\text{NE}} \left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \leq p \quad \forall r \text{ s.t. } 0 < s_r^{\text{NE}} \leq 1 \tag{34}$$

$$\beta_r s_r^{\text{NE}} \left(1 + \frac{s_r^{\text{NE}}}{R-1}\right) \geq p \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{NE}} < 1 \tag{35}$$

$$\beta_r s_r^{\text{OPT}} \leq 1 \quad \forall r \text{ s.t. } 0 < s_r^{\text{OPT}} \leq 1 \tag{36}$$

$$\beta_r s_r^{\text{OPT}} \geq 1 \quad \forall r \text{ s.t. } 0 \leq s_r^{\text{OPT}} < 1 \tag{37}$$

$$0 \leq s_r^{\text{NE}}, s_r^{\text{OPT}} \leq 1 \quad \forall r \tag{38}$$

$$\beta_r > 0 \quad \forall r \tag{39}$$

$$0 \leq p < 1 \tag{40}$$

with the new variables β_r , $r = 1, \dots, R$. From $s_r^{\text{OPT}} \geq s_r^{\text{NE}}$, we can then deduce that (34) and (35) hold with equality. Suppose they don't for some r . Then $s_r^{\text{NE}} = s_r^{\text{OPT}} = 1$. Constraint (34) is $\beta_r < \frac{p}{1+1/(R-1)} < p < 1$. Hence, $\beta_r = \frac{p}{1+1/(R-1)}$ will be a feasible solution (i.e. constraint (36) will still be satisfied). Furthermore, increasing β_r to $\frac{p}{1+1/(R-1)}$ will only decrease the objective function since this is equivalent to subtracting a positive number from the numerator and the denominator. We can further simplify the system by replacing constraints (36) and (37) with $s_r^{\text{OPT}} = \min(1/\beta_r, 1)$. It is easy to see that s_r^{OPT} and β_r satisfy the equation above if and only if they satisfy (36) and (37). The reduced optimization problem now becomes:

$$\min \frac{(1-p)^2 R + p \sum_{r=1}^R s_r^{\text{NE}} - \frac{1}{2} \sum_{r=1}^R \beta_r (s_r^{\text{NE}})^2}{\sum_{r=1}^R s_r^{\text{OPT}} - \frac{1}{2} \sum_{r=1}^R \beta_r (s_r^{\text{OPT}})^2} \tag{41}$$

$$\text{s.t. } \beta_r s_r^{\text{NE}} \left(1 + \frac{s_r^{\text{NE}}}{R-1} \right) = p \quad \forall r \tag{42}$$

$$s_r^{\text{OPT}} = \min(1/\beta_r, 1) \quad \forall r \tag{43}$$

$$0 < s_r^{\text{NE}}, s_r^{\text{OPT}} \leq 1 \quad \forall r \tag{44}$$

$$\beta_r > 0 \quad \forall r \tag{45}$$

$$0 \leq p < 1 \tag{46}$$

We can insert the equality constraints (42) and (43) into the objective function (41) to obtain:

$$\min \frac{(1-p)^2 R + p \sum_{r=1}^R s_r^{\text{NE}} - \frac{p}{2} \sum_{r=1}^R \frac{s_r^{\text{NE}}}{1+s_r^{\text{NE}}/(R-1)}}{\sum_{r=1}^R \min(1/\beta_r, 1) - \frac{p}{2} \sum_{r=1}^R \frac{\min(1/\beta_r, 1)^2}{s_r^{\text{NE}}(1+s_r^{\text{NE}}/(R-1))}} \tag{47}$$

$$\text{s.t. } 0 < s_r^{\text{NE}} \leq 1 \quad \forall r \tag{48}$$

$$\beta_r = \frac{p}{s_r^{\text{NE}} (1 + s_r^{\text{NE}}/(R-1))} \quad \forall r \tag{49}$$

$$0 \leq p < 1 \tag{50}$$

The objective function (47) can be rewritten as:

$$\frac{\sum_{r=1}^R \left((1-p)^2 + p s_r^{\text{NE}} - \frac{p}{2} \frac{s_r^{\text{NE}}}{1+s_r^{\text{NE}}/(R-1)} \right)}{\sum_{r=1}^R \left(\min(1/\beta_r, 1) - \frac{p}{2} \frac{\min(1/\beta_r, 1)^2}{s_r^{\text{NE}}(1+s_r^{\text{NE}}/(R-1))} \right)}$$

Consequently, the minimum of the optimization problem (47)-(50) is greater than or equal to

$$\min \frac{(1-p)^2 + ps - \frac{p}{2} \frac{s}{1+s/(R-1)}}{\min\left(\frac{s(1+s/(R-1))}{p}, 1\right) - \frac{p}{2s(1+s/(R-1))} \min\left(\frac{s(1+s/(R-1))}{p}, 1\right)^2} \tag{51}$$

$$\text{s.t. } 0 < s \leq 1 \tag{52}$$

$$0 \leq p < 1 \tag{53}$$

We have now reduced the system (33)-(40) to a two-dimensional minimization problem. The next step is to try to explicitly find the minimum.

Step VIII: Computing the Worst Case Welfare Ratio. To obtain Theorem 3 we need to solve the optimization problem (51)-(53) with R as a parameter. We show how to do this in the full version of the paper. Thus we have proved our main result. It has several ramifications. Firstly, the worst case welfare ratio occurs with duopolies, that is when $R = 2$. There we obtain $s = 0.566812 \dots$ which gives a worst case welfare ratio of $0.588727 \dots$. Moreover, observe that this bound is tight. Our proof is essentially constructive; costs and valuations can be defined to create an instance that produces the bound. Secondly, the welfare

ratio improves as the number of supplies increases. Specifically as $R \rightarrow \infty$, the bound tends to $\frac{16}{25}$. Thus we obtain Corollaries [1](#) and [2](#).

So, as supply-side competition increases, the welfare ratio does improves. The opposite occurs as demand-side competition increases. Specifically, adapting our approach gives Corollary [3](#).

5 Concave Marginal Cost Functions

The welfare ratio tends to zero if the cost function is linear, that is if the marginal cost function is a constant; for an example see the full version of the paper. We can get some idea of how the welfare ratio tends to zero for concave marginal cost functions by considering a class of polynomial cost functions with degree $1 + \frac{1}{d}$. These functions give a welfare ratio of $\Omega(\frac{1}{d^2})$, for any constant d . A proof of this (Corollary [4](#)) is given in the full version of the paper. See Neumayer [10](#) for another example of inefficiency in the presence of linear cost functions.

6 Extensions to Networks and Arbitrary Markets

We can generalize our results for bandwidth markets over a single network connection to the case where bandwidth is shared over an entire network. In that model, each consumer q is associated with a source-sink pair, and providers at associated with edges of the network at which they can offer bandwidth. A consumer's payoff is a function of the maximum (s_q, t_q) -flow it can obtain using the bandwidth it has purchased in the network.

The welfare guarantees for the network model are the same as for the single-link case. A formal description of the network model and a proof of Theorem [4](#) is given in the full version of the paper. Moreover, if we identify links $e \in E$ with arbitrary resources, then our results extend to a general class of markets with any number of resources. The exact definition of these markets and a proof of Theorem [5](#) are also given in the full version of the paper.

7 Smooth Market-Clearing Mechanisms

It was shown in [3](#) and [8](#) that in one-sided markets, the proportional allocation mechanism uniquely achieves the best possible welfare ratio within a broad class of so-called *smooth market-clearing mechanisms*. This family has a natural extension to the case of two-sided mechanisms, and we show that, given a symmetry condition, the two-sided proportional allocation mechanism is optimal amongst that class of single-parameter mechanisms. A description of smooth market-clearing mechanisms and a proof of Theorem [6](#) is given in the full version of the paper.

References

1. Gupta, A., Stahl, D.O., Whinston, A.B.: The economics of network management. *ACM Commun.* 42(9), 57–63 (1999)
2. Harks, T., Miller, K.: Efficiency and stability of nash equilibria in resource allocation games. In: *GameNets 2009: Proceedings of The First ICST International Conference on Game Theory for Networks*, pp. 393–402. IEEE Press, Piscataway (2009)
3. Johari, R.: Efficiency loss in market mechanisms for resource allocation. Ph.D. thesis, MIT (2004)
4. Johari, R., Mannor, S., Tsitsiklis, J.N.: Efficiency loss in a network resource allocation game: The case of elastic supply. *Mathematics of Operations Research* 29, 407–435 (2005)
5. Johari, R., Tsitsiklis, J.N.: Efficiency Loss in a Network Resource Allocation Game. *Mathematics of Operations Research* 29(3), 407–435 (2004)
6. Johari, R., Tsitsiklis, J.N.: A scalable network resource allocation mechanism with bounded efficiency loss. *IEEE Journal on Selected Areas in Communications* 24(5), 992–999 (2005)
7. Johari, R., Tsitsiklis, J.N.: Parameterized supply function bidding: Equilibrium and Welfare (2006)
8. Johari, R., Tsitsiklis, J.N.: Efficiency of scalar-parameterized mechanisms. *Oper. Res.* 57(4), 823–839 (2009)
9. Kelly, F.: Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8, 33–37 (1997)
10. Neumayer, S.: Efficiency Loss in a Class of Two-Sided Mechanisms. Master’s thesis, MIT (2007)
11. Roughgarden, T.: Potential functions and the inefficiency of equilibria. In: *Proceedings of the International Congress of Mathematicians (ICM)*, vol. 3, pp. 1071–1094 (2006)
12. Semret, N.: Market mechanisms for network resource sharing. Ph.D. thesis, Columbia University, New York, NY, USA (1999)
13. Shapley, L.S., Shubik, M.: Trade using one commodity as a means of payment. *Journal of Political Economy* 85(5), 937–968 (1977)
14. Shenker, S., Clark, D., Estrin, D., Herzog, S.: Pricing in computer networks: re-shaping the research agenda. *SIGCOMM Comput. Commun. Rev.* 26(2), 19–43 (1996)
15. Songhurst, D.J.: Charging communication networks: from theory to practice. Elsevier Science Inc., New York (1999)
16. Stoft, S.: Power system economics: designing markets for electricity. IEEE Press, Piscataway (2002)
17. Tirole, J.: *The Theory of Industrial Organization*, 1st edn., vol. 1. The MIT Press, Cambridge (1988)

Braess's Paradox for Flows over Time

Martin Macko^{1,*}, Kate Larson², and Luboš Steskal^{1,**}

¹ Dept. of Computer Science, Comenius University, Slovak Republic

`{macko,steskal}@dcs.fmph.uniba.sk`

² Cheriton School of Computer Science, University of Waterloo, Canada

`klarson@cs.uwaterloo.ca`

Abstract. We study the properties of Braess's paradox in the context of the model of congestion games with flow over time introduced by Koch and Skutella. We compare them to the well known properties of Braess's paradox for Wardrop's model of games with static flows. We show that there are networks which do not admit Braess's paradox in Wardrop's model, but which admit it in the model with flow over time. Moreover, there is a topology that admits a much more severe Braess's ratio for this model. Further, despite its symmetry for games with static flow, we show that Braess's paradox is not symmetric for flows over time. We illustrate that there are network topologies which exhibit Braess's paradox, but for which the transpose does not. Finally, we conjecture a necessary and sufficient condition of existence of Braess's paradox in a network, and prove the condition of existence of the paradox either in the network or in its transpose.

Keywords: Flows over time, Braess's paradox, Dynamic flows, Selfish routing, Congestion games.

1 Introduction

Selfish routing and congestion games on networks have been analyzed mainly with respect to only static flows. The most prevalent model of congestion games with static flows is Wardrop's model [7,18] extensively studied by Roughgarden and Tardos [15,16]. A game in Wardrop's model is played by an infinite set of players each of which is selfishly routing only a negligible amount of his traffic.

In various applications, e.g. road traffic control and communication networks, however, flow variation over time is a crucial feature. Flow congestion on links and the time to traverse them may change over time in such applications and the flow does not reach its destination instantaneously, but it travels through the network at a certain speed determined by link transit times. We model these

* Supported in part by the Comenius University grant No. UK/358/2008, the National Scholarship Programme of the Slovak Republic, and the Tatra Banka Foundation grant No. 2009sds057.

** Supported in part by the National Scholarship Programme of the Slovak Republic, and the Tatra Banka Foundation grant No. 2009sds056.

phenomena by *flows over time* (also known as *dynamic flows*) introduced by Ford and Fulkerson [5,6].

Nash equilibria for flows over time were introduced by Vickrey [17] and Yagar [19] and mainly studied within the traffic community. For a survey see, e.g. [13]. In 2009, Koch and Skutella [9], defined a new variant of flows over time and introduced the notion of the price of anarchy for them. This model is based on the deterministic queueing model introduced by Vickrey [17], in which if at some point in time, more flow tries to enter a link than its capacity allows, the flow queues up at the link tail and waits until it may actually enter the link. The total time spent by a flow particle to traverse a single link is then the sum of the waiting time in the link queue and the actual time to traverse the link. In the model with flow over time introduced by Koch and Skutella, every link of a given network has a fixed capacity and a fixed free flow transit time. The link capacity bounds the maximal rate at which the flow may traverse the link and the link free flow transit time expresses the time a flow particle spends traveling from the link tail to its head.

It is a well known property of selfish routing with static flows that adding a new link to a network does not necessarily decrease the congestion in Nash equilibrium, but, paradoxically, it may even increase it, and so increase the cost of routing through the network. This phenomenon, discovered by Braess [2], is called Braess's paradox. For a survey see, e.g. [14].

For Wardrop's model of static flows, it is known [10,14] that the ratio by which the efficiency of a network may improve by removing any number of its links, i.e. Braess's ratio, is at most $\lfloor n/2 \rfloor$, where n is the number of network nodes. This bound is tight [8,14] for Wardrop's model. In this paper, we will prove that this bound does not generalize for the model of games with flow over time, and that there is a topology which admits a much more severe Braess's ratio. For flows over time nothing was known in this respect. Akamatsu and Heydecker [1] considered similar paradoxes for flows over time from a different point of view.

In principle, the only kind of topology that admits Braess's paradox in static flows is the Wheatstone network, see Fig. 1(a), also known as the θ -network. For Wardrop's model of static flows, it is known [12] that Braess's paradox may arise on and only on networks which contain the Wheatstone network as a topological minor. Recall, that a network H is called a topological minor of a network G if a subdivision of H is isomorphic to a subgraph of G . The networks that do not contain the Wheatstone network as a topological minor are usually called series-parallel [4], as they can be inductively composed by a number of series and parallel compositions from an edge. In other words, a network admits Braess's paradox in Wardrop's model if and only if it is not series-parallel. We will show that there are networks which do not admit Braess's paradox in Wardrop's model, but which admit it in the model with flow over time. Thus, we will introduce a class of series-parallel networks, for which there are instances of the model with flow over time that admit Braess's paradox. An example of such a network is shown in Fig. 1(b). Actually, all these networks are even extension-parallel.




Fig. 1. (a) The Wheatstone network, in principle the only topology that admits Braess's paradox in static flows; (b) The new topology for Braess's paradox in flows over time

As the Wheatstone network is symmetric, or more precisely, as it is isomorphic to its transpose, a network G admits Braess's paradox in Wardrop's model if and only if its transpose G^T admits Braess's paradox as well. Moreover, we know that an instance of Wardrop's model on a network G admits Braess's paradox if and only if the instance on the transpose network G^T with the same latency functions and the same flow supply admits it as well.

We will illustrate that in the model with flow over time there exist network topologies which exhibit Braess's paradox, but for which the transpose does not. Also, we will show that there is an infinite set of instances of the model with flow over time, which admit Braess's paradox, but none of the corresponding instances on their transpose networks with the same traffic supply and capacities and the same free flow transit times admits it any more. At the end of this paper, we conjecture a necessary and sufficient condition of existence of Braess's paradox in a network, and prove the condition of existence of the paradox either in the network or in its transpose.

The structure of this paper is as follows: In Section 2, we provide the formal definition of Koch's and Skutella's model of games with flow over time and define all notations we use later in this paper. Then, in Section 3, we prove the lower bound on Braess's ratio for the model of games with flow over time, and we show that there are networks which admit Braess's paradox in this model, but which do not admit it in Wardrop's model. In Section 4, we show that Braess's paradox is not symmetric in this model. And finally, in Section 5, we provide the necessary and sufficient conditions of existence of Braess's paradox in a network.

Due to space limitations, we omit all proofs in this extended abstract and refer to the full version¹ of the paper. More details can also be found in Macko's PhD thesis [11].

2 The Model

In this section, following its original definition introduced by Koch and Skutella [9], we define the model of games with flow over time.

¹ The full version of the paper with all proofs included can be found at <http://arxiv.org/pdf/1007.4864>

An instance of a game with flow over time is given by a tuple (G, c, τ, s, t, d) , where G is a network modeled by a directed graph $G = (V, E)$, $s \in V$ and $t \in V$ are source and sink nodes of G , respectively, $c = \{c_e\}_{e \in E}$ is a vector of *link capacities* with all $c_e > 0$, $\tau = \{\tau_e\}_{e \in E}$ is a vector of *link free flow transit times* with all $\tau_e \geq 0$, and $d > 0$ is an amount of the game input supply. We assume that there is at least one path from s to t in G . Let $P_{s,t}$ denote the set of all s - t -paths in G . Note that by the term path we mean a simple path.

The fundamental concept to this model are *waiting queues* that accumulate at the tails of network links if more flow wants to traverse a link than its capacity allows. Therefore, the link capacity bounds its outflow, that is the rate at which the flow leaves the link. The total *transit time* of a flow particle through a network link e at time θ is the sum of the *waiting time* $q_e(\theta)$ in the link queue at time θ and the link free flow transit time τ_e . The link free flow transit time determines the time the flow particle needs to traverse the link after leaving its waiting queue. The term *flow particle* represents an infinitesimally small flow unit that traverses the network along a single path.

For a given link e , the *actual link inflow* $f_e^+(\theta)$ is a function that determines the flow rate at which the flow enters the link e at its tail at time $\theta \geq 0$. Similarly, the *actual link outflow* $f_e^-(\theta)$ is a function that determines the flow rate at which the flow leaves the link e at its head at time $\theta \geq 0$. We have $f_e^+(\theta) \geq 0$ and $f_e^-(\theta) \geq 0$ for all $\theta \geq 0$. We usually omit the word actual and write just the link in- and outflow. Further, the *cumulative link inflow* $F_e^+(\theta)$ is a function that determines the total amount of flow that entered the link e until $\theta \geq 0$, and the *cumulative link outflow* $F_e^-(\theta)$ is a function that determines the total amount of flow that left the link e until $\theta \geq 0$. Thus, $F_e^+(\theta) = \int_0^\theta f_e^+(\vartheta) d\vartheta$ and $F_e^-(\theta) = \int_0^\theta f_e^-(\vartheta) d\vartheta$ for all $\theta \geq 0$. Note that all cumulative in- and outflows are continuous and nondecreasing. A *flow over time* is a vector $f = \{(f_e^+, f_e^-)\}_{e \in E}$ of pairs of the in- and outflows of all network links.

We say that a flow over time f is feasible if it satisfies the following conditions. The outflow of every link $e \in E$ is upper bounded by its capacity, therefore

$$f_e^-(\theta) \leq c_e, \quad (1)$$

for all $\theta \geq 0$. The flow leaves a link e after and only after it waits in its waiting queue and then it traverses the whole link, so for all $e \in E$ and $\theta \geq 0$ we have

$$F_e^+(\theta) - F_e^-(\theta + q_e(\theta) + \tau_e) = 0. \quad (2)$$

All flow that enters a node v continues immediately into node v out-links, and obviously, only flow that just entered the node v may continue into its out-links. This condition has two exceptions, namely the source node and the sink node. The amount of flow that leaves the source node through its out-links is always larger than the amount of flow that enters it through its in-links. This difference is exactly the network supply d . Reciprocally, the amount of flow that enters the sink node through its in-links may be larger than the amount of flow that leaves it through its out-links. This difference at time $\theta \geq 0$ is called the *actual*

sink flow and expresses the amount of flow that successfully finished its route from the source to the sink node at time θ . We denote it by $\gamma(\theta)$, and again, we usually omit the word actual and write just the sink flow. The *cumulative sink flow* $\Gamma(\theta)$ is a function that determines the total amount of flow that finished its route until $\theta \geq 0$, hence $\Gamma(\theta) = \int_0^\theta \gamma(\vartheta) d\vartheta$. Therefore, for all $\theta \geq 0$ the following condition must hold:

$$\sum_{e \in \delta_v^-} f_e^-(\theta) - \sum_{e \in \delta_v^+} f_e^+(\theta) = \begin{cases} 0 & \text{for } v \in V \setminus \{s, t\}, \\ -d & \text{if } v = s, \\ \gamma(\theta) & \text{if } v = t, \end{cases} \tag{3}$$

where δ_v^- and δ_v^+ are the sets of all in- and out-links of the node v , respectively.

Finally, the waiting time on any link e may not be negative, and if there is a nonempty waiting queue on the tail of the link e , the rate, at which the flow leaves the queue, must utilise the entire link capacity. If a flow particle leaves the waiting queue on the link e at time θ , then by condition (2), it will leave the link e at time $\theta + \tau_e$. Therefore, if a flow particle enters the link e at time θ , the cumulative amount of flow that entered the waiting queue until now is $F_e^+(\theta)$ and the cumulative amount of flow that left the waiting queue until now is $F_e^-(\theta + \tau_e)$. Hence, $F_e^+(\theta) - F_e^-(\theta + \tau_e)$ is the amount of flow waiting in the queue at time θ . So, the current waiting time in the queue of the link $e \in E$ at time $\theta \geq 0$ is

$$q_e(\theta) = \frac{F_e^+(\theta) - F_e^-(\theta + \tau_e)}{c_e},$$

and for all $e \in E$ and all $\theta \geq 0$ the following condition must hold:

$$q_e(\theta) \geq 0 \quad \text{and} \quad q_e(\theta) > 0 \Rightarrow f_e^-(\theta + \tau_e) = c_e. \tag{4}$$

Note that the function $\theta \rightarrow \theta + q_e(\theta)$ is increasing and continuous. This means that no flow particle may overtake any other flow particle in the link waiting queue. Similarly, it may not overtake any other flow particle in the rest of the link as well, as the link free flow transit time is constant for all flow particles. This means that every link in a feasible flow is FIFO.

A game with flow over time is a strategic game, in which every flow particle is an independent player with an infinitesimally small amount of traffic. The flows of particular players enter the network at the source node such that the total amount of flow that entered the network at any point of time $\theta \geq 0$ is equal to the game input supply d . Every player, before his piece of flow enters the network, independently chooses his *strategy*, that is a path from the source to the sink node his flow will route along. Then his flow enters the network and follows this path as quickly as possible. For every network s - t -path p , let $f^p(\theta)$ denote the amount of flow that entered the network at time $\theta \geq 0$ and is going to follow the path p . We know that $\sum_{p \in \mathcal{P}_{s,t}} f^p(\theta) = d$ for all $\theta \geq 0$.

Koch and Skutella [9] defined a feasible flow over time to be a Nash equilibrium, or in other words to be a Nash flow over time, if and only if the flow is

sent only over currently shortest paths, or equivalently, if and only if no flow overtakes any other flow.

For a fixed flow over time, let $\ell_v(\theta)$ denote the earliest point in time when a flow particle that entered the network at time θ may arrive at the node v . Then

$$\begin{aligned} \ell_s(\theta) &= \theta & \text{and} \\ \ell_w(\theta) &= \min\{\ell_v(\theta) + q_e(\ell_v(\theta)) + \tau_e \mid e = vw \in \delta_w^-\} \end{aligned}$$

for every node $w \in V \setminus \{s\}$ and all $\theta \geq 0$. We call these functions *label functions*. Note that the label functions are nondecreasing and continuous.

We say that a flow is *sent only over currently shortest paths* if for every link $e = vw \in E$ and all $\theta \geq 0$ we have

$$\ell_w(\theta) < \ell_v(\theta) + q_e(\ell_v(\theta)) + \tau_e \Rightarrow f_e^+(\ell_v(\theta)) = 0.$$

Similarly, we say that *no flow overtakes any other flow* if, for every flow particle, the amount of flow that entered the network before this flow particle equals the amount of flow that left the network before this flow particle. That is, if $d \cdot \theta = \Gamma(\ell_t(\theta))$ for all $\theta \geq 0$.

Koch and Skutella showed that flow over time is sent only over currently shortest paths if and only if no flow overtakes any other flow. This gives us a pair of handy characterizations of Nash equilibria for flows over time. Finally, they showed that for every instance of the model of games with flow over time there exists a flow in a Nash equilibrium.

For convenience in the rest of this paper, let $\lambda_p(\theta)$ denote the time spent by a flow particle traveling along a network path p if the flow particle entered the queue at the first link of p at time $\theta \geq 0$. We know that if the path p consists only of one link $e \in E$, then $\lambda_e(\theta) = q_e(\theta) + \tau_e$. If p contains more links, let e_1 denote its first link and p' the rest of the path p , that is $p = e_1p'$. Then $\lambda_p(\theta) = \lambda_{p'}(\theta + q_{e_1}(\theta) + \tau_{e_1})$. We call these functions *latency functions* or *latencies*.

Further, let $\lambda_v(\theta)$ denote the shortest time in which a flow particle may get to a node $v \in V$ if the flow particle entered the network source node at time $\theta \geq 0$. That is, $\lambda_v(\theta) = \min_{p \in P_{s,v}} \{\lambda_p(\theta)\}$. Notice that $\lambda_v(\theta) = \ell_v(\theta) - \theta$.

In this paper we investigate Braess's paradox with respect to the social cost function SC, which expresses the maximum experienced latency of a flow particle. For a feasible flow f of an instance A of a game with flow over time, it is defined as follows:

$$SC(f) = \sup_{\theta \geq 0} \max_{p \in P_{s,t}} \left([f^p(\theta) > 0] \cdot \lambda_p(\theta) \right).$$

So, we are taking into account the supremum of the latencies of all paths over the points in time in which a non-negligible amount of flow used the particular path. By a non-negligible amount of flow we mean a strictly positive amount of flow.

For every instance of a game with flow over time, we believe, that all its Nash flows are in principle equivalent, in the sense that their social costs are equal.

However, this has not yet been proven. So we define the Braess’s ratio with respect to the worst case Nash flow, where by the worst Nash flow we mean the Nash flow with the highest social cost. Nevertheless, all instances of games with flow over time we use in our proofs have all their Nash flows provably equivalent.

Let $A = (G, c, \tau, s, t, d)$ be an instance of a game with flow over time on a network G and f^* its worst Nash flow. We say that $BR(A)$ is Braess’s ratio of the instance A with respect to the social cost function SC , and define it as follows:

$$BR(A) = \max \left\{ \frac{SC(f^*)}{SC(f_H^*)} \mid H \subseteq G \right\},$$

where f_H^* is the worst Nash flow of the instance (H, c, τ, s, t, d) on the subgraph H . Braess’s ratio of a nonempty class of instances of games with flow over time is the supremum of Braess’s ratios of particular instances.

We say that an instance admits Braess’s paradox if its Braess’s ratio is strictly greater than one. Similarly, we say that a network admits Braess’s paradox if there is an instance on this network with its Braess’s ratio strictly greater than one.

For simplicity, we will write the value of the social cost $SC(f^*)$ of the worst Nash flow f^* of an instance A as $SC^*(A)$.

3 Lower Bound on Braess’s Ratio

In this section we provide a lower bound on Braess’s ratio for the model of games with flow over time, and we show that there is a topology which admits Braess’s paradox in this model, but which does not admit it in games with static flows.

Let’s consider an instance $A_n := (M_n, c, \tau, s, t, d)$ of a game with flow over time on the network M_n , for $n \geq 2$, as shown in Fig. 2 with the source node $s = v_1$ and the sink node $t = v_n$. The input supply of the network is $d = \alpha_0$. The free flow transit times and the capacities of the network links are defined as follows: $\tau_{e_k} = 0$, $\tau_{f_k} = T$ and $c_{e_k} = \alpha_k$, for $1 \leq k \leq n - 1$, $c_{f_k} = \alpha_{k-1} - \alpha_k$, for $1 \leq k \leq n - 2$ and $c_{f_{n-1}} = \alpha_{n-2}$, where $T > 0$ and $0 < \alpha_{n-1} < \dots < \alpha_2 < \alpha_1 < \alpha_0 = d$. Let p_1 denote the s - t -path consisting of the single link f_1 , p_k the s - t -path $e_1 e_2 \dots e_{k-1} f_k$, for $2 \leq k < n$, consisting of several e -links and the link f_k , and finally let p_0 denote the path $e_1 e_2 \dots e_{n-1}$ that uses only e -links, but no f -link. We will show that Braess’s ratio of such an instance may be arbitrarily close to $n - 1$, depending only on α ’s we choose.

In every Nash equilibrium at time zero, all first flow particles follow only the path p_0 , as it is the only s - t -path with zero free flow transit time and so the only s - t -path with zero latency. As $d > c_{e_1} > c_{e_2} > \dots > c_{e_{n-1}}$, a linearly increasing waiting queue accumulates on every e -link and the total transit times on the path p_0 and all paths p_k ($k \geq 2$) linearly increase with the time when a flow particle entered the network. Since the queue accumulates on every e -link and the free flow transit times of all f -links are equal, the latency on every path p_k is strictly greater than the latency on the path p_{k-1} at any positive time θ when a flow particle entered the network, until time θ_1 when the latency on the path




Fig. 2. The network M_n . The network topology which admits a severe Braess's Paradox for congestion games with flows over time, but which does not admit the paradox for Wardrop's model. All these networks are series-parallel, and even extension-parallel.

p_0 reaches T . At this time, the latency on the path p_0 is equal to the latency on the path p_1 . Therefore, any flow particle that enters the network at time θ_1 may follow either the path p_0 or the path p_1 , but not any other path.

After time θ_1 , if $n \geq 3$, the network supply splits between the paths p_0 and p_1 in such a way that the latency on the path p_1 begins to increase uniformly with the latency on the path p_0 . In particular, the path p_0 will gain the amount of $\alpha_{n-1}/(\alpha_0 - \alpha_1 + \alpha_{n-1})$ and the path p_1 the amount of $(\alpha_0 - \alpha_1)/(\alpha_0 - \alpha_1 + \alpha_{n-1})$ portions of the supply d . As $d \cdot \alpha_{n-1}/(\alpha_0 - \alpha_1 + \alpha_{n-1}) < \alpha_1$, the latency on the link e_1 decreases and the waiting queue on the link shortens after time θ_1 . If we are able to choose α 's such that the link e_1 never drains, the link outflow stays constant and equal to α_1 forever, and the flow on the network induced by the nodes v_2 to v_n behaves the same way as a flow on the network M_{n-1} of the instance A_{n-1} with an input supply equal to α_1 .

Indeed, if we choose α 's such that none of the e -links ever drains, the latency on the link e_{n-1} will increase up to T , eventually, with the latencies on all other e -links positive. Therefore, the maximum experienced transit time of a flow particle will be strictly greater than T . The next lemma shows that there are α 's such that the maximum experienced transit time of a flow particle in the instance A_n is almost $(n - 1) \cdot T$.

Lemma 1. *Let $n \geq 2, 0 < \varepsilon < 1/2n, j \geq 1$ and $\alpha_k = 1 + \varepsilon^{j+k}$, for $0 \leq k \leq n-1$. In every Nash equilibrium, the transit time $\lambda_t(\theta)$ of a flow particle that entered the instance A_n at time $\theta > T/\varepsilon^{j+n}$ is:*

$$\lambda_t(\theta) > (1 - 2n\varepsilon) \cdot (n - 1) \cdot T.$$

The previous lemma shows that the maximum experienced transit time of a flow particle, in a Nash equilibrium of A_n , may be arbitrarily close to $(n - 1) \cdot T$, where n is the number of network nodes. The following theorem shows that

there is a subgraph of M_n , for which the maximum experienced transit time of a flow particle in every Nash equilibrium is almost $n - 1$ times better than in the original graph.

Theorem 2. *For every $\varepsilon > 0$ and $n \geq 3$, the network M_n has a subgraph H such that, for the instance $A_n = (M_n, c, \tau, s, t, d)$ we have:*

$$\text{SC}^*(M_n, c, \tau, s, t, d) > (1 - \varepsilon) \cdot (n - 1) \cdot \text{SC}^*(H, c, \tau, s, t, d).$$

Therefore, if we choose a sufficiently small ε , Braess’s ratio of the instance A_n gets arbitrarily close to $n - 1$.

Corollary 3 (Lower bound on Braess’s ratio). *For every $n \geq 3$, Braess’s ratio of the class \mathcal{I}_n of all instances of the game with flow over time on networks with n nodes is $\text{BR}(\mathcal{I}_n) \geq n - 1$.*

Corollary 4 (A new topology for Braess’s paradox). *For every $n \geq 3$, there is a network with n nodes, which admits Braess’s paradox in the model of games with flow over time, but which does not admit it in Wardrop’s model. In particular, it is the network M_n .*

The construction in Lemma 1 also works if we restrict the model to instances with only integer link capacities. For a given ε from the lemma, take the smallest integer a such that $1/2^a \leq \varepsilon$, and let $\alpha_k = 2^{a(n+j)} + 2^{a(n-k)}$. By a proof similar to the proof of Lemma 1, we can show that in a Nash equilibrium the transit time of a flow particle that entered the network at time $\theta > T \cdot 2^{a(j+n)}$ is more than $(1 - n/2^{a-1}) \cdot (n - 1) \cdot T$.

If we restrict the model only to instances with unit link capacities, the lower bound on Braess’s ratio as a function of the number of network nodes still holds. We only need to replace every network link with integer capacity c by a set of c parallel links with unit capacities. However, in this case, the number of network links grows exponentially with n and polynomially with $1/\varepsilon$.

4 Asymmetry of Braess’s Paradox

For every $n \geq 3$, we have shown that the instance A_n as defined in the previous section has Braess’s ratio arbitrarily close to $n - 1$ for sufficiently small ε , and so it admits Braess’s paradox. Now, we will show that the instance on the transpose network with the same traffic supply and the same link capacities and free flow times has Braess’s ratio equal to 1.

Consider an instance $A_n^T := (M_n^T, c, \tau, s', t', d)$ on the network M_n^T , for $n \geq 2$, as shown in Fig. 3, with an input supply $d = \alpha_0$. The network M_n^T is a transpose of the network M_n , that is the network M_n with all its links reversed and its source and sink nodes swapped. Therefore, the M_n^T source and sink nodes are $s' = v_n$ and $t' = v_1$, respectively, and the free flow transit times and the capacities of its links are defined as follows: $\tau_{e_k} = 0$, $\tau_{f_k} = T$ and $c_{e_k} = \alpha_k$, for $1 \leq k \leq n - 1$, $c_{f_k} = \alpha_{k-1} - \alpha_k$, for $1 \leq k \leq n - 2$ and $c_{f_{n-1}} = \alpha_{n-2}$, where $T > 0$ and




Fig. 3. The transpose M_n^T of the network M_n , which admits no Braess's paradox in the model of games with flow over time

$0 < \alpha_{n-1} < \dots < \alpha_2 < \alpha_1 < \alpha_0 = d$. Similarly, denote p_k^T as the reverse of the path p_k . Therefore, p_0^T is the path $e_{n-1}e_{n-2}\dots e_1$ and p_k^T is the path $f_k e_{k-1} e_{k-2} \dots e_1$, for $k \geq 1$.

We will show, that for all $n \geq 2$ the instance A_n^T does not admit Braess's paradox. That is, there is no subgraph H of the network M_n^T , for which the maximum experienced transit time of a flow particle in any Nash equilibrium of the instance (H, c, τ, s', t', d) would be smaller than the maximum experienced transit time of a flow particle in any Nash equilibrium of the instance A_n^T .

Lemma 5. *For every $n \geq 2$, Braess's ratio of the instance A_n^T is $BR(A_n^T) = 1$.*

We have proved, that none of the instances A_n^T admits Braess's paradox. In fact, it is possible to show that there is no Braess's paradox even for the instance $(M_n^T, c, \tau, s', t', d)$ with any $d \leq \alpha_0$, not just $d = \alpha_0$.

By Theorem 2, for every $n \geq 3$, we know that Braess's ratio of the instance A_n is arbitrarily close to $n - 1$, and so the instance A_n admits a rather severe Braess's paradox. Thus, together with the previous lemma, we get the following theorem:

Theorem 6 (Braess's paradox asymmetry). *For every $n \geq 3$, there is an instance of a game with flow over time on a network G with n nodes, which admits Braess's paradox, but for which the corresponding instance on the transpose network G^T does not admit it. In particular, it is the instance A_n .*

So, there is an instance on the network M_3 , see Fig. 4(a), namely A_3 , that admits Braess's paradox, and we have shown that the instance A_3^T on the transpose network M_3^T , see Fig. 4(b), with the same flow supply and the same link capacities and free flow times does not admit it. In fact, for this particular transpose network, there is no instance which would have Braess's ratio strictly larger than 1, and so, which would admit Braess's paradox.




Fig. 4. (a) M_3 ; (b) M_3^T ; (c) M_3' ; (d) M_3''

Theorem 7 (Braess’s paradox asymmetry for networks). *There is a network G with an instance of the game with flow over time that admits Braess’s paradox, for which there is no instance of the game with flow over time on its transpose G^T that would admit the paradox. In particular, it is the network M_3 .*

5 Necessary and Sufficient Condition for Braess’s Paradox

In this section, we would like to answer the question, which topologies in general admit Braess’s paradox in the model of games with flow over time. So, we would like to characterize the class of all such networks in this model. Foremost, we show that every network which contains either the network M_3 , or its variations, the network M_3' (see Fig. 4(c)) or the network M_3'' (see Fig. 4(d)) as a topological minor admits Braess’s paradox in this model, and then we conjecture that these three networks are essentially the only topologies that admit Braess’s paradox in this model in general.

The networks M_3' and M_3'' are very similar to M_3 . If we set the free flow transit times and the link capacities in these two networks the same way as in the network M_3 with the only difference that $\tau_g = 0$ and $c_g = d$, where d is the network supply, the instances on the networks M_3' , M_3'' and M_3 act the same way, and their social costs and Braess’s ratios are equal. Therefore, both M_3' and M_3'' admit Braess’s paradox, since M_3 admits it. Notice that despite their similarity, the networks M_3 , M_3' and M_3'' are not topological minors of each other. The conjecture that these three networks are essentially the only topologies that admit Braess’s paradox is motivated by the result, we investigate at the end of this section, that these three networks together with the network M_3^T are the only topologies that admit Braess’s paradox if we use the network both ways, the forward and also the reverse.

Theorem 8 (Sufficient condition for Braess’s paradox). *If a network G contains either the network M_3 , M_3' or M_3'' as a topological minor, then it admits Braess’s paradox in the model of games with flow over time.*

Conjecture 9 (Necessary and sufficient condition for Braess’s paradox). A network G admits Braess’s paradox in the model of games with flow over time if and

only if the network G contains either the network M_3 , M'_3 or M''_3 as a topological minor.

Another natural question to ask is, which topologies admit Braess's paradox if we would like to use the networks in both directions. That is, if every network could be used in the way it is defined to route the traffic from its source to its sink, and also if we could transpose it and use it to route the traffic in the opposite direction from the original sink to the original source traveling along reversed links.

A network in Wardrop's model admitted Braess's paradox if and only if its transpose admitted it. However, as we have shown, this is not the case for the model of games with flow over time. So, for the model of games with flow over time, we would like to characterize the class of networks which admit Braess's paradox either in their original or in their reverse direction.

By Theorem 8, for every network G , we know that either the network G or its transpose G^T admits Braess's paradox in the model of games with flow over time if the network G contains either the network M_3 , M'_3 or M''_3 as a topological minor or its transpose G^T contains a transpose of any of these three networks as a topological minor. This is equivalent to the condition that the network G contains either the network M_3 , M_3^T , M'_3 or M''_3 as a topological minor, since the networks M'_3 and M''_3 are symmetric, i.e., they are isomorphic to their transposes.

Call a network a *chain of parallel paths* if it can be constructed from a chain of parallel links by a number of link subdivisions, see Fig. 5 for illustration. We say that two nodes u and v of a network *use a chain of parallel paths* if the union of all paths from u to v is a chain of parallel paths, or there is no path from u to v in the network. Further, we say that a network *uses only chains of parallel paths* if every pair of the network nodes uses a chain of parallel paths.




Fig. 5. (a) Example of a chain of parallel paths; (b) and the corresponding chain of parallel links

We will show that the networks that use only chains of parallel paths are the only networks that contain neither the network M_3 , M_3^T , M'_3 nor M''_3 as a topological minor. Then we will show that no network that uses only chains of parallel paths admits Braess's paradox in the model of games with flow over time. This will give us a necessary and sufficient condition of existence of Braess's paradox either in a network or in its transpose.

Lemma 10. *A network G uses only chains of parallel paths if and only if it does not contain any of the networks M_3 , M_3^T , M'_3 and M''_3 as a topological minor.*

Lemma 11. *If a network G uses only chains of parallel paths, then it does not admit Braess's paradox in the model of games with flow over time.*

Thus, from the previous lemmas and the fact that a network admits Braess's paradox in the model of games with flow over time if it contains either M_3 , M'_3 or M''_3 as a topological minor, we get the following theorem:

Theorem 12 (Necessary and sufficient condition for Braess's paradox both-ways). *For any network G , the following statements are equivalent:*

- (i) *Either the network G or its transpose G^T admits Braess's paradox in the model of games with flow over time.*
- (ii) *The network G contains either M_3 , M_3^T , M'_3 or M''_3 as a topological minor.*
- (iii) *The network G does not use only chains of parallel paths.*

6 Conclusion

We have proved several new properties of Braess's paradox for congestion games with flow over time. However, a number of questions have been left open.

We showed that there are networks which do not admit Braess's paradox in games with static flows, but which admit it in the model with flow over time. We showed that these networks admit a much more severe Braess's ratio for this model. In particular, we showed that Braess's ratio of the class of all instances of games with flow over time on networks with n nodes is at least $n - 1$. What is the upper bound on Braess's ratio for this model?

Then, we illustrated that Braess's paradox is not symmetric for flows over time, although it is symmetric for the case of static flows. We showed that there are network topologies which exhibit Braess's paradox, but for which the transpose does not. Is this asymmetry of Braess's paradox inherent for flows over time? What are the properties of Braess's paradox for different models of games with flows over time?

Finally, we conjectured a necessary and sufficient condition of existence of Braess's paradox in a network, and proved the condition of existence of the paradox either in the network or in its transpose. Is this conjecture valid in general?

References

1. Akamatsu, T., Heydecker, B.: Detecting dynamic traffic assignment capacity paradoxes in saturated networks. *Transportation Science* 37(2), 123–138 (2003)
2. Braess, D.: Uber ein paradoxon aus der verkehrsplanung. *Unternehmensforschung* 12, 258–268 (1968); English translation in [3]

3. Braess, D., Nagurney, A., Wakolbinger, T.: On a paradox of traffic planning. *Transportation Science* 39(4), 446–450 (2005)
4. Duffin, R.J.: Topology of series-parallel networks. *J. Math. Anal. Applications* 10, 303–318 (1965)
5. Ford, L.R., Fulkerson, D.R.: Constructing maximal dynamic flows from static flows. *Operations Research* 6(3), 419–433 (1958)
6. Ford, L.R., Fulkerson, D.R.: *Flows in Networks*. Princeton University Press, Princeton (1962)
7. Haurie, A.B., Marcotte, P.: On the relationship between nash-cournot and wardrop equilibria. *Networks* 15, 295–308 (1985)
8. Kameda, H.: How harmful the paradox can be in the Braess/Cohen-Kelly-Jeffries networks. In: *IEEE INFOCOM*, vol. 1, pp. 437–445 (2002)
9. Koch, R., Skutella, M.: Nash equilibria and the price of anarchy for flows over time. In: *Algorithmic Game Theory*, pp. 323–334 (2009)
10. Lin, H., Roughgarden, T., Tardos, É.: A stronger bound on braess's paradox. In: *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2004)*, pp. 340–341. Society for Industrial and Applied Mathematics, Philadelphia (2004)
11. Macko, M.: *The Price of Anarchy in Network Congestion Games*. PhD thesis, Faculty of Mathematics, Physics and Informatics, Comenius University, Bratislava, Slovakia (2010)
12. Milchtaich, I.: Network topology and the efficiency of equilibrium. *Games and Economic Behavior* 57(2), 321–346 (2006)
13. Peeta, S., Ziliaskopoulos, A.K.: Foundations of dynamic traffic assignment: The past, the present and the future. *Networks and Spatial Economics* 1(3-4), 233–265 (2001)
14. Roughgarden, T.: On the severity of braess's paradox: designing networks for selfish users is hard. *Journal Computer System Sciences* 72(5), 922–953 (2006)
15. Roughgarden, T.: Selfish routing and the price of anarchy. *Optima Mathematical Programming Society Newsletter* 74, 1–14 (2006)
16. Roughgarden, T., Tardos, É.: How bad is selfish routing? *Journal of the ACM* 49(2), 236–259 (2002)
17. Vickrey, W.S.: Congestion theory and transport investment. *The American Economic Review* 59(2), 251–260 (1969)
18. Wardrop, J.G.: Some theoretical aspects of road traffic research. In: *Proceedings of the Institute of Civil Engineers, Pt. II*, vol. 1, pp. 325–378 (1952)
19. Yagar, S.: Dynamic traffic assignment by individual path minimization and queuing. *Transportation Research* 5(3), 179–196 (1971)

The Price of Anarchy in Network Creation Games Is (Mostly) Constant

Matúš Mihalák and Jan Christoph Schlegel

Institute of Theoretical Computer Science, ETH Zurich, Switzerland

Abstract. We study the price of anarchy and the structure of equilibria in network creation games. A network creation game (first defined and studied by Fabrikant et al. [4]) is played by n players $\{1, 2, \dots, n\}$, each identified with a vertex of a graph (network), where the strategy of player i , $i = 1, \dots, n$, is to build some edges adjacent to i . The cost of building an edge is $\alpha > 0$, a fixed parameter of the game. The goal of every player is to minimize its *creation cost* plus its *usage cost*. The creation cost of player i is α times the number of built edges. In the SUMGAME (the original variant of Fabrikant et al. [4]) the usage cost of player i is the sum of distances from i to every node of the resulting graph. In the MAXGAME (variant defined and studied by Demaine et al. [3]) the usage cost is the eccentricity of i in the resulting graph of the game. In this paper we improve previously known bounds on the price of anarchy of the game (of both variants) for various ranges of α , and give new insights into the structure of equilibria for various values of α . The two main results of the paper show that for $\alpha > 273 \cdot n$ all equilibria in SUMGAME are trees and thus the price of anarchy is constant, and that for $\alpha > 129$ all equilibria in MAXGAME are trees and the price of anarchy is constant. For SUMGAME this (almost) answers one of the basic open problems in the field – is price of anarchy of the network creation game constant for all values of α ? – in an affirmative way, up to a tiny range of α .

1 Introduction

Network creation game, as defined and introduced by Fabrikant et al. in [4], is a game that models the process of building large autonomous computer and communication networks (such as the Internet). In this game, as in the reality, these networks are built and maintained by entities (*players* in the game-theoretic jargon) that pursue their own goals that may be different from the goals of other players – the players do not necessarily cooperate, they are *selfish* (we leave the meaning of this on an intuitive level). Network creation games is a well-studied and well-known research topic which is covered by many lectures and courses on algorithmic game theory and related subjects.

Network creation game is a strategic game with n players where each player is identified with a vertex (of a to be built graph/network). Every player i has to decide what edges incident to i the player *creates* (or *buys*, or *builds*). Building one edge costs the player $\alpha > 0$, which is a fixed parameter of the game. The

edges that the players buy form a graph (network) which is the result of the game. The players pursue two incompatible goals: pay as little as possible (minimize the *creation cost*), and have a good connection to other nodes of the network (maximize the *usage utility*). The usage utility of player i has been originally expressed as the following usage cost: the sum of distances to all other players in the resulting network [4] (where naturally players want to minimize this sum). Recently, the game where the usage cost of player i is expressed as the maximum distance of i to any node of G has been studied [3]. In this paper we consider both variants.

The central question that motivated the study of network creation games is: what do we lose in terms of quality of a network, if the communication network is built autonomously by selfish agents, as opposed to a communication network that is centrally planned and built? The *price of anarchy* of a game is a way to express this in that one compares the cost of a worst Nash equilibrium¹ (worst in the sense of the cost of the network) with the cost of an optimum network – the ratio of these two values is the *price of anarchy* of the game.

Definition of the game and related concepts. Let $G = (V, E)$ be an undirected graph (and we shall only consider undirected graphs in the following). For $u, v \in V$ we denote by $d_G(u, v)$ the length of a shortest u - v -path in G , and by $D_G(v)$ the *eccentricity* of the vertex v , i.e., the maximum distance between v and any other vertex of G . If G is not connected we define $d_G(u, v) := \infty$. We denote the degree of vertex $v \in V$ in G by $\deg_G(v)$. The *average degree* of G is $\deg(G) := \frac{1}{|V|} \sum_{v \in V} \deg_G(v) = \frac{2|E|}{|V|}$. We sometimes omit the index G and write simply $d(u, v)$, $D(v)$, or $\deg(v)$ if the underlying graph G is clear from the context. For $k \in \mathbb{N}$ we define the k -*neighborhood* of a vertex $v \in V$ as the set $N_k(v) := \{w \in V : d(v, w) \leq k\}$ (observe that v belongs to N_k), and the *boundary of the k -neighborhood* as the set $N_k^-(v) := \{w \in V : d(v, w) = k\}$. Furthermore we define the set of all eccentric vertices of v by $\mathcal{E}(v) := N_{D(v)}^-(v)$. We denote the *diameter* of G by $\text{diam}(G)$, and the *radius* of G by $\text{rad}(G)$. Recall that $\text{diam}(G) = \max_{u, v \in V} d(u, v)$ and $\text{rad}(G) = \min_v D(v)$. A *central vertex* is a vertex v for which $D(v) = \text{rad}(G)$. Graph G is a *star* if it is a tree and all edges of G are incident to one vertex. Recall that a *biconnected graph* is a graph that does not contain a *cut vertex*, i.e. a vertex whose removal makes the graph disconnected, and that a *biconnected component* (or *block*) of a graph G is a maximal biconnected subgraph of G .

We consider n players $N = \{1, \dots, n\}$ in our setting. Let $\alpha > 0$ be a real number which we shall call the *edge price*. The set of *strategies* of player $i \in N$ is the set $S_i = 2^{N \setminus \{i\}}$ (i.e., S_i contains all subsets of the set $N \setminus \{i\}$). A strategy $s_i \in S_i$ corresponds to a set of players to which i *buys* (or *builds*) an edge. We define $S := S_1 \times S_2 \times \dots \times S_n$ and call the elements of S the *strategy profiles*. For every strategy profile $s \in S$ we define the graph $G(s) := (N, \bigcup_{i=1}^n \bigcup_{j \in s_i} \{\{i, j\}\})$, and a cost function $c_i(s)$ for every player i (to be specified later). The triple (N, S, c) , where $c : S \rightarrow \mathbb{R}^n$ is given by $c(s) := (c_1(s), \dots, c_n(s))$, naturally

¹ In Nash equilibrium no player can unilaterally change its strategy and improve.

defines a non-cooperative n -player strategic game. Depending on the form of $c_i(\cdot)$ we distinguish two games. *Sum-Unilateral Network Creation Game*, or shortly SUMGAME, is the game given by (N, S, c) where for $s \in S, i \in N$,

$$c_i(s) = \alpha \cdot |s_i| + \sum_{j=1, \dots, n} d_{G(s)}(i, j).$$

Max-Unilateral Network Creation Game, or shortly MAXGAME, is the game given by (N, S, c) where for $s \in S, i \in N$,

$$c_i(s) = \alpha \cdot |s_i| + \max_{j=1, \dots, n} d_{G(s)}(i, j).$$

We call the term $\alpha \cdot |s_i|$ in both cost functions the *creation cost*, and the term $\sum_{j=1, \dots, n} d_{G(s)}(i, j)$ or $\max_{j=1, \dots, n} d_{G(s)}(i, j)$ in the respective cost function the *usage cost* of player i . A *Nash equilibrium* (NE for short) of the game (N, S, c) is a strategy-profile $s \in S$ such that for every player $i \in N$ and every strategy $\tilde{s}_i \in S_i$ we have $c_i(s) \leq c_i(s_1, \dots, s_{i-1}, \tilde{s}_i, s_{i+1}, \dots, s_n)$, i.e. no player can lower her cost by changing her strategy when all other players keep their strategies unchanged. Observe therefore that for every finite α every NE is a connected graph, and in every NE any edge is bought by at most one player. If $s \in S$ is a Nash equilibrium of the game (N, S, c) , we call $G(s)$ an *equilibrium graph* or sometimes a *stable graph*. The *social cost* C of a strategy-profile $s \in S$ is defined, for the respective cost function of player i , as the sum of the individual costs of the players under this strategy-profile, i.e.:

$$C(s) = \sum_{i=1}^n c_i(s) = \begin{cases} \alpha \cdot \sum_{i=1}^n |s_i| + \sum_{i=1}^n \sum_{j=1}^n d_{G(s)}(i, j) & \text{in SUMGAME,} \\ \alpha \cdot \sum_{i=1}^n |s_i| + \sum_{i=1}^n \max_{j=1, \dots, n} d_{G(s)}(i, j) & \text{in MAXGAME.} \end{cases}$$

Since for every graph $G = (V, E)$ on n vertices there is a strategy-profile inducing this graph, the social cost function generalizes for any graph G on n vertices: $C(G) = \alpha \cdot |E| + \sum_{v \in V} D_G(v)$, or $C(G) = \alpha \cdot |E| + \sum_{v \in V} \sum_{w \in V} d_G(v, w)$ for the respective cost function. We call a graph G_{OPT} minimizing the respective social cost function a *social optimum*. The *price of anarchy* (PoA for short) of a game (SUMGAME or MAXGAME) is defined as $\max_{s \in S; s \text{ is NE}} \frac{C(G(s))}{C(G_{OPT})}$.

Related work. Networks have been an important research topic in the economical and social sciences, as networks naturally model relationships between interacting entities. As such, a link between two entities is usually created upon mutual consensus (“if entity A knows entity B then entity B knows entity A ” is a common assumption). For an overview of economical and social studies from this perspective we refer to the book by Jackson [5] and to the references therein. Strategic network formation in this framework has been studied with tools from cooperative game theory. The trade-off between efficiency and stability for these kind of networks has been studied in Jackson and Wolinski [6].

We study networks where links can be created unilaterally (i.e., without an explicit agreement of both players at the ends of the respective edges) and where the payoff of the players reflects the cost for building the edges as well as the

quality of the resulting network in terms of the players’ distances in the network. The first game of this nature studied in the literature is SUMGAME.

Fabrikant et al. introduced and defined SUMGAME in [4]. They proved an upper bound $\mathcal{O}(\sqrt{\alpha})$ on PoA (by showing that PoA is bounded by the diameter of the equilibrium graph), and showed that every NE which is a tree has constant PoA (we will use this result later on). Albers et al. [1] showed that PoA is constant for $\alpha = \mathcal{O}(\sqrt{n})$ (this was also independently and earlier discovered by Lin [7]) and for $\alpha \geq 12n \lg n$. The latter result is achieved by showing that for $\alpha \geq 12n \lg n$ all NE are trees. Albers et al. also show the general bound $15(1 + (\min\{\alpha^2/n, n^2/\alpha\})^{1/3})$ for all α , which shows that PoA is $\mathcal{O}(n^{1/3})$ for all α . Demaine et al. [3] show that PoA is constant already for $\alpha = \mathcal{O}(n^{1-\varepsilon})$ for any fixed ε , and show the general bound $2^{\mathcal{O}(\sqrt{\lg n})}$ on PoA for all α .

Demaine et al. introduced and defined MAXGAME in [3] as a natural variant of the network creation games, and showed that PoA is at most 2 for $\alpha \geq n$, $\mathcal{O}(\min\{4^{\sqrt{\lg n}}, (n/\alpha)^{1/3}\})$ for $2\sqrt{\lg n} \leq \alpha \leq n$, and $\mathcal{O}(n^{2/\alpha})$ for $\alpha < 2\sqrt{\lg n}$.

Recently, a related model has been introduced by Alon et al. in [2] where players do not buy edges, but only swap the endpoints of existing edges. Alon et al. claim some implications of their model to the models studied in this paper; this, however, seems not to be the case in the claimed extent (but it is not easy to argue as Alon et al. do not state any such claim formally). We refer to [8], the full version of this paper, for a detailed discussion.

Our results. For MAXGAME we show that PoA is constant for $\alpha > 129$ and $\alpha = \mathcal{O}(n^{-1/2})$, and also prove that PoA is $2^{\mathcal{O}(\sqrt{\log n})}$ for any $\alpha > 0$ in Section 2. The result for $\alpha > 129$ is obtained as a corollary of the more general result (proved in Section 2.1) showing that in MAXGAME for $\alpha > 129$ all equilibrium graphs are trees. This is proved by new techniques which establish and use estimates on the average degree of biconnected components of equilibrium graphs. In Section 3 we adopt the new techniques for SUMGAME to prove that for $\alpha > 273n$ all equilibrium graphs are trees. This result implies a constant upper bound on PoA for $\alpha > 273n$ which shrinks the range of edge-prices for which we do not know a constant upper bound to $\alpha = \Theta(n)$. A comparison and overview of the previously known bounds and the new bounds on PoA in both game variants are summarized in Table 1 and Table 2.

2 Bounding the Price of Anarchy in MaxGame

In this section we consider MAXGAME. First we classify social optima. This is rather a folklore and resembles in many aspects the previously shown characterization of social optima in SUMGAME. We use this to bound PoA in MAXGAME for small values of α .

Proposition 1. *For $\alpha \leq \frac{2}{n-2}$ the complete graph is a social optimum. For $\alpha \geq \frac{2}{n-2}$ the star is a social optimum.*

Table 1. Comparison of the previously known bounds for the price of anarchy in MAXGAME (due to [3]) and the bounds proved in this paper. The abbreviations T. and C. stand for Theorem and Corollary, respectively.

$\alpha = 0$	$\frac{1}{n-2}$	$\mathcal{O}(n^{-\frac{1}{2}})$	129	$2\sqrt{\log n}$	$n \infty$
new	1 (T. [1])	$\Theta(1)$ (C. [2])	$2^{\mathcal{O}(\sqrt{\log n})}$ (T. [3])	< 4 (C. [4])	≤ 2
old	$\mathcal{O}(n^{2/\alpha})$			$\mathcal{O}(\min\{4^{\sqrt{\log n}}, (n/\alpha)^{1/3}\}) \leq 2$	

Table 2. Summary of the best known bounds for the price of anarchy in SUMGAME

$\alpha = 0$	1	2	$\sqrt[3]{n/2}$	$\sqrt{n/2}$	$\mathcal{O}(n^{1-\epsilon})$	273n	$12n \lg n$	∞
PoA	$1 \leq \frac{4}{3}$ ([4])	≤ 4 ([3])	≤ 6 ([3])	$\Theta(1)$ ([3])	$2^{\mathcal{O}(\sqrt{\log n})}$ ([3])	< 5 (T. [5])	≤ 1.5 ([1])	

Theorem 1. For $\alpha < \frac{1}{n-2}$ the price of anarchy is 1. For $\alpha < \frac{2}{n-2}$ the price of anarchy is at most 2.

The proofs of the two statements can be found in the full version of this paper [8]. Next we relate the diameter of an equilibrium graph with PoA of the game, where the following lemma is the key ingredient. The lemma exploits that a breadth-first search tree (BFS-tree) of an equilibrium graph already contains much information about the whole graph. For SUMGAME a similar result with a similar proof is known [1].

Lemma 1. If $G = (V, E)$ is an equilibrium graph then $C(G) \leq (2\alpha + 1)(n - 1) + n \cdot \text{rad}(G)$.

Proof. Let T be a BFS-tree of G rooted in a central vertex v_0 of G . Let $v \in V \setminus \{v_0\}$. Let E_v be the edges built by v in T . Consider the following strategy of v : Buy all edges of E_v plus buy the edge to v_0 . The creation cost of v in this strategy is at most $\alpha(|E_v| + 1)$ and the usage cost is at most $D(v_0) + 1$. As G is an equilibrium, every vertex (player) achieves in G the best possible cost, given what other players do. Thus, the above mentioned strategy upper-bounds the cost of v in equilibrium, i.e., $c_v(G) \leq \alpha(|E_v| + 1) + \text{rad}(G) + 1$. For vertex v_0 we have $c_{v_0}(G) = \alpha|E_{v_0}| + \text{rad}(G)$. Summing the obtained inequalities for every vertex of G yields the claimed inequality. \square

Corollary 1. Let G be a worst NE for $\alpha \geq \frac{2}{n-2}$. The price of anarchy is $\mathcal{O}\left(1 + \frac{\text{diam}(G)}{\alpha+1}\right)$.

Proof. By Proposition [1] and Lemma [1] we get

$$\text{PoA} \leq \frac{(2\alpha + 1)(n - 1) + n \cdot \text{rad}(G)}{(\alpha + 2)(n - 1) + 1} \leq \frac{2\alpha + 1}{\alpha + 2} + \frac{n \cdot \text{rad}(G)}{(n - 1)(\alpha + 2)} \leq 2 + \frac{2 \cdot \text{rad}(G)}{\alpha + 2}$$

\square

Demaine et al. showed in [3] that the diameter of equilibrium graphs is bounded by $\mathcal{O}(1 + \alpha 4^{\sqrt{\lg n}})$ and by $\mathcal{O}(1 + (n\alpha^2)^{1/3})$.² Combining these results with Corollary 1 yields an improved bound for the price of anarchy:

Lemma 2 ([3]). *The diameter of an equilibrium graph is $\mathcal{O}(1 + (n\alpha^2)^{1/3})$.*

Theorem 2. *For $\alpha = \mathcal{O}(1)$ the price of anarchy is $\mathcal{O}(1 + (n\alpha^2)^{1/3})$.*

Corollary 2. *For $\alpha = \mathcal{O}(n^{-1/2})$ the price of anarchy is constant.*

Lemma 3 ([3]). *The diameter of an equilibrium graph is $\mathcal{O}(1 + \alpha \cdot 4^{\sqrt{\lg(n)}})$.*

Theorem 3. *The price of anarchy is $2^{\mathcal{O}(\sqrt{\log(n)})}$.*

In the following we show that equilibrium graphs that are trees have cost at most a constant times bigger than the cost of a social optimum. Thus if for given α all equilibrium graphs are trees then PoA is constant. We note that a similar result for SUMGAME has been shown by Fabrikant et al. in [4]. We show in Section 2.1 that for $\alpha > 129$ all equilibrium graphs are trees which shows that PoA for this range of α is constant.

Theorem 4. *The cost of an equilibrium graph that is a tree is less than 4 times the cost of a social optimum.*

Proof. Observe that the claim is trivial when $n \leq 2$, or when $\alpha < 2/(n - 2)$ (as then, by Theorem 1, PoA is at most 2). We therefore assume that $n \geq 3$ and $\alpha \geq 2/(n - 2)$. Let $T = (V, E)$ be a tree on $n \geq 3$ vertices that is NE. We first show that $\text{diam}(T) \leq 2\alpha + 3$. Let $v \in V$ be a vertex with $D(v) = \lceil \text{diam}(T)/2 \rceil$ (observe that there exists such a vertex). Consider T rooted at v . Let l be a leaf of T at depth $D(v)$. Consider the strategy of l where l buys, additionally to what it does in the equilibrium strategy profile, an edge to v . The usage cost of l is at most $1 + D(v)$ using the new strategy. Its usage cost in the equilibrium strategy is $D(l) = \text{diam}(T)$. As T is NE we can conclude that buying the edge to v is not beneficial and therefore $\alpha \geq D(l) - (D(v) + 1) \geq \lceil \text{diam}(T)/2 \rceil - 1$. Hence, $\text{diam}(T) \leq 2\alpha + 3$. We now compare the cost of T with the cost of a social optimum G_{OPT} . As $\alpha \geq 2/(n - 2)$, a star is a social optimum (Proposition 1). Hence, as $C(G_{\text{OPT}}) = (\alpha + 2)(n - 1) + 1$,

$$\frac{C(T)}{C(G_{\text{OPT}})} \leq \frac{\alpha(n - 1) + \text{diam}(T) \cdot n}{(\alpha + 2)(n - 1)} \leq \frac{\alpha}{\alpha + 2} + \frac{(2\alpha + 3) \cdot n}{(\alpha + 2)(n - 1)} < 1 + 2 \cdot \frac{3}{2} = 4,$$

which proves the claim. □

² In fact, [3] claims a bound of $\mathcal{O}(\alpha 4^{\sqrt{\lg n}})$ resp. $\mathcal{O}((n\alpha^2)^{1/3})$ on the diameter, which does not make sense for very small α . The arguments given in [3] show a bound of $\mathcal{O}(1 + \alpha 4^{\sqrt{\lg n}})$ resp. $\mathcal{O}(1 + (n\alpha^2)^{1/3})$ on the diameter.

2.1 For $\alpha > 129$ Every Equilibrium Graph Is a Tree

In this section we present the main result for MAXGAME, namely, we show that for $\alpha > 129$ every equilibrium graph is a tree. This, together with Theorem 4, shows that PoA is smaller than 4 for this range of α . The main idea is to show that an arbitrary (non trivial) biconnected component of an equilibrium graph has average degree $c > 2$ and at the same time smaller than $2 + \frac{c'}{\alpha}$ for some constants c, c' . For big enough α these inequalities become contradicting and thus we know that this cannot happen, i.e., every NE for such α contains no biconnected component other than bridges and therefore no cycle – it has to be a tree.

For the entire section let $G = (V, E)$ be a graph on n vertices that contains at least one cycle and let $H \subseteq G$ be an (arbitrary) biconnected component of G of size $|H| \geq 3$. Furthermore we use the following definitions. For a vertex $v \in V$ and a set $X \subseteq V$ we call a path starting in v and ending in a vertex in X a v - X -path. For every vertex v in H we define $S(v)$ to be the set of all vertices $x \in V$ such that a shortest x - H -path ends in v . Note that by definition: $S(v) \neq \emptyset$ since $v \in S(v)$; v is the only vertex from H in $S(v)$; $S(u) \cap S(v) = \emptyset$ for $u \in V(H), u \neq v$; for every $w \in S(v)$ every shortest u - w -path contains v . We start with the observation of Demaine et al. [3] stating that there are no “short” cycles in equilibrium graphs.

Lemma 4 ([3]). *Every equilibrium graph has no cycle of length less than $\alpha + 2$.*

The following lemma shows that the usage cost of vertices in H differ by at most 4 and “tends to be lower” for a vertex that buys an edge in H .

Lemma 5. *If G is an equilibrium graph and $v \in V(H)$ then $D_G(v) \leq \text{rad}(G) + 3$ if v buys an edge in H and $D_G(v) \leq \text{rad}(G) + 4$ otherwise.*

Proof. We show that for every edge $\{u, v\} \in E(H)$ bought by u we have $D_G(u) \leq \text{rad}(G) + 3$ and $D_G(v) \leq \text{rad}(G) + 4$. The claim then follows. Consider a BFS-tree T rooted in some central vertex v_0 of G . First we consider the case that $\{u, v\} \in E(H) \setminus E(T)$. Trivially, $D_G(u) \leq \text{rad}(G) + 1$ (as otherwise u could buy an edge $\{u, v_0\}$ instead of $\{u, v\}$ and thus improve its cost) and therefore $D_G(v) \leq \text{rad}(G) + 2$. Next we consider the case that $\{u, v\} \in E(T) \cap E(H)$. We note that the edge either leads “up” the tree to v_0 , or it leads “down” the tree such that there is a vertex $s \in V(H)$ below or at v which is incident to an edge in $E(H) \setminus E(T)$ (if not then u would be a cut vertex of H). In the first case we have $D_G(u) \leq \text{rad}(G) + 1$ (as otherwise u could buy an edge $\{u, v_0\}$ instead of $\{u, v\}$ and thus improve its cost). In the second case we have, as shown before, $D_G(s) \leq \text{rad}(G) + 2$ and therefore $D_G(u) \leq \text{rad}(G) + 3$ (as otherwise u could buy an edge $\{u, s\}$ instead of $\{u, v\}$ and thus improve its cost). So in general we have $D_G(u) \leq \text{rad}(G) + 3$ and therefore $D_G(v) \leq \text{rad}(G) + 4$. □

In the following lemmata we show that for every vertex in a biconnected component H of an equilibrium graph G there is a vertex of degree at least 3 in H in a constant-size neighborhood of v .

Lemma 6. *If G is an equilibrium graph for $\alpha > 0$ then for every vertex v in H and every vertex $w \in S(v)$: $d_G(v, w) \leq \text{rad}(G) + \frac{7-\alpha}{2}$.*

Proof. By Lemma 4, H has no cycle of length less than $\alpha + 2$. Thus, as every vertex of H is contained in at least one cycle, there is a vertex $u \in V(H)$ with $d_G(u, v) = d_H(u, v) \geq \lfloor \frac{\alpha+2}{2} \rfloor \geq \frac{\alpha+1}{2}$. Every shortest u - w -path contains vertex v (by definition of $S(v)$). Therefore $d_G(u, w) = d_G(u, v) + d_G(v, w) \geq \frac{\alpha+1}{2} + d_G(v, w)$. By Lemma 5 we have $d_G(u, w) \leq D_G(u) \leq \text{rad}(G) + 4$. Hence $d_G(v, w) \leq \text{rad}(G) + \frac{7-\alpha}{2}$. \square

Lemma 7. *If G is an equilibrium graph for $\alpha > 11$, then for every vertex v in H that buys at least two edges in H there is a vertex $w \in N_1(v)$ with $\text{deg}_H(w) \geq 3$.*

Proof. Let us refer to v by x_2 and let x_1 and x_3 be two vertices to which x_2 buys edges in H . Assume for contradiction that $\text{deg}_H(x_i) = 2$ for $i = 1, 2, 3$. Denote the x_1 's other neighbor in H by x_0 and the x_3 's other neighbor in H by x_4 . Note that, as $\alpha > 11$, the girth of H is at least 14 (Lemma 4) and therefore $x_i \neq x_j$ for $i \neq j$. Also by Lemma 6 we have $d_G(x_2, w) < \text{rad}(G) - 1 \leq D_G(x_2) - 1$ for $w \in \bigcup_{i=1,2,3} S(x_i)$. Thus, all shortest x_2 - $\mathcal{E}(x_2)$ -paths contain either x_0 or x_4 . Hence, by buying edges to x_0 and x_4 instead of x_1 and x_3 , x_2 would decrease its distance to the vertices in $\mathcal{E}(x_2)$, increase its distance to the vertices in $S(x_1)$ and $S(x_3)$ by at most 1 and it would not increase its distance to any other vertex. Therefore (as $d_G(x_2, w) < D_G(x_2) - 1$ for $w \in S(x_1) \cup S(x_3)$), by changing its strategy x_2 could improve. But this contradicts equilibrium and hence we have $\text{deg}_H(x_i) \geq 3$ for some $i \in \{1, 2, 3\}$. \square

Lemma 8. *If G is an equilibrium graph for $\alpha > 13$ then any path x_0, x_1, \dots, x_k in H with $\text{deg}_H(x_i) = 2$ for $0 \leq i \leq k$ such that for $0 \leq i < k$, $\{x_i, x_{i+1}\}$ is bought by x_i , has length at most $k \leq 4$.*

Proof. Consider a maximal path x_0, x_1, \dots, x_k in H of the form from the statement and assume for contradiction $k \geq 5$. By Lemma 5 we have $|D_G(x_i) - D_G(x_j)| \leq 3$ for $0 \leq i, j \leq k - 1$ and therefore, by the pigeonhole principle, there is $0 \leq i_0 \leq 3$ such that $D_G(x_{i_0}) \geq D_G(x_{i_0+1})$. Denote the x_{i_0+2} 's other neighbor in H by x_{i_0+3} (if not already so denoted). For every vertex $w \in S(x_{i_0+j})$, $j = 0, 1, 2$, we have (using Lemma 6) $d_G(x_{i_0+j}, w) < \text{rad}(G) - 3$, and therefore $\mathcal{E}(x_{i_0}) \cap S(x_{i_0+j}) = \emptyset$ for $j = 0, 1, 2$.

We consider the strategy where x_{i_0} buys an edge to x_{i_0+3} instead of the edge to x_{i_0+1} and show that x_{i_0} improves in this strategy, which is a contradiction. We split the vertices of $\mathcal{E}(x_{i_0})$ into two parts: set S where for every $z \in S$ no shortest x_{i_0} - z -path contains x_{i_0+1} , and set $\mathcal{E}(x_{i_0}) \setminus S$ where for every vertex $z \in \mathcal{E}(x_{i_0}) \setminus S$ there is a shortest x_{i_0} - z -path that contains x_{i_0+1} (and therefore also x_{i_0+2} and x_{i_0+3}). Observe that in the new strategy x_{i_0} decreases its distance to vertices in $\mathcal{E}(x_{i_0}) \setminus S$ by 2, and increases its distance to vertices in $S(x_{i_0+1})$ by at most 2, and does not increase its distance to any other vertex of V but perhaps to those in S . We show that x_{i_0} actually decreases its distance to every vertex in S by at least one, which shows that x_{i_0} improves in the new strategy

(recall that $d_G(x_{i_0}, y) < D_G(x_{i_0}) - 2$ for every $y \in S(x_{i_0+1})$). To show that x_{i_0} improves its distance to every vertex $z \in S$, we first observe that because $D_G(x_{i_0}) \geq D_G(x_{i_0+1})$ no shortest x_{i_0+1} - z -path contains x_{i_0} . Thus, all shortest x_{i_0+1} - z -paths contain x_{i_0+3} . Hence, in the new strategy, x_{i_0} decreases its distance to z , which finishes the proof. \square

Lemma 9. *If G is an equilibrium graph for $\alpha > 13$ then for every vertex v in H there is a vertex $w \in N_5(v)$ with $\deg_H(w) \geq 3$.*

Proof. Let $\{u, v\}$ be an arbitrary edge in H and assume without loss of generality that u bought the edge. Let C be a cycle containing $\{u, v\}$ and note that by Lemma 4 it has at least 16 vertices. Denote the vertices after v and u (in that order) in C by x_0, x_1, x_2, \dots . We distinguish two cases. Assume first that there is a vertex $y \in \{u, x_0, x_1, x_2\}$ that buys both its edges in C . Then, by Lemma 7, there is vertex $w \in N_1(y) \subseteq N_4(u) \subseteq N_5(v)$ with $\deg_H(w) \geq 3$. Assume now that there is no vertex $y \in \{u, x_0, x_1, x_2\}$ that buys both its edges in C . But then, as u buys an edge to v , we have a path x_3, x_2, x_1, x_0, u, v of length 5 where one vertex buys the edge to the next one. Thus, by Lemma 8, the vertices of the path cannot have all degree 2 in H , and the lemma follows. \square

Corollary 3. *If G is an equilibrium graph for $\alpha > 13$ then $\deg(H) \geq 2 + \frac{1}{16}$.*

Proof. We assign every vertex $v \in H$ to its closest vertex $c \in H$ with $\deg_H(c) \geq 3$ (thus, c is assigned to itself), breaking ties arbitrarily (by Lemma 9 we know that there is a vertex of degree at least 3 in H). Consider the subgraph of H formed by a vertex c of degree at least 3 and by vertices assigned to it. Observe that these subgraphs form a partition of H . We show that the average degree of every such subgraph is at least $2 + \frac{1}{16}$ which proves the claim. The subgraph consists of $\deg_H(c)$ induced paths $\{p_i(c)\}_{i=1}^{\deg_H(c)}$ that all meet in c . Let $\text{length}(p_i(c))$ denote the length of path $p_i(c)$. By Lemma 9 this length is at most 5. The average degree of the subgraph is then
$$\frac{\deg_H(c) + 2 \sum_{i=1}^{\deg_H(c)} \text{length}(p_i(c))}{1 + \sum_{i=1}^{\deg_H(c)} \text{length}(p_i(c))} = 2 + \frac{\deg_H(c) - 2}{1 + \sum_{i=1}^{\deg_H(c)} \text{length}(p_i(c))} \geq 2 + \frac{\deg_H(c) - 2}{1 + 5 \cdot \deg_H(c)} \geq 2 + \frac{1}{16}. \quad \square$$

Next we prove the last ingredient for our approach – we show an upper bound for $\deg(H)$ involving α :

Lemma 10. *If G is an equilibrium graph for $\alpha > 1$ then $\deg(H) \leq 2 + \frac{8}{\alpha - 1}$.*

Proof. Consider a BFS-tree T of G rooted in a central vertex $v_0 \in V$ and let $\tilde{T} := T \cap H$. Note that \tilde{T} is a spanning tree of H . Then $\deg(H) = \frac{2|E(\tilde{T})| + 2|E(H) \setminus E(\tilde{T})|}{|V(\tilde{T})|} \leq 2 + \frac{2|E(H) \setminus E(\tilde{T})|}{|V(\tilde{T})|}$, and hence we have to bound $|E(H) \setminus E(\tilde{T})|$ (the number of edges outside \tilde{T}). To do that, we consider vertices of H that buy an edge in $E(H) \setminus E(\tilde{T})$. Let us call such a vertex a *shopping vertex*. First observe that every shopping vertex u buys exactly one edge in $E(H) \setminus E(\tilde{T})$, as otherwise u could opt not to buy these edges and buy one edge to v_0 instead, thus saving at

least α on creation cost, and having usage cost at most $D_G(v_0) + 1 \leq D_G(u) + 1$, which (for $\alpha > 1$) would be an improvement, a contradiction. This immediately shows that there are at most $|V(\tilde{T})|$ edges in $E(H) \setminus E(\tilde{T})$. To get a better bound, we bound the number of shopping vertices. We show that the distance in \tilde{T} between any two shopping vertices is at least $\frac{\alpha-1}{2}$. The upper bound on the number of shopping vertices follows: Assign every node v of H to the closest shopping vertex (closest according to the distance in \tilde{T} ; breaking ties arbitrarily); Observe that this assignment forms a partition of H (and that every part contains exactly one shopping vertex); As the distance in \tilde{T} between any two shopping vertices is at least $\frac{\alpha-1}{2}$, the size of every part is at least $\frac{\alpha-1}{4}$. Thus, there are at most $\frac{4|V(\tilde{T})|}{\alpha-1}$ shopping vertices and thus at most that many edges in $E(H) \setminus E(\tilde{T})$; The desired bound $\deg(H) \leq 2 + \frac{8}{\alpha-1}$ now easily follows.

We are left to prove that the distance in \tilde{T} between any two shopping vertices is at least $\frac{\alpha-1}{2}$. Assume for contradiction that there are two shopping vertices $u_1 \neq u_2$ for which $d_{\tilde{T}}(u_1, u_2) < \frac{\alpha-1}{2}$. Let $u_1 = x_1, x_2, \dots, x_k = u_2$ be the shortest u_1 - u_2 -path in \tilde{T} and let us call it P . Let $\{u_1, v_1\}$ and $\{u_2, v_2\}$ be the edges that u_1 and u_2 buy in $E(H) \setminus E(\tilde{T})$. Observe that v_1 and v_2 are not descendant of any vertex $x_i, i = 1, \dots, k$, in P ; If $v_j, j = 1, 2$, is descendant of x_i , then the v_j - x_i -path in \tilde{T} , the x_i - u_j -path in \tilde{T} , and the edge $\{u_j, v_j\}$ form a cycle of length at most $2(d_{\tilde{T}}(u_1, u_2) + 1) < \alpha + 1$ which contradicts Lemma 4. In particular, v_j is not part of P , and therefore $x_0 = v_1, x_1, \dots, x_k, x_{k+1} = v_2$ is a path in H . Also by Lemma 4, $u_j, j = 1, 2$, has distance at least $\frac{\alpha-1}{2}$ from v_0 , and therefore v_0 is not in P . Now, since x_1 buys $\{x_0, x_1\}$ and x_k buys $\{x_k, x_{k+1}\}$, there has to be $1 \leq i^* \leq k$ such that x_{i^*} buys both $\{x_{i^*-1}, x_{i^*}\}$ and $\{x_{i^*}, x_{i^*+1}\}$. Consider the following modification of x_{i^*} 's strategy: Buy edge $\{x_{i^*}, v_0\}$ instead of edges $\{x_{i^*-1}, x_{i^*}\}$ and $\{x_{i^*}, x_{i^*+1}\}$. In this new strategy, x_{i^*} decreases its creation cost by α . We now show that x_{i^*} 's new usage cost is $D_{\text{new}}(x_{i^*}) < D_G(x_{i^*}) + \alpha$ thus implying that the new strategy improves x_{i^*} 's cost, a contradiction.

First note that $D_{\text{new}}(x_{i^*}) \leq 1 + D_{\text{new}}(v_0)$ (where the subscript "new" always corresponds to the situation in a graph where x_{i^*} is using the modified strategy). To bound $D_{\text{new}}(v_0)$ we note that only the vertices in P and their descendants in T can have increased distance to v_0 by the strategy change. Let y be one of these vertices with possibly increased distance and let $1 \leq j \leq k$ be such that x_j is the closest ancestor of y , i.e., an ancestor with $d_G(x_j, y) = \min_{x \in P} d_G(x, y)$. If $j = i^*$ it is easy to see that $d_{\text{new}}(v_0, y) \leq d_G(v_0, y)$ and therefore for such a vertex y there is no increase in usage cost of v_0 . Consider now the case $j \neq i^*$ and assume (without loss of generality, as we shall see) that $j < i^*$. Then $d_{\text{new}}(v_0, y) \leq d_{\text{new}}(v_0, x_0) + d_{\text{new}}(x_0, x_j) + d_{\text{new}}(x_j, y) = d_G(v_0, x_0) + d_G(x_0, x_j) + d_G(x_j, y)$ (since x_0 is not a descendant of a vertex in P and x_0, \dots, x_j is still a path in G_{new}), and $d_G(v_0, y) = d_G(v_0, x_j) + d_G(x_j, y)$. Then the increase of usage cost of v_0 is: $d_{\text{new}}(v_0, y) - d_G(v_0, y) = d_G(v_0, x_0) + d_G(x_0, x_j) - d_G(v_0, x_j) \leq 2 \cdot d_G(x_0, x_j) \leq 2 \cdot d_G(u_1, u_2) \leq 2 \cdot d_{\tilde{T}}(u_1, u_2) < \alpha - 1$, where the last inequality follows from our assumption $d_{\tilde{T}}(u_1, u_2) < \frac{\alpha-1}{2}$. As y was chosen arbitrary,

we have that the increase of usage cost of v_0 is less than $\alpha - 1$ and therefore $D_{\text{new}}(v_0) < D_G(v_0) + \alpha - 1$, which shows $D_{\text{new}}(x_{i^*}) < D_G(x_{i^*}) + \alpha$. \square

Using this result we show that only tree equilibria appear for certain α .

Theorem 5. *For $\alpha > 129$ every equilibrium graph is a tree.*

Proof. If G is a non-tree equilibrium for $\alpha > 129$ and H a block in G with $|H| \geq 3$ then we have by Lemma 10 that $\deg(H) \leq 2 + \frac{8}{\alpha-1} < 2 + \frac{1}{16}$, which contradicts Corollary 3 stating that $\deg(H) \geq 2 + \frac{1}{16}$. \square

This bound is asymptotically tight. Indeed there is a constant $c > 0$ such that for $\alpha < c$ we have non-tree equilibrium graphs. E.g., for $\alpha \leq 1$, the triangle is an equilibrium graph (we can generalize this to any size $n \geq 3$ of vertices: three stars of size $n/3$, where the three centers of the stars are connected in a triangle, form an equilibrium graph, too). Theorems 5 and 4 thus show the following.

Corollary 4. *For $\alpha > 129$ the price of anarchy is smaller than 4.*

3 Bounding the Price of Anarchy in SumGame

In this section we consider SUMGAME. Adapting the methods that we have developed for MAXGAME in Section 2.1 we are able to show that in SUMGAME for $\alpha > 273n$ every equilibrium graph is a tree. This improves the best known bound of $\alpha \geq 12n \log n$ from 11 and is asymptotically the best obtainable bound as for $\alpha < n/2$ there exist non-tree equilibrium graphs 11. As a corollary we obtain constant PoA for $\alpha > 273n$. We omit most of the proofs and refer to 8, the full version of this paper, for missing details. We use the same conventions and notation as in Section 2.1

Similarly to MAXGAME we can show in the following lemmata that in a constant-size neighborhood of every vertex v in a biconnected component H of an equilibrium graph G there is a vertex of degree at least 3 in H . The details of the proofs are for SUMGAME a bit different though.

Lemma 11 (11). *Any equilibrium graph has no cycle of length less than $\frac{\alpha}{n} + 2$.*

Lemma 12. *If G is an equilibrium graph and $u, v \in V(H)$ are two vertices in H with $d(u, v) \geq 3$ such that u buys the edge to its adjacent vertex x in a shortest u - v -path and v buys the edge to its adjacent vertex y in that path then either $\deg_H(x) \geq 3$ or $\deg_H(y) \geq 3$.*

Proof. Assume for contradiction that $\deg_H(x) = 2 = \deg_H(y)$. Assume without loss of generality that $|S(x)| \leq |S(y)|$. Let z be the other vertex in H adjacent to x . Consider a modified strategy of u where u buys an edge to z instead of the edge to x . In this strategy u shortens its distance to the vertices in $S(y)$ and $S(v)$ by at least 1 and increases its distance to the vertices in $S(x)$ by 1. Furthermore it does not increase its distance to any other vertex in the graph. Since $|S(x)| < |S(v) \cup S(y)|$ ($S(v) \neq \emptyset$ by definition), we conclude that u decreases its cost in the modified strategy, a contradiction. \square

The proof of the following lemma is relatively technical (and most different from the techniques used for MAXGAME) and it has been omitted due to space reasons.

Lemma 13. *If G is an equilibrium graph then any path x_0, x_1, \dots, x_k in H , where $\deg_H(x_i) = 2$ for $0 \leq i \leq k$ and x_i buys $\{x_i, x_{i+1}\}$ for $0 \leq i \leq k-1$, has length at most $k \leq 8$.*

Using the previous two lemmas we can show the following.

Lemma 14. *If G is an equilibrium graph for $\alpha > 19n$ then for every vertex v in H there is a vertex $w \in N_{11}(v)$ with $\deg_H(w) \geq 3$.*

Now, quite in the same way as for MAXGAME, we can prove the claims that show the main result of the section.

Corollary 5. *If G is an equilibrium graph for $\alpha > 19n$ then $\deg(H) \geq 2 + \frac{1}{34}$.*

Lemma 15. *If G is an equilibrium graph for $\alpha > n$ then $\deg(H) \leq 2 + \frac{8n}{\alpha-n}$.*

Theorem 6. *For $\alpha > 273n$ every equilibrium graph is a tree.*

Theorem 7 ([4]). *The cost of an equilibrium graph that is a tree is less than 5 times the cost of a social optimum.*

Corollary 6. *For $\alpha > 273n$ the price of anarchy is smaller than 5.*

References

1. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., Roditty, L.: On nash equilibria for a network creation game. In: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 89–98. ACM, New York (2006)
2. Alon, N., Demaine, E.D., Hajiaghayi, M., Leighton, T.: Basic network creation games. In: Proceedings of the 22nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), pp. 106–113. ACM, New York (2010)
3. Demaine, E.D., Hajiaghayi, M., Mahini, H., Zadimoghaddam, M.: The price of anarchy in network creation games. In: Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 292–298. ACM, New York (2007)
4. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C.H., Shenker, S.: On a network creation game. In: Proceedings of the Twenty-Second Annual Symposium on Principles of Distributed Computing (PODC), pp. 347–351. ACM, New York (2003)
5. Jackson, M.O.: Social and Economic Networks. Princeton University Press, Princeton (2008)
6. Jackson, M.O., Wolinski, A.: A strategic model of social and economic networks. *Journal of Economic Theory* 71(1), 44–74 (1996)
7. Lin, H.: On the price of anarchy of a network creation game (December 2003), final class-project
8. Mihalák, M., Schlegel, J.C.: The price of anarchy in network creation games is (mostly) constant. Tech. Rep. 675, Institute of Theoretical Computer Science, ETH Zurich (July 2010), <http://www.inf.ethz.ch/research/disstechreps/techreports/>

Truthful Fair Division

Elchanan Mossel^{1,2,*} and Omer Tamuz^{2,**}

¹ U.C. Berkeley

² Weizmann Institute

Abstract. We address the problem of fair division, or cake cutting, with the goal of finding truthful mechanisms. In the case of a general measure space (“cake”) and non-atomic, additive individual preference measures - or utilities - we show that there exists a truthful “mechanism” which ensures that each of the k players gets at least $1/k$ of the cake. This mechanism also minimizes risk for truthful players. Furthermore, in the case where there exist at least two different measures we present a different truthful mechanism which ensures that each of the players gets more than $1/k$ of the cake.

We then turn our attention to partitions of indivisible goods with bounded utilities and a large number of goods. Here we provide similar mechanisms, but with slightly weaker guarantees. These guarantees converge to those obtained in the non-atomic case as the number of goods goes to infinity.

Keywords: Fair division, cake cutting, truthful mechanisms.

1 Introduction

The basic setting of fair partition problems includes a “cake” - some divisible resource - and a number of players, each with different preferences with regards to the different “pieces” of the cake. The task is to divide the cake among the players in a way that would be “fair” according to some fairness notion. This applies to numerous situations, from divorce trials to nuclear arms reduction, see e.g. [1].

This problem has been studied extensively both in the non-atomic case and in the case of indivisible goods. In this paper we take a fresh look at fair division, by approaching it from the mechanism design point of view, and in particular in search of truthful partition mechanisms.

1.1 Background

The problem of fair division, or cake cutting, is a central and classical problem in economics. A *fair* partition is one where each of k players receives at least

* Supported by NSF Career Award (DMS 054829) by ISF grant 1300/08 and by a Minerva grant.

** Supported by ISF grant 1300/08.

$1/k$ of the cake, each according to her own measure (this is also known as a *proportional* partition). Fair partition into two parts has long been known to be possible using a “cut-and-choose” procedure: a cake is guaranteed to be fairly divided in two if one player cuts it and the other picks a piece.

Fair partition to more than two players requires non-trivial mathematics. The founding work in this field was done by Steinhaus, Knaster and Banach in the forties (e.g. [7], [12], [13]). They and others (e.g. [10]) proved existence theorems for fair division on general measure spaces and non-atomic measures, where “fairness” was again, in general, taken to mean that each of k players receives at least $1/k$ of the cake, each according to her own measure. These proofs are not constructive, but nevertheless were useful in generalizing “cut-and-choose” to more than two players.

Another natural approach is “moving knife algorithms”, first described by Dubins and Spanier [4]. These are not algorithms in the modern Turing Machine sense of the word (as “cut-and-choose” and its generalizations aren’t), but still provide a “practical” way to cut actual cakes, if nothing else, by the same fairness criterion mentioned above.

A stronger concept of fairness is that of “envy-free” partitions ([6], [14]). In an envy free partition there is no player who would trade her allocation with one given to another player. Such partitions have been studied extensively (e.g., see [1]).

Partitioning indivisible goods is a more recently studied variant, which places the problem in a standard algorithmic setup. However, it is easy to see that fairness or envy-freeness could not be achieved in every setup as the example of one good demonstrates. Still, the results of Lipton et al. [8] show that almost envy free partitions exist. These are partitions where no player envies another player by more than the value of a single good.

Classical cake cutting mechanisms such as “cut-and-choose” usually require the players to choose a piece of a subset of the cake or make a cut in it, according to their preferences. Alternatively, the mechanism queries the player about her valuation of a particular piece of the cake. We take an approach which is more prevalent in the modern mechanism design world: each player declares their preference (i.e., their entire measure on the cake) to some “third party”, which then proceeds to divide the cake according to a predetermined algorithm.

While circulating a draft of this paper, it was brought to our attention that similar questions are discussed in a recent working paper by Chen et al. [3]. They restrict themselves to a particular class of measures: “the case where the agents hold piecewise uniform valuation functions, that is, each agent is interested in a collection of subintervals of $[0, 1]$ with the same marginal value for each fractional piece in each subinterval.” In this setting, they show a truthful, deterministic, polynomial-time, fair and envy-free mechanism. For the more general setting which we study they independently derive our mechanism [1].

Truthfulness. The notion of truthfulness is very natural in partitioning problems. Why would a player declare her true measure if declaring a different measure would result in a better partition for her?

The “cut-and-choose” method is in some sense truthful: the players don’t have to trust each other to be guaranteed a half of the cake. However, it is *not* truthful in the sense that players have incentive to “strategize” in order to increase the value of the piece they receive. For example, assume that the players’ preferences differ, and that the first player (the “cutter”) knows, or guesses, the second player’s (the “chooser”) preference. She may then cut the cake into two pieces such that the first piece is worth much more than one half to herself, and $\frac{1}{2} - \epsilon$ to the chooser. This would, perhaps, secure her a large piece of the cake, and leave the chooser with the feeling that dealings were not completely *fair*. To further confound matters we note that the chooser may realize, in this setting, that she could gain a larger piece by manipulating the cutter’s perception of her preferences.

As noted in [1], many of the partition mechanisms which were discovered in the non-atomic case have this same weaker truthfulness property; for fair partitions - they have the property that truthful players are guaranteed to receive at least $1/k$ of their total value of the cake. Similarly for envy free partitions every truthful player is guaranteed not to envy any of the other players.

Still the question remains as to why is it beneficial for a player to declare her true value?

A work which addresses this question is a paper by Lipton et al. [8] who analyze a truthful mechanism for allocating a set of indivisible goods: simply give each good to each of k players with probability $1/k$. This mechanism is further analyzed in [2], who showed that this partition results, with high probability, in $O(\alpha\sqrt{n \ln k})$ envy, where n is the number of goods, k is the number of players, and α is the maximum utility over all goods and players.

The strongest possible notion of truthfulness is the following: a mechanism is *truthful* if it is always the case that a player’s utility when declaring the truth is as high as her utility when lying. The notion of [8,2] is weaker. They focus, as we do, on *truthfulness in expectation*: the expected value of a player’s utility is maximal when telling the truth.

1.2 Our Results

Non Atomic Measures. Although our “third party” approach is one of modern mechanism design, our work in the non-atomic setup lies in the realm of classical fair division, in that it does not consider the computational aspects of the mechanisms: in some cases, an infinite amount of information would have to be conveyed to the third party. Even when not, the calculations required may be intractable or even not recursive.

We first consider the problem of a truthful fair partition with non-atomic measures. For this problem, using a result of Steinhaus [13], we show that such a partition always exists. Furthermore, our mechanism particularly incentivizes risk averse players to play truthfully.

We next ask if it is possible to devise a truthful mechanism that guarantees that each player gets more than $1/k$ of the cake. Obviously the answer is negative if all the measures are the same. However, as mentioned by Steinhaus [12] and

proved by Dubins and Spanier [4], if the measures are not all identical, then there exists a partition where each player gets more than $1/k$ of the entire cake.

Our main result is a *truthful in expectation* mechanism which guarantees that each player gets at least $1/k$ of the cake, and furthermore the expected size of the piece each player gets is strictly larger than $1/k$.

We further show that there exists no deterministic and truthful mechanism that gives these guarantees and so only in the weaker notion of truthfulness it is possible to obtain such results.

An additional argument in favor of randomized mechanisms is that a deterministic mechanism cannot be symmetric: when all the players declare the same preferences, it must arbitrarily break the symmetry and assign different players different slices of the cake. A randomized algorithm can avoid this.

Indivisible Goods. Our results presented in the case of non-atomic measures are existential; we do not provide protocols for implementing them. To address the computational aspect of the problem we consider partitions of a large number of indivisible goods where the number of goods n is exponential in the number of players, and the utility of a single good is bounded. We give efficient versions of all of our mechanisms in this setup, with guarantees that are slightly weaker than those provided in the continuous case. The guarantees hold in expectation, and moreover there is a deterministic guarantee for each player to receive a share of the goods which is at least $(1 - \epsilon)$ of the expected value, where $\epsilon = O(M^k/n)$ where k is the number of players, n is the number of goods and M is the maximum utility of a single good.

More generally, we prove a discrete analogue to a theorem of Dubins and Spanier [4]. They show that the space of partition utilities is convex. We show that the same is true in the discrete case, again up to a factor of $(1 - \epsilon)$, as defined above.

2 Continuous Truthful Mechanisms

2.1 Existence Theorems for Fair Divisions

Dubins and Spanier [4], rephrasing Fisher [5], provide the following description of what they call “The problem of the Nile”:

“Each year the Nile would flood, thereby irrigating or perhaps devastating parts of the agricultural land of a predynastic Egyptian village. The value of different portions of the land would depend upon the height of the flood. In question was the possibility of giving to each of the k residents a piece of land whose value would be $1/k$ of the total land value no matter what the height of the flood.”

Neyman [10] showed that this is possible, given that there are a finite number of levels that the Nile can rise to.

Let \mathcal{C} be a “cake” (a set), and \mathfrak{C} a set of “slices” (a σ -algebra of subsets of \mathcal{C}). Let there be k players, and let μ_1, \dots, μ_k be non-atomic probability (additive) measures on $(\mathcal{C}, \mathfrak{C})$, so that the value of a slice $C \in \mathfrak{C}$ to player i is $\mu_i(C)$. Then

Neyman's theorem establishes that there exists a partition of the cake C_1, \dots, C_k such that for all players i and slices j it holds that $\mu_i(C_j) = 1/k$. Hence all the slices are equal, by all the player's measures.

Dubins and Spanier [4] show that a better partition is always possible when at least two of the players have different measures. Their theorem implies that in this case a partition is possible for which, for all players and slices i , it holds that $\mu_i(C_i) > 1/k$, and so each player gets strictly more than $1/k$ of the cake, by his or her own measure.

2.2 Truthful Mechanisms

Fair division We present a simple truthful “mechanism” for distributing the cake among k players, which assures that each player gets precisely $1/k$ of the cake, by all the players' measures. It is a “mechanism” in quotes since it is as constructive as Neyman's theorem, which is not constructive. Note that this mechanism also appears in [3].

Mechanism 1. *Assume the players' true measures are μ_1, \dots, μ_k , and that they each declare some measure ν_i . Find a partition C_1, \dots, C_k such that $\forall i, j : \nu_i(C_j) = 1/k$. Then choose a random permutation τ of size k , from the uniform distribution, and give $C_{\tau(i)}$ to player i .*

Proposition 1. *Mechanism [1] is truthful in the following sense: No player can increase her expected utility by playing non-truthfully. Further, a player who plays truthfully minimizes the risk/variance of the measure of the piece she gets.*

Proof. The expected size of the slice for player i is $\sum_j \mu_i(C_j) \mathbf{P}[\tau(i) = j] = \sum_j \mu_i(C_j)/k = \mu_i(\cup C_j)/k = 1/k$. Since it is independent of ν_i then player i has no incentive to declare untruthfully. Furthermore, a player that declares $\nu_i = \mu_i$ is guaranteed a slice of size $1/k$, and so the truth minimizes the variance (or risk), to zero.

Super-Fair Division. For this result we set $\mathcal{C} = [0, 1) \in \mathbb{R}$ and let \mathfrak{C} be the Borel σ -algebra. While this result can be extended to more general classes of spaces and algebras, we present it in this restricted form for clarity. We consider the case where at least one pair of measures are not equal, i.e. the case in which “super-fair” partitions exist — partitions in which $\mu_i(C_i) > 1/k$.

We first provide motivation for our usage of “truthfulness in expectation”, by showing that deterministic “super-fair” mechanisms cannot be truthful:

Theorem 2. *Any deterministic mechanism that gives each player $1/k$ of the cake when all declared measures are equal and more than $1/k$ of the cake otherwise is not truthful.*

Proof. Consider the case where all players declare the same arbitrary measure μ . Then they receive slices C_i such that $\forall i : \mu(C_i) = 1/k$. Now, consider the case where player 1's true measure ν is such that $\nu(C_1) = 1$. Then player 1's

utility for declaring μ is 1. We propose that her utility for declaring ν (i.e., being truthful) is less than one, and therefore the mechanism is not truthful.

Assume by way of contradiction that her utility for declaring ν is one. Then in this case she must also receive slice C_1 , since that is the only slice worth one to her. But if player 1 receives C_1 then the rest of the players have, by their measure μ , exactly $(k-1)/k$ of the cake left to share, and so it is impossible that they all receive more than $1/k$ of it. This contradicts the hypothesis, since $\mu(C_1) = 1/k$ and $\nu(C_1) = 1$, and therefore $\mu \neq \nu$ and all players must receive more than $1/k$.

We now describe a randomized mechanism that is “super-fair” and truthful in expectation.

Mechanism 3. *Assume again that μ_1, \dots, μ_k are the players’ true measures, and that they each declare some measure ν_i . To distribute the cake, pick a partition C_1, \dots, C_k from a distribution D over partitions, which we describe below. If it so happens that $\nu_i(C_i) > 1/k$ for all i , then distribute the slices accordingly. Otherwise distribute by mechanism 1, that is, give a slice of value $1/k$ to all players.*

Proposition 2. *In mechanism 3 the expected utility of a truthful player is larger than $1/k$ if super-fair partitions are picked with positive probability.*

Proof. If C_1, \dots, C_k is super-fair and the players are truthful, then this partition is recognized as super-fair, and the players each get strictly more than $1/k$. In the event that the picked partition is not super-fair, and the players are truthful, then the mechanism reverts to giving the players precisely $1/k$ of the cake. Thus the expectation for truthful players is more than $1/k$.

Proposition 3. *In mechanism 3 playing truthfully maximizes a player’s expected utility.*

Proof. Consider again two cases: the first, in which C_1, \dots, C_k is super-fair, and the second, in which it isn’t.

In the first event, a truthful player’s expected share is more than $1/k$. Playing untruthfully could either have no effect, leaving the utility as is, or else the only other possibility is that the partition is misconstrued not to be super-fair, in which case the player’s utility is reduced to $1/k$.

In the second event, in which the partition picked is not super-fair, playing untruthfully may again have no effect, leaving the utility at $1/k$. However, if, to some player, the share allocated by this partition was worth less than $1/k$, playing untruthfully may make it seem to be valued at more than $1/k$, turning the partition into super-fair by the declared preferences, and resulting in a utility less than $1/k$ for that player.

Thus, for any random choice of C_1, \dots, C_k the truthful player’s expected utility is maximal, and the proposition follows.

To assure that this mechanism results in a slice of expected size strictly greater than $1/k$, we must find a distribution D (from which we draw the partition) such

that for any set of measures, where at least one pair is not equal, with positive probability $\mu_i(C_i) > 1/k$ for all i . To this end we make the following definition:

Definition 1. Denote by \mathcal{Q} the set of partitions C_1, \dots, C_k of $[0, 1] \in \mathbb{R}$ for which each C_i is a finite union of half-open intervals with rational endpoints.

We note that \mathcal{Q} is countable. D now need only be some distribution with support \mathcal{Q} :

Theorem 4. Let μ_1, \dots, μ_k be non-atomic probability measures on $[0, 1] \in \mathbb{R}$ and the Borel σ -algebra, such that there exist i, j for which $\mu_i \neq \mu_j$. Let D be a distribution over the partitions C_1, \dots, C_k of $[0, 1]$ into k sets, such that the support of D is \mathcal{Q} . Then

$$\mathbf{P}_D[\forall i : \mu_i(C_i) > 1/k] > 0.$$

The proof appears in Appendix [A](#).

3 Indivisible Goods

Let $\mathcal{C} = \{a_1, \dots, a_n\}$ be a finite set of indivisible goods (“discrete cake”). Let there be k players, where each has an additive bounded measure (utility) on the algebra of subsets of \mathcal{C} , μ_i , such that for all i, j it holds that $\mu_i(\{a_j\}) \in \{1, 2, \dots, M\}$.

We focus on the regime where the number of players is small, so that $n \gg M^k$, and in particular $Mk \cdot M^k/n < \epsilon$ for some ϵ . Then it also holds that $Mk \cdot M^k/\mu_i(\mathcal{C}) < \epsilon$.

3.1 Truthful Fair Division

Let ν_1, \dots, ν_k be the set of declared measures.

Definition 2. Let $S = (s_1, \dots, s_k) \in \{1, 2, \dots, M\}^k$ be some vector. Let the bin $B_S \subseteq \mathcal{C}$ be the set of goods a for which, for each player i , it holds that $a \in B_S$ iff $\nu_i(a) = s_i$:

$$B_S = \{a \in \mathcal{C} \text{ s.t. } \forall i : \nu_i(a) = s_i\} \tag{1}$$

Let \mathcal{B} be the set of bins.

We propose the following mechanism:

Mechanism 5. For each bin B_S , pick from the uniform distribution a partition of it into k parts of equal size B_S^1, \dots, B_S^k , with perhaps some left over elements which number at most $k - 1$. Let $C'_i = \bigcup_{B_S \in \mathcal{B}} B_S^i$ and give C'_i to player i . Then, give each leftover good to some player, uniformly at random.

Denote by C_i the set that player i got, i.e. C'_i union any leftovers given to player i . Then it is easy to see that this mechanism is truthful, since player i 's expectation is $\mu_i(\mathcal{C})/k$, independently of her declared measure ν_i :

Proposition 4.

$$\mathbf{E}[\mu_i(C_j)] = \mu_i(\mathcal{C})/k$$

Proof. This follows from the fact that every a_l ends up in C_j with probability $1/k$.

Truthfulness, however, could have been more simply achieved by, for example, giving each player the entire set \mathcal{C} with probability $1/k$. This mechanism’s merit is that it ensures low risk for truthful players:

Theorem 6. *When $\nu_i = \mu_i$ then for all j it holds that $\mu_i(C_j) \geq (1 - \epsilon)\mu_i(\mathcal{C})/k$ when $i = j$ this implies low risk for truthful players.*

Proof. By definition of the C_j ’s

$$\mu_i(C_j) \geq \mu_i(C'_j) = \mu_i\left(\bigcup_{B_S \in \mathcal{B}} B_S^j\right).$$

Since the different B_S^i ’s are disjoint, and by the definition of B_S

$$\mu_i(C_j) \geq \sum_{B_S \in \mathcal{B}} \mu_i(B_S^j) \geq \sum_{B_S \in \mathcal{B}} s_i |B_S^j|.$$

We denote the number of left over elements r_S , so that $|B_S| = k|B_S^i| + r_S$ for all i . Then

$$\mu_i(C_j) \geq \frac{\mu_i(\mathcal{C})}{k} - \frac{1}{k} \sum_{B_S \in \mathcal{B}} s_i r_S,$$

since $r_S < k$ and $s_i \leq M$, and by the definition of ϵ we finally have that

$$\mu_i(C_j) \geq \frac{\mu_i(\mathcal{C})}{k} - \frac{Mk \cdot M^k}{k} \geq (1 - \epsilon)\mu_i(\mathcal{C})/k. \tag{2}$$

We conclude that assuming players are averse to risk, they may find actual advantage in playing truthfully, since that will result in a utility that is with probability one greater than $(1 - \epsilon)\mu_i(\mathcal{C})/k$. Other strategies, on the other hand, may run the risk of resulting in lower utility.

3.2 Truthful Super-Fair Division

We can naturally adapt mechanism [3](#) to the discrete case, by letting D be the uniform (for example) distribution over the partitions of $\mathcal{C} = \{a_1, \dots, a_n\}$ into k subsets. We then use what is essentially the same mechanism:

Mechanism 7. *Pick a random partition from D , keep it iff everyone was allocated strictly more than $1/k$, and otherwise give everyone $1/k$ using mechanism [5](#).*

In this discrete case it is easy to see that if super-fair divisions exist then they are picked with positive probability. The proofs that this mechanism results in expected utility larger than $1/k$, and that it is truthful, are identical to the ones for the continuous case, [2](#) and [3](#).

3.3 Extending Continuous Fair Division Existence Results

Let \mathcal{M} be the space of k -by- n matrices M such that for measures μ_1, \dots, μ_k and some division C_1, \dots, C_n it holds that $M_{ij} = \mu_i(C_j)$. Dubins and Spanier [4], using a theorem of Lyapunov [9], prove that \mathcal{M} is compact and convex when the measures μ_i are non-atomic. From this follow a plethora of existence theorems for partitions of different characteristics. For example, as mentioned above, this fact can be used to show that there exists a division where each of the k players gets a share worth $1/k$ by everyone’s measure (for probability measures). It can also be used to show that some players have different measures then a division exists in which every player gets more than $1/k$, by her own measure.

This result obviously does not apply to the discrete case; the set of partitions is finite and it is difficult to speak of convexity. Accordingly, in the general case no fair partition exists, and a super-fair partition may not exist even when the preferences are different. A simple example of two goods and three players suffices to illustrate this point.

One could, however, imagine that all this *could* be achieved if the players were somehow able to share the goods. In fact, if we allow, for example, that one player has a third of a good and another two thirds of it (with appropriate utilities), then Dubins and Spanier’s results apply again, and a wealth of partitions with different qualities is possible again. We refer to this as the continuous extension of the discrete problem.

However, indivisible goods must by nature remain indivisible. To overcome this, we propose a randomized partition, similar to the one used in mechanism [5], that makes possible, *in expectation*, any partition values possible in the continuous extension. Moreover, it assures that each player not only receives the correct utility in expectation, but that in the worst case she will not receive less than $1 - \epsilon$ of what she expects.

We thus consider again a set of indivisible goods $\mathcal{C} = \{a_1, \dots, a_l\}$, k players, and their additive measures μ_i , where $\mu_i(\{a_j\}) \in \{1, 2, \dots, M\}$. We now imagine that each good can be continuously subdivided, and so define a_j^* to be copy of $[0, 1] \in \mathbb{R}$, and let $\mathcal{C}^* = \{a_1^*, \dots, a_l^*\}$. Let \mathfrak{F} be the standard Borel σ -algebra on \mathcal{C}^* , and let ν be the Lebesgue measure on \mathfrak{F} . Define μ_i^* , a measure on \mathfrak{F} , as a continuous extension of μ_i by

$$\mu_i^*(A) = \sum_j \mu_i(\{a_j\}) \cdot \nu(A \cap a_j^*),$$

for any $A \in \mathfrak{F}$. We refer to \mathcal{C}^* and μ_i^* as the *continuous extension* of \mathcal{C} and μ_i .

We are now ready to state the main result of this section:

Theorem 8. *Consider the problem of partitioning indivisible goods as defined above, and its continuous extension. Let \mathcal{M} be the space of k -by- n matrices M such that for some C_1^*, \dots, C_n^* it holds that $M_{ij} = \mu_i^*(C_j^*)$.*

Then for every element $M \in \mathcal{M}$ there exists a randomized partition C_1, \dots, C_n satisfying $\mathbf{E}[\mu_i(C_j)] = M_{ij}$, and moreover

$$\frac{\mu_i^*(C_j^*)}{\mu_i^*(\mathcal{C}^*)} + k\epsilon \geq \frac{\mu_i(C_j)}{\mu_i^*(\mathcal{C}^*)} \geq \frac{\mu_i^*(C_j^*)}{\mu_i^*(\mathcal{C}^*)} - \epsilon,$$

where ϵ , as before, is $O(M^k/n)$.

Proof. Given a division C_1^*, \dots, C_n^* of the divisible \mathcal{C}^* , we would like to divide the discrete \mathcal{C} in a way that approximates this division as closely as possible. That is, we would like to find a division C_1, \dots, C_n such that $\mu_i(C_j) \approx \mu_i^*(C_j^*)$. We propose two schemes to do this: the random scheme and the binned scheme. For both of them, we define an n -by- l matrix D (l being the number of indivisible goods), where D_{ij} is the fraction of good a_j^* that belongs to C_i^* : $D_{ij} = \nu(C_i^* \cap a_j^*)$.

The random scheme. In the random scheme, we simply give player i good a_j with probability D_{ij} (note that by the definition of D , $D_{i\cdot}$ is a distribution). Then

$$\mathbf{E}[\mu_i(C_j)] = \sum_m \mu_i(\{a_m\}) \cdot \mathbf{P}[a_m \in C_j] = \sum_m \mu_i(\{a_m\}) \cdot D_{jm} = \mu_i^*(C_j^*)$$

and its standard deviation is $O\left(\sqrt{\mu_i^*(C_j^*)}\right)$.

The binned scheme. In the binned scheme, we bin the elements of \mathcal{C} into bins $\{B_S\}$ as above. Without loss of generality, let there be, for each player i and bin B_S , a single value D_{iS} such that $D_{im} = D_{iS}$ for all m . No generality is indeed lost: because all elements of a bin are equivalent to all the players, then for any partition C_j^* there exists an equivalent partition, in the sense of the utilities of the players, for which such D_{iS} 's exist.

Let n_S be the number of elements in bin B_S . From each bin B_S , we give player i a number of goods equal to $\lfloor n_S D_{iS} \rfloor$, picked from the uniform distribution over such partitions. Any leftover a_m we give according to the random scheme, i.e. to player i with probability D_{iS} .

The expectation for $\mu_i(C_j)$ clearly remains $\mu_i^*(C_j^*)$. However, here we can bound it from below:

$$\frac{\mu_i(C_j)}{\mu_i^*(\mathcal{C}^*)} \geq \frac{1}{\mu_i^*(\mathcal{C}^*)} \sum_{B_S \in \mathcal{B}} (n_S D_{iS} - 1) s_i \geq \frac{\mu_i^*(C_j^*)}{\mu_i^*(\mathcal{C}^*)} - \epsilon.$$

We can also bound it from above, by noting that the most a player can get beyond $\mu_i(C_j)$ is what's lost by the rest of the players:

$$\frac{\mu_i^*(C_j^*)}{\mu_i^*(\mathcal{C}^*)} + k\epsilon \geq \frac{\mu_i(C_j)}{\mu_i^*(\mathcal{C}^*)}$$

References

1. Brams, S., Taylor, A.: Fair Division: from cake cutting to dispute resolution. Cambridge University Press, Cambridge (1996)
2. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: On low-envy truthful allocations. In: ADT, pp. 111–119 (2009)
3. Chen, Y., Lai, J.K., Parkes, D.C., Procaccia, A.D.: Truth, justice and cake cutting (2010), working paper (to appear in AAAI 2010)
4. Dubins, L.E., Spanier, E.H.: How to cut a cake fairly. The American Mathematical Monthly 68, 1–17 (1961)
5. Fisher, R.: Quelques remarques sur l'estimation en statistique. Byotypologie 6, 153–159 (1938)
6. Foley, D.: Resource allocation and the public sector. Yale Economics Essays 7, 45–98 (1967)
7. Knaster, B.: Sur le problème du partage pragmatique de h. steinhaus. Ann. Soc. Polonaise Math. 19, 228–231 (1944)
8. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: EC 2004: Proceedings of the 5th ACM Conference on Electronic Commerce, pp. 125–131. ACM, New York (2004)
9. Lyapunov, A.: Sur les fonctions-vecteurs complètement additives. Bulletin of the Academy of Sciences of the USSR 4, 465–478 (1940)
10. Neyman, J.: Un théorème d'existence. C.R. Acad. Sci. Paris 222, 843–845 (1946)
11. Rudin, W.: Real and Complex Analysis. McGraw-Hill, New York (1987)
12. Steinhaus, H.: The problem of fair division. Econometrica 16, 101–104 (1948)
13. Steinhaus, H.: Sur la division pragmatique. Econometrica 17(Suppl.), 315–319 (1949)
14. Varian, H.: Equity, envy and efficiency. Journal of Economic Theory 9, 63–91 (1974)

A Existence of a Distribution on Partitions with Positive Probability for Super-Fair Division

Let \mathcal{R} be the set of finite unions of half-open intervals of $[0, 1) \in \mathbb{R}$ with rational endpoints. Let \mathcal{Q} be defined to be the class of partitions of $[0, 1) \in \mathbb{R}$ such that each part of the partition is in \mathcal{R} .

The main lemma we want to prove is the following:

Lemma 1. *Let μ_1, \dots, μ_k be non-atomic Borel measures on $[0, 1) \in \mathbb{R}$ and the Borel σ -algebra \mathcal{B} . Let B_1, \dots, B_k be a partition of the interval, where $B_i \in \mathcal{B}$ and $\delta > 0$. Then there exists a partition Q_1, \dots, Q_k in \mathcal{Q} such that $\mu_i(B_i \Delta Q_i) < \delta$ for all i .*

Proof. The proof uses the fact that \mathcal{Q} is an algebra, i.e., it is closed under finite unions, intersections and taking of complements. Let δ' be chosen later. Since all of the measures are Borel we can find open sets $O_{i,j}$ so that $B_j \subset O_{i,j}$ and $\mu_i(O_{i,j} \Delta B_j) < \delta'$ (see, e.g., theorem 2.17 in [11]). Taking $O_j = \bigcap_i O_{i,j}$, we get open sets such that $B_j \subset O_j$ and $\mu_i(B_j \Delta O_j) < \delta'$ for all i and j .

Fix j and note that $O_j = \bigcup_{n=1}^{\infty} I_{j,n}$, where $I_{j,n}$ are open intervals with rational end-points. Take m sufficiently large so that $\mu_i(O_j \setminus \bigcup_{n=1}^m I_{j,n}) < \delta'$ for all i . Let

$J_{j,n}$ be the same as $I_{j,n}$ except that the left-end point of the interval is added and let $\tilde{P}_j = \cup_{n=1}^m J_{j,n}$. Since the measures are non-atomic we have $\mu_i(\cup_{n=1}^m J_{j,n}) > \mu_i(O_j) - \delta'$. Note that the \tilde{P}_j 's are all unions of half-open intervals with rational end-points. Moreover, for all i and j ,

$$\mu_i(\tilde{P}_j \triangle B_j) \leq \mu_i(\tilde{P}_j \triangle O_j) + \mu_i(O_j \triangle B_j) \leq 2\delta'.$$

Note further that for all i it holds that:

$$\mu_i[0, 1] \geq \mu_i(\cup_j \tilde{P}_j) \geq \mu_i(\cup_j O_j) - \sum_j \mu_i(O_j \setminus \cup_{n=1}^m I_{j,n}) \geq \mu_i[0, 1] - k\delta',$$

so

$$\mu_i([0, 1] \setminus \cup_j \tilde{P}_j) \leq k\delta'.$$

Now take $P_i = \tilde{P}_i$ for $i > 1$ and $P_1 = \tilde{P}_1 \cup ([0, 1] \setminus \cup_j \tilde{P}_j)$. Then $\cup P_i = [0, 1]$ and

$$\mu_i(P_j \triangle B_j) \leq (k + 2)\delta'$$

for all i and j . The P_i are almost the desired partition. They satisfy all the needed properties except that they are not a partition. We now take $Q_j = P_j \setminus \cup_{j' < j} P_{j'}$. Q_j is obviously a partition. Moreover:

$$\begin{aligned} \mu_i(Q_j \triangle B_j) &\leq \mu_i(P_j \triangle B_j) + \sum_{j' \neq j} \mu_i(P_{j'} \triangle B_j) \\ &\leq \mu_i(P_j \triangle B_j) + \sum_{j' \neq j} \mu_i(B_{j'} \triangle B_j) + \sum_{j' \neq j} \mu_i(P_{j'} \triangle B_{j'}) \\ &\leq 2k(k + 2)\delta'. \end{aligned}$$

Taking $\delta' = \delta / (2k(k + 2))$ concludes the proof.

Theorem 9. *Let μ_1, \dots, μ_k be non-atomic probability measures on $[0, 1] \in \mathbb{R}$ and the Borel σ -algebra, such that there exist i, j for which $\mu_i \neq \mu_j$. Let D be a distribution over the partitions C_1, \dots, C_k of $[0, 1]$ into k sets, such that the support of D is \mathcal{Q} . Then*

$$\mathbf{P}_D[\forall i : \mu_i(C_i) > 1/k] > 0.$$

Proof. By Dubins and Spanier's theorem, there exists a partition B_1, \dots, B_k of measurable sets such that for all i it holds that $\mu_i(B_i) > 1/k$. Let $\epsilon > 0$ be such that $\mu_i(B_i) > 1/k + \epsilon$.

By the lemma above there exists a partition Q_1, \dots, Q_k in \mathcal{Q} such that $\forall i, j : \mu_i(Q_j \triangle B_j) < \frac{1}{2}\epsilon$. This, in particular, implies for all i that $\mu_i(Q_i) > 1/k + \frac{1}{2}\epsilon$.

No Regret Learning in Oligopolies: Cournot vs. Bertrand

Uri Nadav¹ and Georgios Piliouras^{2,*}

¹ Department of Computer Science, Stanford University, Stanford CA 94305
urinadav@cs.stanford.edu

² Department of Computer Science, Cornell University, Ithaca NY 14853
gpil@cs.cornell.edu

Abstract. Cournot and Bertrand oligopolies constitute the two most prevalent models of firm competition. The analysis of Nash equilibria in each model reveals a unique prediction about the stable state of the system. Quite alarmingly, despite the similarities of the two models, their projections expose a stark dichotomy. Under the Cournot model, where firms compete by strategically managing their output quantity, firms enjoy positive profits as the resulting market prices exceed that of the marginal costs. On the contrary, the Bertrand model, in which firms compete on price, predicts that a duopoly is enough to push prices down to the marginal cost level. This suggestion that duopoly will result in perfect competition, is commonly referred to in the economics literature as the “Bertrand paradox”.

In this paper, we move away from the safe haven of Nash equilibria as we analyze these models in disequilibrium under minimal behavioral hypotheses. Specifically, we assume that firms adapt their strategies over time, so that in hindsight their average payoffs are not exceeded by any single deviating strategy. Given this no-regret guarantee, we show that in the case of Cournot oligopolies, the unique Nash equilibrium fully captures the emergent behavior. Notably, we prove that under natural assumptions the daily market characteristics converge to the unique Nash. In contrast, in the case of Bertrand oligopolies, a wide range of positive average payoff profiles can be sustained. Hence, under the assumption that firms have no-regret the Bertrand paradox is resolved and both models arrive to the same conclusion that increased competition is necessary in order to achieve perfect pricing.

1 Introduction

Oligopoly theory deals with the fundamental economic problem of competition between two or more firms. In this work we study the conditions under which an oligopoly arrives at stability. We focus on the two most notable models in oligopoly theory: Cournot oligopoly [7], and Bertrand oligopoly [5]. In the Cournot

* Supported by NSF grant CCF-0325453.

model, firms control their production level, which influences the *market price*. In the Bertrand model, firms choose the price to charge for a unit of product, which affects the *market demand*.

Competition among firms in an oligopolistic market is a setting of strategic interaction, and is therefore analyzed within a game theoretic framework. Cournot and Bertrand oligopolies are modeled as strategic games, with continuous action sets (either production levels or prices). In both models the revenues of a firm are the product of the firm's part of the market times the price; In addition, a firm incurs a production cost, which depends on its production level.

In the most simple oligopoly model, the firms play a single game, where they all take actions simultaneously. All the firms produce the same good; the demand for this product is a linear in the total production; the cost of production is fixed per unit of production. In this oligopolistic market, a Nash equilibrium in pure strategies exists in both Cournot and Bertrand models. Interestingly, despite the strong similarity between these models, the Nash equilibrium points are very different: in Bertrand oligopoly, Nash equilibrium drives prices to their competitive levels, that is, the price equals the cost of production, while in Cournot oligopoly, the price in the unique Nash equilibrium is strictly above its competitive level. Liu [14] showed that the uniqueness of equilibrium in the linear demand, linear cost model, carries on to correlated equilibrium. Yi [23] have extended Liu's work to the case of Cournot oligopoly where firms produce different products, that are strategic substitutes, and to the case of weakly convex production cost functions.

Equilibrium analysis alone, however, cannot capture the dynamic nature of markets. In the real world, trading is performed over long periods of time, which gives firms the chance to adjust their actions e.g, their prices or production levels. If we assume that the essential market attributes remain unchanged, then this situation gives rise to a repeated game, obtained by repeated play of the original simultaneous, one shot game.

One approach for analyzing the repeated oligopolistic game, is through studying the Nash equilibrium of the repeated game. This models a situation where the firms "commit" to a strategy, and their joint commitment forms an equilibrium (see [15], Chapter 12.D). In practice however, an important feature of an oligopolistic market is that different firms are not perfectly informed about different aspects of the market, e.g., the attributes of the other participants, and cannot pre-compute, or agree on a Nash equilibrium of the repeated game before they begin interacting.

A more pragmatic approach for studying such repeated interactions is through the analysis of adaptive behavior dynamics (see, [12],[24]). The goal here is to investigate the evolution of the repeated game, when the agents (firms) play in accordance to some "natural" rule of behavior. In the setting of an oligopolistic market, we would want a natural behavior to comply with "rationality" and hopefully give rise to some sort of profit maximization on the side of a firm. Another natural aspiration is that our behavior rules should be "distributed", which means that firms should be able to make their choices in each period based

only on their own payoffs, and independently of other firms (in most markets, a firm cannot tell with certainty what are the payoffs, and costs of other firms). The central question, in such a setting, is whether the behavior dynamics finally converge, as this would imply long term stability of the market.

Dynamic behavior in Cournot and Bertrand oligopolies have been studied before. Cournot [7] considered the simple best response dynamics, where at every step of the repeated game, firms react to what happened in the market on the previous step. Cournot showed that in the case of a duopoly, the simultaneous best response dynamics converges to the unique Nash equilibrium of the one shot game, i.e., after sufficiently many steps the two firms will play their Nash equilibrium strategies on every subsequent step. However, this result does not generalize to an arbitrary number of firms, as shown by Theocharis [20].

Milgrom and Roberts [16,17] were the first to explore connections between Cournot competition and super-modular games as a way to show convergence results for learning dynamics. In their work (as well as in followup papers [2,21]) Cournot duopolies as well as specific models of Cournot oligopolies are shown to exhibit strategic complementarities. This identifying property of supermodular games is shown to imply convergence to Nash equilibrium for a specific class of learning dynamics, known as adaptive choice behavior. This class of learning dynamics encompasses best-response dynamics and Bayesian learning but is generally orthogonal to the class of dynamics that we will be focusing on.

In this paper we are interested in dynamic behaviors where firms minimize their long term *regret*¹. Regret compares the firm's average utility to that of the best fixed constant action (e.g., constant production level in Cournot, and constant price in Bertrand). Having no-regret means that no deviating action would significantly improve the firm's utility (see [6]). Several learning algorithms [25], [13], are known to offer such guarantees, as their average regret bounds are $o(T)$, where T is the number of time steps.

Regret minimization procedures prescribe to some rather desirable requirements in regards to modeling market behavior. Firstly, they are rational, in the sense that an agent is given guarantees on her own utility regardless of how the other agents act. Moreover, they are distributed, since an agent needs to be aware only of her own utility. Many of the no regret procedures [11] are rather intuitive, as they share the idea that agents increase the probability of choosing actions that have been performing well in the past. Several learning procedures are known to be of no-regret, but more importantly, the assumption is not tied to any specific algorithmic procedure, but merely captures successful long-term behavior. Lastly, no-regret guarantees can be achieved even in the "multi-armed bandit" setting [3,10,11], where the input for the algorithm consists only of the payoffs received. This feature is important in the case that firms are not fully aware of the market structure (i.e., demand function), and are maybe even uncertain regarding their own production costs.

In the most relevant result to our work, Even-Dar *et al.* [9] study no regret dynamics in a class of games that includes Cournot competition with linear

¹ Regret is sometimes also referred to as external regret.

inverse market demand function, and convex costs functions. They show that the *average* production level of every firm, as well as its *average* profits, converge to the ones in the unique Nash equilibrium of the one shot game.

1.1 Our Results

In this work we examine the behavior of no regret dynamics in Cournot and Bertrand oligopoly models.

In the classic model of Bertrand oligopoly [5], it is well known that oligopolies with more than two firms exhibit several trivial Nash equilibria but in all of them the prices are equal to the marginal costs and all players make zero profit (Bertrand paradox). This phenomenon has been verified for correlated equilibria only for the case of a duopoly [2], where correlated equilibria are unique [16]. In our work, we show that under no-regret behavior the zero-profit postulate does not hold even in the case of two players. In fact, we show that not only does the market not necessarily converge to zero profit outcomes, but that the players can actually enjoy significant profits. In summary, our main results for Bertrand oligopolies under no-regret have as follows:

1. The Bertrand paradox does not hold anymore; firms enjoy non-zero profits under no-regret behavior.
2. Moreover, the identified profits can be rather significant when the number of players is small (e.g. 17% of optimal profits in the case of a duopoly). Profits however, tend to go to zero quickly as the number of firms increases.

Interestingly, our observations about no-regret behavior in Bertrand oligopolies agree to a large extent both with experimental work [8], as well as with empirical observations about real world oligopolistic markets [19].

The study of correlated equilibria [14,23] as well as of no-regret dynamics in [9] in Cournot oligopolies, has been an area of interest in both economics as well as computer science. In our work, we analyze a model of Cournot equilibria, which is a strict generalization of the models in [14,23,9], under no-regret dynamics. In fact, our results can be extended to all dynamics, in which each player's average payoff dominates the one they would receive if they always deviated to their respective Nash equilibrium strategy. This is a strict generalization of no-regret dynamics, since no-regret dynamics must fare well against all fixed strategies. In a novel approach in this line of work, we consider the evolution of the market not only from the perspective of the firms (individual production levels, profits), but also from the consumers' perspective (aggregate production level, prices) which leads to new insights. As a result, we can prove a single unifying message for all models examined: the daily prices converge to their level at Nash equilibrium.

In summary, our main results for Cournot oligopolies under no-regret have as follows:

² In [22] it is claimed that the correlated equilibria of Bertrand games are unique under some special cases.

1. In Cournot oligopoly with linear inverse demand function, and weakly convex costs, when every firm experiences no-regret, the empirical distribution of the daily overall production level, as well as of the daily prices, converges to a single point that corresponds to the Nash equilibrium of the one shot game.
2. When the firms produce products that are not perfect substitutes i.e., when even the tiniest of *product differentiation* is introduced, the empirical distributions of all market characteristics including the daily production levels of every firm converge to their levels in Nash equilibrium.
3. Some product differentiation is necessary in order to alleviate the nondeterminism of the day-to-day behavior on the side of the firms.

Table 1 summarizes what is known about equilibrium, and no-regret in Cournot and Bertrand oligopolies.

Table 1. Overview of results

	Bertrand	Cournot with perfect substitutes	Cournot with product differentiation
Nash equilibria	Infinite, Unique prices, Unique profits	Unique	Unique
Correlated equilibria	Infinite ²	Unique	Unique
No Regret	Infinite, Different prices, Different profits	Infinite, Unique prices, Different profits	Unique

2 Preliminaries

2.1 Models of Oligopoly

We formally define Cournot oligopoly, and Bertrand oligopoly, as strategic games, with continuous action space.

Definition 1. A Cournot oligopoly is a game between n firms, where the strategy space S_i of firm i is the span of its production level q_i . Typically, S_i is defined to be the interval $[0, \infty)$. The utility function for firm i is $u_i(q_1, \dots, q_n) = P_i(q_1, \dots, q_n)q_i - c_i(q_i)$, where P_i is the market inverse demand function for the good of firm i , which maps the vector of production levels to a market clearing price in \mathbb{R}^+ .

Our focus is on the case of linear inverse demand function. The utility of a firm i as a function of the firms' production levels is $u_i(q_1, \dots, q_n) = (a - bQ)q_i - c_i(q_i)$, where a and b are positive constants, and $Q = \sum_i q_i$ denotes the total product supply. In Section 4 we consider an extension of Cournot oligopoly with perfect substitutes, to the case of product differentiation, where the price of

firm i depends in an asymmetric manner on his own production level, and the production levels of the other firms. In this case the market inverse demand function $P^i(q)$ is given by $P^i(q) = a_i - b_i q_i - b_i \gamma_i Q_{-i} = a_i - b_i(1 - \gamma_i)q_i - b_i \gamma_i Q$, where γ_i denotes the degree of product differentiation between products, $0 < \gamma_i \leq 1$, $b_i > 0$.

Definition 2. A Bertrand oligopoly is a strategic game between n firms, where the strategy space P_i of firm i is its declared price p_i , which lies in the interval of all possible prices $[0, \infty)$, and its utility function is $u_i(p_1, \dots, p_n) = D_i(p_1, \dots, p_n)p_i - c_i(D_i(p_1, \dots, p_n))$, where D_i is the market demand function of firm i , that maps from the vector of firms prices to a demand in \mathbb{R}^+ .

We consider Bertrand oligopoly with a linear demand function, in which the market demand is equally shared among the firms with the least price:

$$D_i(p_1, \dots, p_n) = \begin{cases} 0 & p_i > p_j, \text{ for some } j \\ \frac{a-p_i}{b(m+1)} & p_i \leq p_j \text{ for all } j, \text{ and } m = |\{j \neq i | p_j = p_i\}| \end{cases}$$

Intuitively, this means that the market demand goes down linearly as the minimal announced price increases. If the minimal price has been offered by more than one firms, these firms share the market demand equally.

2.2 Regret Minimization

Having no-regret in an online sequential problem is defined as follows:

Definition 3. An online sequential problem consists of a feasible set $F \in \mathbb{R}^m$, and an infinite sequence of functions $\{f^1, f^2, \dots\}$, where $f^t : \mathbb{R}^m \rightarrow \mathbb{R}$.

At each time step t , an online algorithm selects a vector $x^t \in \mathbb{R}^m$. After the vector is selected, the algorithm receives f^t , and collects a payoff of $f^t(x^t)$. All decisions must be made *online*, in the sense that an algorithm does not know f^t before selecting x^t , i.e., at each time t , a (possibly randomized) algorithm can be thought of as a mapping from a history of functions up to time t , f^1, \dots, f^{t-1} , to the set F .

Given an algorithm \mathcal{A} and an online sequential problem $(F, \{f^1, f^2, \dots\})$, if $\{x^1, x^2, \dots\}$ are the vectors selected by \mathcal{A} , then the payoff of \mathcal{A} until time T is $\sum_{t=1}^T f^t(x^t)$. The payoff of a static feasible vector $x \in F$, is $\sum_{t=1}^T f^t(x)$. Regret compares the performance of an algorithm with the best static action in hindsight:

Definition 4. The external regret of algorithm \mathcal{A} , at time T is defined as

$$\mathcal{R}(T) = \max_{x \in F} \sum_{t=1}^T f^t(x) - \sum_{t=1}^T f^t(x^t).$$

An algorithm is said to have no-external regret, if for every online sequential problem, its regret at time T is $o(T)$.

The regret of a firm in a repeated oligopoly game: Consider the case of n firms that engage in a repeated Cournot (alternatively Bertrand) oligopoly game, and suppose that $\{x^t\}_{t=1}^\infty$ is a sequence of vectors, where x^t represents the production levels (alternatively, prices), set by the firms at time t . The regret of firm i at time T is defined as $\mathcal{R}_i(T) = \max_{y \in S_i} \sum_{t=1}^T u_i(y, x_{-i}^t) - \sum_{t=1}^T (u_i(x^t))$, where u_i is the utility function of firm i , and S_i is the strategy set of i .

3 Bertrand Oligopolies

We will be focusing on the case where all firms share the same linear cost function (i.e. $C_i(x) = cx$ for all i). The set of Nash equilibria of this game consists of all price vectors such that the prices of at least two firms are equal to c , whereas all others are greater than c . Although there exist multiple Nash equilibria, all of them imply the same market prices where the firms sell at marginal cost and hence no profit is being made. On the contrary, we will show that firms can achieve positive payoffs while experiencing no-regret. Moreover, we will show that infinitely many positive profit vectors are sustainable under no-regret guarantees.

We will show that by producing a probability distribution on outcomes of Bertrand oligopolies such that when the market outcomes are chosen according to this distribution, then each player’s expected payoff is at least as large as the expected payoff of her best deviating strategy, given that all other players follow the distribution. More formally, we will produce a probability measure F on (P, Σ) ⁴ such that for all $i, p'_i \int_P [u_i(p_i, p_{-i}) - u_i(p'_i, p_{-i})] dF(p) \geq 0$. Such probability distributions are referred to as *coarse correlated equilibria (CCE)* [24]. It is straightforward to check that, any market history whose empirical distribution of outcomes converges to a CCE imposes no regret on the involved players. Indeed, the average profits of the players, will converge to their expected values, which by definition of the CCE exhibit no-regret. Conversely, any CCE can give rise to such a history, merely by infinitely choosing outcomes according to it. Therefore it suffices to prove that we can achieve positive payoffs in a CCE. Our constructions are inspired by observations regarding the structure of Nash in Bertrand games made in [4].

Theorem 1. *All symmetric linear Bertrand games exhibit coarse correlated equilibria (CCE) in which all players exhibit positive profits.*

Proof. We denote $(p - c)(a - p)/b$ by $\pi(p)$, which is equal to the utility function when the winning player is unique. This function is strictly increasing in $[c, (a+c)/2]$. As a result, we can define the following distribution:

$$F_0(p) = \begin{cases} 0 & p \leq \beta \\ 1 - \left(\frac{\pi(\beta)}{\pi(p)}\right)^{\frac{1}{n-1}} & \beta < p < \gamma \\ 1 & p \geq \gamma \end{cases} \tag{1}$$

⁴ P is the set of all strategy (price) profiles and Σ is the Borel σ -algebra on it.

where $\beta > c$ and $\gamma \leq (a + c)/2$. Before, we construct the CCE, we will examine some properties of the mixed strategy profile where each player chooses a strategy according to $F_0(p)$. We will show that each action in the support of the mixed strategy $F_0(p)$ is optimal [3](#) except from β .

The probability distribution $F_0(p)$ sets $p = \gamma$ with probability $(\frac{\pi(\beta)}{\pi(\gamma)})^{\frac{1}{n-1}}$. The rest of the probability distribution is atomless, that is $Pr(p = x | x < \gamma) = 0$. Suppose that the rest $n-1$ players play according to this distribution. The expected payoff for playing price $\beta \leq p < \gamma$ would be equal to:

$$E[u] = [1 - F_0(p)]^{n-1} \pi(p) = \pi(\beta)$$

Next, we will compute the expected payoff for playing γ . The only way for someone to win when playing γ is for everyone else to be playing γ . However, in this case they share the pot. So,

$$E[u] = [(\frac{\pi(\beta)}{\pi(\gamma)})^{\frac{1}{n-1}}]^{n-1} \frac{\pi(\gamma)}{n} = \frac{\pi(\beta)}{n}$$

Also, just to complete the picture, the expected cost for playing $p > \gamma$ is 0 and the expected profit for playing $p < \beta$ is less than $\pi(\beta)$. Lastly, let us compute the expected utility of the players when all of them play according to this strategy distribution. In this case and if we denote $(\frac{\pi(\beta)}{\pi(\gamma)})^{\frac{1}{n-1}}$ as ρ , we have

$$E[u] = (1 - \rho)\pi(\beta) + \rho \frac{\pi(\beta)}{n} = (1 - \frac{n-1}{n}\rho)\pi(\beta).$$

Now, we will define a probability distribution over outcomes of the Bertrand games and we will prove that it is a CCE. We will be using three prices α, β, γ such that $c < \alpha < \beta < \gamma \leq (a + c)/2$. With probability 1/2 all players play α and with probability 1/2 all players play according to F_0 . Regarding the expected payoff for each player, we have that with probability 1/2 they all share the profit at price α and with probability 1/2 they gain the precomputed payoff of the defined mixed strategy profile. Specifically,

$$E[u] = 1/2 \frac{\pi(\alpha)}{n} + 1/2 [(1 - \frac{n-1}{n}\rho)\pi(\beta)]$$

In order for this to be a CCE it must be the case any deviating player cannot increase his payoff by deviating to a single strategy given that the rest of the players keep playing according to this distribution. Let us examine what are the best deviating strategies for a player. First a player can deviate and play $\alpha - \epsilon$ for some small $\epsilon > 0$. Her expected payoff in that case is essentially $\pi(\alpha)$ since she will always be winning the competition. It is obvious that any strategy less than that is clearly worse for him since π is increasing in the range $[0, \alpha] \subset [0, (a+c)/2]$. Another good deviating strategy for the player is to play a strategy in $[\beta, \gamma)$ since this is a best response to the second probability distribution. Actually, given that a player deviates to a price which is greater than α her best choice is to deviate

³ Given that all players play according to $F_0(p)$.

to any price in the $[\beta, \gamma)$ range. This is true because the only way to incur payoff at this point is to maximize her payoff when her opponents play according to $F_0(p)$. As we have seen, the player achieves a maximum expected payoff of $\pi(\beta)$ when playing within that range. So, the best deviating strategy is either $\alpha - \epsilon$ or something in the range $[\beta, \gamma)$. If our current expected payoff exceeds the payoffs at these points then our distribution is a CCE. So, we wish to have:

$$1/2 \frac{\pi(\alpha)}{n} + 1/2 [(1 - \frac{n-1}{n} \rho) \pi(\beta)] \geq \pi(\alpha),$$

and

$$1/2 \frac{\pi(\alpha)}{n} + 1/2 [(1 - \frac{n-1}{n} \rho) \pi(\beta)] \geq 1/2 \pi(\beta) .$$

Let us try to analyze each relation separately:

$$\begin{aligned} 1/2 \frac{\pi(\alpha)}{n} + 1/2 [(1 - \frac{n-1}{n} \rho) \pi(\beta)] &\geq \pi(\alpha) \Leftrightarrow \\ 1/2 [(1 - \frac{n-1}{n} \rho) \pi(\beta)] &\geq (1 - \frac{1}{2n}) \pi(\alpha) \Leftrightarrow \\ (1 - \frac{n-1}{n} \rho) \frac{n}{2n-1} &\geq \frac{\pi(\alpha)}{\pi(\beta)} \end{aligned}$$

Similarly, from the second inequality we have:

$$\begin{aligned} 1/2 \frac{\pi(\alpha)}{n} + 1/2 [(1 - \frac{n-1}{n} \rho) \pi(\beta)] &\geq 1/2 \pi(\beta) \Leftrightarrow \\ \frac{\pi(\alpha)}{n} + [(1 - \frac{n-1}{n} \rho) \pi(\beta)] &\geq \pi(\beta) \Leftrightarrow \\ \frac{\pi(\alpha)}{\pi(\beta)} &\geq (n-1) \rho \end{aligned}$$

So, in order for our probability distribution over outcomes to be a CCE, it suffices that we choose α, β so that:

$$(n-1) \rho \leq \frac{\pi(\alpha)}{\pi(\beta)} \leq (1 - \frac{n-1}{n} \rho) \frac{n}{2n-1}$$

However $\pi(\alpha), \pi(\beta)$ are positive payoffs in the range $(0, \frac{(a-c)^2}{4b}]$ with $\pi(\alpha) < \pi(\beta)$. By choosing proper α, β we can reproduce any number in the range $(0, 1)$. Hence, all we have to do is show that we can choose ρ appropriately such that:

$$(n-1) \rho \leq (1 - \frac{n-1}{n} \rho) \frac{n}{2n-1}$$

as well as $(n-1) \rho < 1$ and $0 < (1 - \frac{n-1}{n} \rho) \frac{n}{2n-1}$. Again, by manipulating the given inequality we get:

$$(n-1) \rho \leq (1 - \frac{n-1}{n} \rho) \frac{n}{2n-1} \Leftrightarrow (n-1 + \frac{n-1}{2n-1}) \rho \leq \frac{n}{2n-1}$$

It suffices to choose $\rho = \frac{1}{2n-1}$ and $\frac{\pi(\alpha)}{\pi(\beta)} = \frac{n-1}{2n-1}$ to satisfy all inequalities. However, $\rho = (\frac{\pi(\beta)}{\pi(\gamma)})^{\frac{1}{n-1}}$. So, we have that we need to choose β and γ such that

$\frac{\pi(\beta)}{\pi(\gamma)} = (\frac{1}{2n-1})^{n-1}$. So, given any $\pi(\gamma) \in (0, \pi(\frac{a+c}{2})] = (0, \frac{(a-c)^2}{4b}]$, we can define β, α such that the distribution we have defined is a CCE. The expected payoffs of all players are positive in this CCE and can vary widely. Hence, no regret behavior can support infinitely many different positive average payoff profiles, in contrast to Bertrand’s paradox. Finally, this construction establishes that increased competition is necessary for converging to marginal cost pricing.

The profitability of the families of Bertrand no-regret histories we have identified, decreases much faster than the profitability of the no-regret histories in the Cournot oligopolies as the number of agents (firms) increases. In fact, for $n = 4$ players we see that essentially the prices reach the level of marginal costs as profitability drops to zero. This theoretical projection is in perfect agreement both with experimental work in the case of Bertrand games [8], as well as with empirical observations about real world oligopolistic markets [19]. Specifically, “the rule of three”, as is presented in [19], states that in most markets three major players will emerge (e.g. ExxonMobil, Texaco and Chevron in petroleum). In order for the smaller companies to be successful they need to specialize and address niche markets. Our works suggests a possible quantitative explanation behind this phenomenon, as a result of the steep drop in profitability in the case of Bertrand markets.

4 Cournot Oligopolies

We will be analyzing a generalization of the Cournot model with product differentiation that was introduced by Yi [23]. By exploring ideas from that work, we will show how we can generalize its results and prove tight convergence guarantees in the case of no-regret algorithms. Our model will be the Cournot competition in the case of linear demand functions with symmetric product differentiation, where the inverse demand function $P^i(q)$ is given by $P^i(q) = a_i - b_i q_i - b_i \gamma_i Q_{-i} = a_i - b_i(1 - \gamma_i)q_i - b_i \gamma_i Q$, where γ_i denotes the degree of product differentiation between products, $0 < \gamma_i \leq 1$, $b_i > 0$ and $Q = \sum_i q_i$ denotes the total product supply. We will assume that the cost functions are convex and twice continuously differentiable. We denote by $q^* = (q_1^*, \dots, q_n^*)$ a pure Nash equilibrium of the one-shot game, which is known to exist by [18]. Finally, Q^* denotes the aggregate production level at the Nash.

Lemma 1. *Let q_i^τ, Q^τ denote respectively the production level of company i and the aggregate production level in period τ of a differentiated Cournot market with differentiation levels γ_i for each product. If each player’s regret converges to zero, then*

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=1}^t \left(\frac{\gamma_i - 1}{\gamma_i} \sum_i (q_i^\tau - q_i^*)^2 - (Q^\tau - Q^*)^2 \right) = 0$$

Depending on the details of the Cournot model, we have the following cases:

A) Perfect Substitutes

This is to the simplest case of Cournot competition and was the model analyzed by Even-Dar et. al. in [9]. We have that $\gamma_i = 1$ and $C''_i(q_i) \geq 0$ for all i, q_i .

Theorem 2. *Suppose that n firms participate in a homogeneous Cournot oligopoly game of perfect substitutes with linear demand ($P^i(q) = a_i - b_i Q$,) and convex cost functions. If all firms experience no-regret as t grows to infinity, then given any $\epsilon > 0$, for all but $o(t)$ periods τ in $[1, t]$ we have that $|Q^\tau - Q^*| < \epsilon$.*

We should stress here that this is a statement about the day-to-day behavior (i.e. aggregate production levels) instead of average behavior as in [9] (Theorem 3.1.). In particular, this statement implies that the average action vector and the average utility of each player converge to their respective levels at the Nash equilibrium, a result that has been shown in [9]. Given the convergence of the day-to-day characteristics of the market prices and total supply, it is rather tempting to try to prove a similar statement about the convergence of the action vector and utilities of the firms and not merely of their averages. Here, we show that this cannot be the case by providing sufficient conditions for a market history to be of no-regret.

This is essentially a negative result, so it suffices to prove that this holds for as simple a model as possible. Therefore, we will focus on the special case of the fully symmetric Cournot oligopoly ($a_i = a$ and $b_i = b$) with linear cost functions. It is well known that these games exhibit a unique Nash $q^* = (q_1^*, q_2^*, \dots, q_n^*)$ where $q_i^* = (a - (n+1)c_i - \sum_{j \in N} c_j) / ((n+1)b)$.

Theorem 3. *Suppose that n firms participate in a homogeneous Cournot oligopoly game with linear demand ($P^i(q) = a - bQ$,) and linear cost functions and let q^* denote the unique Nash of this game. Any market history, where for all time periods τ , $Q^\tau = Q^*$ and where the time average \hat{q}_i of each player's actions converges to her Nash strategy q_i^* does not induce regret to any player.*

An immediate corollary of the above theorem is that one cannot hope to prove convergence of the day-to-day action profiles in any model that generalizes the basic linear Cournot model. Surprisingly, if we introduce product differentiation in the market, then we can actually prove convergence of all attributes (i.e. action profiles, profits, prices e.t.c) of the market.

B) Symmetric Product Differentiation

In the case that $0 < \gamma_i < 1$ we have convergence of the daily quantities of each player to Nash equilibrium.

Theorem 4. *Suppose that n firms participate in a differentiated good Cournot oligopoly game with linear demand ($P^i(q) = a_i - q_i - \gamma Q$, $0 < \gamma < 1$). If all firms experience no-regret, then given any $\epsilon > 0$, as t grows to infinity, for all but $o(t)$ periods τ in $[1, t]$ we have that $|q_i^\tau - q_i^*| < \epsilon$, where q^* is the unique Nash equilibrium.*

References

1. Abernethy, J., Hazan, E., Rakhlin, A.: Competing in the dark: An efficient algorithm for bandit linear optimization. In: COLT (2008)

2. Amir, R.: Cournot oligopoly and the theory of supermodular games. *Games and Economic Behavior* (1996)
3. Auer, P., Cesa-Bianchi, N., Freund, Y., Schapire, R.E.: The nonstochastic multi-armed bandit problem. *SIAM Journal on Computing* (2002)
4. Baye, M., Morgan, J.: A folk theorem for one-shot bertrand games. *Economic Letters* 65(1), 59–65 (1999)
5. Bertrand, J.: *Theorie mathematique de la richesse sociale*. *Journal des Savants* 67, 499–508 (1883)
6. Cesa-Bianchi, N., Lugosi, G.: *Prediction, Learning and Games*. Cambridge University Press, Cambridge (2006)
7. Cournot, A.A.: *Recherches sur les principes mathmatiques de la thorie des richesses* (1838)
8. Dufwenberg, M., Gneezy, U.: Price competition and market concentration: An experimental study. *International Journal of Industrial Organization* 18 (2000)
9. Even-Dar, E., Mansour, Y., Nadav, U.: On the convergence of regret minimization dynamics in concave games. In: 41st ACM Symposium on Theory of Computing, STOC (2009)
10. Flaxman, A., Kalai, A.T., McMahan, B.: Online convex optimization in the bandit setting: Gradient descent without a gradient. In: SODA (2005)
11. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55(1), 119–139 (1997)
12. Fudenberg, D., Levine, D.K.: *The theory of learning in games*. MIT Press, Cambridge (1998)
13. Hazan, E., Agarwal, A., Kale, S.: Logarithmic regret algorithms for online convex optimization. *Machine Learning* 69(2-3), 169–192 (2007)
14. Liu, L.: Correlated equilibrium of cournot oligopoly competition. *Journal of Economic Theory* (1996)
15. Mas-Colell, A., Winston, M., Green, J.: *Microeconomic Theory*. Oxford University Press, Oxford (1995)
16. Milgrom, P., Roberts, J.: Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* 58(6), 1255–1277 (1990)
17. Milgrom, P., Roberts, J.: Adaptive and sophisticated learning in repeated normal form games. In: *Games and Economic Behavior*, pp. 82–100 (February 1991)
18. Rosen, J.: Existence and uniqueness of equilibrium points for concave n-person games
19. Sheth, J.N., Sisodia, R.S.: *The Rule of Three: Surviving and Thriving in Competitive Markets*
20. Theocharis, R.D.: On the stability of the cournot solution on the oligopoly problem. *The Review of Economic Studies* 27(2), 133–134 (1960)
21. Vives, X.: *Oligopoly Pricing: Old Ideas and New Tools*. The MIT Press, Cambridge (2001)
22. Wu, J.: Correlated equilibrium of bertrand competition. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 166–177. Springer, Heidelberg (2008)
23. Yi, S.-S.: On the existence of a unique correlated equilibrium in cournot oligopoly. *Economic Letters* 54, 235–239 (1997)
24. Young, H.: *Strategic Learning and Its Limits*. Oxford University Press, Oxford (2004)
25. Zinkevich, M.: Online convex programming and generalized infinitesimal gradient ascent. In: ICML (2003)

On the Complexity of Pareto-optimal Nash and Strong Equilibria

Martin Hoefer* and Alexander Skopalik**

Dept. of Computer Science, RWTH Aachen University, Germany
{mhoefer,skopalik}@cs.rwth-aachen.de

Abstract. We consider the computational complexity of coalitional solution concepts in scenarios related to load balancing such as anonymous and congestion games. In congestion games, Pareto-optimal Nash and strong equilibria, which are resilient to coalitional deviations, have recently been shown to yield significantly smaller inefficiency. Unfortunately, we show that several problems regarding existence, recognition, and computation of these concepts are hard, even in seemingly special classes of games. In anonymous games with constant number of strategies, we can efficiently recognize a state as Pareto-optimal Nash or strong equilibrium, but deciding existence for a game remains hard. In the case of player-specific singleton congestion games, we show that recognition and computation of both concepts can be done efficiently. In addition, in these games there are always short sequences of coalitional improvement moves to Pareto-optimal Nash and strong equilibria that can be computed efficiently.

1 Introduction

A central theme of (algorithmic) game theory is the study and analysis of equilibria to predict the outcomes of interacting rational agents. Insights about the nature of equilibria yield numerous benefits, e.g., for the design and implementation of regulations such as laws in society or protocols in distributed systems. In strategic games the most frequently studied concept of stability is the Nash equilibrium (NE) – a state, in which no agent has an incentive to unilaterally deviate. The analysis of Nash equilibrium has occupied a central place in game theory since its beginning. More recently, the computational complexity of Nash equilibrium has been analyzed to determine whether the concept is reasonable from a computational point of view [1, 11].

Much of the attractiveness of Nash equilibrium stems from its elegance and simplicity and (in the mixed case) from guaranteed existence. However, Nash equilibrium is only resilient against *unilateral* deviations. It neglects the aspect of *cooperation* or *coordination* between agents. Obviously, in many scenarios

* Supported by DFG through UMIC Research Center at RWTH Aachen University and grant Ho 3831/3-1.

** Supported in part by the German Israeli Foundation (GIF) under contract 877/05.

agents have an incentive to cooperate, as cooperation often allows to dramatically improve the situation of every participant. In these cases, the negligence of cooperation in Nash equilibrium significantly hurts the predictive value of the concept in practice.

This shortcoming of Nash equilibrium has been addressed already in the 1950s, most notably by Aumann [3] who introduced the strong equilibrium (SE) – a state, from which no coalition of agents can jointly deviate and thereby strictly improve all members of the coalition. Strong equilibria include the consideration of cooperation, but this comes at the expense of guaranteed existence. Hence, using strong equilibria we can make better predictions about the outcome in many but not all games. In addition, strong equilibria have recently been shown to exhibit a significantly smaller inefficiency in congestion and load balancing games [2,7]. Similar results have been obtained for a weaker concept of Pareto-optimal Nash equilibria (PoNE) [17], in which only unilateral deviations or deviations of the whole player set are allowed. From a designer perspective, it thus appears attractive to design (distributed) algorithms for cooperation between agents that allow to reach these states if they exist. The analysis of the computational complexity of SE and PoNE has been posed as an open problem in [7] and is the subject of this paper.

Related Work and New Results. In this paper, we examine the computational complexity of SE and PoNE in games related to congestion and load balancing. In particular, we consider problems of the following types. *Existence:* Does a given game have a SE? *Recognition:* Is a given state of a game a SE? *Computation:* If a game has a SE, can we compute it in polynomial time? We consider these problems for SE and PoNE and other related variants. In general, our results shed light on the inherent complexity of cooperation. While in some cases, we can give efficient algorithms, most of our insights turn out to be hardness results.

In Section 3 we study anonymous games [4,5], in which the cost of a player does not depend on the identity of the other players. A notable case are games with a constant number of strategies, in which the existence of pure NE can be decided efficiently [6], and for mixed NE there exists an FPTAS [8,9]. In this case, we can decide the recognition problem efficiently for SE and PoNE. Our algorithm uses computation of perfect matchings together with careful enumeration to find a coalition and a profitable deviation if they exist. Deciding the existence problem for SE and PoNE for a given anonymous game, however, is strongly NP-complete, even for a small constant number of strategies. Note that this is in contrast to general graphical games, where the existence problem is even Σ_2^P -complete and thus at the second level of the polynomial hierarchy [12].

An important class of anonymous games are cases of load balancing, i.e., player-specific singleton congestion games [19]. Previous work has shown existence [16] for such games with non-decreasing cost functions. However, we are not aware of any result providing efficient algorithms to compute SE or PoNE. We show in Section 4 how to obtain a SE in polynomial time and how to recognize a given state as a SE or PoNE. Interestingly, our results imply that there

always exist sequences of coalitional improvement moves to SE and PoNE that are of polynomial length. We show how to obtain these moves for the players efficiently.

In Section 5 we consider standard congestion games [20] with special structure. In congestion games it has been shown that SE can be absent [15], and a characterization result has been given that describes structures of strategy spaces that always allow SE for any set of non-decreasing latency functions. An extension of SE to correlated strategies has been considered in [21]. More recently, it has been shown that in a bottleneck variant of congestion games SE exist [14], and that SE in symmetric network and matroid games can be computed in polynomial time [13]. In standard matroid and symmetric network games NE can be computed efficiently [1,11]. In addition, there is a plethora of work on the complexity of NE in standard, weighted or integer-splittable congestion games [10,18], or local-effect games [18].

We here treat standard congestion games and aim to draw a more detailed picture beyond the characterization of [15]. Unfortunately, even when the strategy space has simultaneously a symmetric network and matroid structure, the existence problem for SE is strongly Co-NP-hard. This is particularly interesting in light of the positive results in related work mentioned above. Additionally, we can even show weak NP-hardness for such games that have only 2 players. This directly implies the hardness result also for PoNE, and k -SE (in which only coalitions of size at most k are allowed), for any $k \geq 2$.

All proofs missing from this extended abstract are deferred to the full version of the paper.

2 Definitions

Strategic games. A strategic game $\Gamma = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$ has a finite set $N = \{1, \dots, n\}$ of players. Player $i \in N$ has a set S_i of strategies. A state $s \in S = S_1 \times \dots \times S_n$ is sometimes referred to as a *strategy profile* or *profile*. The *cost function* of player i is $c_i : S \rightarrow \mathbb{R}$, which maps each state $s \in S$ to a real number. We here denote by $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$. A state $s \in S$ is a *k -strong equilibrium* (k -SE) if no subset of the players $I \subseteq N$ with $|I| \leq k$ can benefit from jointly deviating from their strategies. Formally, there is no tuple $(s', I) \in S \times 2^N$ with $s' \neq s$ and $|I| \leq k$ such that $\forall i \in I$ we have $c_i(s') < c_i(s)$ and $\forall i \in N \setminus I$ it holds $s_i = s'_i$. A n -SE is called *strong equilibrium* (SE), and a 1-SE is a *Nash equilibrium* (NE). A *Pareto-optimal Nash equilibrium* (PoNE) is a NE s , in which there is no other state s' with $c_i(s') < c_i(s)$ for every $i \in N$.

Anonymous games. An anonymous game is a tuple $(N, E, (c_i)_{i \in N})$, where E is a set of resources and the strategy space of every player, i.e., $S_i = E$ for all $i \in N$. The cost function of player i depends only on the *numbers* of players that have chosen the strategies, but not their identities. More formally, for a state $s = (s_1, \dots, s_n)$, we define the *load* $l_e(s)$ on resource e by $l_e(s) = |\{i | e \in s_i\}|$, that is $l_e(s)$ is the number of players that selected resource e as strategy in s . We call the tuple $(l_e(s))_{e \in E}$ the *load profile* of s . Let L be the set of all load

profiles. The cost function of player i is $c_i : E \times L \rightarrow \mathbb{R}$, which maps the strategy of player i and the load profile of s to a real number. The function of player i depends only on numbers of other players but not on their identity (i.e., which set of other players he shares his resource with). However, two players i and j choosing the same resource e can suffer a different cost, as c_i and c_j might map $l(s)$ to a different cost value. An interesting subclass are anonymous games with a *constant-size strategy set* in which the size of E is a fixed constant, which we study in this paper. Another subclass of anonymous games that we study are **player-specific singleton congestion games**. While we assume that these games can have an arbitrary number of strategies, the crucial adjustment is that the cost function of player i is $c_i(l_{s_i}(s)) \in \mathbb{R}$. Thus, it depends only on $l_{s_i}(s)$ of the resource chosen by player i . We assume that cost functions are non-decreasing.

Congestion Games. A congestion game is a tuple $(N, E, (S_i)_{i \in N}, (d_e)_{e \in E})$, where E is a set of resources, $S_i \subseteq 2^E$ is the strategy space of player $i \in N$, and $d_e : \mathbb{N} \rightarrow \mathbb{Z}$ is a delay function associated with resource $e \in E$. As above, we define the *load* on $e \in E$ in state s as $l_e(s) = |\{i | e \in s_i\}|$. The cost (or delay) $c_i(s)$ of player i in s is $c_i(s) = \sum_{e \in s_i} d_e(l_e(s))$. Note that symmetric congestion games can be seen as another subclass of anonymous games (albeit with a number of strategies that is possibly exponential in $|E|$), in which the cost functions have a special structure. In terms of SE and PoNE we can also equivalently view asymmetric congestion games as anonymous games with strategy set $\bigcup_{i \in N} S_i$, where player i has a prohibitively large cost when he chooses a strategy $s_i \notin S_i$.

3 Anonymous Games

In this section we start by considering the case of anonymous games with a constant number of strategies. In this case we can decide efficiently if a given state s is a SE.

Theorem 1. *A state s of an anonymous game with a constant number of strategies can be recognized as a Pareto-optimal Nash or a strong equilibrium in polynomial time.*

Proof. We show here that we can efficiently compute a profitable deviation if it exists. For the given state s let $l(s)$ denote the load profile of s . For a given (possibly different) load profile l we present an algorithm that checks in polynomial time if there exists a profitable joint deviation from s to a state that has load profile l . The algorithm repeatedly tries to compute a perfect matching in a bipartite graph. By running this algorithm for s and all polynomially many load profiles l , the theorem follows.

For a given state s and a given load profile l , we construct a bipartite *deviation graph* G and search for a perfect matching. The vertex set of graph $G = (A \cup B, F)$ is defined by $A = \{v_i \mid \text{for all } 1 \leq i \leq n\}$ and $B = \{v_{e,j} \mid \text{for all } e \in E \text{ and } 1 \leq j \leq l_e\}$. For each $1 \leq i \leq n$ and resource $e \in E$, we add

all the edges $(v_i, v_{e,j})$ for all $1 \leq j \leq l_e$ if and only if $c_i(e, l) < c_i(s_i, l(s))$. In addition, there are edges $(v_i, v_{s_i,j})$ for every player $1 \leq i \leq n$ and his current strategy s_i , for every $1 \leq j \leq l_{s_i}$.

Note that a perfect matching in this graph yields an assignment of players to strategies. From this we can derive a new state s' by setting $s'_i = e$ iff $(v_i, v_{e,j})$ is in the matching for some $1 \leq j \leq l_e$. s' represents an improvement for all players i with $s_i \neq s'_i$. Therefore, if there is a profitable coalitional deviation from s to a state s' with $l(s) \neq l(s')$, the algorithm finds at least one such deviation. Observe, that for $l = l(s)$ the algorithm may return s itself. To check if there is a deviation to a strategy $s' \neq s$ with $l(s) = l(s')$, we run the algorithm n times with s and l as input. However, in the i -th run, we force player i to change his strategy by removing all edges $(v_i, v_{s_i,j})$. Thus, if there is a profitable deviation to a state $s' \neq s$ with $l(s) = l(s')$, then there will exist a perfect matching in at least one of the runs, and thereby we will find such a deviation. This proves the result for SE.

For Pareto-optimal Nash equilibria we first check if all unilateral deviations are unprofitable. For deviations of the complete set of players to a state with load profile l we use the above construction, but we add edges $(v_i, v_{s_i,j})$ if and only if they represent a strict improvement for player i , i.e., $c_i(s_i, l) < c_i(s_i, l(s))$. This implies that for each load profile l we only have to examine exactly one graph for a perfect matching. There is a perfect matching for some load profile l if and only if s is not a PoNE. This proves the result for PoNE. \square

We can decide for a given state whether it is a SE or not, which implies that the existence problem for PoNE and SE is in NP. In fact, deciding the existence of SE and PoNE is strongly NP-complete.

Theorem 2. *It is strongly NP-complete to decide if an anonymous game with a constant number of strategies has a Pareto-optimal Nash or strong equilibrium.*

Proof. We first prove the result for SE and present a reduction from 3SAT. Given a formula φ with the variables x_1, \dots, x_n and clauses c_1, \dots, c_m , we construct an anonymous game Γ_φ with players X_i^0, X_i^1 (for $1 \leq i \leq n$), C_j^k (for $1 \leq j \leq m$ and $1 \leq k \leq 10j$), V_i^k (for $1 \leq i \leq n$ and $1 \leq k \leq 10i + 10m$), Prisoner₁, and Prisoner₂. The set of strategies is {On, Off, Verify, False, Wait, Cooperate, Defect}, costs are shown in Fig. [11](#).

If φ is satisfiable, let b_1, \dots, b_n be a satisfying assignment. The following state is a SE. For each $1 \leq i \leq n$, the player $X_i^{b_i}$ plays On and the player $X_i^{1-b_i}$ plays Off. All players C_j^k and V_i^k play Verify and players Prisoner₁ and Prisoner₂ play Cooperate.

We show that there is no coalition that can improve by jointly deviating to another state. The players $X_i^{1-b_i}$ are playing Off and have the minimal possible cost of 1. Thus, they cannot be part of a deviating coalition. The players Prisoner₁ and Prisoner₂ can improve only if some other players move to False. We will show, this cannot happen.

For the remaining players, i.e., $X_i^{b_i}, C_j^k, V_i^k$, the only possible profitable deviation is to deviate to False. Clearly, if there is a deviation of a subset of these

Player	Strategy	Load profile	Cost
X_i^b $1 \leq i \leq n$ $b \in \{0, 1\}$	On	$ \text{On} = n$ and $ \text{Off} = n$	2
	On	Otherwise	3
	Off	$ \text{On} = n$ and $ \text{Off} = n$	1
	Off	Otherwise	3
	False	$ \text{False} \in \{10j + 3 \mid x_i \text{ appears in clause } c_j \text{ and } x_i = b \text{ does not satisfy } c_j\}$	1
	False	$ \text{False} = 10m + 10i + 2$	1
	False	Otherwise	4
C_j^k $1 \leq j \leq m$ $1 \leq k \leq 10j$	Verify		2
	False	$ \text{False} = 10j + 3$	1
	False	Otherwise	3
V_i^k $1 \leq i \leq n$ $1 \leq k \leq 10i + 10m$	Verify		2
	False	$ \text{False} = 10m + 10i + 2$	1
	False	Otherwise	3
Prisoner ₁ , Prisoner ₂	Cooperate	$ \text{False} = 0$	5
	Cooperate	$ \text{False} \neq 0$ and $ \text{Cooperate} = 2$	2
	Cooperate	$ \text{False} \neq 0$ and $ \text{Cooperate} \neq 2$	4
	Defect	$ \text{False} = 0$	5
	Defect	$ \text{False} \neq 0$ and $ \text{Defect} = 2$	3
	Defect	$ \text{False} \neq 0$ and $ \text{Cooperate} \neq 2$	1

Fig. 1. Description of the cost functions in the game Γ_φ . Strategies that are not listed here have cost of 6 and, therefore, are never played in equilibrium.

players, it must result in $10j + 3$ (for $1 \leq j \leq m$) or $10m + 10i + 2$ (for $1 \leq i \leq n$) players on False. We consider the former case. Assume there is a deviation of a coalition of some of the players that results in $10j' + 3$ many player on False. The coalition must contain the players $C_{j'}^1, \dots, C_{j'}^{10j'}$ and the three players $X_i^{b_i}$ with x_i appearing in clause $c_{j'}$. However, let x_{i^*} be a variable that satisfies $c_{j'}$ with $x_{i^*} = b_{i^*}$. Player $X_{i^*}^{b_{i^*}}$ does not improve by deviation to False. Therefore, no such deviation can exist. Similarly, there is no deviation of a coalition that yields $10m + 10i' + 2$ (for $1 \leq i' \leq n$) players on False. This is only possible if both players $X_{i'}^0$ and $X_{i'}^1$ are on strategy On.

Now, assume φ is not satisfiable, and there is a strategy profile s that is a SE. We first show that in s no player is on False. If some player is on False, the players Prisoner₁ and Prisoner₂ play a game corresponding to the prisoners dilemma. This game does not admit a SE and implies that in s no player can be on False.

Now since s is an equilibrium, there are exactly n players X_i^b on On and exactly n players X_i^b on Off because otherwise they would have cost of 3. There is no $1 \leq i' \leq n$ with both players $X_{i'}^0$ and $X_{i'}^1$ being on strategy On. Otherwise, those two players and the players $V_{i'}^1, \dots, V_{i'}^{10m+10i'}$ could jointly change to False and decrease their costs. Now, let $X_1^{b_1}, \dots, X_n^{b_n}$ be the players on On. Since φ is not satisfiable, the assignment b_1, \dots, b_n implied by the players on On creates at least one clause $c_{j'}$ that is not satisfied. Let $x_{i'}, x_{i''}$, and $x_{i'''}$ be

the three variables of this clause. Then, the players $X_{i'}^{b_{i'}}$, $X_{i''}^{b_{i''}}$, $X_{i'''}^{b_{i'''}}$, and the players $C_{j'}^1, \dots, C_{j'}^{10j'}$ could jointly change to False and decrease their costs. This is a contradiction to the assumption that s is a SE and completes our reduction.

We can easily modify the arguments and obtain a similar result for the complexity of PoNE. \square

Note that this implies that further restrictions on the games are necessary in order to decide existence or compute a SE or PoNE efficiently. We consider games with a constant number of player types, i.e., where each player has one out of a constant number of different cost functions.

Corollary 1. *In anonymous games with constant number of strategies that are (1) symmetric or (2) have only a constant number of different player types we can decide efficiently if Pareto-optimal Nash or strong equilibria exist and compute one efficiently if it exists.*

Note that for symmetric games the assignment of players in a load profile is irrelevant, hence we can use our algorithm from Theorem [1](#) above to check each of the polynomial number of profiles for being a SE or PoNE. For a constant number of player types, the number of essentially different assignments that can be derived from a single load profile is a polynomial number. Again, by enumeration and application of our algorithm we can decide existence and compute SE and PoNE efficiently.

4 Player-Specific Singleton Congestion Games

In this section we treat player-specific singleton congestion games. For games with non-decreasing cost functions it is known that SE always exist [\[16\]](#). Here we provide efficient algorithms to compute a SE and decide whether a given state is a SE or PoNE. To the best of our knowledge these results have not been described in the literature before.

Theorem 3. *In player-specific singleton congestion games with non-decreasing cost functions we can in polynomial time (1) decide whether a given state is a Pareto-optimal Nash or strong equilibrium and (2) compute a strong equilibrium in polynomial time.*

Proof. Obviously, a state s that is a SE or PoNE must be a NE. Consider a NE s and the corresponding load profile $l(s)$. Because cost functions are non-decreasing, every profitable coalitional deviation must result in a state s' with the same load profile $l(s)$. In particular, if the load profile changes to $l(s') \neq l(s)$, there must be a resource e with higher load $l_e(s') > l_e(s)$. Consider a player moving to e . Any player moving to e does not make a strict improvement, because otherwise he could move there unilaterally – a contradiction to s being a NE. Hence, whenever we have a NE, there must be a SE with the same load profile, a fact that was observed in [\[16\]](#). Every profitable coalitional deviation represents a circular switch of players and thereby decreases the sum of player costs.

We use our algorithm presented for anonymous games in Theorem 1 to decide for a given state s whether it is a SE or PoNE. Note that due to the arbitrary number of strategies, there is a possibly exponential number of load profiles. We make sure that s is a NE, then we only have to check one load profile – namely $l(s)$ – to verify that no coalitional deviation exists. In this way, we can efficiently check whether a state is a SE or PoNE.

For the task of computing a SE, we note that there are efficient algorithms to compute a NE in these games [19]. This allows us to obtain a NE s and load profile $l(s)$ in polynomial time. To compute a SE, we construct a bipartite deviation graph G for state s and target profile $l(s)$ as in the proof of Theorem 1. Here we also add costs to the edges, and let the cost of edge $(v_i, v_{e,j})$ be $c_i(l_e(s))$, for all $1 \leq j \leq l_e(s)$. Now consider any other state s' with $l(s)$, in which $c_i(s) = c_i(s')$ for every player i with $s_i = s'_i$ and $c_i(s) > c_i(s')$ for every i with $s_i \neq s'_i$. For every such state we can find a corresponding perfect matching in G . In particular, we construct a *minimum cost* perfect matching. This matching yields a state s' and we now argue that s' is indeed a SE.

Suppose for contradiction that there is a coalitional deviation from s' to a state s'' . s is a NE and $c_i(s) \geq c_i(s') \geq c_i(s'')$ for every $i \in N$, with at least one inequality for a moving player. s'' must also have load profile $l(s)$, and a deviation from s' is a circular switch of players. This switch does not increase the cost of any player but decreases the cost of the moving players. Therefore, the assignment s'' is such that $c_i(s) = c_i(s'')$ for every player i with $s_i = s''_i$ and $c_i(s) > c_i(s'')$ for every i with $s_i \neq s''_i$. Note that s'' corresponds to a perfect matching in G , and the sum of costs $\sum_{i \in N} c_i(s'') < \sum_{i \in N} c_i(s')$. This is a contradiction to s' being derived from a minimum cost perfect matching in G . \square

Interestingly, our proof shows that for every NE there is a single coalitional deviation that turns the state into a SE. Milchtaich [19] proved that from every state s there is a sequence of unilateral deviations with length at most $|E| \cdot \binom{n+1}{2}$ that leads to a NE. Our result implies that even SE can be reached via short sequences of improvement moves from every state of the game. In these sequences we only need one coalitional move which is efficiently computable. A similar result can be derived for PoNE, where we adjust the deviation graph to allow only coalitional improvement moves where all players strictly improve.

Corollary 2. *For every state s of a player-specific singleton congestion game with non-decreasing cost functions there is a sequence of coalitional improvement moves that leads to a strong equilibrium. Each move can be computed in polynomial time. The length of the sequence is at most $|E| \cdot \binom{n+1}{2} + 1$.*

5 Congestion Games

In this section, we consider the complexity of computing SE and PoNE in general congestion games. The class of singleton congestion games is a special case of the games we treated in the previous section, and for which we could establish a variety of positive results. Here we extend the combinatorial structure of strategy




Fig. 2. Construction that proves hardness of the existence and recognition problems of SE in congestion games

spaces only slightly to matroids. This allows to obtain a set of quite strong hardness results concerning the existence and recognition of SE and PoNE. Note that all our results in this section hold even for symmetric games, in which strategy spaces are simultaneously matroids and networks.

Theorem 4. *It is strongly Co-NP-hard to decide (1) if a congestion game has a strong equilibrium and (2) if a given state of a game is a strong equilibrium.*

Proof. We reduce from 3-PARTITION. An instance is given by a multiset of integers a_1, \dots, a_{3m} . Let $b = \frac{1}{m} \sum_{i=1}^{3m} a_i$. An instance $I = (a_1, \dots, a_{3m}) \in 3$ -PARTITION if and only if there exists a partition of $A = \{1, \dots, 3m\}$ into m subsets A_1, \dots, A_m such that the sum $\sum_{i \in A_j} a_i = b$ for all $1 \leq j \leq m$. Without loss of generality, we can assume that every integer $b/2 > a_i > b/4$. Therefore, each subset A_i is forced to consist of exactly three elements.

Given an instance I , we construct a symmetric matroid network congestion game Γ_I as follows. The network is $G = (V, E)$ with vertices $V = \{s, v_1, \dots, v_{3m}, t\}$ and a series of parallel edges as depicted in Figure 2. There are $m + 1$ players. Each player's source node is s and his target node is t . The delay functions are defined as follows. Let $M = 2B$ and $1 > \epsilon > 0$. The delay of an edge a_i^- is $M - a_i$ for one player and M for more than one player. Delay of an edge a_i^+ is always $M + a_i$. The delay of an edge a_i^0 is M for at most $m - 1$ players and $2M$ for m or more players. Delay of edge b^- is $M - b - \epsilon$ for at most m players and M for more than m players. Delay of an edge b^+ is always $M + b - \epsilon$.

If $I \in 3$ -PARTITION, we show that no SE exists. Observe that for a single agent it is never optimal to choose one of the edges a_i^+ or b^+ . Thus, no SE exists in which these edges are used. Thus, in every SE every player has delay of (at least) $(m + 1)M$. However, there is a joint deviation of all players which yields delay of $(m + 1)M - \epsilon$ for each of them. Let A_1, \dots, A_m be a solution of the 3-PARTITION-instance I . Player $m + 1$ chooses edge b^+ and edges a_1^-, \dots, a_{3m}^- . Each player $1 \leq j \leq m$ chooses edge b^- and the following edges: For each $1 \leq i \leq 3m$, if $i \in A_j$ player j plays edge a_i^+ otherwise he plays a_i^0 . As argued above, the resulting state is not a SE either. Thus, no SE exists.

If $I \notin 3$ -PARTITION, all players choosing path $b^-, a_1^-, a_2^-, \dots, a_{3m}^-$ is a SE. Details of this argument will be given in the full version of this paper.

Obviously, the above implies that s with all players choosing path $b^-, a_1^-, a_2^-, \dots, a_{3m}^-$ is a SE if and only if $I \notin 3\text{-PARTITION}$. This proves Co-NP-hardness of deciding whether a given state is a SE. \square

Theorem 5. *It is weakly Co-NP-hard to decide (1) if a congestion game with two players has a strong equilibrium and (2) if a given state of a game is a strong equilibrium.*

This implies the same result for PoNE, as for two players SE and PoNE coincide. Additionally, it implies the result for k -SE, for any $k \geq 2$.

Corollary 3. *It is weakly Co-NP-hard to decide for a congestion game (1) if it has a k -strong equilibrium, (2) if it has a Pareto-optimal Nash equilibrium, (3) if a given state is a k -strong equilibrium, (4) if a given state is a Pareto-optimal Nash equilibrium, for any $k \geq 2$.*

References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. *J. ACM* 55(6) (2008)
2. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. *Games Econ. Behav.* 65(2), 289–317 (2009)
3. Aumann, R.: Acceptable points in general cooperative n -person games. In: *Contributions to the Theory of Games IV*. *Annals of Mathematics Study*, vol. 40, pp. 287–324 (1959)
4. Blonski, M.: Anonymous games with binary actions. *Games Econ. Behav.* 28, 171–180 (1999)
5. Blonski, M.: Characterization of pure strategy equilibria in finite anonymous games. *J. Math. Econ.* 34(2), 225–233 (2000)
6. Brandt, F., Fischer, F., Holzer, M.: Symmetries and the complexity of pure Nash equilibrium. *J. Comput. Syst. Sci.* 75(3), 163–177 (2009)
7. Chien, S., Sinclair, A.: Strong and Pareto price of anarchy in congestion games. In: *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, pp. 279–291 (2009)
8. Daskalakis, C., Papadimitriou, C.: Computing equilibria in anonymous games. In: *Proc. 48th Symp. Foundations of Computer Science (FOCS)*, pp. 83–93 (2007)
9. Daskalakis, C., Papadimitriou, C.: Discretized multinomial distributions and Nash equilibria in anonymous games. In: *Proc. 49th Symp. Foundations of Computer Science (FOCS)*, pp. 25–34 (2008)
10. Dunkel, J., Schulz, A.: On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Math. Oper. Res.* 33(4), 851–868 (2008)
11. Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure Nash equilibria. In: *Proc. 36th Symp. Theory of Computing (STOC)*, pp. 604–612 (2004)
12. Gottlob, G., Greco, G., Scarcello, F.: Pure Nash equilibria: Hard and easy games. *J. Artif. Intell. Res.* 24, 195–220 (2005)
13. Harks, T., Hofer, M., Klimm, M., Skopalik, A.: Computing pure Nash and strong equilibria in bottleneck congestion games. In: *Proc. 18th European Symposium on Algorithms, ESA (to appear 2010)*

14. Harks, T., Klimm, M., Möhring, R.H.: Strong Nash equilibria in games with the lexicographical improvement property. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 463–470. Springer, Heidelberg (2009)
15. Holzman, R., Law-Yone, N.: Strong equilibrium in congestion games. *Games Econ. Behav.* 21(1-2), 85–101 (1997)
16. Konishi, H., Breton, M.L., Weber, S.: Equilibria in a model with partial rivalry. *J. Econ. Theory* 72(1), 225–237 (1997)
17. Mas-Colell, A., Whinston, M., Green, J.: *Microeconomic Theory*. Oxford University Press, Oxford (1995)
18. Meyers, C.: *Network Flow Problems and Congestion Games: Complexity and Approximation Results*. PhD thesis, MIT (June 2006)
19. Milchtaich, I.: Congestion games with player-specific payoff functions. *Games Econ. Behav.* 13(1), 111–124 (1996)
20. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* 2, 65–67 (1973)
21. 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)

2-Player Nash and Nonsymmetric Bargaining Games: Algorithms and Structural Properties

Vijay V. Vazirani

Georgia Institute of Technology, Atlanta, GA 30332-0280
vazirani@cc.gatech.edu

Abstract. The solution to a Nash or a nonsymmetric bargaining game is obtained by maximizing a concave function over a convex set, i.e., it is the solution to a convex program. We show that each 2-player game whose convex program has linear constraints, admits a rational solution and such a solution can be found in polynomial time using only an LP solver. If in addition, the game is succinct, i.e., the coefficients in its convex program are “small”, then its solution can be found in strongly polynomial time. We also give non-succinct linear games whose solution can be found in strongly polynomial time.

1 Introduction

In game theory, 2-player games occupy a special place – not only because numerous applications involve two players but also because 2-player games often have remarkable properties that are not possessed by extensions to more players.

For instance, in the case of Nash equilibrium, the 2-player case is the most extensively studied and used, and captures a rich set of possibilities, e.g., those encapsulated in canonical games such as prisoner’s dilemma, battle of the sexes, chicken, and matching pennies. In terms of properties, 2-player Nash equilibrium games always have rational solutions whereas games with three or more players may have only irrational solutions; an example of the latter, called “a three-man poker game,” was given by Nash [14].

From a computational viewpoint, the difference is even more stark. The problem of finding a Nash equilibrium is PPAD-complete for two players [4] and FIXP-complete for three or more players [6]. Note that whereas PPAD is in $NP \cap \text{co-NP}$, the only fact known about FIXP is that $P \subseteq \text{FIXP} \subseteq \text{PSPACE}$. Next, let us restrict to zero-sum games. For two players, von Neumann’s min-max theorem yields a polynomial time algorithm using LP [15]. On the other hand, 3-player zero-sum games are PPAD-hard, since any 2-player non-zero-sum game can be reduced to a 3-player zero-sum game [15].

John Nash’s seminal paper defining the bargaining game dealt only with the case of 2-players [13]. Later, it was observed that his entire setup, and theorem characterizing the bargaining solution, easily generalize to the case of more than 2 players, e.g., see [10]. Today, Nash bargaining is regarded as a central solution concept within game theory for “fair” allocation of utility among competing players in the presence of complete information, e.g., see [11, 18, 16].

Recently, Vazirani [20] initiated a systematic algorithmic study of Nash bargaining games and also carried this program over to solving nonsymmetric bargaining games of Kalai [10]. In this paper we carry the program further, though only for the case of 2-player games. The solution to a Nash or a nonsymmetric bargaining game is obtained by maximizing a concave function over a convex set, i.e., it is the solution to a convex program. As defined in [20], NB is the class of Nash and nonsymmetric bargaining games that can be solved in polynomial time, and LNB is the subclass of NB consisting of games all of whose constraints are linear. Let NB2 and LNB2 be the restrictions of these classes to 2-player games.

We show that for solving any game in LNB2, it is not essential to solve a convex program – an LP solver suffices. As a consequence, all games in LNB2 have rational solutions; this property does not hold for 3-player games in LNB. Following [20], let us say that a nonlinear convex program is *rational* if it always has a rational solution that can be written using polynomially many bits, if all parameters are rational. Rational convex programs form a surprisingly rich class, see [20, 5, 21, 9, 3, 1, 7]. Next, we define a subclass of LNB2 called SLNB2, consisting of *succinct* games, see Section 4. We show that all games in SLNB2 admit strongly polynomial algorithms; however, these algorithms are not combinatorial. This class includes nontrivial and interesting games, e.g., the game **DG2**, which consists of a directed graph with edge capacities and each player is a source-sink pair desiring flow (see Section 4 for definition). This game is derived from Kelly’s *flow markets* [12].

This raises the question of whether there are games in (LNB2 - SLNB2) that admit strongly polynomial time algorithms. The answer turns out to be “yes”. In the full paper, we will show that the 2-player version of the game **ADNB**, for which a combinatorial polynomial time algorithm is given in [20], admits a combinatorial strongly polynomial algorithm. The game **ADNB** was derived from the linear case of the Arrow-Debreu market model.

Building on the Eisenberg-Gale program, [9] gave the notion of *Eisenberg-Gale markets*, see Section 3. In answering an open question of [9] affirmatively, [3] showed that EG(2) markets, i.e., the restriction of Eisenberg-Gale markets to 2 buyers, always admit a rational solution and it can be found using only an LP solver. Our result is obtained by extending their algorithm. The exact relationship of EG(2) and LNB2 is explained in Section 3.

Finally, the class (NB2 - LNB2) needs to be properly explored and understood, both structurally and algorithmically. As an example of interesting games it may contain, we introduce the circle game. This game can be solved in polynomial time, even without a convex program solver; its solution reduces to solving a degree 4 equation. Alternatively, it also admits an elegant geometric solution.

2 Nash and Nonsymmetric Bargaining Games

An n -person Nash bargaining game consists of a pair $(\mathcal{N}, \mathbf{c})$, where $\mathcal{N} \subseteq \mathbf{R}_+^n$ is a compact, convex set and $\mathbf{c} \in \mathcal{N}$. Set \mathcal{N} is the *feasible set* and its elements give utilities that the n players can simultaneously accrue. Point \mathbf{c} is the *disagreement*

point – it gives the utilities that the n players obtain if they decide not to cooperate. The set of n agent will be denoted by B and the agents will be numbered $1, 2, \dots, n$. Game $(\mathcal{N}, \mathbf{c})$ is said to be *feasible* if there is a point $\mathbf{v} \in \mathcal{N}$ such that $\forall i \in B, v_i > c_i$. The solution to a feasible game is the point $\mathbf{v} \in \mathcal{N}$ that satisfies the four axioms: Pareto optimality, Invariance under affine transformations, symmetry, and independence of irrelevant alternatives.

Theorem 1. Nash [13] *If game $(\mathcal{N}, \mathbf{c})$ is feasible then there is a unique point in \mathcal{N} satisfying the axioms stated above. This is also the unique point that maximizes $\prod_{i \in B} (v_i - c_i)$ over all $\mathbf{v} \in \mathcal{N}$.*

Since the log of $\prod_{i \in B} (v_i - c_i)$, i.e., $\sum_{i \in B} \log(v_i - c_i)$, is a concave function, Nash’s solution involves maximizing a concave function over a convex domain, and is therefore the optimal solution to the convex program that maximizes $\sum_{i \in B} \log(v_i - c_i)$ subject to $\mathbf{v} \in \mathcal{N}$. Therefore, if for a specific game, a separation oracle can be implemented in polynomial time, then using the ellipsoid algorithm one can get as good an approximation to the solution as desired [8].

Kalai [10] generalized Nash’s bargaining game by removing the axiom of symmetry and showed that any solution to the resulting game is the unique point that maximizes $\prod_{i \in B} (v_i - c_i)^{p_i}$, over all $\mathbf{v} \in \mathcal{N}$, for some choice of positive numbers p_i , for $i \in B$, such that $\sum_{i \in B} p_i = 1$. Thus, any particular nonsymmetric bargaining solution is specified by giving the p_i ’s satisfying the two conditions, i.e., $\forall i \in B, p_i > 0$ and $\sum_{i \in B} p_i = 1$. For the purposes of computability, we will restrict to nonsymmetric games in which the p_i ’s are rational. Equivalently, let us define the *n -person nonsymmetric bargaining game* as follows. Assume that $B, \mathcal{N}, \mathbf{c}$ are as defined above. In addition, we are given the *clout* of each player: a positive integer w_i for each player i .

Assuming the game is feasible, the solution to this nonsymmetric bargaining game is the unique point that maximizes $\prod_{i \in B} (v_i - c_i)^{w_i}$ over all $\mathbf{v} \in \mathcal{N}$. As before, we will view this as the solution to a convex program by maximizing $\sum_{i \in B} w_i \log(v_i - c_i)$ over all $\mathbf{v} \in \mathcal{N}$. One more remark is in order. As shown by Kalai [10], any nonsymmetric game can be reduced to a Nash bargaining game over a larger number of players. However, this reduction is not useful for our purpose because once the number of players increases, the special properties of 2-player games are lost.

3 Fisher and Eisenberg-Gale Market Models

We will first state Fisher’s market model for the case of linear utility functions [2]. Consider a market consisting of a set of n buyers $B = \{1, 2, \dots, n\}$, and a set of g divisible goods, $G = \{1, 2, \dots, g\}$; we may assume w.l.o.g. that there is a unit amount of each good. Let m_i be the money possessed by buyer $i, i \in B$; w.l.o.g. assume that each $m_i > 0$. Let u_{ij} be the utility derived by buyer i on receiving one unit of good j . Thus, if x_{ij} is the amount of good j that buyer i gets, for $1 \leq j \leq g$, then the total utility derived by i is $v_i(x) = \sum_{j=1}^g u_{ij}x_{ij}$.

The problem is to find prices $\mathbf{p} = \{p_1, p_2, \dots, p_g\}$ for the goods so that when each buyer is given her utility maximizing bundle of goods, the market clears, i.e., each good having a positive price is exactly sold, without there being any deficiency or surplus. Such prices are called *equilibrium prices*.

The following is the Eisenberg-Gale convex program. Using KKT conditions, one can show that its optimal solution is an equilibrium allocation for Fisher's linear market and the Lagrange variables corresponding to the inequalities give equilibrium prices of goods (e.g., see Theorem 5.1 in [19]).

$$\begin{aligned}
 &\text{maximize} && \sum_{i \in B} m_i \log v_i && (1) \\
 &\text{subject to} && \forall i \in B : v_i = \sum_{j \in G} u_{ij} x_{ij} \\
 &&& \forall j \in G : \sum_{i \in B} x_{ij} \leq 1 \\
 &&& \forall i \in B, \forall j \in G : x_{ij} \geq 0
 \end{aligned}$$

Next, we state the definition of Eisenberg-Gale markets as given in [9]. Let us say that a convex program is an *Eisenberg-Gale-type* convex program if its objective function is of the form $\max \sum_i m_i \log v_i$, subject to linear packing constraints, i.e., constraints of the form $\mathbf{Ax} \leq \mathbf{b}$, where matrix \mathbf{A} and vector \mathbf{b} are non-negative and the vector of variables \mathbf{x} is constrained to be non-negative. Let \mathcal{M} be a Fisher market, with an arbitrary utility function, whose set of feasible allocations and buyers' utilities form a polytope Π . We will assume that the linear constraints defining Π are packing constraints. As a result, \mathcal{M} satisfies the *free disposal property* or is *downward closed*, i.e., if \mathbf{v} is a feasible utility vector then so is any vector dominated by \mathbf{v} . We will say that an allocation x_1, \dots, x_n made to the buyers is a *clearing allocation* if it uses up all goods exactly to the extent they are available in \mathcal{M} . Finally, we will say that \mathcal{M} is an *Eisenberg-Gale market* if any clearing allocation x_1, \dots, x_n that maximizes $\max \sum_i m_i \log v_i(x_i)$ is an equilibrium allocation, i.e., there are prices p_1, \dots, p_g for the goods such that for each buyer i , x_i is a utility maximizing bundle for i at these prices. The class EG(2), defined in [3], is essentially the restriction of Eisenberg-Gale markets to the case of 2 buyers; see [3] for the precise definition. Finally, we point out the relationship between EG(2) and LNB2. If a game \mathcal{G} in LNB2 is such that its feasible set, \mathcal{N} , is downward closed, then the instances of this game in which the disagreement utilities are zero form an EG(2) market.

4 The Class LNB2 and Some Basic Procedures

The classes NB2 and LNB2 were defined in the Introduction. We will assume w.l.o.g. that the convex program for a game \mathcal{G} in LNB2 has the following form:

$$\begin{aligned}
 &\text{maximize} && \sum_{i=1,2} w_i \log(v_i - c_i) && (2) \\
 &\text{subject to} && \mathbf{Ax} + \mathbf{b}_1 v_1 + \mathbf{b}_2 v_2 \leq \mathbf{e} \\
 &&& \text{for } i = 1, 2 : \quad v_i \geq c_i \\
 &&& \mathbf{x} \geq 0
 \end{aligned}$$

where \mathbf{A} is an $m \times r$ matrix, \mathbf{x} is a vector consisting of a total of r allocation and auxiliary variables, and $\mathbf{b}_1, \mathbf{b}_2, \mathbf{e}$ are m -dimensional vectors. We will say that \mathcal{G} is *succinct* if all the entries in $\mathbf{A}, \mathbf{b}_1, \mathbf{b}_2$ are polynomially bounded in m and r ; if so, \mathcal{G} lies in SLNB2.

As an example of a game that lies in SLNB2, consider **DG2** : We are given a directed graph $G = (V, E)$, with $c_e \in \mathbf{Q}^+$ specifying the capacity of edge $e \in E$. Two source-sink pairs are also specified, (s_1, t_1) and (s_2, t_2) . Each source-sink pair represents a player in the game and has its own disagreement utility (flow value) c_i , for $i = 1, 2$; c_i can be thought of as a strict lower bound on the amount of flow player i desires (perhaps because of the resources player i has invested in building the network). In the nonsymmetric version, we are also given the clouts w_1 and w_2 of the two players. The object is to find the Nash or nonsymmetric bargaining solution. Let \mathcal{G} denote the given instance of **DG2**.

Next, we give a convex program that captures the solution to \mathcal{G} . The flow going from s_i to t_i will be referred to as commodity i , for $i = 1, 2$, and f_i will denote the total flow of commodity i . For each edge $e \in E$, we have 2 variables, f_e^1 and f_e^2 which denote the amount of each commodity flowing through e . The constraints ensure that the total flow going through an edge does not exceed its capacity and that for each commodity, at each vertex, other than the source-sink pair of this commodity, flow conservation holds. For vertex $v \in V$, $\text{out}(v) = \{(v, u) \mid (v, u) \in E\}$ and $\text{in}(v) = \{(u, v) \mid (u, v) \in E\}$. The constraints of this program are simply ensuring that (f_1, f_2) lies in the feasible set \mathcal{N} .

$$\begin{aligned}
 &\text{maximize} && \sum_{i=1,2} w_i \log(f_i - c_i) && (3) \\
 &\text{subject to} && \text{for } i = 1, 2 : \quad f_i = \sum_{e \in \text{Out}(s_i)} f_e^i \\
 &&& \forall e \in E : \quad f_e^1 + f_e^2 \leq c_e \\
 &&& \text{for } i = 1, 2 : \quad \forall v \in V - \{s_i, t_i\} : \quad \sum_{e \in \text{in}(v)} f_e^i = \sum_{e \in \text{out}(v)} f_e^i \\
 &&& \text{for } i = 1, 2 : \quad \forall e \in E : \quad f_e^i \geq 0
 \end{aligned}$$

Next, assume that \mathcal{G} is a generic game that lies in LNB2 We can test if \mathcal{G} is feasible by solving the following LP:

$$\begin{aligned}
 &\text{maximize} && t && (4) \\
 &\text{subject to} && v_1 \geq c_1 + t \\
 &&& v_2 \geq c_2 + t \\
 &\text{subject to} && \mathbf{Ax} + \mathbf{b}_1 v_1 + \mathbf{b}_2 v_2 \leq \mathbf{e} \\
 &&& \mathbf{x} \geq 0
 \end{aligned}$$

Now, \mathcal{G} is feasible iff the optimal value of $t > 0$. Henceforth, assume that \mathcal{G} is feasible. Next, we make the following change of variables, for $i = 1, 2$: $y_i = v_i - c_i$, hence obtaining the following program which is equivalent to (2).

$$\begin{aligned}
 &\text{maximize} && \sum_{i=1,2} w_i \log y_i && (5) \\
 &\text{subject to} && \mathbf{Ax} + \mathbf{b}_1(y_1 + c_1) + \mathbf{b}_2(y_2 + c_2) \leq \mathbf{e} \\
 &&& \text{for } i = 1, 2 : \quad y_i \geq 0 \\
 &&& \mathbf{x} \geq 0
 \end{aligned}$$


Henceforth, we will denote $(\mathbf{e} - c_1\mathbf{b}_1 - c_2\mathbf{b}_2)$ by \mathbf{d} . We will denote by Π the polyhedron in \mathbf{R}^{n+2} which is defined by the constraints of program (5). In this paper, we will write the constraints of (5) concisely as follows. This notation will also be used for LP’s optimizing over the polytope Π .

$$\begin{aligned}
 &\text{maximize} && \sum_{i=1,2} w_i \log y_i && (6) \\
 &\text{subject to} && (\mathbf{x}, y_1, y_2) \in \Pi
 \end{aligned}$$

The projection of Π onto the coordinates y_1, y_2 gives a polytope, \mathcal{N} in \mathbf{R}^2 , which we will call the *feasible polytope*. In this section, we will describe its *useful faces*, i.e., faces on which the solution to game \mathcal{G} can lie, and we will give some basic procedures for operating on these faces. We first compute the point (l_1, l_2) by first maximizing y_1 over Π to get l_1 and then maximizing y_2 over Π , subject to $y_1 = l_1$, to get l_2 . Similarly, compute the point (h_1, h_2) by first maximizing y_2 over Π to get h_2 and then maximizing y_1 over Π , subject to $y_2 = h_2$, to get h_1 . Clearly, both these points are vertices of \mathcal{N} . Finally, the set of faces encountered in moving, on the boundary of \mathcal{N} , from (l_1, l_2) to (h_1, h_2) , by increasing the second coordinate are the useful faces. If polytope \mathcal{N} is not full dimensional, we will already obtain the vertex or facet on this the solution lies. For the rest of this section, assume that \mathcal{N} is full dimensional.

Each of the useful facets has the form $y_1 + \alpha y_2 \leq \beta$, where $\alpha > 0$ and $\beta > 0$. We will denote the vertex at the intersection of the two facets $y_1 + \alpha_1 y_2 \leq \beta_1$ and $y_1 + \alpha_2 y_2 \leq \beta_2$, by (α_1, α_2) ; we will assume $\alpha_1 < \alpha_2$. Let α^1 and β^1 (α^2 and β^2) be the α and β values of the first (last) facet encountered in moving from (l_1, l_2) to (h_1, h_2) ; clearly, $\alpha^1 < \alpha^2$. Our binary search will be performed

on the interval $[\alpha^1, \alpha^2]$. In Procedure 3 below, we show how to compute α^1 and α^2 . The solution to \mathcal{G} must lie on a face which is either a useful facet or a vertex at the intersection of 2 useful facets. These 2 possibilities give rise to distinct procedures and proofs throughout.



4.1 Procedure 1: Given α , Find the Face It Lies on

We give an algorithm for the following task: Given a number α s.t. $\alpha^1 \leq \alpha \leq \alpha^2$, determine which of the following possibilities holds: **1.** α defines a facet of \mathcal{N} , $y_1 + \alpha y_2 \leq \beta$, for a suitable value of β . **2.** There is a vertex of \mathcal{N} , (α_1, α_2) , such that $\alpha_1 < \alpha < \alpha_2$. If so, find this vertex.

Let β be the optimal objective function value of the following LP and let a and b denote the optimal values of y_1 and y_2 , respectively.

$$\begin{aligned} &\text{maximize} && y_1 + \alpha y_2 && (7) \\ &\text{subject to} && (\mathbf{x}, y_1, y_2) \in \Pi \end{aligned}$$

Next, solve the following LP and let its value be denoted by a_1 .

$$\begin{aligned} &\text{minimize} && y_1 && (8) \\ &\text{subject to} && y_1 + \alpha y_2 = \beta \\ &&& (\mathbf{x}, y_1, y_2) \in \Pi \end{aligned}$$

Next, change the objective in LP (8) to maximize y_1 , and let its optimal objective function value be a_2 . If $a_1 < a_2$, we are in the first case. Define $b_1 = (\beta - a_1)/\alpha$ and $b_2 = (\beta - a_2)/\alpha$. Then, the endpoints of the facet $y_1 + \alpha y_2 = \beta$ are (a_1, b_1) and (a_2, b_2) . Otherwise, $a_1 = a_2 = a$, say, and we are in the second case.

Let b be the value of y_2 computed in LP (8). Then, the vertex has coordinates (a, b) . Next, we need to find α_1 and α_2 for this vertex. Let us begin by writing the dual for LP (7).

$$\begin{aligned}
 &\text{minimize} && \sum_j d_j p_j && (9) \\
 &\text{subject to} && \sum_j b_{1j} p_j \geq 1 \\
 &&& \sum_j b_{2j} p_j \geq \alpha \\
 &&& \text{for } 1 \leq i \leq n : \sum_j A_{ji} p_j \geq 0 \\
 &&& \text{for } 1 \leq j \leq m : p_j \geq 0
 \end{aligned}$$

Let $(\mathbf{x}^*, y_1^*, y_2^*)$ be an optimal solution to LP (7). Since \mathcal{G} has been assumed to be feasible, $y_1^* > 0$ and $y_2^* > 0$. The next LP is derived from LP (9) by adding constraints on p_j which are implied by the complementary slackness conditions of the primal and dual pair of LP's (7) and (9). It is not optimizing any function, since we are only concerned with its feasible solutions.

$$\begin{aligned}
 &\sum_j b_{1j} p_j = 1 && (10) \\
 &\sum_j b_{2j} p_j = r \\
 &\text{for } 1 \leq i \leq n : \sum_j A_{ji} p_j \geq 0 \\
 &\text{for } 1 \leq i \leq n \text{ s.t. } x_i^* > 0 : \sum_j A_{ji} p_j = 0 \\
 &\text{for } 1 \leq j \leq m \text{ s.t. } \sum_i A_{ji} x_i^* + b_{ij} y_1^* + b_{2j} y_2^* < d_j : p_j = 0 \\
 &\text{for } 1 \leq j \leq m : p_j \geq 0
 \end{aligned}$$

The next lemma follows from the complementary slackness conditions of the primal and dual pair of LP's (7) and (9).

Lemma 1. $\{\alpha \mid \text{LP (7) attains its optimal solution at } (a, b)\}$
 $= \{r \mid \exists \text{ a feasible solution to LP (10) in which } \sum_j b_{2j} p_j = r\}$.

By Lemma 1, we can obtain α_1 and α_2 as follows. First, minimize r subject to the constraints of LP (10); this gives α_1 . Next, maximize r subject to the constraints of LP (10); this gives α_2 .

4.2 Procedure 2: Given (a, b) , Find the Face It Lies on

Given a point (a, b) on the boundary of \mathcal{N} , we give a procedure for finding the facet or vertex it lies on. First, solve LP (11):

$$\begin{aligned} y_1 &= a \\ y_2 &= b \\ (\mathbf{x}, y_1, y_2) &\in \Pi \end{aligned} \tag{11}$$

Next, solve the minimization and maximization versions, with objective function r , of LP (10) to find α_1 and α_2 , respectively. If $\alpha_1 = \alpha_2 = \alpha$, (a, b) lies on the facet $y_1 + \alpha y_2 \leq a + \alpha b$. Otherwise, $\alpha_1 < \alpha_2$ and (a, b) lies on the vertex (α_1, α_2) .

4.3 Procedure 3: Computing α^1 and α^2

We now show how to compute α^1 and α^2 , defined at the beginning of this section. As stated there, our binary search will be performed on the interval $[\alpha^1, \alpha^2]$.

First, use Procedure 2 to find the vertex, say (α_1, α_2) , on which (l_1, l_2) lies. Set, $\alpha^1 \leftarrow \alpha_1$. Next, use Procedure 2 to find the vertex, say (α_1, α_2) , on which (h_1, h_2) lies. Set, $\alpha^2 \leftarrow \alpha_2$.

5 Binary Search on Parameter z

We first give some crucial definitions. Let (f_1, f_2) be the solution to game \mathcal{G} . For player i define $\gamma_i = \frac{f_i}{w_i}$. Define parameter z to be $z = \frac{\gamma_1}{\gamma_2}$. The next lemma relates z to the point where the solution lies.

Lemma 2. *If the solution to game \mathcal{G} lies on the facet $y_1 + \alpha y_2 \leq \beta$, then $z = \alpha$, and if it lies on the vertex (α_1, α_2) , then $\alpha_1 < z < \alpha_2$.*

Proof. In the first case, the objective function of the convex program (3), $g = w_1 \log y_1 + w_2 \log y_2$ must be tangent to the facet at the solution point, say (a, b) . Equating the ratio of the partial derivatives of g and the line $y_1 + \alpha y_2 = \beta$ w.r.t. y_2 and y_1 , we get $\frac{a/w_1}{b/w_2} = \alpha$. But the l.h.s. is $\gamma_1/\gamma_2 = z$, thereby giving $z = \alpha$. In the second case, the tangent to g at the solution must be intermediate between the slopes of the adjacent facets, giving $\alpha_1 < z < \alpha_2$.

Our algorithm will conduct a binary search on z , on the interval $[\alpha^1, \alpha^2]$, to find the right face on which the solution lies. The test given in the next lemma helps determine, in each iteration, if the current face is the right one.

Algorithm 2 (Binary Search)

1. **(Initialization:)** $l \leftarrow \alpha^1$ and $h \leftarrow \alpha^2$.

$$\text{Let } r \leftarrow \frac{w_1}{w_1 + w_2}.$$

2. $\alpha \leftarrow \lfloor \frac{l+h}{2} \rfloor_{\kappa}$.

3. Using Procedure 1 (Section 4.1), determine if α lies on:

Case 1: A facet, say $y_1 + \alpha y_2 \leq \beta$, with endpoints (a_1, b_1) and (a_2, b_2) .

If $r < (a_2/\beta)$ then $l \leftarrow \alpha$ and go to step 2

Else if $r > (a_1/\beta)$ then $h \leftarrow \alpha$ and go to step 2

Else if $r \in \left[\frac{a_1}{\beta}, \frac{a_2}{\beta} \right]$, then solve for y_1 and y_2 :

$$y_1 + \alpha y_2 = \beta \quad \text{and} \quad \frac{y_1/w_1}{y_2/w_2} = \alpha.$$

If the solution is $y_1 = a, y_2 = b$, output the solution to game \mathcal{G} :

$v_1 = a + c_1$ and $v_2 = b + c_2$, and HALT.

Case 2: A vertex, say (α_1, α_2) , with coordinates (a, b) ,

If $r \leq (a/\beta_2)$ then $l \leftarrow \alpha_2$ and go to step 2

Else if $r \geq (a/\beta_1)$ then $h \leftarrow \alpha_1$ and go to step 2

Else if $r \in \left(\frac{a}{\beta_2}, \frac{a}{\beta_1} \right)$,

then output the solution to game \mathcal{G} : $v_1 = a + c_1$ and $v_2 = b + c_2$, and HALT.

4. End.

Our algorithm will conduct a binary search on z , on the interval $[\alpha^1, \alpha^2]$, to find the right face on which the solution lies. The test given in the next lemma helps determine, in each iteration, if the current face is the right one.

Lemma 3. 1. The solution to game \mathcal{G} lies on the facet $y_1 + \alpha y_2 \leq \beta$, having endpoints (a_1, b_1) and (a_2, b_2) , with $a_1 < a_2$, iff $\frac{w_1}{w_1+w_2} \in \left[\frac{a_1}{\beta}, \frac{a_2}{\beta} \right]$.

2. The solution to game \mathcal{G} lies on the vertex (α_1, α_2) , having coordinates (a, b) , which is at the intersection of facets $y_1 + \alpha_1 y_2 \leq \beta_1$ and $y_1 + \alpha_2 y_2 \leq \beta_2$, iff $\frac{w_1}{w_1+w_2} \in \left(\frac{a}{\beta_2}, \frac{a}{\beta_1} \right)$.

Proof. In the first case, substituting $w_i = f_i/\gamma_i$ and $\alpha = \gamma_1/\gamma_2$ (this follows from Lemma 2), we get $\frac{w_1}{w_1+w_2} = \frac{f_1}{f_1+\alpha f_2} = \frac{f_1}{\beta} \in \left[\frac{a_1}{\beta}, \frac{a_2}{\beta}\right]$. For the other direction, if $w_1/(w_1 + w_2)$ lies in the interval given, then by the equation given above, $f_1 \in [a_1, a_2]$, thereby showing that the solution lies on the facet $y_1 + \alpha y_2 \leq \beta$. In the second case, by Lemma 2, $\gamma_1/\gamma_2 \in (\alpha_1, \alpha_2)$, and this leads to the interval in which $w_1/(w_1 + w_2)$ lies. The proof of the other direction also follows in the same manner.

Theorem 3. *Every game in LNB2 has a rational solution; moreover, such a solution can be found in polynomial time using only an LP solver.*

Next assume that the coefficients in the constraints of convex program (2) are “small”, i.e., polynomially bounded in n . Then all LP’s that need to be solved will also have “small” coefficients (the objective function and right hand side don’t need to be “small”). By [17] we get:

Theorem 4. *Every game in SLNB2 can be solved in strongly polynomial time.*

In particular, the game **DG2**, which lies in SLNB2, can be solved in strongly polynomial time. [9] give examples of Eisenberg-Gale markets with 3 buyers which do not have rational solutions. In particular, let **DG3** be the extension of **DG2** to three players, with 3 source-sink pairs. Consider instances of **DG3** in which each player’s disagreement utility is zero. As stated in Section 3, these instances correspond to an Eisenberg-Gale market. [9] show that this market, with 3 buyers, does not have rational solutions. Hence the game **DG3**, which is in (LNB - LNB2), does not admit a rational convex program.

6 The Circle Game

The circle game lies in (NB2 - LNB2). Its feasible set is the intersection of the unit disk with the positive orthant. We will consider only its Nash bargaining version. Its convex program is:

$$\begin{aligned}
 &\text{maximize} && \sum_{i=1,2} \log(v_i - c_i) && (12) \\
 &\text{subject to} && v_1^2 + v_2^2 \leq 1 \\
 &&& \forall i = 1, 2 : v_i \geq 0
 \end{aligned}$$

Using the KKT conditions for this program, it is easy to see that the Nash bargaining solution (x, y) satisfies the following equations: $(2y^2 - c_2y - 1)^2 = c_1^2(1 - y^2)$ and $x^2 + y^2 = 1$. On the other hand, the problem also has a simple geometric solution. Let Q be a point on the unit circle in the positive orthant. Let O denote the origin and P denote the point (c_1, c_2) . Let θ_1 be the angle made by PQ with the x -axis and θ_2 be the angle made by OQ with the y -axis. Then one can show that Q is the Nash bargaining solution iff $\theta_1 = \theta_2$.

References

- [1] Birnbaum, B., Devanur, N., Xiao, L.: New convex programs and distributed algorithms for Fisher markets with linear and spending constraint utilities (2010) (unpublished manuscript)
- [2] Brainard, W.C., Scarf, H.E.: How to compute equilibrium prices in 1891. Cowles Foundation Discussion Paper (1270) (2000)
- [3] Chakrabarty, D., Devanur, N., Vazirani, V.V.: New results on rationality and strongly polynomial solvability in Eisenberg-Gale markets. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) WINE 2006. LNCS, vol. 4286, pp. 239–250. Springer, Heidelberg (2006) (to appear in SIDMA)
- [4] Chen, X., Deng, X., Teng, S.H.: Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM* 56(3) (2009)
- [5] Devanur, N., Papadimitriou, C.H., Saberi, A., Vazirani, V.V.: Market equilibrium via a primal-dual-type algorithm. *JACM* 55(5) (2008)
- [6] Etessami, K., Yannakakis, M.: On the complexity of nash equilibria and other fixed points. In: FOCS (2007)
- [7] Goel, G., Vazirani, V.V.: A perfect price discrimination market model with production and a (rational) convex program for it. In: SAGT (2010)
- [8] Grotschel, M., Lovasz, L., Schrijver, A.: *Geometric Algorithms and Combinatorial Optimization*. Springer, Heidelberg (1988)
- [9] Jain, K., Vazirani, V.V.: Eisenberg-Gale markets: Algorithms and game-theoretic properties. In: *Games and Economic Behavior* (2008) (to appear)
- [10] Kalai, E.: Nonsymmetric Nash solutions and replications of 2-person bargaining. *International Journal of Game Theory* 6(3), 129–133 (1977)
- [11] Kalai, E.: Solutions to the bargaining problem. In: Hurwicz, L., Schmeidler, D., Sonnenschein, H. (eds.) *Social Goals and Social Organization*, pp. 75–105. Cambridge University Press, Cambridge (1985)
- [12] Kelly, F.P.: Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8, 33–37 (1997)
- [13] Nash, J.F.: The bargaining problem. *Econometrica* 18, 155–162 (1950)
- [14] Nash, J.F.: Equilibrium points in n -person games. *Proceedings National Academy of Sciences* 36, 48–49 (1950)
- [15] von Neumann, J., Morgenstern, O.: *Theory of Games and Economic Behavior*. Princeton University Press, Princeton (1944)
- [16] Osborne, M., Rubinstein, A.: *A Course in Game Theory*. The MIT Press, Cambridge (1994)
- [17] Tardos, E.: A strongly polynomial algorithm to solve combinatorial linear programs. *Operations Research* 34(2), 250–256 (1986)
- [18] Thomson, W., Lensberg, T.: *Axiomatic Theory of Bargaining With a Variable Population*. Cambridge University Press, Cambridge (1989)
- [19] Vazirani, V.V.: Combinatorial algorithms for market equilibria. In: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (eds.) *Algorithmic Game Theory*, pp. 103–134. Cambridge University Press, Cambridge (2007)
- [20] Vazirani, V.V.: Nash bargaining via flexible budget markets (2009) (submitted)
- [21] Vazirani, V.V.: Spending constraint utilities, with applications to the Adwords market. *Mathematics of Operations Research* 35(2) (2010)

On the Inefficiency of Equilibria in Linear Bottleneck Congestion Games

Bart de Keijzer, Guido Schäfer, and Orestis A. Telelis*

Center for Mathematics and Computer Science (CWI)
Amsterdam, The Netherlands

{B.de.Keijzer,G.Schaefer,O.Telelis}@cwi.nl

Abstract. We study the inefficiency of equilibrium outcomes in *bottleneck congestion games*. These games model situations in which strategic players compete for a limited number of facilities. Each player allocates his weight to a (feasible) subset of the facilities with the goal to minimize the maximum (weight-dependent) latency that he experiences on any of these facilities. We derive upper and (asymptotically) matching lower bounds on the (strong) price of anarchy of linear bottleneck congestion games for a natural load balancing social cost objective (i.e., minimize the maximum latency of a facility). We restrict our studies to linear latency functions. Linear bottleneck congestion games still constitute a rich class of games and generalize, for example, load balancing games with identical or uniformly related machines with or without restricted assignments.

1 Introduction

Load balancing games constitute an important class of strategic games that capture many applications of practical relevance. These games model situations in which a set of strategically acting players (or jobs) compete for a limited number of resources (or machines). Every player chooses one of the resources available to him and assigns his weight (or load) to this resource. The latency of a resource depends on the total weight of the players using it. The goal of each player is to select a resource such that the latency that he experiences on this resource is minimized.

The study of load balancing games is motivated by the need for quantifying the inefficiency caused by selfish behavior of a set of autonomous players that utilize distributed processors upon which a system is built. The social cost objective of an assignment of loads to processors is measured by the *makespan*, i.e., the completion time of the most loaded machine, which reflects the distance from equi-distribution (balancing) of the load to the machines. Load balancing games have recently been studied extensively for a variety of different machine

* This work was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme.

environments, including identical [15], uniformly related [9,11,14,15], restricted assignment [5,11], and unrelated machines [2].

A natural extension of load balancing games are the *bottleneck congestion games (BCGs)* [6,12]. Here, every player chooses a subset of the resources (also called *facilities* in this context) from a set of feasible facility allocations and assigns his weight to each of these facilities. The goal of each player is to select a subset of the facilities such that the maximum latency over the chosen facilities is minimized. Bottleneck congestion games generalize, for example, load balancing games and network routing games, and have several applications in practice. Despite their importance, bottleneck congestion games have received only very little attention in the literature and are far from being well-understood. In this paper we study the inefficiency of stable outcomes in bottleneck congestion games.

Bottleneck congestion games essentially generalize the context of load balancing games by modeling the activity of each selfish player upon complexes of interrelated resources. This generalization brings the model closer to practice, as in most large scale computing systems the workload of a player occupies different components of the system simultaneously. For example, instantiations of such games emerge if the components form paths in networks, or if they correspond to parallel processors, etc. It is natural to assume that each player wants to balance his load across the different components available to him and hence attempts to minimize the maximum latency of a facility that he uses.

One of the most prominent solution concepts for the prediction of outcomes of rational behavior in strategic games is the *Nash equilibrium* concept. It describes outcomes that are resilient to unilateral player deviations. Throughout this paper we will focus exclusively on pure Nash equilibria. A more general solution concept is the *strong equilibrium* concept introduced by Aumann [3]. It describes outcomes of strategic games that are stable with respect to pure deviations of player subsets (also called *coalitions*). More precisely, an outcome of a strategic game is a strong equilibrium if no coalition of the players can deviate such that every member of the coalition strictly benefits. An outcome is said to be a *k-strong equilibrium* if this property holds for all coalitions of size at most k . Strong equilibria thus generalize the pure Nash equilibrium concept ($k = 1$). Very recently, Harks, Klimm and Möhring [12] showed that (under rather general assumptions) bottleneck congestion games always admit strong equilibria.

It is well known that equilibrium outcomes might be inefficient in the sense that they are suboptimal with respect to some socially desirable objective function. The *price of anarchy (PoA)* [15,16,17] has become the standard measure to assess the inefficiency of equilibrium outcomes. It is defined as the worst-case ratio (over all instances) of the maximum cost of a Nash equilibrium outcome and the cost of a socially optimal outcome. The *strong price of anarchy (SPoA)* and the *k-strong price of anarchy (k-SPoA)* [2] refer to the natural adaptations of this measure to strong and k -strong equilibrium outcomes, respectively.

Contribution. We study the inefficiency of both pure Nash equilibria and strong equilibria of BCGs, under the natural assumption that the *social cost* of an outcome refers to the maximum latency of a facility. We restrict our studies to

Table 1. Summary of the bounds obtained for the SPoA and the k -SPoA of linear BCGs. The PoA of linear BCGs is at most $2m - 1$ and there is an asymptotically matching lower bound showing $\text{SPoA} \geq m - 1$.

	id. facilities		arb. facilities (SPoA)	
	k -SPoA (lower)	SPoA	id. players	arb. players
symmetric	$\max \{2, \lfloor \frac{m}{2k} \rfloor + 1\}$	2	2	$O(m)$
asymmetric	$\max \left\{ \sqrt{2m + \frac{1}{4} - \frac{1}{2}}, \lceil \frac{m}{k-1} \rceil - 1 \right\}$	$\Theta(\sqrt{m})$	$O(\sqrt{n})$	$\Theta(m)$

linear bottleneck congestion games, where the latency of each facility is a linear function of the total weight assigned to it. These games still constitute a rich class of games and generalize, for example, load balancing games with identical or uniformly related machines with or without restricted assignments. We provide upper and lower bounds on the (strong) price of anarchy for symmetric and asymmetric linear BCGs (definitions will be given below). A summary of the results that we obtain in this paper is given in Table 1. Here, we use n and m to refer to the number of players and facilities, respectively.

1. We show that both the PoA and the SPoA of linear BCGs is $\Theta(m)$. More precisely, we show that $m \leq \text{PoA} \leq 2m - 1$ and $m - 1 \leq \text{SPoA} \leq m$.
2. We derive better bounds for identically weighted players. We prove that $\text{SPoA} = 2$ for symmetric linear BCGs and at most $O(\sqrt{n})$ and $O(\sqrt{m\gamma^*})$ for asymmetric linear BCGs, where γ^* refers to the cost of a socially optimal outcome.
3. We consider the case of identical facilities, i.e., all facilities have identical linear latency functions, and show that $\text{SPoA} = \Theta(\sqrt{m})$.
4. We also give elaborate lower bounds on the k -SPoA for symmetric and asymmetric BCGs with identical facilities (see Table 1).

We remark that we also provide asymptotically tight worst-case examples for (directed) network congestion games (definitions will be given below).

Related Work. Network BCGs were considered first by Banner and Orda in [6]. The authors showed existence of pure Nash equilibria and provided an $\Theta(m)$ bound on the PoA for identical network links. Busch and Magdon-Ismael studied in [7] the PoA of network BCGs for identically weighted players. Very recently, Harks, Klimm and Möhring introduced general bottleneck congestion games and showed that strong equilibria are guaranteed to exist in these games.

As mentioned above, bottleneck congestion games generalize load balancing games, which have been studied intensively in recent years. Load balancing games were first studied by Koutsoupias and Papadimitriou [15]. Among other results, the authors provided a lower bound on the PoA of *mixed* Nash equilibria for the case of identical machines. Koutsoupias, Mavronicolas and Spirakis [14] and, independently, Czumaj and Vöcking [9], proved a matching upper bound. Czumaj and Vöcking also proved that $\text{PoA} = \Theta(\log m / \log \log m)$ for pure Nash

equilibria. The same bound on the PoA was shown by Awerbuch et al. [5] for restricted assignments and identical machines. Gairing et al. [11] obtained independently the same bounds and proved $m - 1 \leq \text{PoA} \leq m$ for restricted assignments and uniformly related machines.

Andelman, Feldman and Mansour [2] were the first to study strong and k -strong equilibria in the context of load balancing games. They proved that $m \leq \text{SPoA} \leq 2m - 1$ for the case of unrelated machines, which was tightened to exactly m by Fiat et al. [10]. In this latter work it was also shown that the SPoA of strong equilibria for uniformly related machines is exactly $\Theta(\log m / (\log \log m)^2)$. For results in the context of more general scheduling games and associated scheduling policies (termed *coordination mechanisms*), the interested reader is referred to [13] and the references therein.

Bottleneck congestion games owe their name to their similarity to *congestion games*, which were introduced by Rosenthal [18]. In these games, the latency on each facility depends on the number of players using it (i.e., players have unit weights). The goal of each player is to minimize his cost which is defined as the sum (as opposed to the maximum for BCGs) of the latencies over the facilities used by the player. Rosenthal [18] proved the existence of pure Nash equilibria in congestion games. The price of anarchy of pure Nash equilibria for congestion games was resolved by Christodoulou and Koutsoupias [8] and, independently, by Awerbuch, Azar and Epstein [4]. It is shown in [8] that $\text{PoA} = \Theta(\sqrt{n})$ for asymmetric linear congestion games and the social cost being the maximum over the players' cost, and $\text{PoA} = \frac{5}{2}$ for (symmetric and asymmetric) linear congestion games and the social cost being the sum of the players' costs. Bounds for polynomial latency functions were also derived in [8]. Exact bounds for polynomial latencies and also for weighted players were developed in [1].

2 Preliminaries

In a *bottleneck congestion game*, we are given a set $N = [n]$ of n players that want to utilize (non-cooperatively) a set $E = [m]$ of m resources, which we also call *facilities*.¹ Every player $i \in N$ has a positive weight (or load) $w_i > 0$ and a strategy set $\Sigma_i \subseteq 2^E$ of feasible facility subsets from which he can choose. If player i chooses facility subset $S_i \in \Sigma_i$, he allocates his entire weight w_i to each facility $e \in S_i$. Let $\Sigma = (\Sigma_1, \dots, \Sigma_n)$ be the set of all possible strategy choices of the players. A *strategy profile* $S = (S_1, \dots, S_n) \in \Sigma$ specifies for each player $i \in N$ a strategy $S_i \in \Sigma_i$ that he has chosen. We define $N_e(S)$ as the set of players that have chosen facility $e \in E$ under S , i.e., $N_e(S) = \{i \in N \mid e \in S_i\}$. The total weight of facility $e \in E$ with respect to S is defined as $w_e(S) = \sum_{i \in N_e(S)} w_i$.

Every facility $e \in E$ has a latency function $l_e : \Sigma \rightarrow \mathcal{R}^+$ which satisfies the following three properties (see also [12]):

1. **Non-negativity:** $l_e(S) \geq 0$ for all $S \in \Sigma$.

¹ We use notation $[k]$ to refer to the set $\{1, \dots, k\}$ for some positive integer k .

- 2. **Independence of irrelevant alternatives:** $l_e(S) = l_e(S')$ for all $S, S' \in \Sigma$ with $N_e(S) = N_e(S')$.
- 3. **Monotonicity:** $l_e(S) \geq l_e(S')$ for all $S, S' \in \Sigma$ with $N_e(S) \supseteq N_e(S')$.

Given a strategy profile $S \in \Sigma$, every player $i \in N$ experiences an individual cost $c_i(S)$ equal to the latency of the most loaded facility that he uses, i.e., $c_i(S) = \max_{e \in S_i} l_e(S)$. We assume that every player $i \in N$ acts strategically and chooses his strategy $S_i \in \Sigma_i$ in order to minimize his own individual cost $c_i(S)$.

Aumann [3] introduced the notion of a *strong equilibrium*. Here we consider the refined notion of *k-strong equilibrium*. We use the standard notation S_{-i} to refer to $(S_1, \dots, S_{i-1}, S_{i+1}, \dots, S_n)$. Similarly, we use S_I and S_{-I} to refer to the strategy profiles of S induced by the players in I and $N \setminus I$, respectively.

Definition 1. *A strategy profile $S \in \Sigma$ is a k-strong equilibrium if for every non-empty player set $I \subseteq N$ with $|I| \leq k$ and every possible joint deviation S'_I of I there is at least one player $i \in I$ whose cost with respect to $S' = (S_{-I}, S'_I)$ is not better than with respect to S , i.e., $c_i(S_{-I}, S'_I) \geq c_i(S)$.*

With this definition, a *strong equilibrium* is a *k-strong equilibrium* with $k = n$, and a *pure Nash equilibrium* is a *k-strong equilibrium* with $k = 1$. Very recently, Harks, Klimm and Möhring [12] showed that strong equilibria always exist in BCGs satisfying Properties 1–3 above.

We are interested in characterizing the inefficiency of *k-strong equilibria* for BCGs. We assess the efficiency of a strategy profile S by the maximum load of a facility under S . That is, the *social cost* $C(S)$ of a strategy profile $S \in \Sigma$ is defined as the maximum latency over all facilities, which is equivalent to the maximum cost over all players, i.e., $C(S) = \max_{e \in E} l_e(S) = \max_{i \in N} c_i(S)$. We will use S^* to refer to an optimal strategy profile that minimizes $C(S)$ and denote its cost by $\gamma^* = C(S^*)$.

The *k-strong price of anarchy (k-SPoA)* [2,15] refers to the worst-case ratio over all possible input instances of the maximum cost of a *k-strong equilibrium* and the cost γ^* of the social optimum. We will simply refer to the *price of anarchy (PoA)* and *strong price of anarchy (SPoA)* for the 1-SPoA and the *n-SPoA*, respectively. One can easily make an example to show that the SPoA is unbounded in general. This motivates our studies of *linear BCGs*: We assume that the latency function l_e of each facility $e \in E$ is a linear function of the total weight assigned to it, i.e., $l_e(S) = a_e w_e(S)$ for some $a_e \geq 0$. Linear BCGs constitute an important class of BCGs because they generalize, for example, various load balancing games as outlined in the Introduction.

A BCG is called a *network BCG* if there exists a directed graph $G = (V, E)$ such that every player $i \in N$ is associated with a source $s_i \in V$ and a sink $t_i \in V$ and i 's strategy set Σ_i refers to the set of all directed paths from s_i to t_i in G . We call a game *symmetric* if all players have the same strategy set, i.e., $\Sigma_i = \Sigma_j$ for all $i, j \in N$; we call a game *asymmetric* otherwise. Observe that the above example corresponds to a network BCG, but is not symmetric.

Unless stated otherwise, we assume subsequently that all player weights are at least one, i.e., $w_i \geq 1$ for every $i \in N$, and that the coefficient of each latency function is at least one, i.e., $l_e(S) = a_e w_e(S)$ with $a_e \geq 1$ for every $e \in E$. These assumptions are without loss of generality as we can always enforce them by scaling the weights and coefficients appropriately.

3 Arbitrary Facilities

In this section, we derive bounds on the PoA and SPoA of linear BCGs. We consider both the general and the identical player case.

3.1 Arbitrary Players

We first consider the most general case of arbitrary linear latency functions and arbitrary player weights. We show that the PoA is at most $2m - 1$ in this case. We obtain a better bound of m on the SPoA and present an almost tight lower bound.

Theorem 1. *The price of anarchy of linear BCGs is at most $2m - 1$ and at least m .*

Proof. Let S be a pure Nash equilibrium with cost $C(S) = \alpha\gamma^*$ for some $\alpha \geq 1$. We prove by induction that for every integer k , $1 \leq k < \frac{\alpha+1}{2} + 1$, there is a set E_k of k distinct facilities such that for every $e \in E_k$, $l_e(S) \geq (\alpha - k + 1)\gamma^*$.

The claim holds true for $k = 1$ because there must exist a facility $e \in E$ with latency $l_e(S) = \alpha\gamma^*$. Suppose that the induction hypothesis holds true for $k < \frac{\alpha+1}{2}$. We will prove that there exists a set E_{k+1} of $k + 1$ distinct facilities such that $l_e(S) \geq (\alpha - k)\gamma^*$ for every $e \in E_{k+1}$. Choose from E_k a facility \hat{e} with smallest a_e , i.e., $\hat{e} = \arg \min_{e \in E_k} a_e$. By the induction hypothesis, we have $l_{\hat{e}}(S) \geq (\alpha - k + 1)\gamma^* > k\gamma^*$. Let $I_{\hat{e}} = N_{\hat{e}}(S)$ be the set of players choosing \hat{e} under S . Note that $w_{\hat{e}}(S) \geq l_{\hat{e}}(S)/a_{\hat{e}} > k\gamma^*/a_{\hat{e}}$. Consider the strategies that the players in $I_{\hat{e}}$ choose under S^* and suppose for the sake of a contradiction that for every $i \in I_{\hat{e}}$, $S_i^* \cap E_k \neq \emptyset$. Then there is a facility $e \in E_k$ with $w_e(S^*) \geq w_{\hat{e}}(S)/k > \gamma^*/a_{\hat{e}}$. By the choice of \hat{e} , we have $l_e(S^*) = a_e w_e(S^*) > \gamma^*$, which is a contradiction to the definition of γ^* . Thus there is a player $j \in I_{\hat{e}}$ that chooses a strategy S_j^* that is disjoint from E_k . Note that for every $e \in S_j^*$ we have $a_e w_j \leq \gamma^*$. Since S is a pure Nash equilibrium, player j cannot decrease his cost by deviating to S_j^* and thus there is some facility $e' \in S_j^*$ such that:

$$l_{e'}(S) = (a_{e'} w_{e'}(S) + a_{e'} w_j) - a_{e'} w_j \geq c_i(S) - a_{e'} w_j \geq l_{\hat{e}}(S) - \gamma^* \geq (\alpha - k)\gamma^*$$

The inductive step follows by setting $E_{k+1} = E_k \cup \{e'\}$. By choosing $k = \lceil \frac{\alpha+1}{2} \rceil < \frac{\alpha+1}{2} + 1$, we obtain that there is a set $E_k \subseteq E$ with $|E_k| \geq k$ and thus $m \geq |E_k| \geq k \geq \frac{\alpha+1}{2}$. We conclude that $\text{PoA} = \alpha \leq 2m - 1$.

The following instance shows that $\text{PoA} \geq m$, even for symmetric BCGs with identical facilities and identical players. Consider a BCG with player set $N = [n]$

and facility set $E = [m]$ with $m = n$. Every player $i \in N$ has unit weight $w_i = 1$ and the latency function $l_e(S)$ of every $e \in E$ is the identity function, i.e., $l_e(S) = w_e(S)$. Suppose that each player $i \in N$ has strategy set $\Sigma_i = 2^E$. If every player chooses a distinct facility we obtain an optimal strategy profile S^* with $\gamma^* = 1$. On the other hand, consider the strategy profile S in which every player allocates all facilities in E . This is a pure Nash equilibrium of cost $C(S) = m$. \square

We derive a better upper bound on the SPoA for linear BCGs. The following key lemma will be used several times in the paper.

Lemma 1. *Let S be a strong equilibrium and let $I_\lambda \subseteq I$ be a non-empty subset of the players such that for every $i \in I_\lambda$ we have $c_i(S) \geq \lambda\gamma^*$, for some $\lambda \geq 1$.*

1. *Then there is a player $i \in I_\lambda$ and a facility $e \in S_i^*$ such that $l_e(S_{-I_\lambda}) \geq (\lambda - 1)\gamma^*$.*
2. *Suppose that I_λ is maximal. Then there is a player set $T_\lambda \subseteq N \setminus I_\lambda$ with $w(T_\lambda) \geq \lambda - 1$ and for every $i \in T_\lambda$ we have $(\lambda - 1)\gamma^* \leq c_i(S) < \lambda\gamma^*$.*

Proof. We first prove the first part of the lemma. Note that for every player $i \in I_\lambda$ and every $e \in S_i^*$ we have

$$l_e(S_{I_\lambda}^*) \leq l_e(S^*) \leq \gamma^*. \tag{1}$$

Suppose for the sake of a contradiction that for every player $i \in I_\lambda$ and for every $e \in S_i^*$ it holds that $l_e(S_{-I_\lambda}) < (\lambda - 1)\gamma^*$. Consider the strategy profile $S' = (S_{-I_\lambda}, S_{I_\lambda}^*)$ in which the players in I_λ deviate to their optimal strategies in S^* . Using \square , we obtain for every $i \in I_\lambda$ and for every $e \in S_i^*$:

$$l_e(S') = l_e(S_{I_\lambda}^*) + l_e(S_{-I_\lambda}) < \gamma^* + (\lambda - 1)\gamma^* = \lambda\gamma^*. \tag{2}$$

Thus, for every $i \in I_\lambda$, $c_i(S') = \max_{e \in S_i^*} l_e(S') < \lambda\gamma^*$, which is a contradiction to S being a strong equilibrium.

We next prove the second part of the lemma. Let $i \in I_\lambda$ be a player and $e \in S_i^*$ be a facility satisfying $l_e(S_{-I_\lambda}) \geq (\lambda - 1)\gamma^*$. Define T_λ as the set of players that choose e under S but are not contained in I_λ , i.e., $T_\lambda = N_e(S) \setminus I_\lambda \subseteq N \setminus I_\lambda$. We have

$$a_e w(T_\lambda) = l_e(S_{T_\lambda}) = l_e(S_{-I_\lambda}) \geq (\lambda - 1)\gamma^*. \tag{3}$$

Since $e \in S_i^*$ and $w_i \geq 1$ for every $i \in N$, we have $a_e \leq \gamma^*$. Thus, $w(T_\lambda) \geq \lambda - 1$. Consider an arbitrary player $i \in T_\lambda$. By the above we have, $c_i(S) \geq l_e(S) \geq l_e(S_{T_\lambda}) \geq (\lambda - 1)\gamma^*$. Moreover, by the maximality of I_λ and since $i \notin I_\lambda$, we have $c_i(S) < \lambda\gamma^*$. \square

Remark 1. Observe that in the above proof we exploit the linearity of the latency functions only in \square . In fact, we can draw exactly the same conclusion if all latency functions are *sub-additive*, i.e., for every $e \in E$, $l_e(x + y) \leq l_e(x) + l_e(y)$ for every $x, y \in \mathcal{R}^+$. As a consequence, all our upper bounds on the SPoA (which exploit Lemma \square) hold for sub-additive latency functions.

Theorem 2. *The strong price of anarchy of linear BCGs is at most m .*

Proof. Let S be a strong equilibrium with cost $C(S) = \alpha\gamma^*$ for some $\alpha > 1$. For an arbitrary real value $1 < \lambda \leq \alpha$, let I_λ be the maximal non-empty set of players $I_\lambda = \{i \in N \mid c_i(S) \geq \lambda\gamma^*\}$. Applying Lemma 1, we obtain a player set T_λ such that for every $i \in T_\lambda$ we have $(\lambda - 1)\gamma^* \leq c_i(S) < \lambda\gamma^*$. Moreover, $w(T_\lambda) \geq \lambda - 1 > 0$ because $\lambda > 1$ and thus T_λ is non-empty. We can thus identify a family $F = \{T_\alpha, T_{\alpha-1}, \dots, T_{\alpha-k}\}$ of $k + 1$ player sets that are non-empty and pairwise disjoint, where k is the largest integer satisfying $\alpha - k > 1$. Every set $T_\lambda \in F$ identifies at least one distinct facility $e \in E$ with $(\lambda - 1)\gamma^* \leq l_e(S) < \lambda\gamma^*$. Moreover, there is one facility $e \in E$ with $l_e(S) = \alpha\gamma^*$. We conclude that $m \geq |F| + 1 = k + 2 \geq \alpha$ and thus $\text{SPoA} = \alpha \leq m$. \square

Theorem 3. *The strong price of anarchy is at least $m - 1$ in general linear BCGs and at least $\frac{m+1}{3}$ in single-sink linear network BCGs.*

The proof of this result is deferred to the full version. The lower bound of $m - 1$ can also be derived by a construction in [11].

3.2 Identical Players

We next derive an upper bound on the SPoA for linear BCGs if the weights of all players are identical. In this subsection, we assume without loss of generality that the weight of each player $i \in N$ is $w_i = 1$.

Theorem 4. *The strong price of anarchy is at most $O(\min\{\sqrt{n}, \sqrt{m\gamma^*}\})$ for linear BCGs with identical players and 2 for linear symmetric BCGs with identical players.*

Proof. We prove the first part of the theorem. Let S be a strong equilibrium with cost $C(S) = \alpha\gamma^*$ for some $\alpha > 1$. As in the proof of Theorem 2, we can apply Lemma 1 to identify a family $F = \{T_\alpha, T_{\alpha-1}, \dots, T_{\alpha-k}\}$ of $k + 1$ player sets that are non-empty and pairwise disjoint, where k is the largest integer satisfying $\alpha - k > 1$. Each such set $T_\lambda \in F$ contains at least $\lambda - 1$ players, i.e., $|T_\lambda| \geq \lceil \lambda - 1 \rceil$ for every $\alpha - k \leq \lambda \leq \alpha$. Moreover, there is at least one player that experiences a congestion of $\alpha\gamma^*$. Thus

$$n \geq 1 + \sum_{\lambda=1}^{\lceil \alpha-1 \rceil} \lambda \geq 1 + \frac{\alpha(\alpha - 1)}{2}.$$

Solving for α we obtain $\alpha \leq \frac{1}{2} + \sqrt{2n - 3/2}$. Recall that we assume without loss of generality that $a_e \geq 1$ for every $e \in E$ and thus $\gamma^* \geq n/m$. We therefore also obtain $\alpha \leq \frac{1}{2} + \sqrt{m\gamma^* - 3/2}$. Thus $\text{SPoA} \leq \alpha = O(\min\{\sqrt{n}, \sqrt{m\gamma^*}\})$.

We next prove the second part of the theorem. In a strong equilibrium S , at least one player $i \in N$ must have cost $c_i(s) \leq \gamma^*$ since otherwise the grand coalition could deviate to the socially optimal strategy profile. Suppose there is a player $j \in N$ whose cost is more than two times larger than the cost of i .

Consider the deviation $S' = (S_{-j}, S_j)$ where player j deviates to the strategy of player i . Then $c_j(S') \leq \max_{e \in S_i} a_e(w_e(S) + 1) \leq \max_{e \in S_i} 2a_e w_e(S) \leq 2c_i(S)$, which is a contradiction to S being a strong equilibrium.

The following example establishes the tightness of this bound: Let $N = [3]$ and $E = [6]$. The strategy set of every player is $\{\sigma_1 = \{1\}, \sigma_2 = \{2, 3\}, \sigma_3 = \{4, 5\}, \sigma_4 = \{2, 5, 6\}\}$. The social optimum is $S_i^* = \sigma_i$ for every player $i \in [3]$ with $\gamma^* = 1$. A strong equilibrium is given by $S_1 = \sigma_4$ and $S_2 = S_3 = \sigma_1$. The cost of S is $C(S) = 2$. It is easy to see that this example is a network BCG. \square

4 Identical Facilities

In this section, we study the SPoA for the case of linear BCGs with identical facilities, i.e., the latency function of every facility $e \in E$ is $l_e(S) = w_e(S)$.

Theorem 5. *The strong price of anarchy of linear BCGs with identical facilities is at most $-\frac{1}{2} + \sqrt{2m + \frac{1}{4}}$ in general and exactly 2 in case of symmetric games.*

Proof. For the symmetric case we claim that in any strong equilibrium configuration S , there is at least one player i_0 with $c_{i_0}(S) \leq \gamma^*$. Indeed, if $c_i(S) > \gamma^*$ for all players, then the grand coalition would deviate to S^* . Now for any player i we have $\gamma^* \geq w_i$. Let i be any player with $e \in S_i$ such that $c_i(S) = l_e(S) = C(S)$. Consider unilateral deviation $S'_i = S_{i_0}$ of i . Then, because S is also a pure Nash equilibrium, $C(S) = c_i(S) \leq c_{i_0}(S) + w_i \leq 2\gamma^*$. A tight lower bound has already been presented in Theorem 4. \square

For the asymmetric case let the cost of a strong equilibrium S be $C(S) = \alpha\gamma^*$, for some $\alpha > 1$. Similar to the proof of Theorem 2, let I_λ be the maximal non-empty set of players $I_\lambda = \{i \in N \mid c_i(S) \geq \lambda\gamma^*\}$ for some $1 < \lambda \leq \alpha$. By Lemma 1, we obtain a player set T_λ such that for every $i \in T_\lambda$ we have $(\lambda - 1)\gamma^* \leq c_i(S) < \lambda\gamma^*$. We can refine the argument given in the proof of Lemma 1 to bound the weight of T_λ for identical facilities as follows: By inequality (3), we have $w(T_\lambda) \geq (\lambda - 1)\gamma^*/a_e = (\lambda - 1)\gamma^*$, where the last equality holds because for identical facilities $a_e = 1$ for every $e \in E$. Moreover, $w(T_\lambda) \geq (\lambda - 1)\gamma^* > 0$ because $\lambda > 1$ and thus T_λ is non-empty. That is, we can identify a family $F = \{T_\alpha, T_{\alpha-1}, \dots, T_{\alpha-k}\}$ of $k + 1$ player sets that are non-empty and pairwise disjoint, where k is the largest integer satisfying $\alpha - k > 1$. Moreover, by construction we have $I_\alpha \cap T_\lambda = \emptyset$ for every $T_\lambda \in F$ and $w(I_\alpha) \geq \alpha\gamma^*$ since facilities are identical. The total weight $w(N)$ is then:

$$w(N) \geq \alpha\gamma^* + \sum_{\lambda=\alpha-k}^{\alpha} w(T_\lambda) \geq \alpha\gamma^* + \sum_{\lambda=\alpha-k}^{\alpha} (\lambda - 1)\gamma^* \geq \alpha\gamma^* + \sum_{\lambda=0}^{\alpha-1} \lambda\gamma^*$$

The latter equals $\frac{1}{2}\alpha\gamma^*(1 + \alpha)$. Observe that $\gamma^* \geq w(N)/m$ because facilities are identical. We obtain $2m \geq \alpha(1 + \alpha)$ or equivalently $\alpha \leq -\frac{1}{2} + \sqrt{2m + 1/4}$. Since $\text{SPoA} \leq \alpha$ the claim follows. \square

Theorem 6. *The strong price of anarchy of linear BCGs with identical players and identical facilities is at least $-\frac{1}{2} + \sqrt{2m + \frac{1}{4}}$ in general and at least $-\frac{1}{4} + \frac{1}{2}\sqrt{2 + 2m}$ in single-sink network BCGs .*

Proof. We give a family of instances with m facilities and $n = \Theta(m)$ unweighted players, which we turn into a family of network instances subsequently. Consider a partition of the set of players N into q subsets, $N = \bigcup_{j=1}^q P_j$, where $|P_j| = j$, $j \in [q]$. Denote players in P_j by p_{ji} , $i \in [j]$. For each subset P_j make a new set of j distinct facilities $E_j = \{e_1^j, \dots, e_j^j\}$. Define $E_{q+1} = E_1$. For every player $p_{ji} \in P_j$, $i \in [j]$, set the strategy space of p_{ji} to:

$$\Sigma_{p_{ji}} = \left\{ \{e\} \mid e \in E_j \right\} \cup \{E_{j+1}\}$$


For the socially optimal configuration set $S_{p_{ji}}^* = \{e_i^j\}$. Then $C(s^*) = 1$. Now consider the configuration S where $S_{p_{ji}} = E_{j+1}$ for $i \in [j]$, $j \in [q]$. The cost of S is defined by the latency of the unique facility $e = e_1^1 \in E_1$ and is $C(S) = l_e(S) = |P_q| = q$. For every player $p \in P_j$, we have $c_p(S) = j$. We claim that S is a strong equilibrium. Consider any deviation of any coalition $I \subseteq N$. Denote by S'_p the novel strategy that any player $p \in I$ adopts and let S' denote the resulting configuration. Notice that for the unique player $p \in P_1$ we have $c_p(S) = 1$, hence no deviation may lessen his cost and $P_1 \cap I = \emptyset$.

Let $j = \min\{j' \mid P_{j'} \cap I \neq \emptyset\}$; then $j \geq 2$, and $S'_j \cap E_j \neq \emptyset$. For all $j - 1$ players $p_{j-1,i} \in P_{j-1}$ it holds that $S_{p_{j-1,i}} = E_j$, because $I \cap P_{j-1} = \emptyset$. Hence, $c_j(S') = j - 1 + 1 = j = c_j(S)$. In any deviation of any coalition I , at least one player does not have incentive to deviate jointly with I and hence $\text{SPoA} \geq q$. Now for q we have $m = |\cup_j E_j| = \sum_{j=1}^q j = \frac{q(q+1)}{2}$, which yields $q \geq -\frac{1}{2} + \sqrt{2m + 1/4}$.

We convert the example into a network BCG . To grant access to players in P_{j-1} to facilities in E_j , we make a path of length 3, $\{(s_j, u_{ji}), (u_{ji}, v_{ji}), (v_{ji}, t)\}$, for every facility $e_i^j \in E_j$, $i \leq j - 1$ and a length-2 path $\{(s_j, u_{jj}), (u_{jj}, t)\}$ for e_j^j . Let A_j be the set of arcs in these paths. Node s_j is the source of all players in P_j and t is a common sink for all players. Now we add auxiliary arcs $A'_j = \{(v_{ji}, u_{j,i+1}) \mid i \in [j - 1]\}$. And, finally, an arc (s_{j-1}, u_{j1}) , $j \in \{2, \dots, q\}$, by which players P_{j-1} gain access to A_j . For the last group of players we add an arc (s_q, t) . Let us illustrate the analog of configuration S on the constructed network. All players in $p_{ji} \in P_j$, $i \in [j]$, play the same path strategy:

$$\begin{aligned} S_{ji} &= \{(s_j, u_{j+1,1})\} \\ &\cup \{(u_{j+1,r}, v_{j+1,r}), (v_{j+1,r}, u_{j+1,r+1}) \mid r \in [j - 1]\} \\ &\cup \{(u_{j+1,j}, v_{j+1,j}), (v_{j+1,j}, t)\} \end{aligned}$$

and $S_{iq} = (s_q, t)$ for $i \in [q]$. See Fig. [1a](#) for an example with $q = 4$. The proof that S is strong is analogous to the proof given for the non-network example. For the optimal configuration we set $S_{ji}^* = \{(s_j, u_{ji}), (u_{ji}, v_{ji}), (v_{ji}, t)\}$, for each player $p_{ji} \in P_j$, $i < j$, and $S_{jj} = \{(s_j, u_{jj}), (u_{jj}, t)\}$. The number of links m is:



(a) SPoA on identical links.

(b) 2-SPoA for 6 identical players on identical links; player indices mark links used by each player.

Fig. 1. Lower bound constructions for Strong and 2-Strong Equilibria on identical links

$$m = \sum_{j=1}^q (|A_j| + |A'_j|) + q = \sum_{j=1}^q (3j - 1 + (j - 1)) + q - 1 = 2q^2 + q - 1$$

which yields $q \geq -\frac{1}{4} + \frac{1}{2}\sqrt{2 + 2m}$. □

4.1 Lower Bounds On k -Strong Equilibria

For the k -SPoA of symmetric and general BCGs with identical facilities we show:

Theorem 7. *The k -strong price of anarchy of linear BCGs is at least:*

1. $\lfloor \frac{m}{2k} \rfloor + 1$ for symmetric BCGs and $\lceil \frac{m+2}{6k} \rceil$ for symmetric network BCGs, when $2 \leq k \leq \frac{m}{2}$.
2. $\lceil \frac{m}{k-1} \rceil - 1$ in general, when $2 \leq k \leq \frac{3}{4} + \frac{1}{2}\sqrt{\frac{1}{4} + 2m}$.

The proofs of these results are deferred to the full version. Figure 1b presents a 2-strong equilibrium for 6 identical players and 34 identical links. The maximum latency over all links under this configuration is 3. The social optimum has cost 1 and emerges when all players use link-disjoint paths to reach t from s .

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. Andelman, N., Feldman, M., Mansour, Y.: Strong price of anarchy. In: Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 189–198. ACM Press, New York (2007)

3. Aumann, R.J.: Acceptable points in games of perfect information. *Pacific Journal of Mathematics* 10, 381–417 (1960)
4. Awerbuch, B., Azar, Y., Epstein, A.: Large the price of routing unsplittable flow. In: *Proceedings of the 37th ACM Symposium on Theory of Computing (STOC)*, pp. 57–66 (2005)
5. Awerbuch, B., Azar, Y., Richter, Y., Tsur, D.: Tradeoffs in worst-case equilibria. *Theoretical Computer Science* 361(2-3), 200–209 (2006)
6. Banner, R., Orda, A.: Bottleneck Routing Games in Communication Networks. *IEEE Journal on Selected Areas in Communication* 25(6), 1173–1179 (2007)
7. Busch, C., Magdon-Ismail, M.: Atomic Routing Games on Maximum Congestion. *Theoretical Computer Science* 410(36), 3337–3347 (2009)
8. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: *Proceedings of the 37th ACM Symposium on Theory of Computing (STOC)*, pp. 67–73. ACM Press, New York (2005)
9. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms* 3(1) (2007)
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. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B.: The price of Anarchy for Restricted Parallel Links. *Parallel Processing Letters* 16(1), 117–132 (2006)
12. Harks, T., Klimm, M., Möhring, R.: Strong nash equilibria in games with the lexicographical improvement property. In: Leonardi, S. (ed.) *WINE 2009*. LNCS, vol. 5929, pp. 463–470. Springer, Heidelberg (2009)
13. Immorlica, N., Li, E., Mirrokni, V.S., Schulz, A.S.: Coordination mechanisms for selfish scheduling. *Theoretical Computer Science* 410(17), 1589–1598 (2009)
14. Koutsoupias, E., Mavronicolas, M., Spirakis, P.G.: Approximate Equilibria and Ball Fusion. *Theory of Computing Systems* 36(6), 683–693 (2003)
15. 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)
16. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. *Computer Science Review* 3(2), 65–69 (2009)
17. Papadimitriou, C.H.: Algorithms, Games, and the Internet. In: *Proceedings of the 33rd ACM Symposium on Theory of Computing (STOC)*, pp. 749–753. ACM Press, New York (2001)
18. Rosenthal, R.W.: A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2(1), 65–67 (1973)

Minimal Subsidies in Expense Sharing Games

Reshef Meir¹, Yoram Bachrach², and Jeffrey S. Rosenschein¹

¹ Hebrew University, Jerusalem, Israel
{reshef24, jeff}@cs.huji.ac.il

² Microsoft Research, Cambridge, UK
yobach@microsoft.com

Abstract. A key solution concept in cooperative game theory is the core. The core of an expense sharing game contains stable allocations of the total cost to the participating players, such that each subset of players pays at most what it would pay if acting on its own. Unfortunately, some expense sharing games have an empty core, meaning that the total cost is too high to be divided in a stable manner. In such cases, an external entity could choose to induce stability using an external subsidy. We call the minimal subsidy required to make the core of a game non-empty the *Cost of Stability (CoS)*, adopting a recently coined term for surplus sharing games.

We provide bounds on the CoS for general, subadditive and anonymous games, discuss the special case of *Facility Games*, as well as consider the complexity of computing the CoS of the grand coalition and of coalitional structures.

1 Introduction

We begin with a motivating example. Three hospitals plan to purchase an X-ray machine. A standard machine costs \$5 million, and can fulfill the needs of up to two hospitals. An advanced machine capable of serving all three hospitals costs \$9 million. The hospital managers understand that the right thing to do is to buy the more expensive machine, which can serve all three hospitals and costs less than two standard machines, but cannot agree on how to allocate the cost of the expensive machine among the hospitals. There will always be a pair of hospitals that together need to pay at least \$6 million, and would then rather split off and buy the cheaper machine for themselves. The generous mayor solves the problem by subsidizing the expensive machine: she contributes \$3 million, and lets each hospital add \$2 million. Pairs of hospitals now have no incentive to buy the less efficient machine, as each pair together pays only \$4 million.

The example shows how external monetary funding can increase cooperation among self-interested parties. Clearly, a high enough subsidy can always induce cooperation. For example, if the mayor would decide to have the city pay for the entire expensive X-ray machine on its own, then the hospitals' (zero) costs would be irrelevant. However, we would like to consider the minimal external intervention needed to induce cooperation; in the scenario above, for example, a subsidy of \$1.5 million would suffice.

The concepts of stable payoffs, cost allocation, and subsidies have received much attention in economics, decision-making, and recently also computer science. While some papers concentrate on the fair allocation of payments, strategyproofness, and

other requirements, we focus on finding the minimal subsidy that guarantees cooperation among all parties. We model situations such as the example above as Transferable Utility (TU) Expense games. In such games, every subset of agents has a fixed cost. An *imputation* is an allocation of the cost of the grand coalition containing all agents, and it is stable if no coalition can do better (i.e., pay less) on its own. The set of all stable imputations is known as the *core*, and unfortunately it may be empty (as the example above illustrates). The *Cost of Stability (CoS)* of an expense game is the minimal external payment, or subsidy, required to stabilize a game with an empty core.

Related Work. The term “Cost of Stability” for TU games was coined by Bachrach et al. [34]. They defined it as the minimal monetary infusion required to stabilize a surplus game (where agents try to distribute a positive surplus, rather than a negative cost), focusing on computational problems in Weighted Voting games. Resnick et al. [23] extended the results to Threshold Network Flow games (suggested in [6]). The CoS in *expense games*, which are the complementary class of surplus games, has a more natural interpretation as the necessary *proportion of subsidy*, or “how much of the total expense should be subsidized?” The answer ranges between 0% (when the core is non-empty) and 100% (when no agent is willing to contribute). The relative part of the cost covered by the agents, i.e., the complement of the CoS, is called the *cost recovery ratio*.

Using different terms, several other researchers studied subsidies in expense games, and in facility games in particular (described below), sometimes adding requirements on top of the minimization of subsidies. A common assumption is that players gain some *private utility* from their participation.¹ Devanur et al. [12] suggested a mechanism that covers at least a fraction of $\frac{1}{\ln(n)+1}$ in facility games, and a constant fraction of 0.462 in Metric Facility Location games—with the additional requirement of strategyproofness.

An application that has drawn much attention is routing in networks, which was initially formulated as a Minimum Spanning Tree game [11]. In this game, agents are nodes on a graph, and each edge is a connection that has a fixed price. The cost of a coalition is the price of the cheapest tree that connects all participating nodes to the source node. The CoS in this particular game is always 0, as its core is nonempty [13]. However, there is a more realistic variant of routing scenarios known as the Steiner Tree game, where nodes are allowed to route through nodes that are not part of their coalition. Meggido [19] showed that the core of the Steiner Tree game may be empty, and therefore its CoS is nontrivial. Jain and Vazirani [16] proposed a mechanism for the Steiner Tree game with a cost recovery ratio of $1/2$, under the stronger requirements of group strategyproofness.² Other research [24] suggested a cost sharing mechanism for Steiner Trees that does not consider strategyproofness, and showed *empirically* that it allocates at least 92% of the cost on all tested instances.

Other cost sharing mechanisms for many different games have been suggested; see Pal and Tardos [21], and Immorlica et al. [15] for an overview. Some proofs in this paper use similar techniques. Some of the proposed mechanisms pose strong requirements

¹ Our model drops this assumption, which is equivalent to assuming that participation is mandatory, or that the utility is sufficiently high to guarantee participation at any cost.

² To be exact, Jain and Vazirani demanded full cost recovery, and relaxed stability constraints. The bound on the CoS is achieved if we divide their proposed payments by 2.

such as group strategyproofness, in addition to stability. Therefore it is quite likely that tighter bounds on the CoS can be derived once these requirements are relaxed.

While we focus mainly on bounds, there has also been interest in the complexity of *computing* the CoS [32]. Stability bounds regarding the ratio between the optimal social welfare and a core-stable social welfare were examined in affinity games [9]. External subsidies have also been suggested as a means of stabilizing normal form games [20], including a normal form version of the facility game [10].

Our Contribution. We analyze bounds on the CoS in the general and the anonymous case, and compare them to surplus games. We then focus on a particular class of expense games called *Facility Games* (more widely known as *Set-Cover Games*). We provide a tight bound on the CoS in facility games based on a known relation between the CoS and combinatorial properties of the Set-Cover problem, and discuss some related computational issues. We show that the bounds on facility games apply to all games whose cost function is subadditive. Interestingly, subadditivity can be further exploited to bound the subsidy even in games that are *not* subadditive. We conclude with an efficient algorithm for stabilizing coalitional structures in anonymous games.

2 Preliminaries

We denote by I the set of n agents, and by \mathcal{S} the set of all possible coalitions, i.e., $\mathcal{S} = 2^I \setminus \{\emptyset\}$. $I \in \mathcal{S}$ is referred to as the *grand coalition*. An *expense game* $G = \langle I, c \rangle$ is characterized by a cost function $c : \mathcal{S} \rightarrow \mathbb{R}_+$, where $c(S)$ is the cost of coalition S . An *imputation* $\mathbf{p} \in \mathbb{R}_+^n$ is a vector whose sum is the total cost $c(I)$, which defines the amount each agent pays when the grand coalition is formed. A coalition S *blocks* imputation \mathbf{p} if it can guarantee a lower payment for itself, i.e., if $c(S) < p(S) = \sum_{i \in S} p_i$. The *core* is the set of all imputations that are *stable*, i.e., not blocked by any coalition. We emphasize that the concept of the core refers only to stability of games in which only a single coalition is allowed. We discuss cases where several coalitions may exist in Section 6. For a detailed background treatment of coalitional games, see [22].

Monotonicity. We say that an expense sharing game $\langle I, c \rangle$ is *monotone* if $c(S) \geq c(T)$ for all $T \subset S$. Monotonicity means that adding agents to a coalition can only increase its expenses. We will limit our attention to monotone games.

Subadditivity. The game $\langle I, c \rangle$ is *subadditive*, if $c(S \cup T) \leq c(S) + c(T)$ for all S, T s.t. $S \cap T = \emptyset$. Intuitively, in subadditive expense games larger coalitions are better off, and should therefore be easier to stabilize. Like *superadditivity* in surplus sharing games [3], subadditivity means it is always best to form the grand coalition.

The Cost of Stability. In many games, the core is empty. However an external subsidy lowers the cost of the grand coalition, thus creating a different game, possibly with a nonempty core. Formally, the *adjusted game* is an expense sharing game $G(\Delta) = \langle I, c' \rangle$ where $c'(I) = c(I) - \Delta$, and $c'(S) = c(S)$ for all $S \neq I$. A payment vector whose sum is less than $c(I)$ is referred to as a *subimputation*. Thus, imputations in the adjusted game are subimputations of the original game. Also, the imputation \mathbf{p} in $G(\Delta)$

is blocked by coalition S iff the subimputation \mathbf{p} in G is blocked by S . In the example given in the introduction, $c(I)$ was reduced from 9 (million) to $c'(I) = 6$, while for all $S \subsetneq I$, $c'(S) = c(S) = 5$. Naturally, we would like the external payment to be as small as possible. The *Cost of Stability* (CoS) in an expense sharing game is the minimal non-negative payment Δ s.t. the game $G(\Delta)$ has a non-empty core.³ Thus the CoS can be formulated as the optimal solution of a linear program with exponentially many constraints, similarly to the approach in surplus sharing games; see [3] for details.

From the definition, we have that $0 \leq CoS(G) \leq c(I)$. The worst case ($CoS(G) = c(I)$) occurs when the cost of any coalition (other than I) is 0. We can derive a closed form for the Cost of Stability using *balanced collections*. Let $\delta_S \in \mathbb{R}_+$ be the coefficient of coalition S . $\delta = \{\delta_S\}_{S \in \mathcal{S}}$ is called balanced if for every agent i , $\sum_{S:i \in S} \delta_S = 1$.

Theorem 1 (Bondareva-Shapley theorem). *The game $G = \langle I, c \rangle$ has a non-empty core iff any balanced collection of coefficients satisfies: $\sum_{S \in \mathcal{S}} \delta_S \cdot c(S) \geq c(I)$.*

For a proof and more detailed discussion, see for example [22]. By applying the theorem on the adjusted game, we can write the CoS of games with empty cores as follows:

$$CoS(G) = c(I) - \min_{\text{balanced } \delta} \sum_{S \in \mathcal{S}} \delta_S \cdot c(S) \tag{1}$$

If the righthand-side term is negative (nonempty core), then $CoS(G) = 0$. Unfortunately, in the general case (1) does not provide an efficient way to compute the CoS.

3 Expense Games vs. Surplus Games

We refer to transferable utility games with positive utilities as *surplus games*. In a *surplus game* $G_v = \langle I, v \rangle$, $v(S) \in \mathbb{R}_+$ is the utility that coalition S can generate, and an imputation is a division of $v(I)$ among all agents. There has been much interest in solution concepts for surplus games, as well as in their relation to expense games.

Duality. The *dual* (not to be confused with linear duality) of the surplus game $G_v = \langle I, v \rangle$ is the expense game G_c defined as:

$$c(S) = \begin{cases} v(S) & , S = I \\ v(I) - v(I \setminus S) & , S \neq I \end{cases}$$

G_v 's core is empty iff G_c 's core is empty [8]. Duality also preserves monotonicity.

As in expense games, the CoS of a surplus game is the minimal amount that needs to be added to $v(I)$, so that the adjusted game has a non-empty core. Surplus games have already been studied [3,23,2], so one might ask whether expense games deserve special treatment. Further, we might conjecture that it is possible to derive the CoS of an expense game by analyzing its dual, i.e., that there is some function f such that $CoS(G_c) = f(CoS(G_v))$. Unfortunately, some important properties *are not preserved*

³ Following [3], we define the CoS w.r.t. the *additive difference* between $c(I)$ and $c'(I)$. Related papers use the *multiplicative ratio* between the costs; the transformation is straightforward.

in the dual. For example, a game G_v can be superadditive, while its dual G_c is not sub-additive nor superadditive. Furthermore, although $CoS(G_c) = 0$ implies $CoS(G_v) = 0$ (and vice versa), the following example shows that the CoS of one problem does not reveal much information about the CoS of its dual. Consider G_v, G_c s.t. $c(I) = v(I) = 1$, and the cost of all singletons in G_c is $c(\{i\}) = 0$. This means that $CoS(G_c) = 1$, as no agent has any incentive to contribute anything. This only constrains the value of coalitions of size $n - 1$ in the dual game G_v to 1. $CoS(G_v)$ can still be as low as $\frac{1}{n-1}$ (if all other values are 0), or as high as $n - 1$ (if $v(\{i\}) = 1$ for all agents).

4 Anonymous Games

An anonymous expense sharing game is characterized by a cost function $c : [n] \rightarrow \mathbb{R}_+$, i.e., the cost of S is $c(|S|)$. In the anonymous case, Equation (1) can be simplified:

Theorem 2. *Let G be an anonymous game. $CoS(G) = c(n) - n \cdot \min_{k \leq n} \frac{c(k)}{k}$.*

Proof. We first show that there is an optimal subimputation (i.e., whose sum is minimal) in which all agents pay the same amount. Let \mathbf{p}^* be an optimal subimputation, and define a new subimputation \mathbf{q} , with $q_j = \frac{1}{n} \sum_{i \in I} p_i^*$ for all j . For each coalition S , we have that $q(S) = \frac{|S|}{n} \sum_{i \in I} p_i^*$. Denote by \mathcal{S}_k the set of all coalitions of size k ; then from the stability of \mathbf{p}^* ,

$$\forall S \in \mathcal{S}_k, q(S) \leq \max\{p^*(S') \leq c(k) : S' \in \mathcal{S}_k\},$$

Thus \mathbf{q} is also stable, and therefore a legal subimputation ($q(I) = p^*(I)$).

Since a coalition of size k has to pay at least $c(k)$, every agent i has to pay at least its fair share in the best possible coalition, i.e., $q_i = \min_{k \leq n} \frac{c(k)}{k}$. □

Without further assumptions, the CoS of an anonymous game can still reach the trivial upper bound of $c(n)$, for example if $n = 2, c(1) = 0; c(2) = 1$.

We now consider subadditivity in anonymous games, i.e., assume that $c(s + t) \leq c(s) + c(t)$. The following theorem shows that in such games the subsidy (CoS) will be approximately half of the total cost. This is similar to the corresponding result on superadditive anonymous surplus games, in which the CoS was shown to be roughly twice the value of the grand coalition [3].

Theorem 3. *Let G be a subadditive anonymous expense game.*

$CoS(G) \leq \left(\frac{1}{2} - \frac{1}{2n-2}\right) c(n)$, and this bound is tight. That is, there is a subadditive anonymous expense game for which this is exactly the CoS.

Proof. For $n \leq 2$ the theorem is trivial. Thus assume $n \geq 3$. $c(n) = c(\frac{n}{k} \cdot k) \leq \lceil \frac{n}{k} \rceil c(k)$, which means that $n \frac{c(k)}{k} \geq \frac{n}{k} \lceil \frac{n}{k} \rceil c(n)$ for any k , and in particular for $k^* = \operatorname{argmin} \frac{c(k)}{k}$.

We denote $\frac{n}{k^*}$ by a . Note that $a \geq \frac{n}{n-1} > 1$, thus $\lceil a \rceil \geq 2$. We first look at the case $\lceil a \rceil \geq 3$. This means that $a > 2$, and thus (for $n \geq 4$)

$$\frac{a}{\lceil a \rceil} \geq \frac{a}{a+1} \geq \frac{2}{3} \geq \frac{n}{2n-2}.$$

The alternative case is $\lceil a \rceil = 2$. Here, $a = \frac{n}{n-1}$ minimizes the expression $\frac{a}{\lceil a \rceil}$ (since the denominator is fixed), and we get that $\frac{a}{\lceil a \rceil} \geq \frac{n/(n-1)}{2} = \frac{n}{2n-2}$. Note that for $n = 3$ we are either in the second case, or $k^* = 1$, and thus $\frac{a}{\lceil a \rceil} = \frac{3}{3} = 1 > \frac{n}{2n-2}$ also holds.

We showed that in any case $\frac{a}{\lceil a \rceil} \geq \frac{n}{2n-2}$, thus:

$$n \frac{c(k^*)}{k^*} \geq \frac{n}{k^*} \frac{1}{\lceil \frac{n}{k^*} \rceil} c(n) = \frac{a}{\lceil a \rceil} c(n) \geq \frac{n}{2n-2} c(n) \Rightarrow$$

$$CoS(G) = c(n) - n \frac{c(k^*)}{k^*} \leq \left(1 - \frac{n}{2n-2}\right) c(n) = \left(\frac{1}{2} - \frac{1}{2n-2}\right) c(n).$$

For tightness, consider a game where $c(n) = 2$, and $c(k) = 1$ for any $k < n$. In this game $k^* = n - 1$, and by using Theorem 2

$$CoS(G) = c(n) - n \frac{1}{n-1} = c(n) - c(n) \left(\frac{n}{2(n-1)}\right) = c(n) \left(1 - \frac{n}{2n-2}\right) \quad \square$$

5 Facility Games

We now describe a specific domain on which we demonstrate our approach. Later, we use results for this domain to derive general results for expense games.

Facility Games (also known as Set-Cover Games) are closely related to the MinSet-Cover problem. In the MinSetCover problem, we are given a set $I = \{1, \dots, n\}$, a family of subsets $F = \{A_1, \dots, A_m\} \subseteq \mathcal{S}$, and a weight function $w : F \rightarrow \mathbb{R}_+$. We are asked to find the lightest group J s.t. $\bigcup_{j \in J} A_j = I$. We denote the the optimal set cover by $F^*(I)$ and its value by $opt(I, F) = w(F^*(I)) = \sum_{j \in F^*(I)} w_j$. We assume that each element is contained in at least one set, so $opt(I, F)$ is well-defined.

This algorithmic problem has a natural variant as an expense game: the agents are the elements, and each set represents a *facility* capable of giving service to the agents (corresponding to the elements in the set). The expense of a coalition is the minimal total price of facilities it must buy so that all of its members are served.

Formally, a facility game is a tuple $G = \langle I, F, w \rangle$. The cost function is defined as $c(S) = opt(S, F|_S)$, where $F|_S = \{A_j \cap S \text{ s.t. } A_j \in F\}$. We also denote by $F^*(S) \subseteq F$ the optimal cover of S ; thus $w(F^*(S)) = c(S)$.

The hospital example given in the introduction is a facility game with three agents (the hospitals). As the following lemma shows, facility games are highly expressive.

Lemma 1. *Facility games are subadditive. Furthermore, any subadditive expense game can be described as a facility game.*

Proof. We first prove subadditivity. Let $S, T \in \mathcal{S}$ be distinct coalitions; then $F^*(S) \cup F^*(T)$ is a cover of $S \cup T$. Thus $c(S \cup T) \leq w(F^*(S) \cup F^*(T)) \leq w(F^*(S)) + w(F^*(T)) = c(S) + c(T)$.

In the other direction, every subadditive game has a naïve formulation as a facility game with an exponential number of facilities: we add a facility for each coalition, whose price is the cost of the coalition. As the original game is subadditive, the cost of a coalition in the new game is exactly the price of its corresponding facility. \square

The CoS is tightly coupled with the key concept of the *integrality gap*. Consider the cost of the grand coalition $c(I)$. This is the optimal solution of the MinSetCover problem $\langle I, F, \mathbf{w} \rangle$, which can be written as the following integer linear program, over the variables $\{y_j\}_{A_j \in F}$:

$$\begin{aligned} \min \sum_{A_j \in F} w_j y_j & \quad \text{subject to:} \\ \sum_{j:i \in A_j} y_j & \geq 1 & \quad \text{for each } i \in I, \\ y_j & \in \{0, 1\} & \quad \text{for each } A_j \in F. \end{aligned}$$

In any returned solution, the facility A_j is part of the cover $F^*(I)$ if $y_j = 1$. The linear relaxation of this program is obtained by relaxing the last condition and allowing $y_j \in [0, 1]$. The difference between the optimal integer solution and the optimal fractional solution is known as the *integrality gap* of the problem.⁴

Formally, we denote by $ILP(G)$ ($= c(I)$) and $LP(G)$ the value of the optimal integer and fractional solutions of the linear program corresponding to the facility game $G = \langle I, F, \mathbf{w} \rangle$, and define the integrality gap of G as $IG(G) = ILP(G) - LP(G)$. We use the following equality, which is a known folk theorem; and also supply a simple proof, demonstrating how the optimal subimputation can be computed efficiently.

Theorem 4. *Let G be a facility game. $CoS(G) = IG(G)$.*

Proof. We define the following linear program over the variables $\{p_i\}_{i \in I}$, which is the dual program of $LP(G)$:

$$\begin{aligned} \max \sum_{i \in I} p_i & \quad \text{subject to:} \\ p_i & \in [0, 1] & \quad \text{for each } i \in I, \\ \sum_{i \in A_j} p_i & \leq w_j & \quad \text{for each } A_j \in F. \end{aligned}$$

Denote by \mathbf{p}^* the optimal assignment to the dual variables $\{p_i\}_{i \in I}$, and their sum (which is the optimal value of the dual) by $\hat{LP}(G)$. From strong duality we have that:

$$\sum_{i \in I} p_i^* = \hat{LP}(G) = LP(G) = \sum_{A_j \in F} w_j y_j^*, \tag{2}$$

where \mathbf{y}^* is the optimal solution vector of the primal linear (fractional) program. We can read the dual LP as “the maximal sum of payments, such that all agents belonging to

⁴ We use the term “integrality gap” to denote the *difference* between the solutions, rather than their *ratio*. See also Footnote [3](#)

a set A_j pay at most w_j together⁵. Consider a coalition S with cost $c(S)$. By definition of the cost function, there is a partial cover $F'(S) = \{A_1, \dots, A_k\}$ whose cost is $c(S) = w(F'(S)) = \sum_{r=1}^k w_r$. Note that:

$$\sum_{i \in S} p_i^* \leq \sum_{r=1}^k \sum_{i \in A_r} p_i^* \leq \sum_{r=1}^k w_r = c(S).$$

That is, the vector \mathbf{p}^* is a legal subimputation in G , as it is not blocked by any coalition S . Furthermore, any other subimputation \mathbf{p} must obey the constraints of the dual program (otherwise there is a coalition $S = A_j$ that pays more than the cost of its corresponding facility), and therefore $\sum_{i \in I} p_i \geq \sum_{i \in I} p_i^*$ must hold, so it is not possible that there are better solutions (subimputations).

By definition, the CoS is the gap between the $c(I)$ and the maximal payment. Thus, from (2): $CoS(G) = c(I) - \sum_{i \in I} p_i^* = IG(G)$. □

The proof also reveals the connection with the Bondareva-Shapley theorem: if all agents pay something, then due to complementary slackness, the optimal solution vector \mathbf{y} is a balanced collection of coefficients (if there is an agent that pays 0, we can remove him and obtain a solution \mathbf{y}').

The integrality gap of integer programs, and of MinSetCover in particular, is well-studied in the literature. We can therefore use known bounds on the IG, and apply them to the CoS. The following, for example, it due to Lovász (17):

$$IG(G) \leq c(I) \left(1 - \frac{1}{\ln(d) + 1} \right), \tag{3}$$

where d is the size of the largest set in F (5).

By joining (3) and Lemma 1 to Theorem 4, we have the first part of the following corollary:

Theorem 5. *For any subadditive game $G = \langle I, c \rangle$, $CoS(G) \leq c(I) \left(1 - \frac{1}{\ln(n)+1} \right)$, and this bound is tight, up to a constant.*

The tightness is due to an example by Vazirani (25), showing that the integrality gap of MinSetCover can be as high as $\frac{\log_2(n)}{2}$.

It is interesting to compare this bound to the corresponding bound for superadditive surplus games, which depends on the square root of n , rather than on the logarithm (3).

6 Coalition Structures

Although the CoS of subadditive games is bounded, not all expense games are subadditive. Furthermore, in such games it is not guaranteed that forming the grand coalition is optimal (cheapest). For example, we can think of variants of facility/routing games, where there is an additional cost for using fewer facilities, or for constructing networks with a high branching factor due to increased congestion.

⁵ The greedy cost-sharing scheme of (12) obtains a similar cost-recovery ratio in the worst case (and is also strategyproof), but is inferior to the result of the dual in other cases.

In such cases, the external authority may be interested in stabilizing the structure that minimizes the *social cost*, i.e., the total expenses. It is not hard to see that the same structure also minimizes the subsidy, as stability constraints for each coalition remain the same. See [3] for further discussion of this point.

Formally, a *coalition structure* is a partition of I to distinct coalitions $CS = \{T_j\}_{j=1}^k$ s.t. $\bigcup_{T_j \in CS} T_j = I$. The set $\mathcal{T}(A)$ contains all partitions of the set A , thus for $A = I$, we get the set of all coalition structures $\mathcal{CS} = \mathcal{T}(I)$. The cost of a coalition structure $CS \in \mathcal{CS}(I)$ in the game $G = \langle I, c \rangle$ is $c(CS) = \sum_{T_j \in CS} c(T_j)$. The CS -core of G (denoted by $CORE(G, CS)$) contains all subimputations such that: (a) $c(S) \geq p(S)$ for all coalitions (i.e., \mathbf{p} is stable); and (b) for each $T_j \in CS$, $p(T_j) = c(T_j)$ (i.e., no transfer of payments between coalitions). In the adjusted game $G(CS, \Delta)$ the cost of each $T_j \in CS$ is subsidized by $\Delta_j \geq 0$, whereas the cost of any other coalition or coalition structure remains the same. As in the case of the grand coalition, $CoS(CS, G)$ is the minimal sum of $\Delta = (\Delta_1, \dots, \Delta_k)$ s.t. the CS -core of $G(CS, \Delta)$ is nonempty.

Finally, let $\hat{CS} \in \mathcal{CS}$ be the structure that minimizes the cost $c(\hat{CS})$. Then the minimal amount required to stabilize the cheapest structure in the game, is defined as $CoS_{CS}(G) = CoS(\hat{CS}, G)$.

Theorem 6. For any expense game $G = \langle I, c \rangle$, $CoS_{CS}(G) \leq c(\hat{CS}) \left(1 - \frac{1}{\ln(n)+1}\right)$.

Proof. We define the *subadditive closure* of a game $G = \langle I, c \rangle$, as a coalitional game $G^* = \langle I, c^* \rangle$, whose cost function is $c^*(S) = \min_{T \in \mathcal{T}(S)} c(T)$. In particular, for $S = I$ we get:

$$c^*(I) = \min_{CS \in \mathcal{T}(I)} c(CS) = c(\hat{CS}). \tag{4}$$

It is easy to see that $\langle I, c^* \rangle$ is a subadditive coalitional game, since any non-overlapping coalitions S_1, S_2 are a partition of the coalition $S_1 \cup S_2$.

Lemma 2 (Aumann and Dréze [1]). Let $G = \langle I, c \rangle$ be a coalitional game, and let $CS \in \mathcal{CS}$ be a coalitional structure.

1. If $c^*(I) = c(CS)$ then the CS -core of G is equal to the core of G^* .
2. Otherwise (i.e., if $c^*(I) < c(CS)$), then the CS -core of G is empty.

Let Δ be some fixed subsidy vector whose sum is $\Delta \geq 0$, and consider the optimal coalition structure \hat{CS} . From [4] we have that $c^*(I) - \Delta = c(\hat{CS}) - \Delta$, and thus from Lemma [2] $CORE(G(\Delta)^*) = CORE(G(\hat{CS}, \Delta), \hat{CS})$. In particular, Δ stabilizes $G(\Delta)^*$ iff Δ stabilizes $G(\hat{CS}, \Delta)$. Finally, from Theorem [5]

$$\begin{aligned} CoS_{CS}(G) &= CoS(G, \hat{CS}) = CoS(G^*) \leq \left(1 - \frac{1}{\ln(n)+1}\right) c^*(I) \\ &= \left(1 - \frac{1}{\ln(n)+1}\right) c(\hat{CS}), \end{aligned} \tag{from [4]}$$

since G^* is a subadditive expense game. □

⁶ Aumann and Dréze treated superadditive *surplus games*. However, a slight variation of their work proves the lemma.

Our final result in this section shows that the subadditivity condition in Theorem 5 can in fact be weakened. Since $\{I\}$ is also a coalition structure, we get the following result for non-subadditive games as a corollary of Theorem 6.

Theorem 7. *If $c(I) \leq c(CS)$ for all $CS \in \mathcal{CS}$, then $CoS(G) \leq c(I) \left(1 - \frac{1}{\ln(n)+1}\right)$.*

7 Some Notes on Computational Complexity

In some seemingly simple TU games, such as weighted voting games and threshold network flow games, finding (or even testing) if a subimputation is stable proved to be \mathcal{NP} -hard, while computing the value of a coalition was trivial. In particular, computing the CoS was hard in these games [3][23].

Facility games present a situation opposite to that in the above mentioned surplus games: it is \mathcal{NP} -hard to compute the cost of a coalition, and in particular to know if some set of facilities $F' \subseteq F$ is optimal for the grand coalition. This follows directly from the hardness of the MinSetCover problem [25]. However, the optimal subimputation \mathbf{p}^* can be computed efficiently by solving $LP(G)$.

The computational complexity results regarding the CoS lead to an interesting trade-off between the computational power of the center and the size of the subsidy. Since the maximal costs the agents can safely pay (\mathbf{p}^*) do not depend on the quality of the selected set of facilities, faster computers can assist the city council in finding cheaper solutions (better F') and thereby save money on subsidies (i.e., lower $w(F') - \sum_i p_i^*$).

It is important to note that the runtime of the mechanism would be polynomial in the *description size* of the game, i.e., in the number of facilities. Therefore we cannot efficiently compute optimal payments for arbitrary subadditive expense games with the dual method, as the linear program might contain an exponential number of constraints.

The optimal coalition structure in anonymous games. Recall that for anonymous games the characteristic function is given by $c : [n] \rightarrow \mathbb{R}_+$, and for every coalition structure, $c(CS) = \sum_{C \in CS} c(|C|)$. Computing the CoS of a given coalition structure is easy. As in general games, \mathbf{p}^* is easy to compute. In the anonymous case, it has the form $p_i^* = \min_{j \leq n} \frac{c(j)}{j}$ for all i . Thus for any given CS we can compute $CoS(G, CS)$ as

$$CoS(G, \hat{CS}) = c(CS) - n \cdot \min_{j \leq n} \frac{c(j)}{j}. \tag{5}$$

The proof is the same as the proof of Theorem 2 which is a special case for $CS = \{I\}$.

However, finding the *optimal* coalition structure might be difficult. That is, we know how much money each agent should pay in total, but we do not know how much they can make by themselves, and therefore we are not sure how much to subsidize.

Proposition 1. *Computing $CoS_{CS}(G)$ for anonymous games is in \mathcal{P} .*

We note here that Proposition 1 also holds for surplus sharing games, as defined in [3].

Proof. From (5) it is sufficient to find the optimal coalition structure \hat{CS} .

Our key observation is that the problem of finding \hat{CS} in an anonymous game is equivalent to solving KNAPSACK with bounded weights.

In the KNAPSACK problem, we are given n pairs $\langle w_i, x_i \rangle$, and a threshold t . We can select any pair $a_i \in \mathbb{N}$ times, in order to minimize the total weight of the sack $\sum_i a_i \cdot w_i$, while maintaining the total value above the threshold, i.e., $\sum_i a_i \cdot x_i \geq t$. While the general KNAPSACK problem is \mathcal{NP} -hard, it can be solved by a simple dynamic algorithm, provided that either the weights (w_i) or the values (x_i) are polynomially bounded (see e.g., [18]).

Consider an anonymous cost game, with a characteristic function c . We construct a KNAPSACK instance KN with the pairs $\{\langle c(i), i \rangle\}_{i=1}^n$, and a threshold $t = n$ (hence the values are bounded). As $w_i = c(i)$, and $x_i = i$, we have that

$$CoS_{CS}(G) = \min \sum_{i \leq n} a_i \cdot w_i \quad s.t. \quad \sum_{i \leq n} a_i \cdot x_i = n = t \quad ,$$

which is the optimal knapsack solution.

Thus for anonymous coalitional games we have a dynamic algorithm that finds \hat{CS} and computes $CoS_{CS}(G)$ efficiently. □

8 Discussion

Our study of minimal subsidies in expense sharing games joins together two lines of work: the ongoing study of cost sharing mechanisms with different requirements for specific game classes, and the clean formulation of bounds on the Cost of Stability that depend on the cost function’s properties.

We have focused on the subadditivity property, and were able to provide tight bounds on the CoS for broad families of games, even in the absence of an efficient cost sharing mechanism. Our work also complements previous work on the Cost of Stability, highlighting the similarities and differences between cost sharing games and surplus games in terms of the magnitude of the required subsidy for achieving stability.

Several issues remain open for future research. First, better bounds on the CoS should be developed for specific classes of games, such as those suggested in [14,5,7,9,2]. It is particularly interesting if dropping all requirements except stability (such as strategyproofness or computational efficiency) can result in mechanisms that are better than the existing cost sharing mechanisms described in the related work section.

Another question for future research is whether the CoS of more coalitional games, other than Set-cover and Steiner-tree games, can also be derived from the integrality gap of their underlying combinatorial problem (when there is one).

Finally, we plan to further investigate the relation between the CoS and other solution concepts, such as the Shapley Value, the nucleolus, and the least core.

Acknowledgments

This work was partially supported by Israel Science Foundation grant #898/05, and Israel Ministry of Science and Technology grant #3-6797.

⁷ The KNAPSACK problem is typically formulated as maximization of the value, rather than minimization of the weights (and this version is in fact used in the symmetric proof for surplus games). However, the problems are clearly computationally equivalent.

References

1. Aumann, R.J., Dréze, J.H.: Cooperative games with coalition structures. *International Journal of Game Theory* 3, 217–237 (1974)
2. Aziz, H., Brandt, F., Harrenstein, P.: Monotone cooperative games and their threshold versions. In: *AAMAS 2010*, pp. 1017–1024 (2010)
3. Bachrach, Y., Elkind, E., Meir, R., Pasechnik, D., Zuckerman, M., Rothe, J., Rosenschein, J.S.: The cost of stability in coalitional games. In: *Mavronicolas, M., Papadopoulou, V.G. (eds.) SAGT 2009*. LNCS, vol. 5814, pp. 122–134. Springer, Heidelberg (2009)
4. Bachrach, Y., Meir, R., Zuckerman, M., Rothe, J., Rosenschein, J.S.: The cost of stability in weighted voting games (extended abstract). In: *AAMAS 2009*, pp. 1289–1290 (2009)
5. Bachrach, Y., Rosenschein, J.S.: Coalitional skill games. In: *AAMAS 2008*, pp. 1023–1030 (2008)
6. Bachrach, Y., Rosenschein, J.S.: Power in threshold network flow games. *Autonomous Agents and Multi-Agent Systems* 18(1), 106–132 (2009)
7. Bachrach, Y., Rosenschein, J.S., Porat, E.: Power and stability in connectivity games. In: *AAMAS 2008*, pp. 999–1006 (2008)
8. Bilbao, J.M.: *Cooperative Games on Combinatorial Structures*. Kluwer Publishers, Dordrecht (2000)
9. Brânzei, S., Larson, K.: Coalitional affinity games and the stability gap. In: *IJCAI 2009*, pp. 79–84 (2009)
10. Buchbinder, N., Lewin-Eytan, L., Naor, J(S.), Orda, A.: Non-cooperative cost sharing games via subsidies. In: *Monien, B., Schroeder, U.-P. (eds.) SAGT 2008*. LNCS, vol. 4997, pp. 337–349. Springer, Heidelberg (2008)
11. Claus, A., Kleitman, D.J.: Cost allocation for a spanning tree. *Networks* 3, 289–304 (1973)
12. Devanur, N.R., Mihail, M., Vazirani, V.V.: Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems* 39, 11–22 (2005)
13. Granot, D., Huberman, G.: Minimum cost spanning tree games. *Mathematical Programming* 21, 1–18 (1981)
14. Ieong, S., Shoham, Y.: Marginal contribution nets: a compact representation scheme for coalitional games. In: *ACM EC 2005*, pp. 193–202 (2005)
15. Immorlica, N., Mahdian, M., Mirrokni, V.S.: Limitations of cross-monotonic cost sharing schemes. In: *SODA 2005*, pp. 602–611 (2005)
16. Jain, K., Vazirani, V.V.: Applications of approximation algorithms to cooperative games. In: *STOC 2001*, pp. 364–372 (2001)
17. Lovász, L.: On the ratio of optimal integral and fractional covers. *Discrete Mathematics* 13, 383–390 (1975)
18. Martello, S., Toth, P.: *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, Chichester (1990)
19. Megiddo, N.: Cost allocation for steiner trees. *Networks* 8, 1–6 (1978)
20. Monderer, D., Tennenholtz, M.: k-implementation. In: *ACM EC 2003*, pp. 19–28 (2003)
21. Pál, M., Tardos, É.: Group strategy proof mechanisms via primal-dual algorithms. In: *FOCS 2003*, pp. 584–593 (2003)
22. Peleg, B., Sudhölter, P.: *Introduction to the Theory of Cooperative Games*. Kluwer Publishers, Dordrecht (2003)
23. Resnick, E., Bachrach, Y., Meir, R., Rosenschein, J.S.: The cost of stability in network flow games. In: *Kráľovič, R., Niwiński, D. (eds.) MFCS 2009*. LNCS, vol. 5734, pp. 636–650. Springer, Heidelberg (2009)
24. Skorin-Kapov, D.: On the core of the minimum cost steiner tree game in networks. *Annals of Operations Research* 57, 233–249 (1995)
25. Vazirani, V.V.: *Approximation Algorithms*. Springer, Heidelberg (2001)

Author Index

- Adsul, Bharat 18, 30
Anshelevich, Elliot 42
Auletta, Vincenzo 54
Aumann, Yonatan 66
Awasthi, Pranjal 78
- Babu, Ch. Sobhan 18, 30
Bachrach, Yoram 347
Balcan, Maria-Florina 78
Bilò, Vittorio 90
Blum, Avrim 78
Brandt, Felix 102
Busch, Costas 222
- Caragiannis, Ioannis 90
Caskurlu, Bugra 42
- Daskalakis, Constantinos 114
de Keijzer, Bart 335
Dombb, Yair 66
- Elkind, Edith 210
- Fabrikant, Alex 126
Fanelli, Angelo 90
Feige, Uriel 138, 150
Ferraioli, Diodato 54
Fiat, Amos 1
Fischer, Felix 102
Fotakis, Dimitris 162
Frongillo, Rafael 114
- Gairing, Martin 174
Garg, Jugal 18, 30
Goel, Gagan 186, 234
Goldberg, Paul W. 15
- Hansen, Kristoffer Arnsfelt 198
Harrenstein, Paul 102
Hate, Ameya 42
Hazon, Noam 210
Hofer, Martin 312
- Jaggard, Aaron D. 126
- Kannan, Rajgopal 222
Karakostas, George 162
- Karande, Chinmay 234
Kolliopoulos, Stavros G. 162
Kuleshov, Volodymyr 246
- Larson, Kate 262
- Macko, Martin 262
Mehta, Ruta 18, 30
Meir, Reshef 347
Mihalák, Matúš 276
Miltersen, Peter Bro 198
Monaco, Gianpiero 90
Mossel, Elchanan 288
- Nadav, Uri 300
- Papadimitriou, Christos H. 1, 114
Pasquale, Francesco 54
Persiano, Giuseppe 54
Pierrakos, George 114
Piliouras, Georgios 300
- Rosenschein, Jeffrey S. 347
- Savani, Rahul 174
Schäfer, Guido 335
Schapira, Michael 126
Schlegel, Jan Christoph 276
Sheffet, Or 78
Skopalik, Alexander 312
Sohoni, Milind 18, 30
Sørensen, Troels Bjerre 198
Steskal, Euboš 262
- Talgam-Cohen, Inbal 138
Tamuz, Omer 288
Telelis, Orestis A. 335
Tennenholtz, Moshe 150
- Valiant, Gregory 114
Vazirani, Vijay V. 186, 323
Vempala, Santosh 78
Vetta, Adrian 246
- Wang, Lei 234