

Yang Cai · Adrian Vetta (Eds.)

ARCOSS

LNCS 10123

# Web and Internet Economics

**12th International Conference, WINE 2016  
Montreal, Canada, December 11–14, 2016  
Proceedings**



Springer

*Commenced Publication in 1973*

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

## Editorial Board

David Hutchison, UK

Josef Kittler, UK

Friedemann Mattern, Switzerland

Moni Naor, Israel

Bernhard Steffen, Germany

Doug Tygar, USA

Takeo Kanade, USA

Jon M. Kleinberg, USA

John C. Mitchell, USA

C. Pandu Rangan, India

Demetri Terzopoulos, USA

Gerhard Weikum, Germany

## Advanced Research in Computing and Software Science

Subline of Lecture Notes in Computer Science

## Subline Series Editors

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

Vladimiro Sassone, *University of Southampton, UK*

## Subline Advisory Board

Susanne Albers, *TU Munich, Germany*

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

Bernhard Steffen, *University of Dortmund, Germany*

Deng Xiaotie, *City University of Hong Kong*

Jeannette M. Wing, *Microsoft Research, Redmond, WA, USA*

More information about this series at <http://www.springer.com/series/7407>

Yang Cai · Adrian Vetta (Eds.)

# Web and Internet Economics

12th International Conference, WINE 2016  
Montreal, Canada, December 11–14, 2016  
Proceedings



*Editors*

Yang Cai  
School of Computer Science  
McGill University  
Montreal, QC  
Canada

Adrian Vetta  
Department of Mathematics and Statistics,  
and School of Computer Science  
McGill University  
Montreal, QC  
Canada

ISSN 0302-9743                      ISSN 1611-3349 (electronic)  
Lecture Notes in Computer Science  
ISBN 978-3-662-54109-8              ISBN 978-3-662-54110-4 (eBook)  
DOI 10.1007/978-3-662-54110-4

Library of Congress Control Number: 2016959799

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© Springer-Verlag GmbH Germany 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature  
The registered company is Springer-Verlag GmbH Germany  
The registered company address is: Heidelberger Platz 3, 14197 Berlin, Germany

# Preface

This volume contains the papers and extended abstracts presented at WINE 2016: the 12th Conference on Web and Internet Economics, held during December 11–14, 2016, in Montreal, Canada.

Over the past decade, researchers in theoretical computer science, artificial intelligence, and microeconomics have joined forces to tackle problems involving incentives and computation. These problems are of particular importance in application areas like the Web and the Internet that involve large and diverse populations. The Conference on Web and Internet Economics (WINE) is an interdisciplinary forum for the exchange of ideas and results on incentives and computation arising from these various fields.

WINE 2016 built on the previous success of the series, held annually from 2005 to 2015 with published archival proceedings.

WINE 2016 accepted 35 papers. All submissions were rigorously peer reviewed and evaluated on the basis of originality, soundness, significance, and exposition. The program also included three invited talks by Kevin Leyton-Brown (University of British Columbia), Christos Papadimitriou (University of California at Berkeley), and Rakesh Vohra (University of Pennsylvania). In addition WINE 2016 featured three tutorials by Hu Fu (University of British Columbia), Brendan Lucier (Microsoft Research at New England), and Ruta Mehta (University of Illinois at Urbana-Champaign).

We would like to thank our sponsors, CRM, Facebook, Microsoft Research, Google, GERAD, and Springer, for their generous financial support. We are very grateful to Louis Pelletier and the administrative staff at the CRM for their assistance in the organization of the event.

We acknowledge the work of the Program Committee for their hard work. Special thanks goes to Vasilis Gkatzelis for chairing the poster session. In addition we would like to acknowledge Springer for their help with the proceedings, and the EasyChair paper management system.

November 2016

Yang Cai  
Adrian Vetta

# Organization

## Program Committee

|                       |   |
|-----------------------|---|
| Moshe Babaioff        | Microsoft Research, Israel                                  |
| Xiaohui Bei           | Nanyang Technological University, Singapore                 |
| Gerardo Berbeglia     | Melbourne Business School, Australia                        |
| Nicolas Bousquet      | Ecole Centrale de Lyon, France                              |
| Yang Cai              | McGill University, Canada                                   |
| Nikhil Devanur        | Microsoft Research, Redmond, USA                            |
| Edith Elkind          | University of Oxford, UK                                    |
| John Fearnley         | University of Liverpool, UK                                 |
| Felix Fischer         | University of Glasgow, UK                                   |
| Hu Fu                 | University of British Columbia, Canada                      |
| Martin Gairing        | University of Liverpool, UK                                 |
| Jugal Garg            | University of Illinois at Urbana-Champaign, USA             |
| Vasilis Gkatzelis     | Drexel University, USA                                      |
| Paul Goldberg         | University of Oxford, UK                                    |
| Nick Gravin           | Massachusetts Institute of Technology, USA                  |
| Nima Haghpanah        | Penn State University, USA                                  |
| Zhiyi Huang           | The University of Hong Kong, SAR China                      |
| Elias Koutsoupias     | University of Oxford, UK                                    |
| Kate Larson           | University of Waterloo, Canada                              |
| Ron Lavi              | Technion – Israel Institute of Technology, Israel           |
| Pinyan Lu             | Shanghai University of Finance and Economics, China         |
| Brendan Lucier        | Microsoft Research, New England, USA                        |
| Yishay Mansour        | Tel Aviv University, Israel                                 |
| Vangelis Markakis     | Athens University of Economics and Business, Greece         |
| Aranyak Mehta         | Google Research, USA  |
| Denis Nekipelov       | University of Virginia, USA                                 |
| Neil Olver            | Vrije Universiteit Amsterdam, The Netherlands               |
| Sigal Oren            | Ben-Gurion University, Israel                               |
| Malleesh Pai          | Rice University, USA  |
| Georgios Piliouras    | Singapore University of Technology and Design,<br>Singapore |
| Guido Schaefer        | CWI Amsterdam, The Netherlands                              |
| Balasubramanian Sivan | Google Research, USA  |
| Eric Sodomka          | Facebook, USA   |
| Nicolas Stier-Moses   | Facebook, USA   |
| Vasilis Syrgkanis     | Microsoft Research, New England, USA                        |
| Pingzhong Tang        | Tsinghua University, China                                  |

|               |                                       |
|---------------|---------------------------------------|
| Laszlo Vegh   | London School of Economics, UK        |
| Adrian Vetta  | McGill University, Canada             |
| Matt Weinberg | Princeton University, USA             |
| Chris Wilkens | Yahoo Research, USA                   |
| Lirong Xia    | Rensselaer Polytechnic Institute, USA |

## Additional Reviewers

|                          |                         |                         |
|--------------------------|-------------------------|-------------------------|
| Abeliuk, Andres          | Goldner, Kira           | Panageas, Ioannis       |
| Alaei, Saeed             | Gonczarowski, Yannai A. | Pountourakis, Emmanouil |
| Amanatidis, Georgios     | Gur, Yonatan            | Psomias,                |
| An, Bo                   | Harks, Tobias           | Christos-Alexandros     |
| Badanidiyuru Varadaraja, | Hauser, Daniel          | Ravi, Nishant           |
| Ashwinkumar              | Ho, Chien-Ju            | Rogers, Ryan            |
| Bilò, Davide             | Igarashi, Ayumi         | Schewior, Kevin         |
| Bilò, Vittorio           | Karimi, Majid           | Schrijvers, Okke        |
| Brill, Markus            | Kesselheim, Thomas      | Sekar, Shreyas          |
| Burq, Maximilien         | Kleer, Pieter           | Simon, Sunil Easaw      |
| Cavallo, Ruggiero        | Kliemann, Lasse         | Swamy, Chaitanya        |
| Chen, Wei                | Klimm, Max              | Telelis, Orestis        |
| Chestnut, Stephen        | Ko, Young Kun           | Todo, Taiki             |
| Cheung, Yun Kuen         | Kotsialou, Grammateia   | Tzamos, Christos        |
| Cooper, Frances          | Kroer, Christian        | Uetz, Marc              |
| Cseh, Ágnes              | Lang, Jérôme            | Vaidya, Tushar          |
| Cummings, Rachel         | Leme, Renato            | von Stengel, Bernhard   |
| De Keijzer, Bart         | Lev, Omer               | Wang, Junxing           |
| Deligkas, Argyrios       | Li, Minming             | Wang, Zihe              |
| Delikas, Argyrios        | Li, Yunan               | Wojtczak, Dominik       |
| Deng, Yuan               | Lianas, Thanasis        | Wright, James           |
| Dobzinski, Shahar        | Liu, Yicheng            | Xiao, Tao               |
| Eden, Alon               | Lund, Troels Bjerre     | Yazdanbod, Sadra        |
| Farczadi, Linda          | Lykouris, Thodoris      | Zampetakis, Emmanouil   |
| Feldmann, Andreas Emil   | Maillé, Patrick         | Zhang, Hanrui           |
| Ferraioli, Diodato       | Mattei, Nicholas        | Zhang, Hong             |
| Filos-Ratsikas, Aris     | Meir, Reshef            | Zhang, Jialin           |
| Freeman, Rupert          | Meng, DeLong            | Zhang, Qiang            |
| Friedler, Ophir          | Miller, Benjamin        | Zubrickas, Robertas     |
| Frongillo, Rafael        | Niazadeh, Rad           | Zuo, Song               |
| Giannakopoulos, Yiannis  | Obraztsova, Svetlana    |                         |

# Contents

|  |     |
|--|-----|
| Computing Equilibria with Partial Commitment . . . . .   | 1   |
| <i>Vincent Conitzer</i>  |     |
| Distributed Methods for Computing Approximate Equilibria. . . . .  | 15  |
| <i>Artur Czumaj, Argyrios Deligkas, Michail Fasoulakis, John Fearnley, Marcin Jurdziński, and Rahul Savani</i> |     |
| Inapproximability Results for Approximate Nash Equilibria . . . . .  | 29  |
| <i>Argyrios Deligkas, John Fearnley, and Rahul Savani</i>  |     |
| Multilinear Games. . . . .   | 44  |
| <i>Hau Chan, Albert Xin Jiang, Kevin Leyton-Brown, and Ruta Mehta</i>  |     |
| Power-Law Distributions in a Two-Sided Market and Net Neutrality. . . . .                                      | 59  |
| <i>Xiaotie Deng, Zhe Feng, and Christos H. Papadimitriou</i>   |     |
| On-Demand or Spot? Selling the Cloud to Risk-Averse Customers . . . . .  | 73  |
| <i>Darrell Hoy, Nicole Immorlica, and Brendan Lucier</i>   |     |
| Buying Data from Privacy-Aware Individuals: The Effect of Negative Payments . . . . .                          | 87  |
| <i>Weina Wang, Lei Ying, and Junshan Zhang</i>   |     |
| Bidding Strategies for Fantasy-Sports Auctions. . . . .  | 102 |
| <i>Aris Anagnostopoulos, Ruggiero Cavallo, Stefano Leonardi, and Maxim Sviridenko</i>                          |     |
| Competitive Equilibria for Non-quasilinear Bidders in Combinatorial Auctions . . . . .                         | 116 |
| <i>Rad Niazadeh and Christopher A. Wilkens</i>   |     |
| Correlated and Coarse Equilibria of Single-Item Auctions . . . . .   | 131 |
| <i>Michal Feldman, Brendan Lucier, and Noam Nisan</i>  |     |
| Pricing to Maximize Revenue and Welfare Simultaneously in Large Markets. . . . .                               | 145 |
| <i>Elliot Anshelevich, Koushik Kar, and Shreyas Sekar</i>  |     |
| A Prior-Independent Revenue-Maximizing Auction for Multiple Additive Bidders. . . . .                          | 160 |
| <i>Kira Goldner and Anna R. Karlin</i>   |     |

|  |     |
|--|-----|
| Optimal Mechanism for Selling Two Items to a Single Buyer Having Uniformly Distributed Valuations . . . . .                    | 174 |
| <i>Thirumulanathan D., Rajesh Sundaresan, and Y. Narahari</i>  |     |
| Anonymous Auctions Maximizing Revenue . . . . .  | 188 |
| <i>Christos Tzamos and Christopher A. Wilkens</i>  |     |
| Revenue Maximizing Envy-Free Pricing in Matching Markets with Budgets. . .   | 207 |
| <i>Riccardo Colini-Baldeschi, Stefano Leonardi, and Qiang Zhang</i>  |     |
| Conference Program Design with Single-Peaked and Single-Crossing Preferences . . . . .   | 221 |
| <i>Dimitris Fotakis, Laurent Gourvès, and Jérôme Monnot</i>  |     |
| Truthful Facility Assignment with Resource Augmentation: An Exact Analysis of Serial Dictatorship . . . . .                    | 236 |
| <i>Ioannis Caragiannis, Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, Kristoffer Arnsfelt Hansen, and Zihan Tan</i> |     |
| Putting Peer Prediction Under the Micro(economic)scope and Making Truth-Telling Focal . . . . .                                | 251 |
| <i>Yuqing Kong, Katrina Ligett, and Grant Schoenebeck</i>  |     |
| Truthful Mechanisms for Matching and Clustering in an Ordinal World. . . . .   | 265 |
| <i>Elliot Anshelevich and Shreyas Sekar</i>  |     |
| Computer-Aided Verification for Mechanism Design . . . . .   | 279 |
| <i>Gilles Barthe, Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, Aaron Roth, and Pierre-Yves Strub</i>                |     |
| Smoothness for Simultaneous Composition of Mechanisms with Admission. . . . .  | 294 |
| <i>Martin Hoefer, Thomas Kesselheim, and Bojana Kodric</i>   |     |
| Motivating Time-Inconsistent Agents: A Computational Approach . . . . .  | 309 |
| <i>Susanne Albers and Dennis Kraft</i>   |     |
| FPT Approximation Schemes for Maximizing Submodular Functions . . . . .  | 324 |
| <i>Piotr Skowron</i>   |     |
| Bounds for the Convergence Time of Local Search in Scheduling Problems. . .  | 339 |
| <i>Tobias Brunsch, Michael Etscheid, and Heiko Röglin</i>  |     |
| On the Price of Stability of Undirected Multicast Games . . . . .  | 354 |
| <i>Rupert Freeman, Samuel Haney, and Debmalya Panigrahi</i>  |     |
| Efficiency and Budget Balance . . . . .  | 369 |
| <i>Swaprava Nath and Tuomas Sandholm</i>   |     |

The Core of the Participatory Budgeting Problem . . . . . 384  
*Brandon Fain, Ashish Goel, and Kamesh Munagala*

Approximating Gains-from-Trade in Bilateral Trading . . . . . 400  
*Liad Blumrosen and Yehonatan Mizrahi*

Coverage, Matching, and Beyond: New Results on Budgeted  
Mechanism Design . . . . . 414  
*Georgios Amanatidis, Georgios Bimpas, and Evangelos Markakis*

Strategic Network Formation with Attack and Immunization . . . . . 429  
*Sanjeev Goyal, Shahin Jabbari, Michael Kearns, Sanjeev Khanna,  
and Jamie Morgenstern*

Opinion Formation Games with Dynamic Social Influences . . . . . 444  
*Vittorio Bilò, Angelo Fanelli, and Luca Moscardelli*

Complex Contagions on Configuration Model Graphs  
with a Power-Law Degree Distribution . . . . . 459  
*Grant Schoenebeck and Fang-Yi Yu*

**Abstracts**

The Magician’s Shuffle: Reusing Lottery Numbers for School  
Seat Redistribution . . . . . 475  
*Itai Feigenbaum, Yash Kanoria, Irene Lo, and Jay Sethuraman*

Near-Efficient Allocation Using Artificial Currency in Repeated Settings  
(Extended Abstract). . . . . 476  
*Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer*

Multi-unit Facility Location Games . . . . . 478  
*Omer Ben-Porat and Moshe Tennenholtz*

**Author Index** . . . . . 481

# Computing Equilibria with Partial Commitment

Vincent Conitzer<sup>(✉)</sup>

Duke University, Durham, NC, USA  
conitzer@cs.duke.edu

**Abstract.** In security games, the solution concept commonly used is that of a Stackelberg equilibrium where the defender gets to commit to a mixed strategy. The motivation for this is that the attacker can repeatedly observe the defender’s actions and learn her distribution over actions, before acting himself. If the actions were not observable, Nash (or perhaps correlated) equilibrium would arguably be a more natural solution concept. But what if some, but not all, aspects of the defender’s actions are observable? In this paper, we introduce solution concepts corresponding to this case, both with and without correlation. We study their basic properties, whether these solutions can be efficiently computed, and the impact of additional observability on the utility obtained.

## 1 Introduction

Algorithms for computing game-theoretic solutions have long been of interest, but were for a long time not deployed in real-world applications (at least if we do not count, e.g., computer poker programs—for an overview of those, see Sandholm [21]—as real-world applications). This changed in 2007 with a series of deployed applications coming out of Milind Tambe’s TEAMCORE research group at the University of Southern California. The games in question are what are now called security games, where a defender has to allocate limited resources to defend certain targets or patrol a certain area, and an attacker chooses a target to attack. The deployed applications include airport protection [20], assigning Federal Air Marshals to flights [22], patrolling in ports [2], fare inspection in transit systems [25], and patrolling to prevent wildlife poaching [11].

While most of the literature on computing game-theoretic solutions has focused on the computation of Nash equilibria—including the breakthrough result that even computing a single Nash equilibrium is PPAD-complete [6, 10]—in the security games applications the focus is instead on computing an optimal mixed strategy to commit to [8]. In this model, one player (in security games, the defender) chooses a mixed strategy, and the other (the attacker) observes this mixed strategy and best-responds to it. This sometimes helps, and never hurts, the former player [23]. Intriguingly, in two-player normal-form games, such a strategy can be computed in polynomial time via linear programming [8, 23]. Another benefit of this model is that it sidesteps issues of equilibrium selection that the approach of computing (say) a Nash equilibrium might face.

---

I dedicate this paper to my sister Jessica, her fiancé Jeremy, and their upcoming full commitment. I wish them a lifetime of happiness.



Such technical conveniences aside, the standard motivation for assuming that the defender in security games can commit to a mixed strategy is as follows. The defender has to choose a course of action every day. The attacker, on the other hand, does not, and can observe the defender’s actions over a period of time. Thus, the defender can establish a reputation for playing any particular mixed strategy. This can be beneficial for the defender: whereas in a simultaneous-move model (say, using Nash equilibrium as the solution concept), she can play only best responses to the attacker’s strategy, in the commitment model she can commit to play something that is not a best response, which may incentivize the attacker to play something that is better for the defender. Of course, for this argument to work, it is crucial that the attacker observes over time which actions the defender takes before taking any action himself. Previous work has questioned this and considered models where there is uncertainty about whether the attacker observes the defender’s actions at all [14, 15], as well as models where the attacker only gets a limited number of observations [1, 19].

In this paper, we consider a different setting where some defender actions are (externally) indistinguishable from each other. This captures, for example, the case where there are both observable and unobservable security measures, as is often the case. Here, two courses of action are indistinguishable if and only if they differ only in the unobservable component. It also captures the case where a guard can be assigned to a visible location (1), or to one of two invisible locations (2 or 3). In this case, the first action is distinguishable from the latter two, but the latter two are indistinguishable from each other. Indistinguishability is an equivalence relation that partitions the player’s strategy space; we call one element of this partition a SIS (subset of indistinguishable strategies). Thus, the defender can establish a reputation for playing a particular distribution over the SISes. However, she cannot establish any reputation for how she plays *within* each SIS, because this is not externally observable. Thus, intuitively, when the defender plays from a particular SIS, she needs to play a strategy that, within that SIS, is a best response; however, if there is another strategy in a *different* SIS that is a better response, that is not a problem, because deviating to that strategy would be observable.

The specific contributions of this paper are as follows. We formalize solution concepts for these settings that generalize both Nash and correlated equilibrium, as well as the basic Stackelberg model with (full) commitment to mixed strategies. Further contributions include illustrative examples of these solutions, basic properties of the concepts, analysis of their computational complexity, and analysis of how the row player (defender)’s utility varies as a function of the amount of commitment power (as measured by observability).

## 2 Definitions and Basic Properties

We are now ready to define some basic concepts. Throughout, the row player (player 1) is the player with (some) commitment power, in the sense of being able to build a reputation.  $R$  denotes the set of rows,  $C$  the set of columns, and  $\sigma_1$  and  $\sigma_2$  denote mixed strategies over these, respectively.

**Definition 1.** A subset of indistinguishable strategies (SIS)  $S$  is a maximal subset of  $R$  such that for any two rows  $r_1, r_2 \in S$ , the column player's observation is identical for  $r_1$  and  $r_2$ . Let  $\mathcal{S}$  denote the set of all SISes, constituting a partition of  $R$ . Given a mixed strategy  $\sigma_1$  for the row player and a SIS  $S$ , let  $\sigma_1(S) = \sum_{r \in S} \sigma_1(r)$  (where  $\sigma_1(r)$  is the probability  $\sigma_1$  puts on  $r$ ).

Since our focus is on games in which one player can build up a reputation and the other cannot, we do not consider SISes for the column player. Equivalently, we consider all the column player's strategies to be in the same SIS.

**Definition 2.** Two mixed strategies  $\sigma_1, \sigma'_1$  are indistinguishable to the column player if for all  $S \in \mathcal{S}$ ,  $\sigma_1(S) = \sigma'_1(S)$ .

**Example.** Consider the following game:

|   | A   | B   |
|---|-----|-----|
| a | 7,0 | 2,1 |
| b | 6,1 | 0,0 |
| c | 5,0 | 0,1 |
| d | 4,1 | 1,0 |

If the players move simultaneously, then  $a$  is a strictly dominant strategy and we obtain  $(a, B)$  as the iterated strict dominance solution (and hence the unique Nash equilibrium), with a utility of 2 for the row player. If the row player gets to commit to a mixed strategy, then she could commit to play  $a$  and  $b$  with probability  $1/2$  each, inducing the column player to play  $A$ ,<sup>1</sup> resulting in a utility of 6.5 for the row player. (Even committing to a pure strategy—namely,  $b$ —would result in a utility of 6.) Now suppose  $\mathcal{S} = \{\{a, b\}, \{c, d\}\}$ , i.e.,  $a$  and  $b$  are indistinguishable and so are  $c$  and  $d$ . In this case, playing  $a$  and  $b$  with probability  $1/2$  each (or playing  $b$  with probability 1) is indistinguishable from playing  $a$  with probability 1. Hence, it is not credible that the row player would ever play  $b$ , given that  $a$  is a strictly dominant strategy. But can the row player still do better than always playing  $a$  (and thereby inducing the column player to play  $B$ )?

We will return to this example shortly, but first we need to formalize the idea of a deviation that cannot be detected by the column player.

**Definition 3.** A profile  $(\sigma_1, \sigma_2)$  has no undetectable beneficial deviations if (1) for all  $\sigma'_2$ ,  $u_2(\sigma_1, \sigma'_2) \leq u_2(\sigma_1, \sigma_2)$ , and (2) for all  $\sigma'_1$  indistinguishable from  $\sigma_1$ ,  $u_1(\sigma'_1, \sigma_2) \leq u_1(\sigma_1, \sigma_2)$ .

The following simple proposition points out that this is equivalent to the column player only putting probability on best responses, and the row player only putting probability on rows that *within their SIS* are best responses.

<sup>1</sup> As is commonly assumed in this model, ties for the column player are broken in the row player's favor; if not, the row player can simply commit to  $1/2 - \epsilon$  on  $a$  and  $1/2 + \epsilon$  on  $b$ .

**Proposition 1.** *A profile  $(\sigma_1, \sigma_2)$  has no undetectable beneficial deviations if and only if (1) for all  $c, c' \in C$  with  $\sigma_2(c) > 0$ ,  $u_2(\sigma_1, c') \leq u_2(\sigma_1, c)$ , and (2) for all  $S \in \mathcal{S}$ , for all  $r, r' \in S$  with  $\sigma_1(r) > 0$ ,  $u_1(r', \sigma_2) \leq u_1(r, \sigma_2)$ .*

**Example Continued.** In the game above, consider the profile

$$(((1/2)c, (1/2)d), ((1/2)A, (1/2)B))$$

This profile has no undetectable deviations: (1) the column player is playing a best response, and (2) the only undetectable deviations for the row player do not put any probability on  $\{a, b\}$ , and  $c$  and  $d$  are both equally good responses.

Note that a profile that has no undetectable beneficial deviations may still not be stable, in the sense that player 1 may prefer to deviate to a mixed strategy that is in fact distinguishable from  $\sigma_1$ , and build up a reputation for playing that strategy instead. But in a sense, these profiles are *feasible* solutions for the row player: *given* that the row player decides to build up a reputation for the distribution over SISes resulting from  $\sigma_1$ , the profile  $(\sigma_1, \sigma_2)$  is stable. This is similar to the sense in which in the regular Stackelberg model, any profile consisting of a mixed strategy for the row player and a best response for the column player is feasible: the row player may not have had good reason to commit to that particular mixed strategy, but *given* that she did, the profile is stable. In fact, this just corresponds to the special case of our model where all rows are distinguishable.

**Proposition 2.** *If  $|\mathcal{S}| = 1$  (all rows are indistinguishable), then a profile has no undetectable beneficial deviations if and only if it is a Nash equilibrium of the game. If  $|\mathcal{S}| = |\mathcal{R}|$  (all rows are distinguishable), then a profile has no undetectable beneficial deviations if and only if the column player is best-responding.*

We can now define an optimal solution.

**Definition 4.** *A profile with no undetectable beneficial deviations is a Stackelberg equilibrium with limited observation (SELO) if among such profiles it maximizes the row player's utility.*

**Example Continued.** In the game above, consider the profile

$$(((1/2)a, (1/2)d), ((1/2)A, (1/2)B))$$

This profile has no undetectable deviations:  $A$  and  $B$  are both best responses for the column player, and the row player strictly prefers  $a$  to  $b$  and is indifferent between  $c$  and  $d$ . It gives the row player utility 3.5. We now argue that it is in fact a SELO. First, note that a SELO must put at least probability 1/2 on  $d$ : for, if it did not, then, because the row player would never play  $b$ , the column player would strictly prefer  $B$ , which would result in lower utility for the row player. Second, the column player must play  $B$  at least half the time, because otherwise, the row player would strictly prefer  $c$  to  $d$ —but if the row player only plays  $a$

and  $c$ , the column player would strictly prefer  $B$ . Under these two constraints, the row player would be best off having as much as possible of the remaining probabilities on  $a$  and  $A$ , and this results in the profile above.

**Proposition 3.** *If  $|\mathcal{S}| = 1$  (all rows are indistinguishable), then a profile is a SELO if and only if it is a Nash equilibrium that maximizes the row player's utility among Nash equilibria. If  $|\mathcal{S}| = |\mathcal{R}|$  (all rows are distinguishable), then a profile is a SELO if and only if it is a Stackelberg equilibrium (with full observation).*

### 3 Computational Results

We now consider the complexity of computing a SELO. We immediately obtain:

**Corollary 1.** *When  $|\mathcal{S}| = 1$ , computing a SELO is NP-hard (and the maximum utility for the row player in a profile with no undetectable beneficial deviations is inapproximable unless  $P = NP$ ).*

*Proof.* By Propositions 2 and 3, these problems are equivalent to maximizing the row player's utility in a Nash equilibrium, which is known to be NP-hard and inapproximable [9, 13].

This still leaves open the question of whether the problem becomes easier if the individual SISes have small size. Unfortunately, the next result shows that the problem remains NP-hard and inapproximable in this case. This motivates extending the model to one that allows correlation, as we will do in Sect. 4.

**Theorem 1.** *Computing a SELO remains NP-hard even when  $|\mathcal{S}| = 2$  for all  $S \in \mathcal{S}$  (and in fact it is NP-hard to check whether there exists a profile with no undetectable beneficial deviations that gives the row player positive utility, even when all payoffs are nonnegative).*

*Proof.* We reduce from the EXACT-COVER-BY-3-SETS problem, in which we are given a set of elements  $T$  ( $|T| = m$ , with  $m$  divisible by 3) and subsets  $T_j \subseteq T$  that each satisfy  $|T_j| = 3$ , and are asked whether there exist  $m/3$  of these subsets that together cover all of  $T$ . For an arbitrary instance of this problem, we construct the following game. For each  $T_j$ , we add a SIS consisting of two rows,  $\{T_j^+, T_j^-\}$ , as well as a column  $T_j$ . For each element  $t \in T$ , we add a column  $t$ . The utility functions are as follows.

- $u_1(T_j^+, T_j) = m/3$  for any  $j$
- $u_1(T_j^+, T_{j'}) = 0$  for any  $j, j'$  with  $j \neq j'$
- $u_1(T_j^-, T_{j'}) = 1$  for any  $j, j'$
- $u_1(r, t) = 0$  for any row  $r$  and element  $t$
- $u_2(r, T_j) = m/3 - 1$  for any row  $r$  and any  $j$
- $u_2(T_j^+, t) = 0$  for any  $j$  and  $t \in T_j$
- $u_2(r, t) = m/3$  for any element  $t$  and row  $r$  that is not some  $T_j^+$  with  $t \in T_j$

First suppose the EXACT-COVER-BY-3-SETS instance has a solution. Let the row player play uniformly over the  $m/3$  corresponding rows  $T_j^+$ , and the column player uniformly over the  $m/3$  corresponding columns  $T_j$ . The row player's expected utility for any of the rows in her support is 1; deviating to the corresponding  $T_j^-$  would still only give her 1. The column player's expected utility is  $m/3 - 1$  for any  $T_j$ ; because the row player plays an exact cover, deviating to any  $t$  gives him expected utility  $(m/3)(m/3 - 1)/(m/3) = m/3 - 1$ . So this profile has no undetectable beneficial deviations (in fact it is a Nash equilibrium) and gives the row player an expected utility of 1.

Now suppose that the game has a SELO in which the row player gets positive utility, which implies that the column player puts total probability  $p > 0$  on his  $T_j$  columns. It follows that for every  $t \in T$ , the total probability that the row player puts on rows  $T_j^+$  with  $t \in T_j$  is at least  $3p/m$ , or otherwise the column player would strictly prefer playing  $t$  to playing any  $T_j$ . However, note that the row player can only put positive probability on rows  $T_j^+$  where the corresponding column  $T_j$  receives probability at least  $3p/m$  (thereby resulting in expected utility at least  $p$  for the row player for playing  $T_j^+$ ), because otherwise the corresponding row  $T_j^-$  (which is indistinguishable) would be strictly preferable (resulting in expected utility  $p$ ). But of course there can be at most  $m/3$  such columns  $T_j$ , and these  $T_j$  must cover all the elements  $t$  by what we said before. Hence the EXACT-COVER-BY-3-SETS instance has a solution.

## 4 Adding Signaling

The notion of correlated equilibrium [4] results from augmenting a game with a trusted mediator that sends correlated signals to the agents. As is well known, without loss of generality, we can assume the signal that an agent receives is simply the action she is to take. This is for the following reason. If a correlated equilibrium relies on an agent randomizing among multiple actions conditional on receiving a particular signal, then we may as well have the mediator do this randomization on behalf of the agent before sending out the signal. It is well known that correlated equilibria can outperform Nash equilibria from all agents' perspectives. For example, consider Shapley's game, which is a version of rock-paper-scissors where choosing the same action as the other counts as a loss.

|     | $A$ | $B$ | $C$ |
|-----|-----|-----|-----|
| $a$ | 0,0 | 1,0 | 0,1 |
| $b$ | 0,1 | 0,0 | 1,0 |
| $c$ | 1,0 | 0,1 | 0,0 |

Whereas the only Nash equilibrium of this game is for both players to randomize uniformly (resulting in 0,0 payoffs 1/3 of the time), there is a correlated equilibrium that only results in the 1,0 and 0,1 outcomes, each 1/6 of the time. That is, if the mediator is set up to draw one of these six entries uniformly at

random, and then tell each agent what she is supposed to play (but not what the other is supposed to play), then each agent has an incentive to follow the recommendation: doing so will result in a win half the time, and it is not possible to do better given what the agent knows.

Correlated equilibria are easier to compute than Nash equilibria: given a game in normal form, there is a linear program formulation for computing even optimal correlated equilibria (say, ones that maximize the row player’s utility). The linear program presented later in Fig. 1 is closely related.

Similar signaling has received attention in the Stackelberg model. One may assume a more powerful leader in this model that can commit not only to taking actions in a particular way, but also to sending signals in a way that is correlated with how she takes actions. (Again, the motivation for using this in real applications might be that over time the leader develops a reputation for sending out signals according to a particular distribution, and playing particular distributions over actions conditional on those signals.) Because the leader can commit to sending signals in a particular way, there is no need to introduce an independent mediator entity in this context. As it turns out, in a two-player normal-form game this additional power does not buy the leader anything, but with more players it does [7]. Such signaling can also help in Bayesian games [24] and stochastic games [18], both from the perspective of increasing the leader’s utility and from the perspective of making the computation easier.

It is straightforward to see that signaling can be useful in our limited commitment model as well. For example, if we just take Shapley’s game with  $|\mathcal{S}| = 1$ , then by Proposition 3 without signaling we are stuck with the Nash equilibrium, but it seems we should be able to obtain the improved correlated equilibrium outcome with some form of signaling. But what is the right model of signaling here? We consider a very powerful model of signaling in this version of the paper. The full version of the paper also contains a discussion of weaker signaling models.

**Definition 5.** *In the trusted mediator model, the row player can design an independent trusted mediator that sends signals privately to each player according to a pre-specified joint distribution. After the round of play has completed, the mediator publicly reveals the signal sent to the row player.*

The after-the-fact public revelation of the signal sent to the row player allows the row player to commit to (i.e., in the long run develop a reputation for) responding to each signal with a particular distribution of play. Specifically, after each completed round, the column player learns the signal sent to, and the SIS played by, the row player.<sup>2</sup> Thus, if the row player according to the signal

<sup>2</sup> It is easy to get confused here—does the column player not learn more in a round purely by virtue of his own payoff from that round? It is important to remember that we are not considering repeated play by the column player. The idea is that the column player can observe over time the signals and how the row player acts *before* the column player ever acts. For discussion of security contexts in which certain types of players can receive messages that are inaccessible to other types, see Xu et al. [24].

that she received was supposed to play an action from a particular SIS, then the column player can verify that she did. However, the row player may have an incentive to deviate *within* a SIS, because this is undetectable.

In the appendix of the full version of the paper, we show that under the trusted mediator model, without loss of generality a signal consists of just an action to play. With this in mind, we now define formally what it means for a correlated profile to have no undetectable beneficial deviations.

**Definition 6.** *A correlated profile  $\sigma$  has no undetectable beneficial deviations if (1) for all  $c, c' \in C$  with  $\sum_{r \in R} \sigma(r, c) > 0$ , we have  $\sum_{r \in R} \sigma(r, c)(u_2(r, c) - u_2(r, c')) \geq 0$ , and (2) for all  $S \in \mathcal{S}$ , for all  $r, r' \in S$  with  $\sum_{c \in C} \sigma(r, c) > 0$ , we have  $\sum_{c \in C} \sigma(r, c)(u_1(r, c) - u_1(r', c)) \geq 0$ .*

Note that, as is well known in the formulation of correlated equilibrium, in the first inequality, we can use  $\sigma(r, c)$  rather than the more cumbersome  $\sigma(r, c) / \sum_{r'' \in R} \sigma(r'', c)$ , which would be the conditional probability of seeing  $r$  given a signal of  $c$ , because the denominator is a constant (similar for the second inequality). As a result, the condition that  $\sum_{r \in R} \sigma(r, c) > 0$  is in fact not necessary because the inequality is vacuously true otherwise. This is what allows the standard linear program formulation of correlated equilibrium, as well as the linear program we present below in Fig. 1.

**Definition 7.** *A correlated profile with no undetectable beneficial deviations is a Stackelberg equilibrium with signaling and limited observation (SESLO) if among such profiles it maximizes the row player's utility.*

**Example.** Consider the following game:

|   | A    | B    | C    | D   |
|---|------|------|------|-----|
| a | 0,0  | 12,0 | 0,1  | 0,0 |
| b | 0,1  | 0,0  | 12,0 | 0,0 |
| c | 12,0 | 0,1  | 0,0  | 0,0 |
| d | 5,0  | 5,0  | 5,0  | 0,1 |
| e | 7,0  | 7,0  | 7,0  | 1,1 |

Suppose  $\mathcal{S} = \{\{a, b, c, d\}, \{e\}\}$ . Then the following correlated profile (in which the signal an agent receives is which action to take) is a SESLO:

$$((1/9)(a, B), (1/9)(a, C), (1/9)(b, A), (1/9)(b, C), (1/9)(c, A), (1/9)(c, B), \\ (1/9)(e, A), (1/9)(e, B), (1/9)(e, C))$$

With this profile, for any signal the column player can receive, following the signal will give him utility 1/3, and so will any deviation. For any signal the row player receives in SIS  $\{a, b, c, d\}$ , following the signal will give her 6; deviating to  $a$ ,  $b$ , or  $c$  will give either 0 or 6, and deviating to  $d$  will give 5.

The row player obtains utility  $19/3$  from this profile.<sup>3</sup> In contrast, without any commitment (if  $|\mathcal{S}|$  had been 1), the outcome  $(e, D)$  would have been a SESLO, giving the row player utility only 1. Also, without signaling (but still with  $\mathcal{S} = \{\{a, b, c, d\}, \{e\}\}$ ), the outcome  $(e, D)$  would have been a SELO. For consider a mixed-strategy profile without any undetectable beneficial deviations, and suppose it puts positive probability on at least one of  $A, B$ , and  $C$ . Then at least one of  $a, b$ , and  $c$  must get positive probability as well, for otherwise the column player would be better off playing  $D$ . Because  $a, b$ , and  $c$  are all in the same SIS and perform equally well against  $D$ , and because  $A, B$ , and  $C$  all perform equally well against  $d$  and  $e$ , if we condition on the players playing from  $a, b, c$  and  $A, B, C$ , the result must be a Nash equilibrium of that  $3 \times 3$  game, which means that all of  $A, B$ , and  $C$  get the same probability. But in that case,  $d$  (which is in the same SIS) is a better response for the row player, and we have a contradiction. Hence any SELO involves the column player always playing  $D$  and the most the row player can obtain is 1.

We next have the following simple proposition that the ability to signal never hurts the row player.

**Proposition 4.** *The row player’s utility from a SESLO is always at least that of a SELO.*

*Proof.* We show that an uncorrelated profile  $(\sigma_1, \sigma_2)$  that has no undetectable deviations (in the sense of Definition 3) also has no undetectable deviations (in the sense of Definition 6) when interpreted as a correlated profile  $\sigma$  (with  $\sigma(r, c) = \sigma_1(r)\sigma_2(c)$ ); the result follows. First, for all  $c, c' \in C$  with  $\sum_{r \in R} \sigma(r, c) > 0$  (which is equivalent to  $\sigma_2(c) > 0$ ), we have  $\sum_{r \in R} \sigma(r, c)(u_2(r, c) - u_2(r, c')) = \sigma_2(c) \sum_{r \in R} \sigma_1(r)(u_2(r, c) - u_2(r, c')) = \sigma_2(c)(u_2(\sigma_1, c) - u_2(\sigma_1, c')) \geq 0$  by the best-response condition of Definition 3. Similarly, for all  $S \in \mathcal{S}$ , for all  $r, r' \in S$  with  $\sum_{c \in C} \sigma(r, c) > 0$  (which is equivalent to  $\sigma_1(r) > 0$ ), we have  $\sum_{c \in C} \sigma(r, c)(u_1(r, c) - u_1(r', c)) = \sigma_1(r) \sum_{c \in C} \sigma_2(c)(u_1(r, c) - u_1(r', c)) = \sigma_1(r)(u_1(r, \sigma_2) - u_1(r', \sigma_2)) \geq 0$  by the best-response-within-a-SIS condition of Definition 3.

**Proposition 5.** *If  $|\mathcal{S}| = 1$  (all rows are indistinguishable), then a profile is a SESLO if and only if it is a correlated equilibrium that maximizes the row player’s utility. If  $|\mathcal{S}| = |R|$  (all rows are distinguishable), then a profile is a SESLO if and only if it is a Stackelberg equilibrium with signaling (which can do no better than a Stackelberg equilibrium without signaling).*

## 5 Computational Results

It turns out that with signaling, we do not face hardness. The linear program in Fig. 1 can be used to compute a SESLO. It is a simple modification of the standard linear program for correlated equilibrium, the differences being that (1) for the row player, only deviations within a SIS are considered, and (2) there

<sup>3</sup> This was verified to be optimal using the linear program in Fig. 1; same for the next case.



$$\boxed{
\begin{array}{ll}
\text{maximize} & \sum_{r \in R, c \in C} u_1(r, c)p(r, c) \\
(\forall S \in \mathcal{S}) (\forall r, r' \in S) & \sum_{c \in C} (u_1(r', c) - u_1(r, c))p(r, c) \leq 0 \\
(\forall c, c' \in C) & \sum_{r \in R} (u_2(r, c') - u_2(r, c))p(r, c) \leq 0 \\
& \sum_{r \in R, c \in C} p(r, c) = 1 \\
(\forall r \in R, c \in C) & p(r, c) \geq 0
\end{array}
}$$

**Fig. 1.** Linear program for computing a SESLO.

is an objective of maximizing the row player’s utility. The special case where  $|\mathcal{S}| = |R|$  has no constraints for the row player, and that special case of the linear program has previously been described by Conitzer and Korzhyk [7].

**Theorem 2.** *A SESLO can be computed in polynomial time.*

## 6 The Value of More Commitment Power

More strategies being distinguishable corresponds to more commitment power for the row player. As commitment power (in this particular sense) increases, does the utility the row player can obtain always increase gradually? (Note that it can never *decrease* the row player’s utility, because all it will do is remove constraints in the optimization.) If she has close to full commitment power, does this guarantee her most of the benefit of full commitment power? Is some non-trivial minimal amount of commitment power necessary to obtain much benefit from it? The next two results demonstrate that the answer to all these questions is “no”: there can be big jumps in the utility that the row player can obtain, both on the side close to full commitment power (Proposition 6) and on the side close to no commitment power (Proposition 7). (For an earlier study comparing the value of being able to commit completely to that of not being able to commit at all, see Letchford et al. [17]; for one assessing the value of correlation without commitment, see Ashlagi et al. [3].)

**Proposition 6.** *For any  $\epsilon > 0$  and any  $n > 1$ , there exists an  $n \times (n + 1)$  game with all payoffs in  $[0, 1]$  such that if  $|\mathcal{S}| = |R| = n$ , the row player can obtain utility  $1 - \epsilon$  (even without signaling), but for any  $\mathcal{S}$  with  $|\mathcal{S}| < |R| = n$ , the row player can obtain utility at most  $\epsilon$  (even with signaling).*

*Proof.* Let  $R = \{1, \dots, n\}$  and  $C = \{1, \dots, n + 1\}$ . Let  $u_1(i, j) = i\epsilon/n$  for  $j \leq n$ , and let  $u_1(i, n + 1) = 1 - (n - i)\epsilon/n$ . Let  $u_2(i, j) = (1 + 1/n)/2$  for  $i \neq j$  and  $j \leq n$ , let  $u_2(i, i) = 0$  (for  $i \leq n$ ), and let  $u_2(i, n + 1) = 1/2$  for all  $i$ .

Suppose  $|\mathcal{S}| = |R| = n$ . Then, by Proposition 3, we are in the regular Stackelberg model, and the row player can commit to a uniform strategy, putting probability  $1/n$  on each  $i$ . As a result the expected utility for the column player for playing some  $j \leq n$  is  $((n - 1)/n)(1 + 1/n)/2 = (n - 1)(n + 1)/(2n^2) = (n^2 - 1)/(2n^2) < 1/2$ , so to best-respond he needs to play  $n + 1$ , resulting in a utility for the row player that is greater than  $1 - (n - 1)\epsilon/n > 1 - \epsilon$ .

On the other hand, suppose that  $|\mathcal{S}| < |R| = n$ . Hence there exists some  $S \in \mathcal{S}$  with  $i, i' \in S$ ,  $i < i'$ . Note that  $i'$  strictly dominates  $i$ , so the row player will never play  $i$  in a SELO or even a SESLO. But then, the column player can obtain  $(1 + 1/n)/2 > 1/2$  by playing  $i$ , and hence will not play  $n + 1$ . As a result the row player obtains at most  $n\epsilon/n = \epsilon$ .


**Proposition 7.** *For any  $\epsilon > 0$  and any  $n > 1$ , there exists an  $n \times (n + 1)$  game with all payoffs in  $[0, 1]$  such that for any  $\mathcal{S}$  with  $|\mathcal{S}| > 1$ , the row player can obtain utility  $1 - \epsilon$  (even without signaling), but if  $|\mathcal{S}| = 1$ , the row player can only obtain utility 0 (even with signaling).*

*Proof.* Let  $R = \{1, \dots, n\}$  and  $C = \{1, \dots, n + 1\}$ . Let  $u_1(i, j) = 1 - \epsilon$  for  $i \neq j$  and  $j \leq n$ , let  $u_1(i, i) = 1$  (for  $i \leq n$ ), and let  $u_1(i, n + 1) = 0$  for all  $i$ . Let  $u_2(i, j) = 1$  for  $i \neq j$  and  $j \leq n$ , let  $u_2(i, i) = 0$  (for  $i \leq n$ ), and let  $u_2(i, n + 1) = (n - 1/2)/n$  for all  $i$ .

Suppose  $|\mathcal{S}| > 1$ . Then, the row player can commit to put 0 probability on some  $S \in \mathcal{S}$ , and therefore, 0 probability on some  $i$ . Hence, this  $i$  is a best response for the column player, and the row player obtains  $1 - \epsilon$ . (The row player may be able to do better yet, but this is a feasible solution.)

On the other hand, suppose  $|\mathcal{S}| = 1$ . By Proposition 3, the row player can only obtain the utility of the best Nash equilibrium of the game for her (or, in the case with signaling, the utility of the best correlated equilibrium, by Proposition 5). We now show that in every Nash equilibrium (or even correlated equilibrium) of the game, the column player puts all his probability on  $n + 1$ , from which the result follows immediately. For suppose the column player sometimes plays some  $j \leq n$ . Then, for the row player to best-respond, she has to maximize the probability of choosing the same  $j$  (conditional on the column player playing some  $j \leq n$ ). (Or, more precisely in the case of correlated equilibrium, conditional on receiving any signal that leaves open the possibility that the column player plays some  $j \leq n$ , the row player has to maximize the probability of picking the same  $j$ .) She can always make this probability at least  $1/n$  by choosing uniformly at random. Hence, the column player's expected utility (conditional on playing  $j \leq n$ ) is at most  $(n - 1)/n$ . But then  $n + 1$  is a strictly better response, so we do not have a Nash (or correlated) equilibrium.

Of course, the above two results are extreme cases. Can we say anything about what happens “typically”? To illustrate this, we present the results for randomly generated games in Fig. 2. For each data point, 1000 games of size  $m \times n$  were generated by choosing utilities uniformly at random. The rows were then evenly (round-robin) spread over a given number of SISes, and the game was solved using the GNU Linear Programming Kit (GLPK) with the linear program from Fig. 1. The leftmost points (1 SIS) correspond to no commitment power (best correlated equilibrium), and the rightmost points (at least when the number of SISes is at least  $m$ ) correspond to full commitment power (best Stackelberg mixed strategy). From this experiment, we can observe that most of the value of commitment is already obtained when moving from one SIS to two.



**Fig. 2.** Utility obtained by the row player as a function of commitment power (number of SISes), for various sizes of  $m \times n$  games.

## 7 Conclusion

The model of the defender being able to commit to a mixed strategy has been popular in security games, motivated by the idea that the attacker can learn the distribution over time. This model has previously been questioned, and limited observability has previously been studied in various senses, including the attacker obtaining only a limited number of observations [1, 19] as well as the attacker observing (perfectly) only with some probability [14, 15]. Here, we considered a different type of limited observability, where certain courses of action are distinguishable from each other, but others are not. As a result, the row player’s pure strategies partition into SISes, and she can commit to a distribution over SISes but not to how she plays within each SIS. We showed that it is NP-hard to compute a Stackelberg equilibrium with limited observation in this context, even when the SISes are small (Theorem 1). We then introduced a modified model with signaling and showed that in it, Stackelberg equilibria can be computed in polynomial time (Theorem 2). Finally, we showed that the cost of introducing a bit of additional unobservability can be large both when close to full observability (Proposition 6) and close to no observability (Proposition 7); however, in simulations, introducing a little bit of observability already gives most of the value of full observability.

Future research may be devoted to the following questions. Are there algorithms for computing a SELO that are efficient for special cases of the problem

or that run fast on “typical” games? Another direction for future work concerns learning in games, which is a topic that has been thoroughly studied in the simultaneous-move case (see, e.g., Fudenberg and Levine [12]), but also already to some extent in the mixed-strategy commitment case [5, 16]. A model of learning in games with partial commitment needs to generalize models for both of these cases. Finally, can we mathematically prove what is suggested by the experiment in Fig. 2, namely that in random games most of the value of commitment is already obtained with only two SISes?

**Acknowledgments.** I am thankful for support from ARO under grants W911NF-12-1-0550 and W911NF-11-1-0332, NSF under awards IIS-1527434, IIS-0953756, CCF-1101659, and CCF-1337215, and a Guggenheim Fellowship.

## References

1. An, B., Kempe, D., Kiekintveld, C., Shieh, E., Singh, S.P., Tambe, M., Vorobeychik, Y.: Security games with limited surveillance. In: Proceedings of the National Conference on Artificial Intelligence (AAAI), Toronto, ON, Canada (2012)
2. An, B., Shieh, E., Tambe, M., Yang, R., Baldwin, C., DiRenzo, J., Maule, B., Meyer, G.: PROTECT - a deployed game theoretic system for strategic security allocation for the United States Coast Guard. *AI Mag.* **33**(4), 96–110 (2012)
3. Ashlagi, I., Monderer, D., Tennenholtz, M.: On the value of correlation. *J. Artif. Intell. Res.* **33**, 575–613 (2008)
4. Aumann, R.: Subjectivity and correlation in randomized strategies. *J. Math. Econ.* **1**, 67–96 (1974)
5. Balcan, M.F., Blum, A., Haghtalab, N., Procaccia, A.D.: Commitment without regrets: online learning in Stackelberg security games. In: Proceedings of the ACM Conference on Economics and Computation (EC), Portland, OR, USA, pp. 61–78 (2015)
6. Chen, X., Deng, X., Teng, S.H.: Settling the complexity of computing two-player Nash equilibria. *J. ACM* **56**(3), 14 (2009)
7. Conitzer, V., Korzhyk, D.: Commitment to correlated strategies. In: Proceedings of the National Conference on Artificial Intelligence (AAAI), San Francisco, CA, USA, pp. 632–637 (2011)
8. Conitzer, V., Sandholm, T.: Computing the optimal strategy to commit to. In: Proceedings of the ACM Conference on Electronic Commerce (EC), Ann Arbor, MI, USA, pp. 82–90 (2006)
9. Conitzer, V., Sandholm, T.: New complexity results about Nash equilibria. *Games Econ. Behav.* **63**(2), 621–641 (2008)
10. Daskalakis, C., Goldberg, P., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. *SIAM J. Comput.* **39**(1), 195–259 (2009)
11. Fang, F., Nguyen, T.H., Pickles, R., Lam, W.Y., Clements, G.R., An, B., Singh, A., Tambe, M., Lemieux, A.: Deploying PAWS: field optimization of the protection assistant for wildlife security. In: Proceedings of the Twenty-Eighth Innovative Applications of Artificial Intelligence Conference (IAAI), Phoenix, AZ, USA (2016)
12. Fudenberg, D., Levine, D.: *The Theory of Learning in Games*. MIT Press, Cambridge (1998)
13. Gilboa, I., Zemel, E.: Nash and correlated equilibria: some complexity considerations. *Games Econ. Behav.* **1**, 80–93 (1989)

14. Korzhyk, D., Conitzer, V., Parr, R.: Solving Stackelberg games with uncertain observability. In: Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Taipei, Taiwan, pp. 1013–1020 (2011)
15. Korzhyk, D., Yin, Z., Kiekintveld, C., Conitzer, V., Tambe, M.: Stackelberg vs. nash in security games: an extended investigation of interchangeability, equivalence, and uniqueness. *J. Artif. Intell. Res.* **41**(2), 297–327 (2011)
16. Letchford, J., Conitzer, V., Munagala, K.: Learning and approximating the optimal strategy to commit to. In: Mavronicolas, M., Papadopoulou, V.G. (eds.) SAGT 2009. LNCS, vol. 5814, pp. 250–262. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-04645-2\\_23](https://doi.org/10.1007/978-3-642-04645-2_23)
17. Letchford, J., Korzhyk, D., Conitzer, V.: On the value of commitment. *Auton. Agent. Multi-Agent Syst.* **28**(6), 986–1016 (2014)
18. Letchford, J., MacDermed, L., Conitzer, V., Parr, R., Isbell, C.: Computing optimal strategies to commit to in stochastic games. In: Proceedings of the National Conference on Artificial Intelligence (AAAI), Toronto, ON, Canada, pp. 1380–1386 (2012)
19. Pita, J., Jain, M., Ordóñez, F., Tambe, M., Kraus, S.: Robust solutions to Stackelberg games: addressing bounded rationality and limited observations in human cognition. *Artif. Intell.* **174**(15), 1142–1171 (2010)
20. Pita, J., Jain, M., Western, C., Portway, C., Tambe, M., Ordóñez, F., Kraus, S., Parachuri, P.: Deployed ARMOR protection: the application of a game-theoretic model for security at the Los Angeles International Airport. In: Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Estoril, Portugal, pp. 125–132 (2008)
21. Sandholm, T.: The state of solving large incomplete-information games, and application to poker. *AI Mag.* **31**(4), 13–32 (2010). Special Issue on Algorithmic Game Theory
22. Tsai, J., Rathi, S., Kiekintveld, C., Ordóñez, F., Tambe, M.: IRIS - a tool for strategic security allocation in transportation networks. In: Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Budapest, Hungary, pp. 37–44 (2009)
23. von Stengel, B., Zamir, S.: Leadership games with convex strategy sets. *Games Econ. Behav.* **69**, 446–457 (2010)
24. Xu, H., Freeman, R., Conitzer, V., Dughmi, S., Tambe, M.: Signaling in Bayesian Stackelberg games. In: Proceedings of the Fifteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Singapore, pp. 150–158 (2016)
25. Yin, Z., Jiang, A.X., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Sullivan, J.P.: TRUSTS: scheduling randomized patrols for fare inspection in transit systems using game theory. *AI Mag.* **33**(4), 59–72 (2012)

# Distributed Methods for Computing Approximate Equilibria

Artur Czumaj<sup>1</sup>, Argyrios Deligkas<sup>2</sup>, Michail Fasoulakis<sup>1</sup>, John Fearnley<sup>2</sup>,  
Marcin Jurdziński<sup>1</sup>, and Rahul Savani<sup>2</sup>(✉)

<sup>1</sup> Department of Computer Science and DIMAP, University of Warwick,  
Coventry, UK

<sup>2</sup> Department of Computer Science, University of Liverpool, Liverpool, UK  
{a.deligkas, rahul.savani}@liverpool.ac.uk

**Abstract.** We present a new, distributed method to compute approximate Nash equilibria in bimatrix games. In contrast to previous approaches that analyze the two payoff matrices at the same time (for example, by solving a single LP that combines the two players' payoffs), our algorithm first solves two independent LPs, each of which is derived from one of the two payoff matrices, and then computes an approximate Nash equilibrium using only limited communication between the players. Our method gives improved bounds on the complexity of computing approximate Nash equilibria in a number of different settings. Firstly, it gives a polynomial-time algorithm for computing *approximate well supported Nash equilibria (WSNE)* that always finds a 0.6528-WSNE, beating the previous best guarantee of 0.6608. Secondly, since our algorithm solves the two LPs separately, it can be applied to give an improved bound in the limited communication setting, giving a randomized expected-polynomial-time algorithm that uses poly-logarithmic communication and finds a 0.6528-WSNE, which beats the previous best known guarantee of 0.732. It can also be applied to the case of *approximate Nash equilibria*, where we obtain a randomized expected-polynomial-time algorithm that uses poly-logarithmic communication and always finds a 0.382-approximate Nash equilibrium, which improves the previous best guarantee of 0.438. Finally, the method can also be applied in the query complexity setting to give an algorithm that makes  $O(n \log n)$  payoff queries and always finds a 0.6528-WSNE, which improves the previous best known guarantee of  $2/3$ .

## 1 Introduction

The problem of finding equilibria in non-cooperative games is a central problem in game theory. Nash's seminal theorem proved that every finite normal-form

---

The first, third and fifth author are partially supported by Research partially supported by the Centre for Discrete Mathematics and its Applications (DIMAP) and by EPSRC award EP/D063191/1. The second, fourth and sixth author are supported by EPSRC grant EP/L011018/1. The second author is also supported by ISF grant #2021296. The full version of this paper, with complete proofs, is available at <http://arxiv.org/abs/1512.03315>.

game has at least one *Nash equilibrium* [17], and this raises the natural question of whether we can find one efficiently. After several years of extensive research, it was shown that finding a Nash equilibrium is PPAD-complete [6] even for two-player *bimatrix games* [2], which is considered to be strong evidence that there is no polynomial-time algorithm for this problem.

**Approximate Equilibria.** The fact that computing an exact Nash equilibrium of a bimatrix game is unlikely to be tractable, has led to the study of *approximate Nash equilibria*. There are two natural notions of approximate equilibrium, both of which will be studied in this paper. An  $\epsilon$ -*approximate Nash equilibrium* ( $\epsilon$ -NE) is a pair of strategies in which neither player can increase their expected payoff by more than  $\epsilon$  by unilaterally deviating from their assigned strategy. An  $\epsilon$ -*well-supported Nash equilibrium* ( $\epsilon$ -WSNE) is a pair of strategies in which both players only place probability on strategies whose payoff is within  $\epsilon$  of the best response payoff. Every  $\epsilon$ -WSNE is an  $\epsilon$ -NE but the converse does not hold, so a WSNE is a more restrictive notion.

Approximate Nash equilibria are the more well studied of the two concepts. A line of work has studied the best guarantee that can be achieved in polynomial time [1, 5, 7]. The best algorithm known so far is the gradient descent method of Tsaknakis and Spirakis [19] that finds a 0.3393-NE in polynomial time, and examples upon which the algorithm finds no better than a 0.3385-NE have been found [11]. On the other hand, progress on computing approximate-well-supported Nash equilibria has been less forthcoming. The first correct algorithm was provided by Kontogiannis and Spirakis [15] (which shall henceforth be referred to as the KS algorithm), who gave a polynomial time algorithm for finding a  $\frac{2}{3}$ -WSNE. This was later slightly improved by Fearnley et al. [9] (whose algorithm we shall refer to as the FGSS-algorithm), who gave a new polynomial-time algorithm that extends the KS algorithm and finds a 0.6608-WSNE; prior to this work, this was the best known approximation guarantee for WSNEs. For the special case of symmetric games, there is a polynomial-time algorithm for finding a  $\frac{1}{2}$ -WSNE [4].

Previously, it was considered a strong possibility that there is a PTAS for this problem (either for finding an  $\epsilon$ -NE or  $\epsilon$ -WSNE, since their complexity is polynomially related). A very recent result of Rubinfeld [18] sheds serious doubt on this possibility. END-OF-THE-LINE is the canonical problem that defines the complexity class PPAD. The “Exponential Time Hypothesis” (ETH) for END-OF-THE-LINE says that any algorithm that solves an END-OF-THE-LINE instance with  $n$ -bit circuits, requires  $2^{\Omega(n)}$  time. Rubinfeld’s result says that, subject to the ETH for END-OF-THE-LINE, there exists a constant, but so far undetermined,  $\epsilon^*$ , such that for  $\epsilon < \epsilon^*$ , every algorithm for finding an  $\epsilon$ -NE takes quasi-polynomial time, so the quasi-PTAS of Lipton et al. [16] is optimal.

**Communication Complexity.** Approximate Nash equilibria can also be studied from the *communication complexity* point of view, which captures the amount of communication the players need to find a good approximate Nash equilibrium. It models a natural scenario where the two players each know their own payoff matrix, but do not know their opponent’s payoff matrix. The players must then follow a communication protocol that eventually produces strategies for both

players. The goal is to design a protocol that produces a sufficiently good  $\epsilon$ -NE or  $\epsilon$ -WSNE while also minimizing the amount of communication between the two players.

Communication complexity of equilibria in games has been studied in previous works [3, 14]. The recent paper of Goldberg and Pastink [12] initiated the study of communication complexity in the bimatrix game setting. There they showed  $\Theta(n^2)$  communication is required to find an exact Nash equilibrium of an  $n \times n$  bimatrix game. Since these games have  $\Theta(n^2)$  payoffs in total, this implies that there is no communication-efficient protocol for finding exact Nash equilibria in bimatrix games. For approximate equilibria, they showed that one can find a  $\frac{3}{4}$ -NE *without any communication*, and that in the no-communication setting, finding a  $\frac{1}{2}$ -NE is impossible. Motivated by these positive and negative results, they focused on the most interesting setting, which allows only a poly-logarithmic (in  $n$ ) number of bits to be exchanged between the players. They showed that one can compute 0.438-NE and 0.732-WSNE in this context.

**Query Complexity.** The payoff query model is motivated by practical applications of game theory. It is often the case that we know that there is a game to be solved, but we do not know what the payoffs are, and in order to discover the payoffs, we would have to play the game. This may be costly, so it is natural to ask whether we can find an equilibrium while minimizing the number of experiments that we must perform.

*Payoff queries* model this situation. In the payoff query model we are told the structure of the game, i.e., the strategy space, but we are not told the payoffs. We can then make payoff queries, where we propose a pure strategy profile, and we are told the payoff of each player under that strategy profile. Our task is to compute an equilibrium of the game while minimizing the number of payoff queries that we make.

The study of query complexity in bimatrix games was initiated by Fearnley et al. [10], who gave a deterministic algorithm for finding a  $\frac{1}{2}$ -NE using  $2n - 1$  payoff queries. A subsequent paper of Fearnley and Savani [8] showed a number of further results. Firstly, they showed an  $\Omega(n^2)$  lower bound on the query complexity of finding an  $\epsilon$ -NE with  $\epsilon < \frac{1}{2}$ , which combined with the result above, gives a complete view of the deterministic query complexity of approximate Nash equilibria in bimatrix games. They then give a randomized algorithm that finds a  $(\frac{3-\sqrt{5}}{2} + \epsilon)$ -NE using  $O(\frac{n \cdot \log n}{\epsilon^2})$  queries, and a randomized algorithm that finds a  $(\frac{2}{3} + \epsilon)$ -WSNE using  $O(\frac{n \cdot \log n}{\epsilon^4})$  queries.

**Our Contribution.** In this paper we introduce a *distributed* technique that allows us to efficiently compute  $\epsilon$ -NE and  $\epsilon$ -WSNE using limited communication between the players.

Traditional methods for computing WSNEs have used an LP based approach that, when used on a bimatrix game  $(R, C)$ , solves the zero-sum game  $(R - C, C - R)$ . The KS algorithm uses the fact that if there is no pure  $\frac{2}{3}$ -WSNE, then the solution to this zero-sum game is a  $\frac{2}{3}$ -WSNE. The slight improvement of the FGSS-algorithm [9] to 0.6608 was obtained by adding two further methods to the KS algorithm: if the KS algorithm does not produce a 0.6608-WSNE, then



either there is a  $2 \times 2$  *matching pennies* sub-game that is 0.6608-WSNE or the strategies from the zero-sum game can be improved by shifting the probabilities of both players within their supports in order to produce a 0.6608-WSNE.

In this paper, we take a different approach. We first show that the bound of  $\frac{2}{3}$  can be matched using a pair of *distributed* LPs. Given a bimatrix game  $(R, C)$ , we solve the two zero-sum games  $(R, -R)$  and  $(-C, C)$ , and then give a simple procedure that we call the *base algorithm*, which uses the solutions to these games to produce a  $\frac{2}{3}$ -WSNE of  $(R, C)$ . Goldberg and Pastink [12] also considered this pair of LPs, but their algorithm only produces a 0.732-WSNE. We then show that the base algorithm can be improved by applying the probability-shifting and matching-pennies ideas from the FGSS-algorithm. That is, if the base algorithm fails to find a 0.6528-WSNE, then a 0.6528-WSNE can be obtained either by shifting the probabilities of one of the two players, or by identifying a  $2 \times 2$  sub-game. This gives a polynomial-time algorithm that computes a 0.6528-WSNE, which provides the best known approximation guarantees for WSNEs.

It is worth pointing out that, while these techniques are thematically similar to the ones used by the FGSS-algorithm, the actual implementation is significantly different. The FGSS-algorithm attempts to improve the strategies by shifting probabilities *within the supports* of the strategies returned by the two player game, with the goal of reducing the other player's payoff. In our algorithm, we shift probabilities *away from bad strategies* in order to improve that player's payoff. This type of analysis is possible because the base algorithm produces a strategy profile in which one of the two players plays a pure strategy, which simplifies the analysis that we need to carry out. On the other hand, the KS-algorithm can produce strategies in which both players play many strategies, and so the analysis used for the FGSS-algorithm is necessarily more complicated.

Since our algorithm solves the two LPs separately, it can be used to improve upon the best known algorithms in the limited communication setting. This is because no communication is required for the row player to solve  $(R, -R)$  and the column player to solve  $(-C, C)$ . The players can then carry out the rest of the algorithm using only poly-logarithmic communication. Hence, we obtain a randomized expected-polynomial-time algorithm that uses poly-logarithmic communication and finds a 0.6528-WSNE. Moreover, the base algorithm can be implemented as a communication efficient algorithm for finding a  $(\frac{1}{2} + \epsilon)$ -WSNE in a *win-lose* bimatrix game, where all payoffs are either 0 or 1.

The algorithm can also be used to beat the best known bound in the query complexity setting. It has already been shown by Goldberg and Roth [13] that an  $\epsilon$ -NE of a *zero-sum game* can be found by a randomized algorithm that uses  $O(\frac{n \log n}{\epsilon^2})$  payoff queries. Since the rest of the steps used by our algorithm can also be carried out using  $O(n \log n)$  payoff queries, this gives us a query efficient algorithm for finding a 0.6528-WSNE.

We also show that the base algorithm can be adapted to find a  $\frac{3-\sqrt{5}}{2}$ -NE in a bimatrix game. Once again, this can be implemented in a communication efficient manner, and so we obtain an algorithm that computes a  $(\frac{3-\sqrt{5}}{2} + \epsilon)$ -NE (i.e., 0.382-NE) using only poly-logarithmic communication. For a summary of our contribution, see Table 1.

**Table 1.** Comparison of our approximation guarantees with the previous best-known guarantees.

| Complexity setting                 | Payoffs | Solution         | Previous best approximation | This paper          |
|------------------------------------|---------|------------------|-----------------------------|---------------------|
| Computational (polynomial)         | [0, 1]  | $\epsilon$ -WSNE | 0.6608 [9]                  | 0.6528              |
| Query ( $n \cdot \log(n)$ queries) | [0, 1]  | $\epsilon$ -WSNE | 0.6667 [8]                  | $0.6528 + \epsilon$ |
| Communication (polylogarithmic)    | [0, 1]  | $\epsilon$ -WSNE | 0.7321 [12]                 | $0.6528 + \epsilon$ |
| Communication (polylogarithmic)    | [0, 1]  | $\epsilon$ -NE   | 0.4384 [12]                 | $0.3820 + \epsilon$ |
| Communication (polylogarithmic)    | {0, 1}  | $\epsilon$ -WSNE | 0.7321 [12]                 | $0.5 + \epsilon$    |

## 2 Preliminaries

**Bimatrix Games.** Throughout, we use  $[n]$  to denote the set of integers  $\{1, 2, \dots, n\}$ . An  $n \times n$  bimatrix game is a pair  $(R, C)$  of two  $n \times n$  matrices:  $R$  gives payoffs for the *row* player and  $C$  gives the payoffs for the *column* player. We make the standard assumption that all payoffs lie in the range  $[0, 1]$ . For simplicity, as in [12], we assume that each payoff has constant bit-length<sup>1</sup>. A *win-lose* bimatrix game has all payoffs in  $\{0, 1\}$ .

Each player has  $n$  *pure* strategies. To play the game, both players simultaneously select a pure strategy: the row player selects a row  $i \in [n]$ , and the column player selects a column  $j \in [n]$ . The row player then receives payoff  $R_{i,j}$ , and the column player receives payoff  $C_{i,j}$ .

A *mixed strategy* is a probability distribution over  $[n]$ . We denote a mixed strategy for the row player as a vector  $\mathbf{x}$  of length  $n$ , such that  $\mathbf{x}_i$  is the probability that the row player assigns to pure strategy  $i$ . A mixed strategy of the column player is a vector  $\mathbf{y}$  of length  $n$ , with the same interpretation. Given a mixed strategy  $\mathbf{x}$  for either player, the *support* of  $\mathbf{x}$  is the set of pure strategies  $i$  with  $\mathbf{x}_i > 0$ . If  $\mathbf{x}$  and  $\mathbf{y}$  are mixed strategies for the row and the column player, respectively, then we call  $(\mathbf{x}, \mathbf{y})$  a *mixed strategy profile*. The expected payoff for the row player under strategy profile  $(\mathbf{x}, \mathbf{y})$  is given by  $\mathbf{x}^T R \mathbf{y}$  and for the column player by  $\mathbf{x}^T C \mathbf{y}$ . We denote the *support* of a strategy  $\mathbf{x}$  as  $\text{supp}(\mathbf{x})$ , which gives the set of pure strategies  $i$  such that  $\mathbf{x}_i > 0$ .

**Nash Equilibria.** Let  $\mathbf{y}$  be a mixed strategy for the column player. The set of *pure best responses* against  $\mathbf{y}$  for the row player is the set of pure strategies that maximize the payoff against  $\mathbf{y}$ . More formally, a pure strategy  $i \in [n]$  is a best response against  $\mathbf{y}$  if, for all pure strategies  $i' \in [n]$  we have:  $\sum_{j \in [n]} \mathbf{y}_j \cdot R_{i,j} \geq \sum_{j \in [n]} \mathbf{y}_j \cdot R_{i',j}$ . Column player best responses are defined analogously.

A mixed strategy profile  $(\mathbf{x}, \mathbf{y})$  is a *mixed Nash equilibrium* if every pure strategy in  $\text{supp}(\mathbf{x})$  is a best response against  $\mathbf{y}$ , and every pure strategy in

<sup>1</sup> The statements of our results can easily be extended to the case where all payoffs can be represented using  $b$  bits by including a factor  $b$  in all our communication complexity bounds.

$\text{supp}(\mathbf{y})$  is a best response against  $\mathbf{x}$ . Nash [17] showed that all bimatrix games have a mixed Nash equilibrium.

**Approximate Equilibria.** There are two commonly studied notions of approximate equilibrium, and we consider both of them in this paper. The first notion is of an  $\epsilon$ -approximate Nash equilibrium ( $\epsilon$ -NE), which weakens the requirement that a player's expected payoff should be equal to their best response payoff. Formally, given a strategy profile  $(\mathbf{x}, \mathbf{y})$ , we define the *regret* suffered by the row player to be the difference between the best response payoff and actual payoff, i.e.,

$$\max_{i \in [n]} ((R \cdot \mathbf{y})_i) - \mathbf{x}^T \cdot R \cdot \mathbf{y}.$$

Regret for the column player is defined analogously. We have that  $(\mathbf{x}, \mathbf{y})$  is an  $\epsilon$ -NE if and only if both players have regret less than or equal to  $\epsilon$ .

The other notion is of an  $\epsilon$ -approximate-well-supported equilibrium ( $\epsilon$ -WSNE), which weakens the requirement that players only place probability on best response strategies. Given a strategy profile  $(\mathbf{x}, \mathbf{y})$  and a pure strategy  $j \in [n]$ , we say that  $j$  is an  $\epsilon$ -best-response for the row player if

$$\max_{i \in [n]} ((R \cdot \mathbf{y})_i) - (R \cdot \mathbf{y})_j \leq \epsilon.$$

An  $\epsilon$ -WSNE requires that both players only place probability on  $\epsilon$ -best-responses. In an  $\epsilon$ -WSNE both players place probability only on  $\epsilon$ -best-responses. Formally, the row player's *pure strategy regret* under  $(\mathbf{x}, \mathbf{y})$  is defined to be

$$\max_{i \in [n]} ((R \cdot \mathbf{y})_i) - \min_{i \in \text{supp}(\mathbf{x})} ((R \cdot \mathbf{y})_i).$$

Pure strategy regret for the column player is defined analogously. A strategy profile  $(\mathbf{x}, \mathbf{y})$  is an  $\epsilon$ -WSNE if both players have pure strategy regret less than or equal to  $\epsilon$ .

**Communication Complexity.** We consider the communication model for bimatrix games introduced by Goldberg and Pastink [12]. In this model, both players know the payoffs in their own payoff matrix, but do not know the payoffs in their opponent's matrix. The players then follow an algorithm that uses a number of communication rounds, where in each round they exchange a single bit of information. Between each communication round, the players are permitted to perform arbitrary randomized computations (although it should be noted that, in this paper, the players will only perform polynomial-time computations) using their payoff matrix, and the bits that they have received so far. At the end of the algorithm, the row player outputs a mixed strategy  $\mathbf{x}$ , and the column player outputs a mixed strategy  $\mathbf{y}$ . The goal is to produce a strategy profile  $(\mathbf{x}, \mathbf{y})$  that is an  $\epsilon$ -NE or  $\epsilon$ -WSNE for a sufficiently small  $\epsilon$  while limiting the number of communication rounds used by the algorithm. The algorithms given in this paper will use at most  $O(\log^2 n)$  communication rounds. In order to achieve this, we use the following result of Goldberg and Pastink [12].

**Lemma 1** [12]. *Given a mixed strategy  $\mathbf{x}$  for the row-player and an  $\epsilon > 0$ , there is a randomized expected-polynomial-time algorithm that uses  $O(\frac{\log^2 n}{\epsilon^2})$ -communication to transmit a strategy  $\mathbf{x}_s$  to the column player where  $|\text{supp}(\mathbf{x}_s)| \in O(\frac{\log n}{\epsilon^2})$  and for every strategy  $i \in [n]$  we have:*

$$|(\mathbf{x}^T \cdot R)_i - (\mathbf{x}_s^T \cdot R)_i| \leq \epsilon.$$

The algorithm uses the well-known sampling technique of Lipton, Markakis, and Mehta to construct the strategy  $\mathbf{x}_s$ , so for this reason we will call the strategy  $\mathbf{x}_s$  the *sampled strategy* from  $\mathbf{x}$ . Since this strategy has a logarithmically sized support, it can be transmitted by sending  $O(\frac{\log n}{\epsilon^2})$  strategy indexes, each of which can be represented using  $\log n$  bits. By symmetry, the algorithm can obviously also be used to transmit approximations of column player strategies to the row player.

**Query Complexity.** In the query complexity setting, the algorithm knows that the players will play an  $n \times n$  game  $(R, C)$ , but it does not know any of the entries of  $R$  or  $C$ . These payoffs are obtained using *payoff queries* in which the algorithm proposes a pure strategy profile  $(i, j)$ , and then it is told the value of  $R_{ij}$  and  $C_{ij}$ . After each payoff query, the algorithm can make arbitrary computations (although, again, in this paper the algorithms that we consider take polynomial time) in order to decide the next pure strategy profile to query. After making a sequence of payoff queries, the algorithm then outputs a mixed strategy profile  $(\mathbf{x}, \mathbf{y})$ . Again, the goal is to ensure that this strategy profile is an  $\epsilon$ -NE or  $\epsilon$ -WSNE, while minimizing the number of queries made overall.

There are two results that we will use for this setting. Goldberg and Roth [13] have given a randomized algorithm that, with high probability, finds an  $\epsilon$ -NE of a zero-sum game using  $O(\frac{n \cdot \log n}{\epsilon^2})$  payoff queries. Given a mixed strategy profile  $(\mathbf{x}, \mathbf{y})$ , an  $\epsilon$ -approximate payoff vector for the row player is a vector  $v$  such that, for all  $i \in [n]$  we have  $|v_i - (R \cdot \mathbf{y})_i| \leq \epsilon$ . Approximate payoff vectors for the column player are defined symmetrically. Fearnley and Savani [8] observed that there is a randomized algorithm that when given the strategy profile  $(\mathbf{x}, \mathbf{y})$ , finds approximate payoff vectors for both players using  $O(\frac{n \cdot \log n}{\epsilon^2})$  payoff queries and that succeeds with high probability. We summarise these two results in the following lemma.

**Lemma 2** [8, 13]. *Given an  $n \times n$  zero-sum bimatrix game, with probability at least  $(1 - n^{-\frac{1}{8}})(1 - \frac{2}{n})^2$ , we can compute an  $\epsilon$ -Nash equilibrium  $(\mathbf{x}, \mathbf{y})$ , and  $\epsilon$ -approximate payoff vectors for both players under  $(\mathbf{x}, \mathbf{y})$ , using  $O(\frac{n \cdot \log n}{\epsilon^2})$  payoff queries.*

### 3 The Base Algorithm

In this section, we introduce the *base algorithm*. This algorithm provides a simple way to find a  $\frac{2}{3}$ -WSNE. We present this algorithm separately for three reasons. Firstly, we believe that the algorithm is interesting in its own right, since it

provides a relatively straightforward method for finding a  $\frac{2}{3}$ -WSNE that is quite different from the technique used in the KS-algorithm. Secondly, our algorithm for finding a 0.6528-WSNE will replace the final step of the algorithm with two more involved procedures, so it is worth understanding this algorithm before we describe how it can be improved. Finally, at the end of this section, we will show that this algorithm can be adapted to provide a communication efficient way to find a  $(0.5 + \epsilon)$ -WSNE in win-lose games.

#### Algorithm 1

1. Solve the zero-sum games  $(R, -R)$  and  $(-C, C)$ .
  - Let  $(\mathbf{x}^*, \mathbf{y}^*)$  be a NE of  $(R, -R)$ , and let  $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$  be a NE of  $(-C, C)$ .
  - Let  $v_r$  be the value secured by  $\mathbf{x}^*$  in  $(R, -R)$ , and let  $v_c$  be the value secured by  $\hat{\mathbf{y}}$  in  $(-C, C)$ . Without loss of generality assume that  $v_c \leq v_r$ .
2. If  $v_r \leq 2/3$ , then return  $(\hat{\mathbf{x}}, \mathbf{y}^*)$ .
3. If for all  $j \in [n]$  it holds that  $C_j^T \cdot \mathbf{x}^* \leq 2/3$ , then return  $(\mathbf{x}^*, \mathbf{y}^*)$ .
4. Otherwise:
  - Let  $j^*$  be a pure best response to  $\mathbf{x}^*$ .
  - Find a row  $i$  such that  $R_{ij^*} > 1/3$  and  $C_{ij^*} > 1/3$ .
  - Return  $(i, j^*)$ .

We argue that this algorithm is correct. For that, we must prove that the row  $i$  used in Step 4 actually exists.

**Lemma 3.** *If Algorithm 1 reaches Step 4, then there exists a row  $i$  such that  $R_{ij^*} > 1/3$  and  $C_{ij^*} > 1/3$ .*

We now argue that the algorithm always produces a  $\frac{2}{3}$ -WSNE. There are three possible strategy profiles that can be returned by the algorithm, which we consider individually.

**The algorithm returns in Step 2:** Since  $v_c \leq v_r$  by assumption, and since  $v_r \leq \frac{2}{3}$ , we have that  $(R \cdot \mathbf{y}^*)_i \leq \frac{2}{3}$  for every row  $i$ , and  $((\hat{\mathbf{x}})^T \cdot C)_j \leq \frac{2}{3}$  for every column  $j$ . So, both players can have pure strategy regret at most  $\frac{2}{3}$  in  $(\hat{\mathbf{x}}, \mathbf{y}^*)$ , and thus this profile is a  $\frac{2}{3}$ -WSNE.

**The algorithm returns in Step 3:** Much like in the previous case, when the column player plays  $\mathbf{y}^*$ , the row player can have pure strategy regret at most  $\frac{2}{3}$ . The requirement that  $C_j^T \cdot \mathbf{x}^* \leq \frac{2}{3}$  also ensures that the column player has pure strategy regret at most  $\frac{2}{3}$ . Thus, we have that  $(\mathbf{x}^*, \mathbf{y}^*)$  is a  $\frac{2}{3}$ -WSNE.

**The algorithm returns in Step 4:** Both players have payoff at least  $\frac{1}{3}$  under  $(i, j^*)$  for the sole strategy in their respective supports. Hence, the maximum pure strategy regret that can be suffered by a player is  $1 - \frac{1}{3} = \frac{2}{3}$ .

Observe that the zero-sum game solved in Step 1 can be solved via linear programming, and so the algorithm runs in polynomial time. Therefore, we have shown the following.

**Theorem 1.** *Algorithm 1 always produces a  $\frac{2}{3}$ -WSNE in polynomial time.*

**Win-Lose Games.** The base algorithm can be adapted to provide a communication efficient method for finding a  $(0.5 + \epsilon)$ -WSNE in win-lose games. In brief, the algorithm can be modified to find a 0.5-WSNE in a win-lose game by making Steps 2 and 3 check against the threshold of 0.5. It can then be shown that if these steps fail, then there exists a pure Nash equilibrium in column  $j^*$ . This can then be implemented as a communication efficient protocol using the algorithm from Lemma 1.

**Theorem 2.** *For every win-lose game and  $\epsilon > 0$ , there is a randomized expected-polynomial-time algorithm that finds a  $(0.5 + \epsilon)$ -WSNE with  $O\left(\frac{\log^2 n}{\epsilon^2}\right)$  communication.*

### 4 An Algorithm for Finding a 0.6528-WSNE

In this section, we show how Algorithm 1 can be modified to produce a 0.6528-WSNE.

**Outline.** Our algorithm replaces Step 4 of Algorithm 1 with a more involved procedure. This procedure uses two techniques, that both find an  $\epsilon$ -WSNE with  $\epsilon < \frac{2}{3}$ . Firstly, we attempt to turn  $(\mathbf{x}^*, j^*)$  into a WSNE by *shifting probabilities*. Observe that, since  $j^*$  is a best response, the column player has a pure strategy regret of 0 in  $(\mathbf{x}^*, j^*)$ . On the other hand, we have no guarantees about the row player since  $\mathbf{x}^*$  might place a small amount of probability strategies with payoff strictly less than  $\frac{1}{3}$ . However, since  $\mathbf{x}^*$  achieves a high *expected* payoff (due to Step 2,) it cannot place too much probability on these low payoff strategies. Thus, the idea is to shift the probability that  $\mathbf{x}^*$  assigns to entries of  $j^*$  with payoff less than or equal to  $\frac{1}{3}$  to entries with payoff strictly greater than  $\frac{1}{3}$ , and thus ensure that the row player’s pure strategy regret is below  $\frac{2}{3}$ . Of course, this procedure will increase the pure strategy regret of the column player, but if it is also below  $\frac{2}{3}$  once all probability has been shifted, then we have found an  $\epsilon$ -WSNE with  $\epsilon < \frac{2}{3}$ .

If shifting probabilities fails to find an  $\epsilon$ -WSNE with  $\epsilon < \frac{2}{3}$ , then we show that the game contains a *matching pennies* sub-game. More precisely, we show that there exists a column  $j'$ , and rows  $b$  and  $s$  such that the  $2 \times 2$  sub-game induced by  $j^*$ ,  $j'$ ,  $b$ , and  $s$  has the following form:

|   |     |             |             |
|---|-----|-------------|-------------|
|   |     | II          |             |
|   |     | $j^*$       | $j'$        |
| I | $b$ | 0           | $\approx 1$ |
|   | $s$ | $\approx 1$ | 0           |
|   |     | 0           | $\approx 1$ |

Thus, if both players play uniformly over their respective pair of strategies, then  $j^*$ ,  $j'$ ,  $b$ , and  $s$  will have payoff  $\approx 0.5$ , and so this yields an  $\epsilon$ -WSNE with  $\epsilon < \frac{2}{3}$ .

**The Algorithm.** We now formalize this approach, and show that it always finds an  $\epsilon$ -WSNE with  $\epsilon < \frac{2}{3}$ . In order to quantify the precise  $\epsilon$  that we obtain, we parametrise the algorithm by a variable  $z$ , which we constrain to be in the range  $0 \leq z < \frac{1}{24}$ . With the exception of the matching pennies step, all other steps of the algorithm will return a  $(\frac{2}{3} - z)$ -WSNE, while the matching pennies step will return a  $(\frac{1}{2} + f(z))$ -WSNE for some increasing function  $f$ . Optimizing the trade off between  $\frac{2}{3} - z$  and  $\frac{1}{2} + f(z)$  then allows us to determine the quality of WSNE found by our algorithm.

The algorithm is displayed as Algorithm 2. Steps 1, 2, and 3 are versions of the corresponding steps from Algorithm 1, which have been adapted to produce a  $(\frac{2}{3} - z)$ -WSNE. Step 4 implements the probability shifting procedure, while Step 5 finds a matching pennies sub-game.

Observe that the probabilities used in  $\mathbf{x}_{\text{mp}}$  and  $\mathbf{y}_{\text{mp}}$  are only well defined when  $z \leq \frac{1}{24}$ , because we have that  $\frac{1-15z}{2-39z} > 1$  whenever  $z > \frac{1}{24}$ , which explains our required upper bound on  $z$ .

**The Correctness of Step 5.** This step of the algorithm relies on the existence of the rows  $b$  and  $s$ , which is shown in the following lemma.

**Lemma 4.** *Suppose that the following conditions hold:*

1.  $\mathbf{x}^*$  has payoff at least  $\frac{2}{3} - z$  against  $j^*$ .
2.  $j^*$  has payoff at least  $\frac{2}{3} - z$  against  $\mathbf{x}^*$ .
3.  $\mathbf{x}^*$  has payoff at least  $\frac{2}{3} - z$  against  $j'$ .
4. Neither  $j^*$  or  $j'$  contains a pure  $(\frac{2}{3} - z)$ -WSNE  $(i, j)$  with  $i \in \text{supp}(\mathbf{x}^*)$ .

Then, both of the following are true:

- There exists a row  $b \in B$  such that  $R_{bj^*} > 1 - \frac{18z}{1+3z}$  and  $C_{bj'} > 1 - \frac{18z}{1+3z}$ .
- There exists a row  $s \in S$  such that  $C_{sj^*} > 1 - \frac{27z}{1+3z}$  and  $R_{sj'} > 1 - \frac{27z}{1+3z}$ .

Observe that the preconditions are indeed true if the Algorithm reaches Step 5. The first and third conditions hold because, due to Step 2, we know that  $\mathbf{x}^*$  is a min-max strategy that secures payoff at least  $v_r > \frac{2}{3} - z$ . The second condition holds because Step 3 ensures that the column player's best response payoff is at least  $\frac{2}{3} - z$ . The fourth condition holds because Step 5 explicitly checks for these pure strategy profiles.

**Quality of Approximation.** We now analyse the quality of WSNE our algorithm produces. Steps 2, 3, 4, 5 each return a strategy profile. Observe that Steps 2 and 3 are the same as the respective steps in the base algorithm, but with the threshold changed from  $\frac{2}{3}$  to  $\frac{2}{3} - z$ . Hence, we can use the same reasoning as we gave for the base algorithm to argue that these steps return  $(\frac{2}{3} - z)$ -WSNE. We now consider the other two steps.

## Algorithm 2

1. Solve the zero-sum games  $(R, -R)$  and  $(-C, C)$ .
  - Let  $(\mathbf{x}^*, \mathbf{y}^*)$  be a NE of  $(R, -R)$ , and let  $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$  be a NE of  $(C, -C)$ .
  - Let  $v_r$  be the value secured by  $\mathbf{x}^*$  in  $(R, -R)$ , and let  $v_c$  be the value secured by  $\hat{\mathbf{y}}$  in  $(-C, C)$ . Without loss of generality assume that  $v_c \leq v_r$ .
2. If  $v_r \leq 2/3 - z$ , then return  $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ .
3. If for all  $j \in [n]$  it holds that  $C_j^T \mathbf{x}^* \leq 2/3 - z$ , then return  $(\mathbf{x}^*, \mathbf{y}^*)$ .
4. Otherwise:
  - Let  $j^*$  be a pure best response against  $\mathbf{x}^*$ . Define:

$$S := \{i \in \text{supp}(\mathbf{x}^*) : R_{ij^*} < 1/3 + z\}$$

$$B := \text{supp}(\mathbf{x}^*) \setminus S$$

- Define the strategy  $\mathbf{x}_b$  as follows. For each  $i \in [n]$  we have:

$$(\mathbf{x}_b)_i = \begin{cases} \frac{1}{\Pr(B)} \cdot \mathbf{x}_i^* & \text{if } i \in B \\ 0 & \text{otherwise.} \end{cases}$$

- If  $(\mathbf{x}_b^T \cdot C)_{j^*} \geq \frac{1}{3} + z$ , then return  $(\mathbf{x}_b, j^*)$ .

5. Otherwise:
  - Let  $j'$  be a pure best response against  $\mathbf{x}_b$ .
  - If there exists an  $i \in \text{supp}(\mathbf{x}^*)$  such that  $(i, j^*)$  or  $(i, j')$  is a pure  $(\frac{2}{3} - z)$ -WSNE, then return it.
  - Find a row  $b \in B$  such that  $R_{bj^*} > 1 - \frac{18z}{1+3z}$  and  $C_{bj'} > 1 - \frac{18z}{1+3z}$ .
  - Find a row  $s \in S$  such that  $C_{sj^*} > 1 - \frac{27z}{1+3z}$  and  $R_{sj'} > 1 - \frac{27z}{1+3z}$ .
  - Define the row player strategy  $\mathbf{x}_{\text{mp}}$  and the column player strategy  $\mathbf{y}_{\text{mp}}$  as follows. For each  $i \in [n]$  we have:

$$\mathbf{x}_{\text{mp}_i} = \begin{cases} \frac{1-24z}{2-39z} & \text{if } i = b, \\ \frac{1-15z}{2-39z} & \text{if } i = s, \\ 0 & \text{otherwise.} \end{cases} \quad \mathbf{y}_{\text{mp}_i} = \begin{cases} \frac{1-24z}{2-39z} & \text{if } i = j^*, \\ \frac{1-15z}{2-39z} & \text{if } i = j', \\ 0 & \text{otherwise.} \end{cases}$$

- Return  $(\mathbf{x}_{\text{mp}}, \mathbf{y}_{\text{mp}})$ .

**The algorithm returns in Step 4:** By definition all rows  $r \in B$  satisfy  $R_{rj^*} \geq \frac{1}{3} + z$ , so since  $\text{supp}(\mathbf{x}_b) \subseteq B$ , the pure strategy regret of the row player can be at most  $1 - (\frac{1}{3} + z) = \frac{2}{3} - z$ . For the same reason, since  $(\mathbf{x}_b^T \cdot C)_{j^*} \geq \frac{1}{3} + z$  holds, the pure strategy regret of the column player can also be at  $\frac{2}{3} - z$ . Thus, the profile  $(\mathbf{x}_b, j^*)$  is a  $(\frac{2}{3} - z)$ -WSNE.



**The algorithm returns in Step 5:** Since  $R_{bj^*} > 1 - \frac{18z}{1+3z}$ , the payoff of  $b$  when the column player plays  $\mathbf{y}_{\mathbf{mp}}$  is at least:

$$\frac{1 - 24z}{2 - 39z} \cdot \left(1 - \frac{18z}{1 + 3z}\right) = \frac{1 - 39z + 360z^2}{2 - 33z - 117z^2}$$

Similarly, since  $R_{sj'} > 1 - \frac{27z}{1+3z}$ , the payoff of  $s$  when the column player plays  $\mathbf{y}_{\mathbf{mp}}$  is at least:

$$\frac{1 - 15z}{2 - 39z} \cdot \left(1 - \frac{27z}{1 + 3z}\right) = \frac{1 - 39z + 360z^2}{2 - 33z - 117z^2}$$

In the same way, one can show that the payoffs of  $j^*$  and  $j'$  are also  $\frac{1-39z+360z^2}{2-33z-117z^2}$  when the row player plays  $\mathbf{x}_{\mathbf{mp}}$ . Thus, we have that  $(\mathbf{x}_{\mathbf{mp}}, \mathbf{y}_{\mathbf{mp}})$  is a  $(1 - \frac{1-39z+360z^2}{2-33z-117z^2})$ -WSNE.

To find the optimal value for  $z$ , we need to find the largest value of  $z$  for which the following inequality holds.

$$1 - \frac{1 - 39z + 360z^2}{2 - 33z - 117z^2} \leq \frac{2}{3} - z.$$

Setting the inequality to an equality and rearranging gives us a cubic polynomial equation:  $117z^3 + 432z^2 - 30z + \frac{1}{3} = 0$ . Since the discriminant of this polynomial is positive, this polynomial has three real roots, which can be found via the trigonometric method. Only one of these roots lies in the range  $0 \leq z < \frac{1}{24}$ , which is the following:

$$z = \frac{1}{117} \sqrt{3} \left( \sqrt{2434} \sqrt{3} \cos \left( \frac{1}{3} \arctan \left( \frac{39}{240073} \sqrt{9749} \sqrt{3} \right) \right) - 3 \sqrt{2434} \sin \left( \frac{1}{3} \arctan \left( \frac{39}{240073} \sqrt{9749} \sqrt{3} \right) \right) - 48 \sqrt{3} \right).$$

Thus, we get  $z \approx 0.013906376$ , and we have found an algorithm that always produces a 0.6528-WSNE. So we have the following theorem.

**Theorem 3.** *There is a polynomial time algorithm that, given a bimatrix game, finds a 0.6528-WSNE.*

**Communication Complexity.** We claim that our algorithm can be adapted for the limited communication setting by making the following modifications. After computing  $\mathbf{x}^*$ ,  $\mathbf{y}^*$ ,  $\hat{\mathbf{x}}$ , and  $\hat{\mathbf{y}}$ , we then use Lemma 1 to construct and communicate the sampled strategies  $\mathbf{x}_s^*$ ,  $\mathbf{y}_s^*$ ,  $\hat{\mathbf{x}}_s$ , and  $\hat{\mathbf{y}}_s$ . These strategies are communicated between the two players using  $4 \cdot (\log n)^2$  bits of communication, and the players also exchange  $v_r = (\mathbf{x}_s^*)^T \cdot R\mathbf{y}_s^*$  and  $v_c = \hat{\mathbf{x}}_s^T C\hat{\mathbf{y}}_s$  using  $\log n$  rounds of communication. The algorithm then continues as before, except the sampled

strategies are used in place of their non-sampled counterparts. Finally, in Steps 2 and 3, we test against the threshold  $\frac{2}{3} - z + \epsilon$  instead of  $\frac{2}{3} - z$ .

Observe that, when sampled strategies are used, all steps of the algorithm can be carried out in at most  $(\log n)^2$  communication. In particular, to implement Step 4, the column player can communicate  $j^*$  to the row player, and then the row player can communicate  $R_{ij^*}$  for all rows  $i \in \text{supp}(\mathbf{x}_s^*)$  using  $(\log n)^2$  bits of communication, which allows the column player to determine  $j'$ . Once  $j'$  has been determined, there are only  $2 \cdot \log n$  payoffs in each matrix that are relevant to the algorithm (the payoffs in rows  $i \in \text{supp}(\mathbf{x}_s^*)$  in columns  $j^*$  and  $j'$ ), and so the two players can communicate all of these payoffs to each other, and then no further communication is necessary.

**Theorem 4.** *For every  $\epsilon > 0$ , there is a randomized expected-polynomial-time algorithm that uses  $O\left(\frac{\log^2 n}{\epsilon^2}\right)$  communication and finds a  $(0.6528 + \epsilon)$ -WSNE.*

**Query Complexity.** We now show that Algorithm 2 can be implemented in a payoff-query efficient manner. Let  $\epsilon > 0$  be a positive constant. We now outline the changes needed in the algorithm.

- In Step 1 we use the algorithm of Lemma 2 to find  $\frac{\epsilon}{2}$ -NEs of  $(R, -R)$ , and  $(C, -C)$ . We denote the mixed strategies found as  $(\mathbf{x}_a^*, \mathbf{y}_a^*)$  and  $(\hat{\mathbf{x}}_a, \hat{\mathbf{y}}_a)$ , respectively, and we use these strategies in place of their original counterparts throughout the rest of the algorithm. We also compute  $\frac{\epsilon}{2}$ -approximate payoff vectors for each of these strategies, and use them whenever we need to know the payoff of a particular strategy under one of these strategies. In particular, we set  $v_r$  to be the payoff of  $\mathbf{x}_a^*$  according to the approximate payoff vector of  $\mathbf{y}_a^*$ , and we set  $v_c$  to be the payoff of  $\hat{\mathbf{y}}_a$  according to the approximate payoff vector for  $\hat{\mathbf{x}}_a$ .
- In Steps 2 and 3 we test against the threshold of  $\frac{2}{3} - z + \epsilon$  rather than  $\frac{2}{3} - z$ .
- In Step 4 we select  $j^*$  to be the column that is maximal in the approximate payoff vector against  $\mathbf{x}_a^*$ . We then spend  $n$  payoff queries to query every row in column  $j^*$ , which allow us to proceed with the rest of this step as before.
- In Step 5 we use the algorithm of Lemma 2 to find an approximate payoff vector  $v$  for the column player against  $\mathbf{x}_b$ . We then select  $j'$  to be a column that maximizes  $v$ , and then spend  $n$  payoff queries to query every row in  $j^*$ , which allows us to proceed with the rest of this step as before.

Observe that the query complexity of the algorithm is  $O\left(\frac{n \cdot \log n}{\epsilon^2}\right)$ , where the dominating term arises due to the use of the algorithm from Lemma 2 to approximate solutions to the zero-sum games.

**Theorem 5.** *There is a randomized algorithm that, with high probability, finds a  $(0.6528 + \epsilon)$ -WSNE using  $O\left(\frac{n \cdot \log n}{\epsilon^2}\right)$  payoff queries.*

## References

1. Bosse, H., Byrka, J., Markakis, E.: New algorithms for approximate Nash equilibria in bimatrix games. *Theoret. Comput. Sci.* **411**(1), 164–173 (2010)
2. Chen, X., Deng, X., Teng, S.-H.: Settling the complexity of computing two-player Nash equilibria. *J. ACM* **56**(3), 14:1–14:57 (2009)
3. Conitzer, V., Sandholm, T.: Communication complexity as a lower bound for learning in games. In: *Proceedings of ICML*, pp. 185–192 (2004)
4. Czumaj, A., Fasoulakis, M., Jurdzinski, M.: Approximate well-supported Nash equilibria in symmetric bimatrix games. In: *Proceedings of SAGT*, pp. 244–254 (2014)
5. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: Progress in approximate Nash equilibria. In: *Proceedings of EC*, pp. 355–358 (2007)
6. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. *SIAM J. Comput.* **39**(1), 195–259 (2009)
7. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: A note on approximate Nash equilibria. *Theoret. Comput. Sci.* **410**(17), 1581–1588 (2009)
8. Fearnley, J., Savani, R.: Finding approximate Nash equilibria of bimatrix games via payoff queries. In: *Proceedings of EC*, pp. 657–674 (2014)
9. Fearnley, J., Goldberg, P.W., Savani, R., Sørensen, T.B.: Approximate well-supported Nash equilibria below two-thirds. In: Serna, M. (ed.) *SAGT 2012*. LNCS, vol. 7615, pp. 108–119. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33996-7\\_10](https://doi.org/10.1007/978-3-642-33996-7_10)
10. Fearnley, J., Gairing, M., Goldberg, P.W., Savani, R.: Learning equilibria of games via payoff queries. In: *Proceedings of EC*, pp. 397–414 (2013)
11. Fearnley, J., Igwe, T.P., Savani, R.: An empirical study of finding approximate equilibria in bimatrix games. In: *Proceedings of SEA*, pp. 339–351 (2015)
12. Goldberg, P.W., Pastink, A.: On the communication complexity of approximate Nash equilibria. *Games Econ. Behav.* **85**, 19–31 (2014)
13. Goldberg, P.W., Roth, A.: Bounds for the query complexity of approximate equilibria. In: *Proceedings of EC*, pp. 639–656 (2014)
14. Hart, S., Mansour, Y.: How long to equilibrium? The communication complexity of uncoupled equilibrium procedures. *Games Econ. Behav.* **69**(1), 107–126 (2010)
15. Kontogiannis, S.C., Spirakis, P.G.: Well supported approximate equilibria in bimatrix games. *Algorithmica* **57**(4), 653–667 (2010)
16. Lipton, R.J., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: *Proceedings of EC*, pp. 36–41 (2003)
17. Nash, J.: Non-cooperative games. *Ann. Math.* **54**(2), 286–295 (1951)
18. Rubinstein, A.: Settling the complexity of computing approximate two-player Nash equilibria. In: *Proceedings of FOCS*, pp. 258–265 (2016)
19. Tsaknakis, H., Spirakis, P.G.: An optimization approach for approximate Nash equilibria. *Internet Math.* **5**(4), 365–382 (2008)

# Inapproximability Results for Approximate Nash Equilibria

Argyrios Deligkas, John Fearnley, and Rahul Savani<sup>(✉)</sup>

Department of Computer Science, University of Liverpool, Liverpool, UK  
{a.deligkas,rahul.savani}@liverpool.ac.uk

**Abstract.** We study the problem of finding approximate Nash equilibria that satisfy certain conditions, such as providing good social welfare. In particular, we study the problem  $\epsilon$ -NE  $\delta$ -SW: find an  $\epsilon$ -approximate Nash equilibrium ( $\epsilon$ -NE) that is within  $\delta$  of the best social welfare achievable by an  $\epsilon$ -NE. Our main result is that, if the randomized exponential-time hypothesis (RETH) is true, then solving  $(\frac{1}{8} - O(\delta))$ -NE  $O(\delta)$ -SW for an  $n \times n$  bimatrix game requires  $n^{\tilde{\Omega}(\delta^A \log n)}$  time, where  $A$  is a constant.

Building on this result, we show similar conditional running time lower bounds on a number of decision problems for approximate Nash equilibria that do not involve social welfare, including maximizing or minimizing a certain player's payoff, or finding approximate equilibria contained in a given pair of supports. We show quasi-polynomial lower bounds for these problems assuming that RETH holds, and these lower bounds apply to  $\epsilon$ -Nash equilibria for all  $\epsilon < \frac{1}{8}$ . The hardness of these other decision problems has so far only been studied in the context of exact equilibria.

## 1 Introduction

One of the most fundamental problems in game theory is to find a Nash equilibrium of a game. Often, we are not interested in finding any Nash equilibrium, but instead we want to find one that also satisfies certain constraints. For example, we may want to find a Nash equilibrium that provides high *social welfare*, which is the sum of the player's payoffs.

In this paper we study such problems for *bimatrix games*, which are two-player strategic-form games. Unfortunately, for bimatrix games, it is known that these problems are hard. Finding any Nash equilibrium of a bimatrix game is PPAD-complete [10], while finding a constrained Nash equilibrium turns out to be even harder. Gilboa and Zemel [16] studied several decision problems related to Nash equilibria. They proved that it is NP-complete to decide whether there exist Nash equilibria in bimatrix games with some “desirable” properties, such as high social welfare. Conitzer and Sandholm [7] extended the list of NP-complete problems of [16] and furthermore proved inapproximability results for some of

---

The authors were supported by EPSRC grant EP/L011018/1. The full version of this paper, with complete proofs, is available at <http://arxiv.org/abs/1608.03574>.

them. Recently, Garg et al. [15] and Bilo and Mavronicolas [4] extended these results to many player games and provided ETR-completeness results for them.

**Approximate Equilibria.** Due to the apparent hardness of finding exact Nash equilibria, focus has shifted to *approximate* equilibria. There are two natural notions of approximate equilibrium, both of which will be studied in this paper. An  $\epsilon$ -*approximate Nash equilibrium* ( $\epsilon$ -NE) requires that each player has an expected payoff that is within  $\epsilon$  of their best response payoff. An  $\epsilon$ -*well-supported Nash equilibrium* ( $\epsilon$ -WSNE) requires that both players only play strategies whose payoff is within  $\epsilon$  of the best response payoff. Every  $\epsilon$ -WSNE is an  $\epsilon$ -NE but the converse does not hold, so a WSNE is a more restrictive notion.

There has been a long line of work on finding approximate equilibria [5, 8, 11–13, 18, 22]. Since we use an additive notion of approximation, it is common to rescale the game so that the payoffs lie in  $[0, 1]$ , which allows different algorithms to be compared. The state of the art for polynomial-time algorithms is the following. There is a polynomial-time algorithm that computes an 0.3393-NE [22], and a polynomial-time algorithm that computes a 0.6528-WSNE [8].

There is also a *quasi-polynomial time approximation scheme* (QPTAS) for finding approximate Nash equilibria. The algorithm of Lipton, Markakis, and Mehta finds an  $\epsilon$ -NE in  $n^{O(\frac{\log n}{\epsilon^2})}$  time [19]. They proved that there is always an  $\epsilon$ -NE with logarithmic support, and then uses a brute-force search over all possible candidates to find one. We will refer to their algorithm as the LMM algorithm.

A recent breakthrough of Rubinfeld implies that we cannot do better than a QPTAS like the LMM algorithm [21]: assuming the ETH for PPAD (PETH), there is a small constant,  $\epsilon^*$ , such that for  $\epsilon < \epsilon^*$ , every algorithm for finding an  $\epsilon$ -NE requires quasi-polynomial time. Briefly, PETH is the conjecture that ENDOFTHELINE, the canonical PPAD-complete problem, cannot be solved faster than exponential time.

**Constrained Approximate Nash Equilibria.** While deciding whether a game has an exact Nash equilibrium that satisfies certain constraints is NP-hard for most interesting constraints, this is not the case for approximate equilibria, because the LMM algorithm can be adapted to provide a QPTAS for them. The question then arises whether these results are tight.

Let the problem  $\epsilon$ -NE  $\delta$ -SW be the problem of finding an  $\epsilon$ -NE whose social welfare is within  $\delta$  of the best social welfare that can be achieved by an  $\epsilon$ -NE. Hazan and Krauthgamer [17] and Austrin et al. [3] proved that there is a small but constant  $\epsilon$  such that  $\epsilon$ -NE  $\epsilon$ -SW is at least as hard as finding a hidden clique of size  $O(\log n)$  in the random graph  $G_{n,1/2}$ . This was further strengthened by Braverman et al. [6] who showed a lower bound based on the *exponential-time hypothesis* (ETH), which is the conjecture that any deterministic algorithm for 3SAT requires  $2^{\Omega(n)}$  time. More precisely, they showed that under the ETH there

is a small constant  $\epsilon$  such that any algorithm for  $O(\epsilon)$ -NE  $O(\epsilon)$ -SW<sup>1</sup> requires  $n^{\text{poly}(\epsilon) \log(n)^{1-o(1)}}$  time<sup>2</sup>. We shall refer to this as the BKW result.

It is worth noting that the Rubinstein’s hardness result [21] almost makes this result redundant. If one is willing to accept that PETH is true, which is a stronger conjecture than ETH, then Rubinstein’s result says that for small  $\epsilon$  we require quasi-polynomial time to find *any*  $\epsilon$ -NE, which obviously implies that the same lower bound applies to  $\epsilon$ -NE  $\delta$ -SW for any  $\delta$ .

**Our Results.** Our first result is a lower bound for the problem of finding  $\epsilon$ -NE  $\delta$ -SW. The *randomized ETH* (RETH) is the conjecture that any randomized algorithm for 3SAT requires  $2^{\Omega(n)}$  time. We show that, assuming RETH, there exists a small constant  $\Delta$  such that for all  $\delta \in [1/n, \Delta]$  the problem  $\left(\frac{1-4g\delta}{8}\right)$ -NE  $\left(\frac{g\delta}{4}\right)$ -SW requires  $n^{\tilde{\Omega}(\delta^A \log n)}$  time<sup>3</sup>, where  $g = \frac{1}{138}$ , and  $A$  is a constant.

To understand this result, let us compare it to the BKW result. First, observe that as  $\delta$  gets smaller, the  $\epsilon$  in our  $\epsilon$ -NE gets larger, whereas the approximate Nash equilibria in the BKW result get smaller. Asymptotically, our  $\epsilon$  approaches  $1/8$ . Moreover, since  $\delta \leq 1$ , our lower bound applies to all  $\epsilon$ -NE with  $\epsilon \leq \frac{1-4g}{8} \approx 0.1214$ . This is orders of magnitude larger than the inapproximability bound given by Rubinstein’s hardness result, and so is not made redundant by that result. In short, our hardness result is about the hardness of obtaining good social welfare, rather than the hardness of simply finding an approximate equilibrium.

Secondly, when compared to the BKW result, we obtain a slightly better lower bound. The exponent in their lower bound is logarithmic only in the limit, while ours is always logarithmic. In particular, we obtain quasi-polynomial lower bounds whenever  $\delta$  is constant.

Finally, our result uses a stronger conjecture when compared to the BKW result. While they assume ETH, our result requires that we assume RETH. This is a stronger conjecture, since even if ETH is true, there may exist randomized sub-exponential algorithms for 3SAT. This means that our result is ultimately incomparable to the BKW result: we obtain a lower bound for larger  $\epsilon$ , and we have a better lower bound on the running time, but we do so by assuming a stronger conjecture.

To prove our result, we reduce from the problem of approximating the value of a *free game*. Aaronson, Impagliazzo, and Moshkovitz showed quasi-polynomial lower bounds for this problem assuming ETH [1]. In fact, they give two different lower bounds: the *high error* result shows a quasi-polynomial lower bound for determining whether the value of the game is 1 or  $1 - \delta$  for small  $\delta$ , while

<sup>1</sup> While the proof in [6] produces a lower bound for  $0.8$ -NE  $(1 - O(\epsilon))$ -SW, this is in a game with maximum payoff  $O(1/\epsilon)$ . Therefore, when the payoffs in this game are rescaled to  $[0, 1]$ , the resulting lower bound only applies to  $\epsilon$ -NE  $\epsilon$ -SW.

<sup>2</sup> Although the paper claims that they obtain a  $n^{\tilde{O}(\log n)}$  lower bound, the proof reduces from the *low error* result from [1] (cf. Theorem 36 in [2]), which gives only the weaker lower bound of  $n^{\text{poly}(\epsilon) \log(n)^{1-o(1)}}$ .

<sup>3</sup> Here  $\tilde{\Omega}(\log n)$  means  $\Omega\left(\frac{\log n}{(\log \log n)^c}\right)$  for some constant  $c$ .

the *low error* result shows a weaker almost-quasi-polynomial lower bound on determining whether the value of the game is 1 or  $\delta$  for small  $\delta$ . The BKW result was proved via a reduction from the low error case, while our result uses the high error case. We reduce the free game to a bimatrix game, and prove that in any  $\epsilon$ -NE of the game, the players must simulate the free game well enough so that we can determine whether the value of the free game is 1 or  $1 - \delta$ . Our reduction is substantially different from the BKW reduction: we use a sub-sampling result for free games to reduce the number of questions in the free game, and then we use a different method to force the players to simulate the free game.

Once we have our lower bound on the problem of finding  $\epsilon$ -NE  $\delta$ -SW, we use it to prove lower bounds for other problems regarding constrained approximate NEs and WSNEs. Table 1 gives a list of the problems that we consider. For each one, we provide a reduction from  $\epsilon$ -NE  $\delta$ -SW to that problem. Ultimately, we prove that if RETH is true, then for every  $\epsilon < \frac{1}{8}$  finding an  $\epsilon$ -NE with the given property requires  $n^{\tilde{\Omega}(\log n)}$  time.

**Table 1.** The decision problems that we consider. All of them take as input a bimatrix games and a quality of approximation  $\epsilon \in (0, 1)$ . Problems 1–6 relate to  $\epsilon$ -NE, and Problems 7–10 relate to  $\epsilon$ -WSNE.

| Problem description   | Problem definition  |
|---|---|
| Problem 1: Large payoffs $u \in (0, 1]$   | Is there an $\epsilon$ -NE $(\mathbf{x}, \mathbf{y})$ such that $\min(\mathbf{x}^T R \mathbf{y}, \mathbf{x}^T C \mathbf{y}) \geq u$ ?     |
| Problem 2: Small total payoff $v \in [0, 2)$  | Is there an $\epsilon$ -NE $(\mathbf{x}, \mathbf{y})$ such that $\mathbf{x}^T R \mathbf{y} + \mathbf{x}^T C \mathbf{y} \leq v$ ?          |
| Problem 3: Small payoff $u \in [0, 1)$  | Is there an $\epsilon$ -NE $(\mathbf{x}, \mathbf{y})$ such that $\mathbf{x}^T R \mathbf{y} \leq u$ ?                                      |
| Problem 4: Restricted support $S \subset [n]$                                       | Is there an $\epsilon$ -NE $(\mathbf{x}, \mathbf{y})$ with $\text{supp}(\mathbf{x}) \subseteq S$ ?  |
| Problem 5: Two $\epsilon$ -NE $d \in (0, 1]$ apart in Total Variation (TV) distance | Are there two $\epsilon$ -NE with TV distance $\geq d$ ?  |
| Problem 6: Small largest probability $p \in (0, 1)$                                 | Is there an $\epsilon$ -NE $(\mathbf{x}, \mathbf{y})$ with $\max_i \mathbf{x}_i \leq p$ ?   |
| Problem 7: Large total support size $k \in [n]$                                     | Is there an $\epsilon$ -WSNE $(\mathbf{x}, \mathbf{y})$ such that $ \text{supp}(\mathbf{x})  +  \text{supp}(\mathbf{y})  \geq 2k$ ?       |
| Problem 8: Large smallest support size $k \in [n]$                                  | Is there an $\epsilon$ -WSNE $(\mathbf{x}, \mathbf{y})$ such that $\min\{ \text{supp}(\mathbf{x}) ,  \text{supp}(\mathbf{y}) \} \geq k$ ? |
| Problem 9: Large support size $k \in [n]$   | Is there an $\epsilon$ -WSNE $(\mathbf{x}, \mathbf{y})$ such that $ \text{supp}(\mathbf{x})  \geq k$ ?                                    |
| Problem 10: Restricted support $S_R \subseteq [n]$                                  | Is there an $\epsilon$ -WSNE $(\mathbf{x}, \mathbf{y})$ with $S_R \subseteq \text{supp}(\mathbf{x})$ ?                                    |

**Other Related Work.** The only positive result for finding  $\epsilon$ -NE with good social welfare that we are aware of was given by Czumaj et al. [9]. They showed that if there is a polynomial-time algorithm for finding an  $\epsilon$ -NE, then for all  $\epsilon' > \epsilon$  there is also a polynomial-time algorithm for finding an  $\epsilon'$ -NE that is within a constant multiplicative approximation of the best social welfare. They also give further results for the case where  $\epsilon > \frac{1}{2}$ .

## 2 Preliminaries

Throughout the paper, we use  $[n]$  to denote the set of integers  $\{1, 2, \dots, n\}$ . An  $n \times n$  bimatrix game is a pair  $(R, C)$  of two  $n \times n$  matrices:  $R$  gives payoffs for the *row* player and  $C$  gives the payoffs for the *column* player.

Each player has  $n$  *pure* strategies. To play the game, both players simultaneously select a pure strategy: the row player selects a row  $i \in [n]$ , and the column player selects a column  $j \in [n]$ . The row player then receives payoff  $R_{i,j}$ , and the column player receives payoff  $C_{i,j}$ .

A *mixed strategy* is a probability distribution over  $[n]$ . We denote a mixed strategy for the row player as a vector  $\mathbf{x}$  of length  $n$ , such that  $\mathbf{x}_i$  is the probability that the row player assigns to pure strategy  $i$ . A mixed strategy of the column player is a vector  $\mathbf{y}$  of length  $n$ , with the same interpretation. If  $\mathbf{x}$  and  $\mathbf{y}$  are mixed strategies for the row and the column player, respectively, then we call  $(\mathbf{x}, \mathbf{y})$  a *mixed strategy profile*. The expected payoff for the row player under strategy profile  $(\mathbf{x}, \mathbf{y})$  is given by  $\mathbf{x}^T R \mathbf{y}$  and for the column player by  $\mathbf{x}^T C \mathbf{y}$ . We denote the *support* of a strategy  $\mathbf{x}$  as  $\text{supp}(\mathbf{x})$ , which gives the set of pure strategies  $i$  such that  $\mathbf{x}_i > 0$ .

**Nash Equilibria.** Let  $\mathbf{y}$  be a mixed strategy for the column player. The set of *pure best responses* against  $\mathbf{y}$  for the row player is the set of pure strategies that maximize the payoff against  $\mathbf{y}$ . More formally, a pure strategy  $i \in [n]$  is a best response against  $\mathbf{y}$  if, for all pure strategies  $i' \in [n]$  we have:  $\sum_{j \in [n]} \mathbf{y}_j \cdot R_{i,j} \geq \sum_{j \in [n]} \mathbf{y}_j \cdot R_{i',j}$ . Column player best responses are defined analogously.

A mixed strategy profile  $(\mathbf{x}, \mathbf{y})$  is a *mixed Nash equilibrium* if every pure strategy in  $\text{supp}(\mathbf{x})$  is a best response against  $\mathbf{y}$ , and every pure strategy in  $\text{supp}(\mathbf{y})$  is a best response against  $\mathbf{x}$ . Nash [20] showed that every bimatrix game has a mixed Nash equilibrium. Observe that in a Nash equilibrium, each player's expected payoff is equal to their best response payoff.

**Approximate Equilibria.** There are two commonly studied notions of approximate equilibrium, and we consider both of them in this paper. The first notion is that of an  $\epsilon$ -*approximate Nash equilibrium* ( $\epsilon$ -NE), which weakens the requirement that a player's expected payoff should be equal to their best response payoff. Formally, given a strategy profile  $(\mathbf{x}, \mathbf{y})$ , we define the *regret* suffered by the row player to be the difference between the best response payoff and the actual payoff:  $\max_{i \in [n]} ((R \cdot \mathbf{y})_i) - \mathbf{x}^T \cdot R \cdot \mathbf{y}$ . Regret for the column player is



defined analogously. We have that  $(\mathbf{x}, \mathbf{y})$  is an  $\epsilon$ -NE if and only if both players have regret less than or equal to  $\epsilon$ .

The other notion is that of an  $\epsilon$ -approximate-well-supported equilibrium ( $\epsilon$ -WSNE), which weakens the requirement that players only place probability on best response strategies. Given a strategy profile  $(\mathbf{x}, \mathbf{y})$  and a pure strategy  $j \in [n]$ , we say that  $j$  is an  $\epsilon$ -best-response for the row player if:  $\max_{i \in [n]} ((R \cdot y)_i) - (R \cdot y)_j \leq \epsilon$ . An  $\epsilon$ -WSNE requires that both players only place probability on  $\epsilon$ -best-responses. Formally, the row player's *pure strategy regret* under  $(\mathbf{x}, \mathbf{y})$  is defined to be:  $\max_{i \in [n]} ((R \cdot y)_i) - \min_{i \in \text{supp}(\mathbf{x})} ((R \cdot y)_i)$ . Pure strategy regret for the column player is defined analogously. A strategy profile  $(\mathbf{x}, \mathbf{y})$  is an  $\epsilon$ -WSNE if both players have pure strategy regret less than or equal to  $\epsilon$ .

Since approximate Nash equilibria use an additive notion of approximation, it is standard practice to rescale the input game so that all payoffs lie in the range  $[0, 1]$ , which allows us to compare different results on this topic. For the most part, we follow this convention. However, for our result in Sect. 3, we will construct a game whose payoffs do not lie in  $[0, 1]$ . In order to simplify the proof, we will prove results about approximate Nash equilibria in the unscaled game, and then rescale the game to  $[0, 1]$  at the very end. To avoid confusion, we will refer to an  $\epsilon$ -approximate Nash equilibrium in this game as an  $\epsilon$ -UNE, to mark that it is an additive approximation in an unscaled game.

**Two-Prover Games.** A two-prover game is defined as follows.

**Definition 1 (Two-prover game).** *A two-prover game  $\mathcal{T}$  is defined by a tuple  $(X, Y, A, B, \mathcal{D}, V)$  where  $X$  and  $Y$  are finite sets of questions,  $A$  and  $B$  are finite sets of answers,  $\mathcal{D}$  is a probability distribution defined over  $X \times Y$ , and  $V$  is a verification function of the form  $V : X \times Y \times A \times B \rightarrow \{0, 1\}$ .*

The game is a co-operative game played between two players, who are called Merlin<sub>1</sub> and Merlin<sub>2</sub>, and an adjudicator called Arthur. At the start of the game, Arthur chooses a question pair  $(x, y) \in X \times Y$  randomly according to  $\mathcal{D}$ . He then sends  $x$  to Merlin<sub>1</sub> and  $y$  to Merlin<sub>2</sub>. Crucially, Merlin<sub>1</sub> does not know the question sent to Merlin<sub>2</sub> and vice versa. Having received  $x$ , Merlin<sub>1</sub> then chooses an answer from  $A$  and sends it back to Arthur. Merlin<sub>2</sub> similarly picks an answer from  $B$  and returns it to Arthur. Arthur then computes  $p = V(x, y, a, b)$  and awards payoff  $p$  to both players. The size of the game, denoted  $|\mathcal{T}| = |X \times Y \times A \times B|$  is the total number of entries needed to represent  $V$  as a table.

A *strategy* for Merlin<sub>1</sub> is a function  $a : X \rightarrow A$  that gives an answer for every possible question, and likewise a strategy for Merlin<sub>2</sub> is a function  $b : Y \rightarrow B$ . We define  $S_i$  to be the set of all strategies for Merlin <sub>$i$</sub> . The *payoff* of the game under a pair of strategies  $(s_1, s_2) \in S_1 \times S_2$  is denoted as  $p(\mathcal{T}, s_1, s_2) = E_{(x,y) \sim \mathcal{D}}[V(x, y, s_1(x), s_2(y))]$ .

The *value* of the game, denoted  $\omega(\mathcal{T})$ , is the maximum expected payoff to the Merlins when they play optimally:  $\omega(\mathcal{T}) = \max_{s_1 \in S_1} \max_{s_2 \in S_2} p(\mathcal{T}, s_1, s_2)$ .

**Free Games.** A two-prover game is called a *free game* if the probability distribution  $\mathcal{D}$  is the uniform distribution  $\mathcal{U}$  over  $X \times Y$ . In particular, this means that

there is no correlation between the question sent to Merlin<sub>1</sub> and the question sent to Merlin<sub>2</sub>. We are interested in the problem of approximating the value of a free game within an additive error of  $\delta$ .

#### FREEGAME $_{\delta}$

Input: A free game  $\mathcal{T}$  and a constant  $\delta > 0$ .

Output: A value  $p$  such that  $|\omega(\mathcal{T}) - p| \leq \delta$ .

The *exponential time hypothesis* (ETH) is the conjecture that any deterministic algorithm for solving 3SAT requires  $2^{\Omega(n)}$  time. The *randomized exponential time hypothesis* (RETH) is the same hypothesis, but for randomized algorithms. Aaronson, Impagliazzo, and Moshkovitz have shown that, if ETH holds, then we have the following inapproximability result [1].

**Theorem 2 (Theorem 32 in [2]).** *If the ETH holds, then there exists a constant  $\Delta$  such that for all  $\delta \in [1/n, \Delta]$  the problem FREEGAME $_{\delta}$  cannot be solved faster than  $n^{\frac{\tilde{O}(\log n)}{\delta}}$ .*

This theorem was proved by providing a family of games such that, each game  $\mathcal{F}$  had either  $\omega(\mathcal{F}) = 1$ , or  $\omega(\mathcal{F}) < 1 - \delta$ , and showing that it is hard to decide which of these is the case. Theorem 2 produces a free game where the size of the question sets  $X$  and  $Y$  is proportional to the size of the answer sets  $A$  and  $B$ . For our proof we would like the size of  $X$  and  $Y$  to be *logarithmic* in the size of  $A$  and  $B$ . Fortunately, this can be achieved by applying the following sub-sampling result from the same paper. Since our results will rely on this sub-sampling lemma, our lower bounds will depend on RETH, rather than ETH.

**Lemma 3 (Corollary 46 in [2]).** *Given a free game  $\mathcal{F} = (X, Y, A, B, \mathcal{U}, V)$  and  $\epsilon > 0$ , we can randomly select a free game  $\mathcal{F}' = (X', Y', A, B, \mathcal{U}, V)$  such that  $|X| = |Y| = 2 \cdot \epsilon^{-\Lambda} \cdot \log(|A| + |B|)$  for some constant  $\Lambda$  such that, with high probability, we have  $|\omega(\mathcal{F}) - \omega(\mathcal{F}')| \leq \epsilon$ .*

### 3 Hardness of Approximating Social Welfare

**Overview.** In this section, we study the following *social welfare* problem. The *social welfare* of a pair of strategies  $(\mathbf{x}, \mathbf{y})$  is denoted by  $\text{SW}(\mathbf{x}, \mathbf{y})$  and is defined to be  $\mathbf{x}^T R \mathbf{y} + \mathbf{x}^T C \mathbf{y}$ . Given an  $\epsilon \geq 0$ , we define the set of all  $\epsilon$  equilibria as  $E^{\epsilon} = \{(\mathbf{x}, \mathbf{y}) : (\mathbf{x}, \mathbf{y}) \text{ is an } \epsilon\text{-NE}\}$ . Then, we define the *best social welfare* achievable by an  $\epsilon$ -NE in  $\mathcal{G}$  as  $\text{BSW}(\mathcal{G}, \epsilon) = \max\{\text{SW}(\mathbf{x}, \mathbf{y}) : (\mathbf{x}, \mathbf{y}) \in E^{\epsilon}\}$ . Using these definitions we now define the main problem that we consider:

#### $\epsilon$ -NE $\delta$ -SW

Input: A game  $\mathcal{G}$ , and two constants  $\epsilon, \delta > 0$ .

Output: An  $\epsilon$ -NE  $(\mathbf{x}, \mathbf{y})$  s.t.  $\text{SW}(\mathbf{x}, \mathbf{y})$  is within  $\delta$  of  $\text{BSW}(\mathcal{G}, \epsilon)$ .

We show a lower bound for this problem by reducing from  $\text{FREEGAME}_\delta$ . Let  $\mathcal{F}$  be a free game of size  $n$  from the family of free games that were used to prove Theorem 2. We have that either  $\omega(\mathcal{F}) = 1$  or  $\omega(\mathcal{F}) < 1 - \delta$  for some fixed constant  $\delta$ , and that it is hard to determine which of these is the case. We will construct a game  $\mathcal{G}$  such that for  $\epsilon = 1 - 4g \cdot \delta$ , where  $g < \frac{5}{12}$  is a fixed constant that we will define at the end of the proof, we have the following properties.

- **(Completeness)** If  $\omega(\mathcal{F}) = 1$ , then the unscaled  $\text{BSW}(\mathcal{G}, \epsilon) = 2$ .
- **(Soundness)** If  $\omega(\mathcal{F}) < 1 - \delta$ , then the unscaled  $\text{BSW}(\mathcal{G}, \epsilon) < 2(1 - g \cdot \delta)$ .

This will allow us to prove our lower bound using Theorem 2.

### 3.1 The Construction

The first step of the proof is to apply Lemma 3 to  $\mathcal{F}$  with  $\epsilon = \delta/2$  to produce a free game  $\mathcal{F}_s = (X, Y, A, B, \mathcal{U}, V)$  that will be fixed for the rest of this section. Since the question sets in  $\mathcal{F}$  have size  $O(|\mathcal{F}|)$ , we have that the question sets  $X$  and  $Y$  in  $\mathcal{F}_s$  have size  $\log(|\mathcal{F}|)$ . Furthermore, with high probability, it is hard to decide whether  $\omega(\mathcal{F}_s) = 1$  or  $\omega(\mathcal{F}_s) = 1 - \delta/2$ . Next, we use  $\mathcal{F}_s$  to construct a bimatrix game, which we will denote as  $\mathcal{G}$  throughout the rest of this section. The game is built out of four subgames, which are arranged and defined as follows.

|                 |                                  |                                  |
|-----------------|----------------------------------|----------------------------------|
|                 | C                                | D <sub>2</sub>                   |
| R               | R, C                             | -D <sub>2</sub> , D <sub>2</sub> |
| -D <sub>1</sub> | D <sub>1</sub>                   | 0                                |
| D <sub>1</sub>  | D <sub>1</sub> , -D <sub>1</sub> | 0, 0                             |

- The game  $(R, C)$  is built from  $\mathcal{F}_s$  in the following way. Each row of the game corresponds to a pair  $(x, a) \in X \times A$  and each column corresponds to a pair  $(y, b) \in Y \times B$ . Since all free games are cooperative, the payoff for each strategy pair  $(x, a), (y, b)$  is defined to be  $R_{(x,a),(y,b)} = C_{(x,a),(y,b)} = V(x, y, a(x), b(y))$ .
- The game  $(D_1, -D_1)$  is a zero-sum game. The game is a slightly modified version of a game devised by Feder et al. [14]. Let  $H$  be the set of all functions of the form  $f : Y \rightarrow \{0, 1\}$  such that  $f(y) = 1$  for exactly half<sup>4</sup> of the elements  $y \in Y$ . The game has  $|Y \times B|$  columns and  $|H|$  rows. For all  $f \in H$  and all  $(y, b) \in Y$  the payoffs are

$$(D_1)_{f,(y,b)} = \begin{cases} \frac{4}{1+4g \cdot \delta} & \text{if } f(y) = 1, \\ 0 & \text{otherwise.} \end{cases}$$

<sup>4</sup> If  $|Y|$  is not even, then we can create a new free game in which each question in  $|Y|$  appears twice. This will not change the value of the free game.

- The game  $(-D_2, D_2)$  is built in the same way as the game  $(D_1, -D_1)$ , but with the roles of the players swapped. That is, each column of  $(-D_2, D_2)$  corresponds to a function that picks half of the elements of  $X$ .
- The game  $(0, 0)$  is a game in which both players have zero matrices.

Observe that the size of  $(R, C)$  is the same as the size of  $\mathcal{F}_s$ , which is at most  $|\mathcal{F}|$ . The game  $(D_1, -D_1)$  has the same number of columns as  $C$ , and the number of rows is at most  $2^{|Y|} \leq 2^{2\delta^{-\Lambda} \log |\mathcal{F}|} = |\mathcal{F}|^{2\delta^{-\Lambda}}$ , where  $\Lambda$  is the constant obtained from Lemma 3. By the same reasoning, the number of columns in  $(-D_2, D_2)$  is at most  $|\mathcal{F}|^{2\delta^{-\Lambda}}$ . Thus, the size of  $\mathcal{G}$  is  $|\mathcal{F}|^{O(\delta^{-\Lambda})}$ , and in particular, for every constant  $\delta > 0$ , this reduction is polynomial.

### 3.2 Completeness

To prove completeness, it suffices to show that, if  $\omega(\mathcal{F}_s) = 1$ , then there exists a  $(1 - 4g \cdot \delta)$ -UNE of  $\mathcal{G}$  that has social welfare 2. To do this, assume that  $\omega(\mathcal{F}_s) = 1$ , and take a pair of optimal strategies  $(s_1, s_2)$  for  $\mathcal{F}_s$  and turn them into strategies for the players in  $\mathcal{G}$ . More precisely, the row player will place probability  $\frac{1}{|X|}$  on each answer chosen by  $s_1$ , and the column player will place probability  $\frac{1}{|Y|}$  on each answer chosen by  $s_2$ . By construction, this gives both players payoff 1, and hence the social welfare is 2. The hard part is to show that this is an approximate equilibrium, and in particular, that neither player can gain by playing a strategy in  $(D_1, -D_1)$  or  $(-D_2, D_2)$ . We prove this in the following lemma.

**Lemma 4.** *If  $\omega(\mathcal{F}_s) = 1$ , then there exists a  $(1 - 4g \cdot \delta)$ -UNE  $(\mathbf{x}, \mathbf{y})$  of  $\mathcal{G}$  with  $\text{SW}(\mathbf{x}, \mathbf{y}) = 2$ .*

### 3.3 Soundness

We now suppose that  $\omega(\mathcal{F}_s) < 1 - \delta/2$ , and we will prove that all  $(1 - 4g \cdot \delta)$ -UNE provide social welfare at most  $2 - 2g \cdot \delta$ . Throughout this subsection, we will fix  $(\mathbf{x}, \mathbf{y})$  to be a  $(1 - 4g \cdot \delta)$ -UNE of  $\mathcal{G}$ . We begin by making a simple observation about the amount of probability that is placed on  $(R, C)$ .

**Lemma 5.** *If  $\text{SW}(\mathbf{x}, \mathbf{y}) > 2 - 2g \cdot \delta$ , then  $\mathbf{x}$  places at least  $(1 - g \cdot \delta)$  probability on rows in  $(R, C)$ , and  $\mathbf{y}$  places at least  $(1 - g \cdot \delta)$  probability on columns in  $(R, C)$ .*

So, for the rest of this subsection, we can assume that both  $\mathbf{x}$  and  $\mathbf{y}$  place at least  $1 - g \cdot \delta$  probability on the subgame  $(R, C)$ . We will ultimately show that, if this is the case, then both players have payoff at most  $1 - \frac{1}{2} \cdot \delta + mg \cdot \delta$  for some constant  $m$  that will be derived during the proof. Choosing  $g = 1/(2m + 2)$  then ensures that both players have payoff at most  $1 - g \cdot \delta$ , and therefore that the social welfare is at most  $2 - 2g \cdot \delta$ .

**A Two-Prover Game.** We use  $(\mathbf{x}, \mathbf{y})$  to create a two-prover game. First, we define two distributions that capture the marginal probability that a question is played by  $\mathbf{x}$  or  $\mathbf{y}$ . Formally, we define a distribution  $\mathbf{x}'$  over  $X$  and a distribution

$\mathbf{y}'$  over  $Y$  such that for all  $x \in X$  and  $y \in Y$  we have  $\mathbf{x}'(x) = \sum_{a \in A} \mathbf{x}(x, a)$ , and  $\mathbf{y}'(y) = \sum_{b \in B} \mathbf{y}(y, b)$ . By Lemma 5, we can assume that  $\|\mathbf{x}'\|_1 \geq 1 - g \cdot \delta$  and  $\|\mathbf{y}'\|_1 \geq 1 - g \cdot \delta$ .

Our two-prover game will have the same question sets, answer sets, and verification function as  $\mathcal{F}_s$ , but a different distribution over the question sets. Let  $\mathcal{T}_{(\mathbf{x}, \mathbf{y})} = (X, Y, A, B, \mathcal{D}, V)$ , where  $\mathcal{D}$  is the product of  $\mathbf{x}'$  and  $\mathbf{y}'$ . Note that we have cheated slightly here, since  $\mathcal{D}$  is not actually a probability distribution. If  $\|\mathcal{D}\|_1 = c < 1$ , then we can think of this as Arthur having a  $1 - c$  probability of not sending any questions to the Merlins and awarding them payoff 0.

The strategies  $\mathbf{x}$  and  $\mathbf{y}$  can also be used to give us a strategy for the Merlins in  $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$ . Without loss of generality, we can assume that for each question  $x \in X$  there is exactly one answer  $a \in A$  such that  $\mathbf{x}(x, a) > 0$ , because if there are two answers  $a_1$  and  $a_2$  such that  $\mathbf{x}(x, a_1) > 0$  and  $\mathbf{x}(x, a_2) > 0$ , then we can shift all probability onto the answer with (weakly) higher payoff, and (weakly) improve the payoff to the row player. Since  $(R, C)$  is cooperative, this can only improve the payoff of the columns in  $(R, C)$ , and since the row player does not move probability between questions, the payoff of the columns in  $(-D_2, D_2)$  does not change either. Thus, after shifting, we arrive at a  $(1 - 4g \cdot \delta)$ -UNE of  $\mathcal{G}$  whose social welfare is at least as good as  $\text{SW}(\mathbf{x}, \mathbf{y})$ . Similarly, we can assume that for each question  $y \in Y$  there is exactly one answer  $b \in B$  such that  $\mathbf{y}(y, b) > 0$ .

So, we can define a strategy  $s_{\mathbf{x}}$  for Merlin<sub>1</sub> in the following way. For each question  $x \in X$ , the strategy  $s_{\mathbf{x}}$  selects the unique answer  $a \in A$  such that  $\mathbf{x}(x, a) > 0$ . The strategy  $s_{\mathbf{y}}$  for Merlin<sub>2</sub> is defined symmetrically.

We will use  $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$  as an intermediary between  $\mathcal{G}$  and  $\mathcal{F}_s$  by showing that the payoff of  $(\mathbf{x}, \mathbf{y})$  in  $\mathcal{G}$  is close to the payoff of  $(s_{\mathbf{x}}, s_{\mathbf{y}})$  in  $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$ , and that the payoff of  $(s_{\mathbf{x}}, s_{\mathbf{y}})$  in  $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$  is close to the payoff of  $(s_{\mathbf{x}}, s_{\mathbf{y}})$  in  $\mathcal{F}_s$ . Since we have a bound on the payoff of any pair of strategies in  $\mathcal{F}_s$ , this will ultimately allow us to bound the payoff to both players when  $(\mathbf{x}, \mathbf{y})$  is played in  $\mathcal{G}$ .

**Relating  $\mathcal{G}$  to  $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$ .** For notational convenience, let us define  $p_r(\mathcal{G}, \mathbf{x}, \mathbf{y})$  and  $p_c(\mathcal{G}, \mathbf{x}, \mathbf{y})$  to be the payoff to the row player and column player, respectively, when  $(\mathbf{x}, \mathbf{y})$  is played in  $\mathcal{G}$ . We begin by showing that the difference between  $p_r(\mathcal{G}, \mathbf{x}, \mathbf{y})$  and  $p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})$  is small. Once again we prove this for the payoff of the row player, but the analogous result also holds for the column player.

**Lemma 6.** *We have  $|p_r(\mathcal{G}, \mathbf{x}, \mathbf{y}) - p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})| \leq 4g \cdot \delta$ .*

**Relating  $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$  to  $\mathcal{F}_s$ .** First we show that if  $(\mathbf{x}, \mathbf{y})$  is indeed a  $(1 - 4g \cdot \delta)$ -UNE, then  $\mathbf{x}'$  and  $\mathbf{y}'$  must be close to uniform over the questions. We prove this for  $\mathbf{y}'$ , but the proof can equally well be applied to  $\mathbf{x}'$ . The idea is that, if  $\mathbf{y}'$  is sufficiently far from uniform, then there is set  $B \subseteq Y$  of  $|Y|/2$  columns where  $\mathbf{y}'$  places significantly more than 0.5 probability. This, in turn, means that the row of  $(D_1, -D_1)$  that corresponds to  $B$ , will have payoff at least 2, while the payoff of  $(\mathbf{x}, \mathbf{y})$  can be at most  $1 + 3g \cdot \delta$ , and so  $(\mathbf{x}, \mathbf{y})$  would not be a  $(1 - 4g \cdot \delta)$ -UNE. We formalise this idea in the following lemma. Define  $\mathbf{u}_X$  to be the uniform distribution over  $X$ , and  $\mathbf{u}_Y$  to be the uniform distribution over  $Y$ .

**Lemma 7.** *We have  $\|\mathbf{u}_Y - \mathbf{y}'\|_1 < 16g \cdot \delta$  and  $\|\mathbf{u}_X - \mathbf{x}'\|_1 < 16g \cdot \delta$ .*

With Lemma 7 at hand, we can now prove that the difference between  $p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})$  and  $p(\mathcal{F}_s, s_{\mathbf{x}}, s_{\mathbf{y}})$  must be small. This is because the question distribution  $\mathcal{D}$  used in  $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$  is a product of two distributions that are close to uniform, while the question distribution  $\mathcal{U}$  used in  $\mathcal{F}_s$  is a product of two uniform distributions. In the following lemma, we show that if we transform  $\mathcal{D}$  into  $\mathcal{U}$ , then we do not change the payoff of  $(s_{\mathbf{x}}, s_{\mathbf{y}})$  very much.

**Lemma 8.** *We have  $|p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}}) - p(\mathcal{F}_s, s_{\mathbf{x}}, s_{\mathbf{y}})| \leq 64g \cdot \delta$ .*

**Completing the Soundness Proof.** The following lemma uses the bounds derived in Lemmas 6 and 8, along with a suitable setting for  $g$ , to bound the payoff of both players when  $(\mathbf{x}, \mathbf{y})$  is played in  $\mathcal{G}$ .

**Lemma 9.** *If  $g = \frac{1}{138}$ , then both players have payoff at most  $1 - g \cdot \delta$  when  $(\mathbf{x}, \mathbf{y})$  is played in  $\mathcal{G}$ .*

Hence, we have proved that  $\text{SW}(\mathbf{x}, \mathbf{y}) \leq 2 - 2g \cdot \delta$ .

### 3.4 The Result

We can now state the theorem that we have proved in this section. We first rescale the game so that it lies in  $[0, 1]$ . The maximum payoff in  $\mathcal{G}$  is  $\frac{4}{1+4g \cdot \delta} \leq 4$ , and the minimum payoff is  $-\frac{4}{1+4g \cdot \delta} \geq -4$ . To rescale this game, we add 4 to all the payoffs, and then divide by 8. Let us refer to the scaled game as  $\mathcal{G}_s$ . Observe that an  $\epsilon$ -UNE in  $\mathcal{G}$  is a  $\frac{\epsilon}{8}$ -NE in  $\mathcal{G}_s$  since adding a constant to all payoffs does not change the approximation guarantee, but dividing all payoffs by a constant does change the approximation guarantee. So, we have the following theorem.

**Theorem 10.** *If RETH holds, then there exists a constant  $\Delta$  such that for all  $\delta \in [1/n, \Delta]$  the problem  $(\frac{1-4g \cdot \delta}{8})$ -NE  $\frac{g}{4} \cdot \delta$ -SW, where  $g = \frac{1}{138}$ , cannot be solved faster than  $n^{\tilde{O}(\delta^\Delta \log n)}$ , for some fixed constant  $\Delta$ .*

## 4 Hardness Results for Other Decision Problems

In this section we study a range of decision problems associated with approximate equilibria. Most are known to be NP-complete for the case of exact Nash equilibria [7, 16]. Table 1 shows all of the decision problems that we consider. For each problem, the input includes a bimatrix game and a quality of approximation  $\epsilon \in (0, 1)$ . We consider decision problems related to both  $\epsilon$ -NE and  $\epsilon$ -WSNE. Since  $\epsilon$ -NE is a weaker concept than  $\epsilon$ -WSNE, the hardness results for  $\epsilon$ -NE imply the same hardness for  $\epsilon$ -WSNE. We consider problems for  $\epsilon$ -WNSE only where the corresponding problem for  $\epsilon$ -NE is trivial. For example, observe that approximate  $\epsilon$ -NE with large support is a trivial problem, since we can always add a tiny amount of probability to each pure strategy without changing our expected payoff very much.

Our conditional quasi-polynomial lower bounds will hold for all  $\epsilon < \frac{1}{8}$ . Thus fix  $\epsilon^* < \frac{1}{8}$  for the rest of this section. We will appeal to Theorem 10, and thus we compute from  $\epsilon^*$  the parameters  $n$  and  $\delta$  that we require to apply this theorem. In particular, compute  $\delta^*$  to solve  $\epsilon^* = (\frac{1-4q \cdot \delta^*}{8})$ , which comes from Theorem 10, and choose  $n^*$  as  $\frac{1}{\delta^*}$ . Then, for  $n > n^*$  and  $\delta = \delta^*$  we can apply Theorem 10 to bound the social welfare achievable if  $\omega(\mathcal{F}_s) < 1 - \delta^*$  as  $\mathbf{u} = \frac{6}{8} - \frac{1}{522}\delta^*$ . Theorem 10 implies that in order to decide whether the game  $\mathcal{G}_s$  possess an  $\epsilon^*$ -NE that yields social welfare strictly greater than  $\mathbf{u}$  requires  $n^{\tilde{O}(\log n)}$  time, where  $\delta$  no longer appears in the exponent since we have fixed it as a constant  $\delta^*$  according to our choice of  $\epsilon^*$ .

The hardness of Problem 1 is a corollary of Theorem 10 when we set  $u = \frac{3}{8}$ . For the other problems in Table 1, we use  $\mathcal{G}_s$  to construct two new games:  $\mathcal{G}'$ , which adds one row and column to  $\mathcal{G}_s$ , is used to show hardness of Problems 2–9, and  $\mathcal{G}''$ , which in turn adds one row and column to  $\mathcal{G}'$ , is used to show hardness of Problem 10. We define  $\mathcal{G}'$  and  $\mathcal{G}''$  using the constants  $\mathbf{u}$  and  $\epsilon^*$  fixed above. The game  $\mathcal{G}'$  extends  $\mathcal{G}_s$  by adding the pure strategy  $i$  for the row player, and the pure strategy  $j$  for the column player, with payoffs as shown in Fig. 1.

$$\mathcal{G}' =$$

|                 |                               |         |                               |        |
|-----------------|-------------------------------|---------|-------------------------------|--------|
| $\mathcal{G}_s$ |                               |         | j                             |        |
|                 |                               |         | $0, \frac{3}{8} + \epsilon^*$ |        |
|                 |                               |         | $\vdots$                      |        |
|                 |                               |         | $0, \frac{3}{8} + \epsilon^*$ |        |
| i               | $\frac{3}{8} + \epsilon^*, 0$ | $\dots$ | $\frac{3}{8} + \epsilon^*, 0$ | $1, 1$ |

**Fig. 1.** The game  $\mathcal{G}'$ .

The payoffs for  $i$  and  $j$  were chosen so that: If the game  $\mathcal{G}_s$  possess an  $\epsilon^*$ -NE with social welfare  $\frac{6}{8}$ , then  $\mathcal{G}'$  posses at least one  $\epsilon^*$ -NE where the players do not play the pure strategies  $i$  and  $j$ ; if every  $\epsilon^*$ -NE of the game  $\mathcal{G}_s$  yields social welfare at most  $\mathbf{u}$ , then in *every*  $\epsilon^*$ -NE of  $\mathcal{G}'$ , the players place almost all of their probability on  $i$  and  $j$  respectively. Lemmas 11 and 12 show further properties that hold in the first case but not the second.

Notice that the expected payoff for the row player from the pure strategy  $i$  is at least  $\frac{3}{8} + \epsilon^*$  irrespective of the strategy the column player chooses. The same holds for the column player and the pure strategy  $j$ , i.e., the expected payoff that the column players gets from the pure strategy  $j$  is at least  $\frac{3}{8} + \epsilon^*$  irrespective of the strategy chosen by the row player. In what follows we will use  $S_R$  ( $S_C$ ), or  $S$  when it is clear from the context, to denote the set of pure strategies available to the row (column) from the  $(R, C)$  part of  $\mathcal{G}_s$  that corresponds to different questions in the free game  $\mathcal{F}_s$ .

First, we derive some properties of the equilibria of  $\mathcal{G}'$  when  $\mathcal{G}_s$  posses an  $\epsilon^*$ -NE with social welfare  $\frac{6}{8}$ .

**Lemma 11.** *If  $\mathcal{G}_s$  possesses an  $\epsilon^*$ -NE  $(\mathbf{x}, \mathbf{y})$  with social welfare  $\frac{6}{8}$ , then  $(\mathbf{x}, \mathbf{y})$  is an  $\epsilon^*$ -WSNE for  $\mathcal{G}'$  such that:*

- (a)  $\mathbf{x}^T R \mathbf{y} = \frac{3}{8}$  and  $\mathbf{x}^T C \mathbf{y} = \frac{3}{8}$ ,
- (b)  $\text{supp}(\mathbf{x}) \subseteq S_R$  and  $\text{supp}(\mathbf{y}) \subseteq S_C$ ,
- (c)  $|\text{supp}(\mathbf{x})| = |S_R|$  and  $|\text{supp}(\mathbf{y})| = |S_C|$ ,
- (d)  $\max_i x_i \leq \frac{1}{|S_R|}$  and  $\max_j y_j \leq \frac{1}{|S_C|}$ .

Next, we prove certain properties that all  $\epsilon^*$ -NE and  $\epsilon^*$ -WSNE of  $\mathcal{G}'$  possess if every  $\epsilon^*$ -NE of  $\mathcal{G}_s$  yields social welfare at most  $u$ .

**Lemma 12.** *If every  $\epsilon^*$ -NE of  $\mathcal{G}_s$  yields social welfare at most  $u$ , then in every  $\epsilon^*$ -NE  $(\mathbf{x}, \mathbf{y})$  of  $\mathcal{G}'$  it holds that:*

- ( $\alpha$ )  $x_i > 1 - \epsilon^*$  and  $y_j > 1 - \epsilon^*$ ,
- ( $\beta$ )  $\mathbf{x}^T R \mathbf{y} > 1 - 2\epsilon^*$  and  $\mathbf{x}^T C \mathbf{y} > 1 - 2\epsilon^*$ .

Furthermore, in every  $\epsilon^*$ -WSNE  $(\mathbf{x}, \mathbf{y})$  of  $\mathcal{G}'$  it holds that

- ( $\gamma$ )  $|\text{supp}(\mathbf{x})| = |\text{supp}(\mathbf{y})| = 1$ .

Observe that the combination of the claims of Lemmas 11 and 12 give the desired hardness results for the Problems 2–9. The combination of claim (a) from Lemma 11 with the claim ( $\beta$ ) from Lemma 12 gives the hardness result for the Problems 2 and 3; the combination of (b) with ( $\alpha$ ) gives the hardness for Problems 4 and 5; the combination of (d) with ( $\alpha$ ) gives the hardness for the Problem 6; and finally that hardness of Problems 7–9 follows from the combination of (c) with ( $\gamma$ ).

For Problem 10, we define a new game  $\mathcal{G}''$  by extending  $\mathcal{G}'$ . We add the new pure strategy  $i'$  for the row player and the new pure strategy  $j'$  for the column player, with payoffs constructed as shown in Fig. 2. We prove that if the game  $\mathcal{G}_s$  possesses an  $\epsilon^*$ -NE with social welfare  $\frac{3}{8}$ , then the game  $\mathcal{G}''$  possesses an  $\epsilon^*$ -WSNE  $(\mathbf{x}, \mathbf{y})$  such that  $i' \in \text{supp}(\mathbf{x})$ . Furthermore, we prove that if all  $\epsilon^*$ -NE of  $\mathcal{G}_s$  yield social welfare at most  $u$ , then for any  $\epsilon^*$ -WSNE  $(\mathbf{x}, \mathbf{y})$  it holds that  $i' \notin \text{supp}(\mathbf{x})$ .

$$\mathcal{G}'' = \begin{array}{c} \begin{array}{c} \phantom{i'} \\ \phantom{i'} \\ \phantom{i'} \\ \phantom{i'} \\ i' \end{array} \begin{array}{|c|c|c|c|} \hline & & & j' \\ \hline & & & \begin{array}{c} \frac{3}{8}, \frac{3}{8} \\ \vdots \\ \frac{3}{8}, \frac{3}{8} \end{array} \\ \hline \begin{array}{c} \frac{3}{8}, \frac{3}{8} \\ \dots \\ \frac{3}{8}, \frac{3}{8} \end{array} & & & 0, 0 \\ \hline \end{array} \end{array}$$

**Fig. 2.** The game  $\mathcal{G}''$ .



**Lemma 13.** *If the game  $\mathcal{G}_s$  possesses an  $\epsilon^*$ -NE with social welfare  $\frac{6}{8}$ , then the game  $\mathcal{G}''$  possesses an  $\epsilon^*$ -WSNE  $(\mathbf{x}, \mathbf{y})$  such that  $i' \in \text{supp}(\mathbf{x})$ .*

**Lemma 14.** *If all the  $\epsilon^*$ -NE of  $\mathcal{G}_s$  yield social welfare at most  $u$ , then for any  $\epsilon^*$ -WSNE  $(\mathbf{x}, \mathbf{y})$  of  $\mathcal{G}''$  it holds that  $i' \notin \text{supp}(\mathbf{x})$ .*

We now summarize the results of this section in the following theorem. Given the game  $\mathcal{G}_s$  we can construct games  $\mathcal{G}'$  and  $\mathcal{G}''$  such that the answer to the Problems 2–10 is “Yes” if  $\mathcal{G}_s$  possess an  $\epsilon^*$ -NE with social welfare  $\frac{3}{8}$  and “No” if every  $\epsilon^*$ -NE of  $\mathcal{G}_s$  has social welfare at most  $u$ .

**Theorem 15.** *Assuming the RETH, any algorithm that solves the Problems 1–10 for any constant  $\epsilon < \frac{1}{8}$  requires  $n^{\bar{\Omega}(\log n)}$  time.*

## References

1. Aaronson, S., Impagliazzo, R., Moshkovitz, D.: AM with multiple Merlins. In: Proceedings of CCC, pp. 44–55 (2014)
2. Aaronson, S., Impagliazzo, R., Moshkovitz, D.: AM with multiple Merlins. CoRR, abs/1401.6848 (2014)
3. Austrin, P., Braverman, M., Chlamtac, E.: Inapproximability of NP-complete variants of Nash equilibrium. *Theory Comput.* **9**, 117–142 (2013)
4. Bilò, V., Mavronicolas, M.: A catalog of  $\exists\mathbb{R}$ -complete decision problems about Nash equilibria in multi-player games. In: Proceedings of STACS, pp. 17:1–17:13 (2016)
5. Bosse, H., Byrka, J., Markakis, E.: New algorithms for approximate Nash equilibria in bimatrix games. *Theoret. Comput. Sci.* **411**(1), 164–173 (2010)
6. Braverman, M., Kun-Ko, Y., Weinstein, O.: Approximating the best Nash equilibrium in  $n^{o(\log n)}$ -time breaks the exponential time hypothesis. In: Proceedings of SODA, pp. 970–982 (2015)
7. Conitzer, V., Sandholm, T.: New complexity results about Nash equilibria. *Games Econ. Behav.* **63**(2), 621–641 (2008)
8. Czumaj, A., Deligkas, A., Fasoulakis, M., Fearnley, J., Jurdziński, M., Savani, R.: Distributed methods for computing approximate equilibria. In: Proceedings of WINE (2016)
9. Czumaj, A., Fasoulakis, M., Jurdziński, M.: Approximate Nash equilibria with near optimal social welfare. In: Proceedings of IJCAI, pp. 504–510 (2015)
10. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. *SIAM J. Comput.* **39**(1), 195–259 (2009)
11. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: Progress in approximate Nash equilibria. In: Proceedings of EC, pp. 355–358 (2007)
12. Daskalakis, C., Mehta, A., Papadimitriou, C.H.: A note on approximate Nash equilibria. *Theoret. Comput. Sci.* **410**(17), 1581–1588 (2009)
13. Fearnley, J., Goldberg, P.W., Savani, R., Sørensen, T.B.: Approximate well-supported Nash equilibria below two-thirds. In: Proceedings of SAGT, pp. 108–119 (2012)
14. Feder, T., Nazerzadeh, H., Saberi, A.: Approximating Nash equilibria using small-support strategies. In: Proceedings of EC, pp. 352–354 (2007)

15. Garg, J., Mehta, R., Vazirani, V.V., Yazdanbod, S.: ETR-completeness for decision versions of multi-player (symmetric) Nash equilibria. In: Proceedings of ICALP, pp. 554–566 (2015)
16. Gilboa, I., Zemel, E.: Nash and correlated equilibria: some complexity considerations. *Games Econ. Behav.* **1**(1), 80–93 (1989)
17. Hazan, E., Krauthgamer, R.: How hard is it to approximate the best Nash equilibrium? *SIAM J. Comput.* **40**(1), 79–91 (2011)
18. Kontogiannis, S.C., Spirakis, P.G.: Well supported approximate equilibria in bimatrix games. *Algorithmica* **57**(4), 653–667 (2010)
19. Lipton, R.J., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: Proceedings of EC, pp. 36–41 (2003)
20. Nash, J.: Non-cooperative games. *Ann. Math.* **54**(2), 286–295 (1951)
21. Rubinfeld, A.: Settling the complexity of computing approximate two-player Nash equilibria. In: Proceedings of FOCS, pp. 258–265 (2016)
22. Tsaknakis, H., Spirakis, P.G.: An optimization approach for approximate Nash equilibria. *Internet Math.* **5**(4), 365–382 (2008)

# Multilinear Games

Hau Chan<sup>1</sup>, Albert Xin Jiang<sup>1(✉)</sup>, Kevin Leyton-Brown<sup>2</sup>, and Ruta Mehta<sup>3</sup>

<sup>1</sup> Department of Computer Science, Trinity University,  
San Antonio, USA

{hchan,xjiang}@trinity.edu

<sup>2</sup> Department of Computer Science, University of British Columbia,  
Vancouver, Canada

kevinlb@cs.ubc.ca

<sup>3</sup> Department of Computer Science, UIUC, Champaign, USA

rutamehta@cs.illinois.edu

**Abstract.** In many games, players' decisions consist of multiple sub-decisions, and hence can give rise to an exponential number of pure strategies. However, this set of pure strategies is often structured, allowing it to be represented compactly, as in network congestion games, security games, and extensive form games. Reduction to the standard normal form generally introduces exponential blow-up in the strategy space and therefore are inefficient for computation purposes. Although individual classes of such games have been studied, there currently exists no general purpose algorithms for computing solutions. equilibrium.

To address this, we define *multilinear games* generalizing all. Informally, a game is multilinear if its utility functions are linear in each player's strategy, while fixing other players' strategies. Thus, if pure strategies, even if they are exponentially many, are vectors in polynomial dimension, then we show that mixed-strategies have an equivalent representation in terms of *marginals* forming a polytope in polynomial dimension.

The canonical representation for multilinear games can still be exponential in the number of players, a typical obstacle in multi-player games. Therefore, it is necessary to assume additional structure that allows computation of certain sub-problems in polynomial time. Towards this, we identify two key subproblems: computation of *utility gradients*, and optimizing linear functions over strategy polytope. Given a multilinear game, with polynomial time subroutines for these two tasks, we show the following: (a) We can construct a polynomially computable and continuous fixed-point formulation, and show that its approximate fixed-points are approximate NE. This gives containment of approximate NE computation in PPAD, and settles its complexity to PPAD-complete. (b) Even though a coarse correlated equilibrium can potentially have exponential representation, through LP duality and a carefully designed separation oracle, we provide a polynomial-time algorithm to compute one with polynomial representation. (c) We show existence of an approximate NE with support-size logarithmic in the strategy polytope dimensions.

# 1 Introduction

The computation of game-theoretic solution concepts is a central problem at the intersection of game theory and computer science. For games with large numbers of players, the standard normal form game representation requires exponential space even if the number of strategies per players is two, and is thus not a practical option as a basis for computation. Most games of practical interest have highly structured *utility functions*, and it is possible to represent them compactly. A line of research thus exists to look for *compact game representations* that are able to succinctly describe structured games, including work on graphical games [16], multi-agent influence diagrams [17] and action-graph games [15].

In many real-world domains, each player needs to make a decision that consists of multiple sub-decisions (e.g., assigning a set of resources or finding a path in a network), and hence the number of pure strategies per player itself can be exponential. The single-player versions of these decision problems have been well studied in the field of combinatorial optimization, with mature general modeling languages such as AMPL and solvers like CPLEX. For the multi-player case, several classes of games studied in the recent literature have structured strategy spaces, including network congestion games [4, 7], simultaneous auctions and other multi-item auctions [23, 25], dueling algorithms [12], integer programming games [18], Blotto games [1], and security games [19, 24]. These papers proposed compact game representations suitable for their specific domains, and corresponding algorithms for computing solution concepts, which take advantage of the specific structure in the representations. However, it is not obvious whether algorithmic techniques developed for one domain can be transferred to another.

A general successful approach in the study of efficient computation for compact representations is the following: identify subtasks that are required for most existing algorithms of these solution concepts, and then speed up these subtasks by exploiting the structure of the compact representation. [7, 21] identified *expected utility computation* given a mixed strategy as the subtask to compute correlated equilibrium efficiently, and to show NE computation is in PPAD. They also showed that games like graphical, polymatrix, and symmetric, this subtask can be done in polynomial time. A crucial assumption behind these results is *polynomial type*: roughly, it is feasible to enumerate pure strategies of all the players. This is not the case for games with structured strategies, in which such explicit strategy enumeration can take exponential time. [7] showed PPAD membership of NE computation for two additional subclasses: network congestion and extensive form games, but the general case remained open.

In this paper, we present a unified algorithmic framework for games with structured polytopal strategy spaces, in which each player’s set of pure strategies is defined to be integer points in a polytope. Our contributions are as follows.

1. We identify *multilinearity* as an important property of games that enables us to represent the players’ mixed strategies compactly. Informally, a game is multilinear if its utility functions are linear in each player’s strategy, while fixing other players’ strategies. We show that many existing game forms, like Bayesian, congestion, security, etc., are multilinear (see [5]).

2. The canonical representation of multilinear games still grows exponentially in the number of players. Therefore, it is necessary to assume additional structure that allows some computation in polynomial time, like done in [7, 21]. Towards this, we identify two key subproblems: computation of *utility gradients*, and optimizing linear functions over strategy polytopes. Given a multilinear game, with polynomial time subroutines for these two tasks, we show the following: (a) computing an approximate Nash equilibrium is in PPAD and (b) a coarse correlated equilibrium can be computed in polynomial time. These results are generalizations of [7, 21], respectively, from games of polynomial type to multilinear games.
3. We prove that given a multilinear game, there exists an approximate NE with support-size logarithmic in the strategy polytope dimensions. This generalizes [2], which gave bounds logarithmic in the number of strategies.

### 1.1 Technical Overview

Our approach is based on a compact representation of mixed strategies as *marginal vectors*, which is a point in the strategy polytope induced by the mixed strategy distributions. When the game is multilinear, all mixed strategies with the same marginal vector are payoff-equivalent (Lemmas 1 and 2). Thus, we can work in the marginal vector space instead of the exponentially higher-dimensional space of mixed strategies. We adapt existing algorithmic approaches such that whenever the algorithm calls for enumeration of pure strategies (e.g., for computing a best response), we instead solve a linear optimization problem in the space of marginal vectors, which can in turn be reduced to the two subproblems, namely computation of utility gradient given a marginal strategy profile, and optimizing a linear function over the polytope of marginal strategies. Given polynomial-time procedures for these two, we show a number of computational results.

Next we analyze complexity of computing an equilibrium. Since normal-form games are subcase of multilinear games, irrationality of NE [20], and PPAD-hardness for NE computation [6, 8] follows. Due to exponentially many pure strategies per player, containment of approximate NE computation in PPAD does not carry forward to multilinear games. Towards this, we design a fixed-point formulation to capture NE in marginal profiles, and show that corresponding approximate fixed-points exactly capture approximate NE. Furthermore, we show polynomial-continuity and polynomial-computability (see [5] or [9] for definitions) of the function by finding its equivalent representation in terms of projection operator, and obtaining a convex quadratic formulation for function evaluation, respectively. Finally, due to a result of [9], all of these together implies containment of finding an approximate NE in PPAD for multilinear games.

For computing CCE (Theorem 2), we adapt the Ellipsoid Against Hope approach of [21] and its refinement [14]. Applied directly to our setting, this approach would involve running the ellipsoid method in a space whose dimension is roughly the total number of pure strategies of all the players, yielding an exponential-time algorithm. We instead use a related but different convex programming

formulation, and then (through use of the multilinear property) transform it into a linear program of polynomial number of variables, which is then amenable to the ellipsoid method. Although the final output is not in terms of mixed strategies or marginal vectors (instead it is a correlated distribution with small support), a crucial intermediate step (the separation oracle of the ellipsoid method) requires linear optimization over the space of marginal vectors.

Finally, we show existence of approximate NE with logarithmic support using the probabilistic method, together with applying concentration inequalities on marginals to avoid union bound on exponentially many terms (Theorem 4).

Due to space constraint next we give an overview of our results, while all the proofs and some of the details can be found in the full paper [5].

## 2 Preliminaries

**Notations.** We use boldface letters, like  $\mathbf{x}$ , to denote vectors, and  $x_i$  to denote its  $i^{\text{th}}$  coordinate. To denote the set of  $\{1, \dots, m\}$  we use  $[m]$ . We use  $\mathbb{Z}_+$  and  $\mathbb{R}_+$  to denote the sets of non-negative integers and reals, respectively.

A game is specified by  $(N, S, u)$ , where  $N = \{1, \dots, n\}$  is the set of players. Each player  $i \in N$  chooses from a finite set of pure strategies  $S_i$ . Denote by  $s_i \in S_i$  a pure strategy of player  $i$ . Then  $S = \prod_i S_i$  is the set of pure-strategy profiles. Moreover,  $u = (u_1, \dots, u_n)$  are the utility functions of the players, where the utility function of player  $i$  is  $u_i : S \rightarrow \mathbb{R}$ .

In normal-form games, strategy sets  $S_i$ s and utility functions  $u_i$ s are specified explicitly. Thus, the size of the representation is of the order of  $n|S| = n \prod_i |S_i|$ .

A mixed strategy  $\sigma_i$  of player  $i$  is a probability distribution over her pure strategies. Let  $\Sigma_i = \Delta(S_i)$  be  $i$ 's set of mixed strategies, where  $\Delta(\cdot)$  denotes the set of probability distributions over a finite set. Denote by  $\sigma = (\sigma_1, \dots, \sigma_n)$  a mixed strategy profile, and  $\Sigma = \prod_i \Sigma_i$  the set of mixed strategy profiles. Denote by  $\sigma_{-i}$  the mixed strategy profile of players other than  $i$ .  $\sigma$  induces a probability distribution over pure strategy profiles. Denote by  $u_i(\sigma)$  the expected utility of player  $i$  under  $\sigma$ :  $u_i(\sigma) = E_{\mathbf{s} \sim \sigma} [u_i(\mathbf{s})] = \sum_{\mathbf{s} \in S} u_i(\mathbf{s}) \prod_{k \in N} \sigma_k(s_k)$ , where  $\sigma_k(s_k)$  is player  $k$ 's probability of playing the pure strategy  $s_k$ .

**Nash Equilibrium (NE).** Player  $i$ 's strategy  $\sigma_i$  is a best response to  $\sigma_{-i}$  if  $\sigma_i \in \arg \max_{\sigma'_i \in \Sigma_i} u_i(\sigma'_i, \sigma_{-i})$ . A mixed strategy profile  $\sigma$  is a Nash equilibrium if for each player  $i \in N$ ,  $\sigma_i$  is a best response to  $\sigma_{-i}$ .

Another important solution concept is Coarse Correlated Equilibrium (CCE). Consider a distribution over the set of pure-strategy profiles. This can be represented by a vector  $\mathbf{x}$ , satisfying  $\mathbf{x} \geq 0$ ,  $\sum_{\mathbf{s} \in S} x_{\mathbf{s}} = 1$ . The expected utility for player  $i$  under  $\mathbf{x}$  is  $u_i(\mathbf{x}) = \sum_{\mathbf{s} \in S} x_{\mathbf{s}} u_i(\mathbf{s})$ . Given  $\mathbf{x}$ , the expected utility for player  $i$  if he deviates to strategy  $s_i$  is:  $u_i^{s_i}(\mathbf{x}) = \sum_{\mathbf{s}_{-i}} x_{\mathbf{s}_{-i}} u_i(s_i, \mathbf{s}_{-i})$ , where  $x_{\mathbf{s}_{-i}} = \sum_{s_i \in S_i} x_{(s_i, \mathbf{s}_{-i})}$  is the marginal probability of  $\mathbf{s}_{-i}$  in distribution  $\mathbf{x}$ . Let

$$g_i(\mathbf{x}) = \max_{s_i \in S_i} u_i^{s_i}(\mathbf{x}), \quad (1)$$

i.e. player  $i$ 's expected utility if he deviates to a best response against  $\mathbf{x}$ .

**Definition 1.** A distribution  $\mathbf{x}$  satisfying  $\mathbf{x} \geq 0$ ,  $\sum_{s \in S} x_s = 1$  is a Coarse Correlated Equilibrium (CCE) if it satisfies the following incentive constraints:  $u_i(\mathbf{x}) \geq g_i(\mathbf{x}), \forall i \in N$ .

A rational polytope,  $P = \{\mathbf{x} \in \mathbb{R}^m \mid D\mathbf{x} \leq \mathbf{f}\}$ , is defined by a set of inequalities with integer coefficients (i.e., matrix  $D$  and vector  $\mathbf{f}$  consist of integers).

### 3 Multilinear Games

#### 3.1 Polytopal Strategy Space

We are interested in games in which a pure strategy has multiple components. Without loss of generality, if each pure strategy of player  $i$  has  $m_i$  components, we can associate each such pure strategy with an  $m_i$ -dimensional nonnegative integer vector. Then the set of pure strategies for each player  $i$  is  $S_i \subset \mathbb{Z}_+^{m_i}$ . In general the number of integer points in  $S_i$  can grow exponentially in  $m_i$ . Thus, we need a compact representation of  $S_i$ .

In most studies of games with structured strategy spaces, each  $S_i$  can be expressed as the set of integer points in a rational polytope  $P_i \subset \mathbb{R}_+^{m_i}$ , i.e.,  $S_i = P_i \cap \mathbb{Z}_+^{m_i}$ . We call such an  $S_i$  a *polytopal pure strategy set*. We assume  $P_i$  is nonempty, bounded and contained in the nonnegative quadrant  $\mathbb{R}_+^{m_i}$ . To represent the strategy space, we only need to specify the set of linear constraints defining  $P_i = \{\mathbf{p} \in \mathbb{R}_+^{m_i} \mid D_i \mathbf{p} \leq \mathbf{f}_i\}$ , with each linear constraint requiring us to store  $O(m_i)$  integers. We call this game a *game with polytopal strategy spaces*.

For example, one common scenario is when there are  $k$  finite sets  $S_i^1, \dots, S_i^k$ , and player  $i$  needs to simultaneously select one action in each of these sets. This happens in Bayesian games in which a player needs to choose an action for each of his type, extensive form games in which a player needs to choose an action in each information set, and simultaneous auctions, among others. The player's pure strategy set  $S_i$  is a polytopal strategy space with  $P_i$  being the product of  $k$  simplices. Second common type of strategy set is a uniform matroid: given a universe  $[m_i]$ , player  $i$ 's pure strategy is a subset of size  $k$ . This can (e.g.) represent security scenarios in which a defender player  $i$  in charge of protecting  $m_i$  target, but due to limited resources can only cover  $k$  of targets [19]. Then player  $i$ 's strategy can be represented as the 0–1 vector encoding the subset, and the strategy set can be represented as a polytopal strategy set with  $P_i = \{\mathbf{p} \in \mathbb{R}^{m_i} \mid \sum_{j \in [m_i]} p_j = k\}$ . Third common type of strategy is to select a path in a network, from a given source to a given destination. This can model routing of data traffic in a network congestion game, or patrol / attack routes in security settings [13, 26]. Here,  $s_i$  can be modeled as a 0–1 vector specifying the subset of edges forming the chosen path.  $S_i$  can be represented as a polytopal strategy space, where  $P_i$  consists of a set of flow constraints, as in [7].

#### 3.2 Mixed Strategies and Multilinearity

In this paper, we are focusing on computation of solution concepts in which players are playing mixed strategies, such as Nash equilibrium. The first challenge we face is the representation of mixed strategies. Recall that a mixed strategy  $\sigma_i$  of

player  $i$  is a probability distribution over the set of pure strategies  $S_i$ . When  $|S_i|$  is exponential, representing  $\sigma_i$  explicitly would take exponential space. Thus we would like a compact representation of mixed strategies, i.e., a way to represent a mixed strategy using only polynomial number of bits. One approach would be to only use mixed strategies of polynomial-sized *support*, where support is the set of pure strategies played with non-zero probability. Such strategies can be stored as sparse vectors requiring polynomial space; however, the space of small-support mixed strategies is not convex, and this is problematic for computation.

We list a set of desirable features for a compact representation of mixed strategies: (1) the expected utilities of the game can be expressed in terms of this compact representation; (2) the space of the resulting compactly-represented strategies is convex; (3) given this compact representation, we can efficiently recover a mixed strategy (e.g., as a mixture over a small number of pure strategies, or by providing a way to efficiently pure strategies from the mixed strategy). We show that such a compact representation is possible if the game is *multilinear*.

**Definition 2.** Consider a game  $\Gamma$  with polytopal strategy sets, with  $S_i = P_i \cap \mathbb{Z}_+^{m_i}$  for each player  $i$ .  $\Gamma$  is a multilinear game if

1. for each player  $i$ , there exists  $U^i \in \mathbb{R}^{\prod_{k \in N} m_k}$  such that for all  $\mathbf{s} \in S$ ,  $u_i(\mathbf{s}) = \sum_{(j_1 \dots j_n) \in \prod_k [m_k]} U_{j_1 \dots j_n}^i \prod_{k \in N} s_{k, j_k}$ , where  $[m_k] = \{1, \dots, m_k\}$ ;
2. The extreme points (i.e. vertices) of  $P_i$  are integer vectors, which implies that  $P_i = \text{conv}(S_i)$ , where  $\text{conv}(S_i)$  is the convex hull of  $S_i$ .

In particular, given a fixed  $\mathbf{s}_{-j}$ ,  $u_i$  is a linear function of  $\mathbf{s}_j$ . In other words, a multilinear game's utility functions are multilinear in the players' strategies.

Condition 2 of Definition 2 is satisfied if  $P_i$ 's constraint matrix  $D_i$  is *totally unimodular*. Total unimodularity is a well-studied property satisfied by the constraint matrices of many polytopal strategy spaces studied in the literature, including the network flow constraint matrix of network congestion games, the uniform matroid constraints of security games [19], and the doubly-stochastic constraints representing rankings in the search engine ranking duel [12]. When Condition 2 is not satisfied, we can redefine  $P_i$  to be  $\text{conv}(S_i)$ , but the new  $P_i$  may have exponentially more constraints. Indeed, dropping Condition 2 would allow us to express various NP-hard single-agent combinatorial optimization problems (e.g. set cover, knapsack). See [5] for examples that demonstrates how security, congestion, extensive-form, and Bayesian games are multilinear.

Given a mixed strategy  $\sigma_i$ , define the *marginal vector*  $\pi_i$  corresponding to  $\sigma_i$  as the expectation over the pure strategy space  $S_i$  induced by the distribution  $\sigma_i$ , i.e.,  $\pi_i = E_{\sigma_i}[\mathbf{s}_i] = \sum_{\mathbf{s}_i \in S_i} \sigma_i(\mathbf{s}_i) \mathbf{s}_i$ . Denote by  $\pi_{ij}$  the  $j$ -th component of  $\pi_i$ . The set of marginal vectors is exactly  $\text{conv}(S_i) = P_i$ . Given a mixed strategy profile  $\sigma$ , we call the corresponding collection of marginal vectors  $\pi = (\pi_1, \dots, \pi_n) \in P = \times_i P_i$  the *marginal strategy profile*. By slight abuse of notation let us denote by

$$u_i(\pi) = \sum_{(j_1 \dots j_n) \in \prod_k [m_k]} U_{j_1 \dots j_n}^i \prod_{k \in N} \pi_{k, j_k} \quad (2)$$

player  $i$ 's expected utility under marginal strategy profile  $\pi$ .



**Lemma 1.** *Given a mixed strategy profile  $\sigma \in \Sigma$  and a marginal vector  $\pi \in P$ , if  $\forall i, \pi_i = \sum_{s_i \in S_i} \sigma_i(s_i) s_i$  then  $\forall i, u_i(\sigma) = u_i(\pi)$ .*

That is, marginal vectors capture all payoff-relevant information about mixed strategies, and thus we can use them to compactly represent the space of mixed strategies. We note that this property does not hold for arbitrary games.

Suppose a mixed strategy profile  $\sigma$  with marginals  $\pi = (\pi_1, \dots, \pi_n)$  is a Nash equilibrium of a multilinear game. By multilinearity any mixed strategy profile having the same marginals are payoff-equivalent to  $\sigma$ , and therefore also a Nash equilibrium. Let us define Nash equilibrium in terms of marginals:

**Marginal NE.**  $\pi \in P$  is a *marginal NE* iff  $\forall i, u_i(\pi) \geq u_i(\pi'_i, \pi_{-i}), \forall \pi'_i \in P_i$ .

The next lemma follows easily using Lemma 1, and the fact that any vector  $\pi_i \in P_i$  can be represented as a convex combination of extreme points of  $P_i$ , and extreme points of  $P_i$  are in  $S_i$ .

**Lemma 2.** *A mixed-strategy profile  $\sigma \in \Sigma$  is a NE iff corresponding marginal strategy profile  $\pi \in P$ , where  $\pi_i = \sum_{s_i \in S_i} \sigma_i(s_i) s_i, \forall i \in N$ , is a marginal NE.*

## 4 Computation with Multilinear Games

We now show that many algorithmic results for computing various solutions for normal form games and other game representations of polynomial type can be adapted to multilinear games, with strategies represented as marginals. We follow a “modular” approach, similar to [7, 21]’s treatment of computation of Nash equilibrium and correlated equilibrium in games of polynomial type: we first identify certain key subproblems, then develop general algorithmic results assuming these subproblems can be efficiently computed. We note that a wide variety of games do has such specific structure (see [5]).

### 4.1 Utility Gradient

Recall that we can express the expected utilities of players using marginal vectors by Eq. (2) (Lemma 1). However, a direct computation of expected utility using (2) would require summing over a number of terms exponential in  $n$ . Also, computing expected utilities may not be enough: consider the task of determining if a mixed strategy profile (as marginals) is a Nash equilibrium. One needs to compute the expected utility for each pure strategy deviation of  $i$  in order to verify that  $i$  is playing a best response, but that would require enumerating all pure strategies. Instead, we identify a related but different computational problem as the key subtask for equilibrium computation for multilinear games.

Due to multilinearity, after fixing the strategies of players  $N \setminus \{k\}$ ,  $u_i(\pi)$  is a linear function of  $\pi_{k1}, \dots, \pi_{km_k}$ . We define the *utility gradient* of player  $i$  with respect to player  $k$ ’s marginal,  $\nabla_k(u_i(\pi_{-k})) \in \mathbb{R}^{m_k}$ , to be the vector of coefficients of this linear function. Formally,  $\forall j_k \in [m_k]$ ,

$$(\nabla_k u_i(\pi_{-i}))_{j_k} \equiv \sum_{(j_1, \dots, j_{k-1}, j_{k+1}, \dots, j_n) \in \prod_{\ell=1}^{N \setminus \{k\}} [m_\ell]} U_{j_1 \dots j_n}^i \prod_{\ell \in N \setminus \{k\}} \pi_{\ell, j_\ell}.$$

*Problem 1 (UtilGradient).* Given a compactly represented game that satisfies multilinearity, given players  $i, k \in N$ , and  $\boldsymbol{\pi}_{-k}$ , compute  $\nabla_k(u_i(\boldsymbol{\pi}_{-k}))$ .

Consider the problem of computing the utility gradients. As with expected utility computation, direct summation would require time exponential in  $n$ . With a compact game representation this problem could be solved in polynomial time.

## 4.2 PolytopeSolve and Decomposing Marginals

The other key subproblem we identify, PolytopeSolve, is the optimization of an arbitrary linear objective in each player's strategy polytope.

*Problem 2 (PolytopeSolve).* Given a compactly represented game with polytopal strategy space, player  $i$ , and a vector  $\mathbf{d} \in \mathbb{R}^{m_i}$ , compute  $\arg \max_{\mathbf{x} \in P_i} \mathbf{d}^T \mathbf{x}$ .

Let us consider the issue of constructing a mixed strategy given a marginal vector. First of all, since we have assumed that the extreme points of the polytope  $P_i$  are integer points, and thus  $P_i = \text{conv}(S_i)$ , this becomes the problem of describing a point in a polytope by a convex combination of extreme points of the polytope. By Caratheodory theorem, given  $\boldsymbol{\pi}_i \in \mathbb{R}^{m_i}$  there exists a mixed strategy of support size at most  $(m_i + 1)$  that matches the marginals. Existing work, such as the Birkhoff-von Neumann theorem and its generalizations [3], provides efficient constructions for different types of polytopes. The most general result by Grostchel et al. [11] reduces the problem to the task of optimizing an arbitrary linear objective over the polytope, i.e., PolytopeSolve.

**Theorem 1 (Grostchel et al. [11]).** *Suppose the PolytopeSolve can be solved in polynomial time. Then, the following problem DECOMPOSE( $P_i$ ) can be solved in polynomial time: Given  $\boldsymbol{\pi}_i \in P_i$ , find a polynomial number of extreme points of  $P_i$  (i.e., pure strategies)  $\mathbf{s}_i^1, \dots, \mathbf{s}_i^K \in S_i$  and weights  $\lambda_1, \dots, \lambda_K \geq 0$  such that  $\sum_{k=1}^K \lambda_k = 1$  and  $\boldsymbol{\pi}_i = \sum_{k=1}^K \lambda_k \mathbf{s}_i^k$ .*

We note that the computational complexity of PolytopeSolve depends only on the strategy polytopes  $P_i$ s of the game, and not on the utility functions. PolytopeSolve can be definitely solved in polynomial time by linear programming if  $P_i$  is given by a polynomial number of linear constraints; this holds for all examples we discussed in this paper. Since the objective is linear,  $\arg \max_{\mathbf{x} \in P_i} \mathbf{d}^T \mathbf{x} = \arg \max_{\mathbf{x} \in S_i} \mathbf{d}^T \mathbf{x}$ , i.e., we can alternatively solve the optimization problem over  $S_i$ , which may be more amenable to combinatorial methods.

For the case when  $P_i$  has exponentially many constraints, Grostchel et al. [11] also showed that PolytopeSolve is equivalent to the SEPARATION problem (also known as a *separation oracle*): Given a vector  $\boldsymbol{\pi}_i \in \mathbb{R}^{m_i}$ , either answers that  $\boldsymbol{\pi}_i \in P_i$ , or produces a hyperplane that separates  $\boldsymbol{\pi}_i$  and  $P_i$ .

## 4.3 Best Response

We observe that by construction,  $u_i(\boldsymbol{\pi}) = \boldsymbol{\pi}_i^T \nabla_i u_i(\boldsymbol{\pi}_{-i})$ . Then given  $\boldsymbol{\pi}$ , the best response for player  $i$  is the solution of the following optimization: maximize

$\pi_i^T \nabla_i u_i(\pi_{-i})$  subject to  $\pi_i \in P_i$ . This is a linear program with feasible region  $P_i$ , which is an instance of the problem PolytopeSolve. The coefficients of the linear objective are exactly the utility gradient  $\nabla_i u_i(\pi_{-i})$ .

**Proposition 1.** *Suppose we have a compact game representation with polynomial-time procedures for both UtilGradient and PolytopeSolve. Then the best response problem can be computed in polynomial time.*

As a corollary, under the same assumptions, we get that checking if a given profile  $\pi$  is a Nash equilibrium can be done in polynomial time.

#### 4.4 Computing Coarse Correlated Equilibrium

**Approximate CCE.** Given a multilinear game, an approximate CCE can be computed by simulating no-regret dynamics (a.k.a. online convex programming) for each player. For example, one such no-regret dynamic is Generalized Infinitesimal Gradient Ascent (GIGA) [27], where in each iteration, for each player  $i$  we move  $\pi_i$  along the direction of the utility gradient  $\nabla_i u_i(\pi_{-i})$ , and then project the resulting point back to  $P_i$ . The projection step is a convex optimization problem on  $P_i$ , and can be solved efficiently given an efficient separation oracle, or equivalently a procedure for PolytopeSolve. Therefore, under the same assumptions as Proposition 1, approximate CCE can be found efficiently.

**Exact CCE.** The above procedure does not guarantee exact CCE in polynomial-time. Next we obtain such a procedure, using LP duality and carefully designed separation oracle to get the following theorem.

**Theorem 2.** *Consider a multilinear game, with polynomial time subroutines for UtilGradient and PolytopeSolve. Then an exact Coarse Correlated Equilibrium (CCE) can be computed in polynomial time.*

Recall that a distribution over the set of pure-strategy profiles can be represented by a vector  $\mathbf{x}$ , satisfying  $\mathbf{x} \geq 0$ ,  $\sum_{\mathbf{s} \in S} x_{\mathbf{s}} = 1$ . Given a multilinear game, the expected utility for player  $i$  under  $\mathbf{x}$  is  $u_i(\mathbf{x}) = \sum_{\mathbf{s}} x_{\mathbf{s}} u_i(\mathbf{s}) = \sum_{\mathbf{s}} \sum_{j_1, \dots, j_n} U_{j_1, \dots, j_n}^i x_{\mathbf{s}} \prod_k s_{k, j_k}$ . Given  $\mathbf{x}$ , the expected utility for player  $i$  if he deviates to strategy  $\mathbf{s}_i$  is:  $u_i^{\mathbf{s}_i}(\mathbf{x}) = \sum_{\mathbf{s}_{-i}} x_{\mathbf{s}_{-i}} u_i(\mathbf{s}_i, \mathbf{s}_{-i}) = \sum_{\mathbf{s}_{-i}} \sum_{j_1, \dots, j_n} U_{j_1, \dots, j_n}^i x_{\mathbf{s}_{-i}} \prod_k s_{k, j_k}$ , where  $x_{\mathbf{s}_{-i}} = \sum_{\mathbf{s}_i \in S_i} x_{(\mathbf{s}_i, \mathbf{s}_{-i})}$  is the marginal probability of  $\mathbf{s}_{-i}$  in distribution  $\mathbf{x}$ . We observe that  $u_i^{\mathbf{s}_i}(\mathbf{x})$  is linear in  $\mathbf{s}_i$ . Specifically,  $u_i^{\mathbf{s}_i}(\mathbf{x}) = \sum_{j_i} s_{i, j_i} \sum_{\mathbf{s}_{-i}} \sum_{j_{-i}} U_{j_1, \dots, j_n}^i x_{\mathbf{s}_{-i}} \prod_{k \neq i} s_{k, j_k}$ . We can extend the definition of  $u_i^{\mathbf{s}_i}(\mathbf{x})$  beyond  $\mathbf{s}_i \in S_i$  to any vector in the convex hull  $P_i$ ; specifically for  $\mathbf{p}_i \in P_i$ ,  $u_i^{\mathbf{p}_i}(\mathbf{x})$  is defined to be  $\sum_{j_i} p_{i, j_i} \sum_{\mathbf{s}_{-i}} \sum_{j_{-i}} U_{j_1, \dots, j_n}^i x_{\mathbf{s}_{-i}} \prod_{k \neq i} s_{k, j_k}$ . Recall from (1) that  $g_i(\mathbf{x}) = \max_{\mathbf{s}_i \in S_i} u_i^{\mathbf{s}_i}(\mathbf{x})$ , i.e. player  $i$ 's expected utility if he deviates to a best response against  $\mathbf{x}$ . Since  $u_i^{\mathbf{s}_i}(\mathbf{x})$  is linear in  $\mathbf{s}_i$ , we can write  $g_i(\mathbf{x}) = \max_{\mathbf{p}_i \in P_i} u_i^{\mathbf{p}_i}(\mathbf{x})$ . Recall that a distribution  $\mathbf{x}$  is a Coarse Correlated Equilibrium (CCE) if it satisfies the *incentive constraints*:  $u_i(\mathbf{x}) \geq g_i(\mathbf{x}), \forall i$ .

Consider the following optimization problem:

$$\max \sum_i z_i \quad (3)$$

$$\mathbf{x} \geq 0, \sum_{\mathbf{s}} x_{\mathbf{s}} = 1, \quad (4)$$

$$u_i(\mathbf{x}) - g_i(\mathbf{x}) - z_i \geq 0, \forall i \quad (5)$$

$$z_i \leq 0, \forall i \quad (6)$$

The feasible region correspond to a relaxation of CCE, due to the introduction of slack variables  $z$ . A feasible solution  $(x, z)$  with  $z = 0$  is an optimal solution of the above problem (since  $z \leq 0$ ); furthermore such a solution corresponds to a CCE  $\mathbf{x}$  by construction.

This optimization problem is convex, but is difficult to handle directly because it has exponential number of variables  $x_{\mathbf{s}}$  for each  $\mathbf{s} \in S$ . Take the dual optimization problem:

$$\min_{\mathbf{y} \geq 0} \max_{\mathbf{x} \in \Delta, z \leq 0} \sum_i z_i + \sum_i y_i (u_i(\mathbf{x}) - g_i(\mathbf{x}) - z_i) \quad (7)$$

$$= \min_{\mathbf{y} \geq 0} \max_{\mathbf{x} \in \Delta, z \leq 0} \sum_i (1 - y_i) z_i + \sum_i \min_{\mathbf{p}_i \in P_i} y_i (u_i(\mathbf{x}) - u_i^{\mathbf{p}_i}(\mathbf{x})) \quad (8)$$

$$= \min_{0 \leq \mathbf{y} \leq 1} \max_{\mathbf{x} \in \Delta} \min_{\mathbf{p}_1 \in P_1, \dots, \mathbf{p}_n \in P_n} \sum_i y_i (u_i(\mathbf{x}) - u_i^{\mathbf{p}_i}(\mathbf{x})) \quad (9)$$

$$= \min_{0 \leq \mathbf{y} \leq 1} \min_{\mathbf{p}_1 \in P_1, \dots, \mathbf{p}_n \in P_n} \max_{\mathbf{x} \in \Delta} \sum_i y_i (u_i(\mathbf{x}) - u_i^{\mathbf{p}_i}(\mathbf{x})) \quad (10)$$

where  $\Delta = \{\mathbf{x} \in \mathbb{R}^{|S|} : \mathbf{x} \geq 0, \mathbf{1}^T \mathbf{x} = 1\}$ . Going from (8) to (9), we used the fact that if  $y_i > 1$ , the maximizer can take  $z_i$  towards  $-\infty$  and get arbitrarily high objective value. Therefore the outer minimizer should keep  $y_i \leq 1$ , in which case it is optimal for the maximizer to set  $z = 0$  and the term  $(1 - y_i)z_i$  disappears. In the last line we used the Minimax Theorem to switch the min and max operators. Since  $\sum_i y_i (u_i(\mathbf{x}) - u_i^{\mathbf{p}_i}(\mathbf{x}))$  is linear in  $\mathbf{x}$ , it attains its maximum at one of the extreme points of  $\Delta$ , i.e., one of the pure strategy profiles. Thus the dual problem is equivalent to

$$\min_{\mathbf{y}, \mathbf{p}_1, \dots, \mathbf{p}_n, t} t \quad (11)$$

$$0 \leq \mathbf{y} \leq 1; \quad \mathbf{p}_i \in P_i \quad \forall i \quad (12)$$

$$t \geq \sum_i y_i (u_i(\mathbf{s}') - u_i(\mathbf{p}_i, \mathbf{s}'_{-i})), \quad \forall \mathbf{s}' \in S \quad (13)$$

This is a nonlinear optimization problem due to the multiplication of  $y_i$  and  $\mathbf{p}_i$  in (13), but can be transformed to a linear optimization problem via the following variable substitution: let  $\mathbf{w}_i = y_i \mathbf{p}_i$ . We now try to express the dual problem in terms of  $y_i$  and  $\mathbf{w}_i$ . Recall that  $P_i = \{p \in \mathbb{R}^{m_i} | D_i p \leq \mathbf{f}_i\} \subset \mathbb{R}_+^{m_i}$ . Then  $\mathbf{w}_i$  satisfies  $D_i \mathbf{w}_i \leq y_i \mathbf{f}_i$ . For positive  $y_i$ , given  $\mathbf{w}_i$  we can recover  $\mathbf{p}_i = \mathbf{w}_i / y_i$ . When

$y_i = 0$ , we need to make sure that  $\mathbf{w}_i$  is also 0. This can be achieved using the constraints  $\mathbf{w}_i \geq 0$  and  $w_{ij} \leq M_{ij}y_i$ , where the constant  $M_{ij} = \max_{\mathbf{p}_i \in P_i} p_{ij}$  for all  $j \in [m_i]$ . Note that this is a valid bound on  $w_{ij}$  when  $y_i > 0$ .  $M_{ij}$  can be computed in polynomial time by calling `PolytopeSolve`, and hence is polynomial-sized. The dual problem is then equivalent to

$$\min_{\mathbf{y}, \mathbf{w}_1 \dots \mathbf{w}_n, t} t \quad (14)$$

$$0 \leq \mathbf{y} \leq 1; \quad D_i \mathbf{w}_i \leq y_i \mathbf{f}_i \quad \forall i \quad (15)$$

$$\mathbf{w}_i \geq 0, w_{ij} \leq M_{ij}y_i \quad \forall i, \forall j \in [m_i] \quad (16)$$

$$t \geq \sum_i y_i u_i(\mathbf{s}') - u_i(\mathbf{w}_i, \mathbf{s}'_{-i}), \quad \forall \mathbf{s}' \in S \quad (17)$$

where  $u_i(\mathbf{w}_i, \mathbf{s}'_{-i})$  is the linear extension of  $u_i(\mathbf{s}_i, \mathbf{s}'_{-i})$ ; i.e.  $u_i(\mathbf{w}_i, \mathbf{s}'_{-i}) = \sum_{j_i} w_{i,j_i} \sum_{j_{-i}} U_{j_1 \dots j_n}^i \prod_{k \neq i} s'_{k,j_k}$ . This is a linear program, with polynomial number of variables and exponential number of constraints. Since we know the primal objective is less or equal to 0, by LP duality, the optimal  $t$  in the dual is less or equal to 0. The following lemma establishes the existence of CCE in a way that does not use the existence of NE.

**Lemma 3.** *The dual LP (and therefore the primal LP) has optimal objective 0.*

This lemma says that for every candidate solution with  $t < 0$ , we can produce a hyperplane that separates it from the feasible set of the dual LP. We can use this lemma as a separation oracle in an algorithm similar to Papadimitriou & Roughgarden's [21] Ellipsoid Against Hope method to compute a CCE. However it would encounter similar numerical precision issues as discussed in [14], essentially due to the use of a convex combination of constraints which has a higher bit complexity than the individual constraints.

On the other hand, if we use a *pure separation oracle* that given  $\mathbf{y}, \mathbf{w}_1 \dots \mathbf{w}_n$ , finds  $\mathbf{s}'$  such that  $\sum_i y_i u_i(\mathbf{s}') - u_i(\mathbf{w}_i, \mathbf{s}'_{-i}) \geq 0$ , we can use the approach as described in [14] to compute a CCE.

**Lemma 4.** *Consider a multilinear game, with polynomial-time subroutines for `UtilGradient` and `PolytopeSolve`. Then there is a polynomial-time algorithm for the following pure separation oracle problem: given  $\mathbf{y}, \mathbf{w}_1 \dots \mathbf{w}_n$ , find pure strategy profile  $\mathbf{s}' \in S$  such that  $\sum_i y_i u_i(\mathbf{s}') - u_i(\mathbf{w}_i, \mathbf{s}'_{-i}) \geq 0$ .*

Using Lemmas 3 and 4, in [5] we extend the approach of approach of [14, 21] and complete the proof of Theorem 2.

## 5 Complexity of Approximate NE: Membership in PPAD

In this section we analyze complexity of computing Nash equilibrium in multilinear games. Existence of a NE in multilinear game follows from [10] makes the problem total. On the other hand, since multilinear games contain normal-form

multi-player games as a subcase, the Nash equilibria may be irrational [20]. In such a case the standard approach is to try approximation.

**$\epsilon$ -approximate NE ( $\epsilon$ -NE).** Given a rational  $\epsilon > 0$  in binary, a mixed strategy profile  $\sigma$  is an  $\epsilon$ -approximate NE iff  $\forall i \in N, u_i(\sigma) \geq \max_{\sigma'_i \in \Sigma_i} u_i(\sigma'_i, \sigma_{-i}) - \epsilon$ . In case of multilinear games, due to Lemma 1, this is iff corresponding marginal strategy profile  $\pi$  satisfy  $u_i(\pi) \geq \max_{\pi'_i \in P_i} u_i(\pi'_i, \pi_{-i}) - \epsilon$ .

It is well known that even in two player normal form games, computing approximate NE is PPAD-complete [6, 8, 22]. Roughly speaking, PPAD captures the class of total search problems that can be reduced to **End-Of-Line** [22], which includes computing approximate fixed-points. Since normal form games are contained in multilinear games, the next corollary follows:

**Corollary 1.** *Given a rational  $\epsilon > 0$  in binary, computing  $\epsilon$ -approximate NE in multilinear games is PPAD-hard.*

Due to exponential size of the strategy spaces, it seems that computing an  $\epsilon$ -NE in multilinear games could be a much harder problem given its generality. However, as we will show, it is no harder than computing a NE in 2-player games.

We note that, there has been recent efforts on showing PPAD membership for different classes of games [7]. However, the techniques are for games with polynomial type property, i.e. polynomial time computation of expected utility given mixed-strategy. Instead, we will use the characterization result (Proposition 2.2) of [9] to show that computing NE in multilinear games is in PPAD. See [5, 9] for relevant definitions, and the proposition statement.

Proposition 2.2 of [9] implies that to show membership of computing  $\epsilon$ -NE in PPAD, it is enough to capture them as approximate fixed-points of a polynomially continuous and polynomially computable function. Next we will construct such a function for multi-linear games.

Consider the following function  $\varphi : \Sigma \rightarrow \Sigma$  from [10] where  $\varphi = (\varphi_1, \dots, \varphi_n)$  and  $\varphi_i : \Sigma \rightarrow \Sigma_i$  such that,  $\varphi_i(\sigma_i, \sigma_{-i}) = \operatorname{argmax}_{\bar{\sigma}_i \in \Sigma_i} [u_i(\bar{\sigma}_i, \sigma_{-i}) - \|\bar{\sigma}_i - \sigma_i\|^2]$ . It was used to show existence of NE in concave games which includes multilinear games. However, notice that for multilinear games, description of mixed strategies is of exponential size, hence the function is not polynomially-computable. Its' polynomial-continuity is unclear. Instead, once again we will use marginal strategies. Moreover, we can compute the expected utilities using the marginal strategies efficiently as long as there is polynomial-time procedure to compute the utility gradient. Let  $P = \prod_{i \in N} P_i$ , we redefine  $\varphi : P \rightarrow P$  where  $\varphi = (\varphi_1, \dots, \varphi_n)$  and  $\varphi_i : P \rightarrow P_i$  is

$$\varphi_i(\pi_i, \pi_{-i}) = \operatorname{argmax}_{\bar{\pi}_i \in P_i} [u_i(\bar{\pi}_i, \pi_{-i}) - \|\bar{\pi}_i - \pi_i\|^2]. \quad (18)$$

Clearly,  $\varphi$  is a continuous function and therefore has a fixed-point. Next we show that its approximate fixed-points give approximate NE of the corresponding game. As the approximation goes to zero in the former we get exact NE in the latter, in other words exact fixed-points of (18) captures exact NE.

**Lemma 5.** *Given a rational  $\epsilon > 0$ , let  $\epsilon' = \frac{\epsilon}{|S|U_{max}H^n}$ , where  $H = \max_{i, \mathbf{p}_i \in P_i} \|\mathbf{p}_i\|_1$  and  $U_{max} = \max_{i, (j_1, \dots, j_n) \in \prod_k [m_k]} |U_{j_1, \dots, j_n}^i|$ . Then if  $\boldsymbol{\pi} \in P$  is an  $\epsilon'$ -approximate fixed-point of (18), i.e.,  $\|\varphi(\boldsymbol{\pi}) - \boldsymbol{\pi}\|_\infty < \epsilon'$  then it is a  $2\epsilon$ -approximate NE of the corresponding multilinear game.*

Lemma 5 implies that, for computation of approximate NE, it is enough to compute approximate fixed-point of function  $\varphi$ . Next we show that this function is polynomially continuous and polynomially computable and therefore computing its approximate fixed-point is contained in PPAD using Proposition 2.2 of [9], and therefore containment of approximate NE computation in PPAD follows. Next lemma shows polynomial-continuity and polynomial-computability by establishing equivalence of  $\varphi_i$  and a projection operator and by establishing connection to convex quadratic programming, respectively.

**Lemma 6.** *The function  $\varphi$  is polynomially continuous and computable.*

Due to the assumption that *PolytopeSolve* has polynomial-time sub-routine, the size of  $\max_{\mathbf{p}_i \in P_i} p_{ij}$ ,  $\forall i, \forall j \in [m_i]$  is polynomial in the description of the game. Furthermore,  $|S|$  is  $2^{\text{poly}(n \sum_i m_i)}$ . Therefore, if  $L$  is the size of the game description, then in Lemma 5 bit-length of  $H$  is polynomially bounded, and hence  $\text{size}(\epsilon') = O(\log(1/\epsilon), \text{poly}(\text{size}(L)))$ . Therefore, next theorem follows using Lemmas 5 and 6, together with Proposition 2.2 of [9], and Corollary 1.

**Theorem 3.** *Given a multilinear game with polynomial-time subroutines for PolytopeSolve and UtilGradient, and  $\epsilon > 0$  in binary, computing an  $\epsilon$ -approximate NE of the game is in PPAD. Furthermore, it is PPAD-complete.*

**Small Support Approximate NE.** Using discussion of Sect. 4.2, given an  $\epsilon$ -approximate NE  $\boldsymbol{\pi} \in P$ , each  $\boldsymbol{\pi}_i$  can be represented as distribution over  $m_i + 1$  pure strategies from  $S_i$ . However, existence of smaller support approximate NE is not clear. In [5], we study the same and obtain the following result.

**Theorem 4.** *Given a multilinear game, and given an  $\epsilon > 0$ , there exists an  $\epsilon$ -approximate NE with support size  $O(M^2 \frac{\log(n) + \log(m) - \log(\epsilon)}{\epsilon^2})$  for each player, where  $m = \max_i m_i$  and  $M = (\max_{i, \boldsymbol{\pi} \in P} \|\nabla_i u_i(\boldsymbol{\pi})\|_\infty) \max_{i, \boldsymbol{\pi}_i \in P_i} \|\boldsymbol{\pi}_i\|_1$ .*

Note that  $M$  upper bounds the magnitude of the game's utilities  $u_i(\mathbf{s}), \forall i, \forall \mathbf{s} \in S$ . Finally we provide discussion in [5].

## References

1. Ahmadinejad, A., Dehghani, S., Hajiaghayi, M., Lucier, B., Mahini, H., Seddighin, S.: From duels to battlefields: computing equilibria of blotto and other games. In: AAAI: Proceedings of the AAAI Conference on Artificial Intelligence (2016)
2. Babichenko, Y., Barman, S., Peretz, R.: Empirical distribution of equilibrium play and its testing application. CoRR abs/1310.7654 (2013). <http://arxiv.org/abs/1310.7654>

3. Budish, E., Che, Y.K., Kojima, F., Milgrom, P.: Designing random allocation mechanisms: theory and applications. *Am. Econ. Rev.* **103**(2), 585–623 (2013)
4. Chan, H., Jiang, A.X.: Congestion games with polytopal strategy spaces. In: Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI, pp. 165–171 (2016)
5. Chan, H., Jiang, A.X., Leyton-Brown, K., Mehta, R.: Multilinear games. Full Version (2016). <http://www.cs.trinity.edu/xjiang/papers/multilinear.pdf>
6. Chen, X., Deng, X., Teng, S.H.: Computing nash equilibria: approximation and smoothed complexity. In: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science, FOCS, pp. 603–612. IEEE Computer Society, Los Alamitos (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). doi:[10.1007/11786986\\_45](https://doi.org/10.1007/11786986_45)
8. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. In: Proceedings of the Annual ACM Symposium on Theory of Computing, STOC, pp. 71–78 (2006)
9. Etessami, K., Yannakakis, M.: On the complexity of nash equilibria and other fixed points (extended abstract). In: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science, FOCS, pp. 113–123 (2007)
10. Geanakoplos, J.: Nash and Walras equilibrium via Brouwer. *Econ. Theory* **21**(2), 585–603 (2003)
11. Grötschel, M., Lovász, L., Schrijver, A.: The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* **1**(2), 169–197 (1981)
12. Immorlica, N., Kalai, A.T., Lucier, B., Moitra, A., Postlewaite, A., Tennenholtz, M.: Dueling algorithms. In: Proceedings of the Annual ACM Symposium on Theory of Computing, STOC (2011)
13. Jain, M., Korzhyk, D., Vanek, O., Conitzer, V., Pechoucek, M., Tambe, M.: A double oracle algorithm for zero-sum security games on graphs. In: AAMAS (2011)
14. Jiang, A.X., Leyton-Brown, K.: Polynomial computation of exact correlated equilibrium in compact games. In: Proceedings of the ACM Conference on Electronic Commerce, EC (2011)
15. Jiang, A.X., Leyton-Brown, K., Bhat, N.: Action-graph games. *Games Econ. Behav.* **71**(1), 141–173 (2011)
16. Kearns, M., Littman, M., Singh, S.: Graphical models for game theory. In: Proceedings of the Conference on Uncertainty in Artificial Intelligence, UAI, pp. 253–260 (2001)
17. Koller, D., Milch, B.: Multi-agent influence diagrams for representing and solving games. In: Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI (2001)
18. Köppe, M., Ryan, C.T., Queyranne, M.: Rational generating functions and integer programming games. *Oper. Res.* **59**(6), 1445–1460 (2011)
19. Korzhyk, D., Conitzer, V., Parr, R.: Complexity of computing optimal stackelberg strategies in security resource allocation games. In: AAAI (2010)
20. Nash, J.: Non-cooperative games. *Ann. Math.* **54**(2), 289–295 (1951)
21. Papadimitriou, C., Roughgarden, T.: Computing correlated equilibria in multi-player games. *J. ACM* **55**(3), 14 (2008)
22. 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)



23. Rabinovich, Z., Gerding, E., Polukarov, M., Jennings, N.R.: Generalised fictitious play for a continuum of anonymous players. In: IJCAI: Proceedings of the International Joint Conference on Artificial Intelligence, pp. 245–250 (2009)
24. Tambe, M.: Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned. Cambridge University Press, Cambridge (2011)
25. Vorobeychik, Y.: Mechanism design and analysis using simulation-based game models. Ph.D. thesis, University of Michigan (2008)
26. Yin, Z., Jiang, A.X., Johnson, M.P., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Sullivan, J.: TRUSTS: Scheduling randomized patrols for fare inspection in transit systems. In: IAAI (2012)
27. Zinkevich, M.: Online convex programming and generalized infinitesimal gradient ascent. In: ICML (2003)

# Power-Law Distributions in a Two-Sided Market and Net Neutrality

Xiaotie Deng<sup>1</sup>, Zhe Feng<sup>2</sup>(✉), and Christos H. Papadimitriou<sup>3</sup>

<sup>1</sup> Department of Computer Science, Shanghai Jiao Tong University, Shanghai, China  
deng-xt@cs.sjtu.edu.cn

<sup>2</sup> School of Engineering and Applied Sciences, Harvard University, Cambridge, USA  
zhe\_feng@g.harvard.edu

<sup>3</sup> Department of Computer Science, University of California, Berkeley, CA, USA  
christos@cs.berkeley.edu

**Abstract.** “Net neutrality” often refers to the policy dictating that an Internet service provider (ISP) cannot charge content providers (CPs) for delivering their content to consumers. Many past quantitative models designed to determine whether net neutrality is a good idea have been rather equivocal in their conclusions. Here we propose a very simple two-sided market model, in which the types of the consumers and the CPs are *power-law distributed* — a kind of distribution known to often arise precisely in connection with Internet-related phenomena. We derive mostly analytical, closed-form results for several regimes: (a) Net neutrality, (b) social optimum, (c) maximum revenue by the ISP, or (d) maximum ISP revenue under quality differentiation. One unexpected conclusion is that (a) and (b) will differ significantly, unless average CP productivity is very high.

## 1 Introduction

The Internet is by far the world’s most crucial technological artifact. A mere quarter century after its beginning, it has emerged to become, through connecting over two billion people, the nexus of all human activity — intellectual, social, economic — and to satisfy, to varying and rapidly evolving degrees, humanity’s thirst for information and access, communication and interaction, education and wisdom, entertainment and excitement, opportunity and publicity, let alone justice, freedom, democracy. The Internet is also a gestalt, complex system; a novel, mysterious, and fascinating scientific object studied intensely by researchers of all colors, including computer scientists and economists.

From the point of view of economics (that is to say, efficiency and scarcity) one useful abstraction of the Internet is that of a *two-sided market* [1, 19]. In such a market, a platform (e.g., a game console, or an operating system, or an Internet service provider (ISP)) brings together two populations of agents: players with game developers, or users with application programmers, or, in the case of ISPs, Internet users with Internet content providers (CPs, such as Google, NYT, or Shtetl-optimized). Two-sided markets are interesting because they can exhibit network effects and other complex externalities. An important question is, if the

two populations are passive price-takers, what is the platform policy (typically, pricing for access from both sides<sup>1</sup>) that maximizes platform revenue, and what is the socially optimum policy? Plus, if these two differ substantially, how should the two-sided market be optimally regulated?

In the case of the Internet, this question has been known as the *net neutrality debate*, see [8, 21] and the related work section for the complex history and diverse and precarious current status. The term “net neutrality” has been used in many different senses. Most fundamentally, and closest to home, net neutrality is the computer science argument that the *end-to-end* principle in networking [18] implies that ISPs have no access to the content or origin of packets (as such information adds nothing to the network’s ability to operate properly). In policy, law, and economics, by “net neutrality” one typically understands two implied consequences of the end-to-end principle, namely that ISPs cannot/should not (a) treat flows differentially depending on the originating CP; or (b) charge CPs for resource use, or for content delivery to consumers.

There is a substantial and growing literature of economic research on net neutrality, and the two subtly different interpretations of the term “net neutrality” (a) and (b) above give rise to divergent threads within it (see the Related Work section). Typically the models include only one ISP (even though interesting analyses of multiple ISPs exist [14]) who charges (or does not) the two sides of the market for access, while the utility of the two populations is modeled in a number of different natural ways. Unsurprisingly, there is no definitive answer in the literature to the key question above (“which ISP policy is socially optimal?”), even though interesting points can be made based on such models (more in the related work section).

*The Model.* In this paper we introduce and analyze a new model of two-sided market motivated by the net neutrality problem. Our goals in defining this model have been these:

- *Keep the model very simple*, with very few and crisp parameters and assumptions, so that general conclusions can be drawn.
- At the same time, adopt assumptions (e.g. about distributions) compatible with the acknowledged reality of the Internet. Our model is the first to assume that the types of both users and CPs are *power-law distributed*.

Power law distributions [7, 13] (see also [9] for their use in economic modeling) are simple distributions outside the exponential family, with one parameter (the exponent) typically ranging between 2 and 3 — thus, they also serve our goal of parametric parsimony. Even though they had been observed in many places since the early 20th century (in city populations, word frequencies, incomes, etc.<sup>2</sup>), it was the Internet that brought them to the center of technical discourse — indeed

<sup>1</sup> Possibly negative prices: recall that in the first two examples the practice includes subsidies.

<sup>2</sup> Power law distributions have been called “the signature of human activity,” even though they also appear in life and the cosmos, and they are easy to confuse with the lognormal distribution.

it seems almost impossible to understand and model any aspect of the Internet and the web without resorting to these types of distributions. It seems natural to suppose that CP type (capturing the CP’s quality, or market share, or size) is so distributed, since power-law distributed firm size is a known characteristic of dynamic industries. It is also natural to accept that consumer types (measuring motivation, interest in the Internet) are power-law distributed — for example, incomes are distributed this way. Type distributions lie at the basis of our model. The product of consumer type  $x$  times CP type  $y$ , times the speed of the net, captures the *matching probability*, the probability that a consumer of type  $x$  will “like” (download content from) a CP of type  $y$ . This, together with a simple assumption on network speed (we take it inversely proportional to total traffic) defines the expected utility of both CPs and consumers: For a CP we assume it is proportional to the number of consumers who like it (modeling advertising income, or else popularity) and for a consumer a concave function, such as the square root, of the number of CPs that s/he likes. Finally, in the appendix we also briefly discuss a simple model of quality-of-service differentiation where an ISP charges CPs for using a privileged channel akin to the so-called “Paris metro pricing” [16].

Naturally, there are many aspects of this complex problem that we do not model: We do not model ISP costs, and, most importantly for the net neutrality debate, ISP technology and investment. However, our work can inform this crucial aspect of the problem, as our analytical results depend explicitly on the total network capacity. Our model of CP cost is simplistic (we assume that it is, in expectation, proportional to its type), but we have obtained similar results under different assumptions. We assume that there is only one ISP (as does most of the literature); however, our results can be used to solve simple models with many ISPs. And we do not model one of the salient characteristics of the Internet, namely its rapid growth; however, our use of power-law distributions in CP size can be seen as taking into account the exquisitely dynamic nature of the Internet market.

*Our Results.* We derive closed-form analytical results for almost all of the questions raised by our model: For the optimum ISP policy, for the optimum ISP policy under net neutrality, as well as for the ISP policy that maximizes social welfare, but also for comparisons between them; for a few points that are hard to answer analytically, we have very clean computational results.

Our most surprising conclusion is that, in this model, *net neutrality is not socially optimal* unless CP costs are very small. That is, there is in general a socially optimum price the ISP should charge the CPs, and this price is zero only if a parameter measuring CP costs (essentially, the average inverse productivity in the CP industry) is below a threshold. Regulation is needed for efficiency, requiring the ISP to charge CPs not necessarily zero, as in net neutrality, but the socially optimum price, typically smaller than what the ISP would like to charge.

The question then, for the regime of large CP costs, becomes: among the two suboptimal extremes (net neutrality or ISP revenue maximization), which is the more efficient? It turns out that the answer varies (see the computational results in Fig. 1): For CP costs just above the neutrality threshold, net neutrality

is better. For larger CP costs, ISP revenue maximization is better. Interestingly, in both cases the differences in social welfare between the three regimes does not seem that great. Overall, our model yields concrete, quantitative, and crisp results for the net neutrality problem, stemming from rather involved analysis, of the kind we believe had not been available in the literature, for a kind of model (consumers and CPs of power-law distributed types) that is arguably especially fit for the problem in hand.

Our results are summarized in Table 1. The parameters shown in this table will be mentioned in next section.

**Table 1.** Summary of results

|                        | CP costs ( $a$ )  | Optimal CP fee ( $b_{opt}$ )  | Optimal membership fee ( $c_{opt}$ )   |
|------------------------|---|---|--|
| Max-Rev<br>( $c = 0$ ) | $a > \frac{\lambda}{2} \frac{x_0^{2-\gamma}}{\gamma-2}$<br>$0 \leq a \leq \frac{\lambda}{2} \frac{x_0^{2-\gamma}}{\gamma-2}$  | $a$<br>$\frac{\lambda}{\gamma-2} x_0^{2-\gamma} - a$  | 0<br>0   |
| Max-Rev<br>( $c > 0$ ) | $a \leq \frac{1}{2} (\frac{\gamma-2}{\gamma-1} \phi'(\bar{Y}x_0) + \lambda) \bar{X}$<br>$a > \frac{1}{2} (\frac{\gamma-2}{\gamma-1} \phi'(\bar{Y}x_0) + \lambda) \bar{X}$           | $\frac{\lambda}{\gamma-2} x_0^{2-\gamma} - a$<br>$(\frac{2\lambda}{\frac{\gamma-1}{\gamma-2} \phi'(\sqrt{\bar{Y} \frac{1}{\beta-2} y^{*2-\beta} x_0} + \lambda)} - 1)a$                                     | $\phi(\frac{1}{\beta-2} y_0^{2-\beta} x_0)$<br>$\phi(\sqrt{\bar{Y} \frac{1}{\beta-2} y^{*2-\beta} x_0})$                                   |
| Socially<br>Optimum    | $a \leq \frac{1}{2} (\int_{x_0}^{\infty} \phi'(x\bar{Y})x^{1-\gamma} dx + \lambda\bar{X})$<br>$a \frac{1}{2} (\int_{x_0}^{\infty} \phi'(x\bar{Y})x^{1-\gamma} dx + \lambda\bar{X})$ | $\leq \frac{\lambda}{\gamma-2} x_0^{2-\gamma} - a$<br>$(\frac{2\lambda}{\frac{1}{\bar{X}} \int_{x_0}^{\infty} \phi'(\sqrt{\bar{Y} \frac{1}{\beta-2} \hat{y}^{2-\beta} x} x^{1-\gamma} dx + \lambda)} - 1)a$ | $\leq \phi(\sqrt{\bar{Y} \frac{1}{\beta-2} \hat{y}^{2-\beta} x_0})$<br>$\leq \phi(\sqrt{\bar{Y} \frac{1}{\beta-2} \hat{y}^{2-\beta} x_0})$ |

## Related Work

For aspects of policy, law, and history of the subject see [8, 17, 20, 21]. [12] is an eloquent advocacy of net neutrality backed by modest quantitative argument, while [2] is an exploration of possible business models in the CP industry and the ways they affect the net neutrality issue; the model involves only one CP. [5] propose a sophisticated and realistic model of CP-consumer interaction, but the complexity of their model prevents definite conclusions about net neutrality; an important monotonicity principle is shown, stating that social welfare is always coterminous with the total content transmitted through the network. In earlier work [6], a simple model in a similar spirit to ours was proposed, albeit with CP and consumer types uniformly distributed. Their results are dependent on parameter value ranges, with CP costs playing an important role, as they do in our results. In the model of [14] there are many regional monopolist ISPs, and deviation from net neutrality leads in a tragedy in the commons situation (the commons being the CP industry). The effect and nature of competition among ISPs is taken on in [10], through a mostly qualitative analysis.

Net neutrality as differentiation in quality of service has also been addressed in the economic literature. In [3, 4] consumers are connected to *two* CPs through a single ISP running a network with realistic (i.e., informed by queueing theory) delays, and two levels of service (a fast lane sold through bidding in [4], a priority service in [3]), and the two CPs choose level of service according to their

profitability. In contrast, [11] models CPs by their tolerance of network delays. Finally, [15] model the network as a sophisticated extensive-form game, in which CPs, ISPs, and consumers interact by setting prices and choosing services; they conclude that net neutrality prevails in several environments, for example in the presence of priority lanes.

## 2 The Model

In our basic model an ISP delivers the content of CPs to a population of consumers:

- The consumers are modeled as a continuum of values for the *consumer type*  $X$ , intuitively, a measure of the value this particular consumer receives from browsing the Internet. Importantly we assume that  $X$  is power-law distributed, that is, the density function is  $p_\gamma(x) = x^{-\gamma}$  for  $x \geq x_0$ , where  $x_0 = (\frac{1}{\gamma-1})^{\frac{1}{\gamma-1}}$  is the minimum type. We denote the expectation of  $X$  by  $\bar{X} = \frac{1}{\gamma-2}(\frac{1}{\gamma-1})^{\frac{2-\gamma}{\gamma-1}}$ .
- Similarly, each CP has a type  $Y$  with density function  $p_\beta(y) = y^{-\beta}$  for all  $y \geq y_0 = (\frac{1}{\beta-1})^{\frac{1}{\beta-1}}$ , a measure of the CPs quality, or size. Again,  $\bar{Y} = \frac{1}{\beta-2}(\frac{1}{\beta-1})^{\frac{2-\beta}{\beta-1}}$ .
- Bandwidth and speed: The ISP provides bandwidth  $B$  ( $B$  is taken to be one for simplicity, even though our results can be rewritten as functions also of  $B$ , for the study of issues of investment and technology innovation by the ISP). The speed of the network is then a decreasing function of the total traffic  $T$ , denoted  $Sp(T)$ , specified next.
- Calculation of  $T$ . Crucially, we assume that the infinitesimal contribution to traffic by consumers of type<sup>3</sup>  $x$  and CPs of type  $y$ , or equivalently, the intensity with which a consumer of type  $x$  will like and download the content of a CP of type  $y$ , is proportional to the product of the three magnitudes  $x$ ,  $y$ , and  $Sp(T)$  (times  $dx \cdot dy$ , of course). Therefore, the total traffic is

$$T = \int_{x_t}^\infty \int_{y_t}^\infty Sp(T) x y p_\gamma(x) p_\beta(y) dx dy$$

Here  $x_t$  and  $y_t$  are the key parameters sought by our analysis, namely the minimum types of consumers and CPs respectively that participate in the market (do not drop out), given the charges imposed by the ISP. The maximum traffic  $T_0$  occurs when  $x_t = x_0$  and  $y_t = y_0$ . We use the relative speed function  $Sp(T) = \frac{T_0}{T}$ . Thus,

$$T_0 = \int_{x_0}^\infty \int_{y_0}^\infty x y p_\gamma(x) p_\beta(y) dx dy = \int_{x_0}^\infty x^{1-\gamma} dx \int_{y_0}^\infty y^{1-\beta} dy = \bar{X} \bar{Y}$$

Since  $T$  only depends on  $x_t, y_t$ ,

$$T = \sqrt{\int_{x_t}^\infty \int_{y_t}^\infty T_0 x y p_\gamma(x) p_\beta(y) dx dy} = \sqrt{\bar{X} \bar{Y} \int_{x_t}^\infty x^{1-\gamma} dx \int_{y_t}^\infty y^{1-\beta} dy}$$

---

<sup>3</sup> More formally, of types between  $x$  and  $x + dx$ , etc.

– Utility functions.

- The utility of a user of type  $x$  is assumed to be  $\phi(N_x) - c$ , where  $N_x$  is the expected number of content providers this user likes,  $c$  is the membership fee imposed on users by the ISP (independent of the traffic), and  $\phi(r)$  is a concave function such as  $\sqrt{r}$ . Therefore, the utility function for a user of type  $x$  is

$$\phi\left(\int_{y_t}^{\infty} \frac{T_0}{T} x y^{1-\beta} dy\right) - c$$

- Finally, we assume that the utility function of a content provider of type  $y$  is  $\lambda N_y - by - ay$  where
  - \*  $\lambda$  is a needed “exchange rate” between the utility of consumers and that of CPs;
  - \*  $N_y$  is the expected number of consumers who like this content provider — notice that we assume advertising income to be proportional to the number of users;
  - \*  $ay$  is the expected costs of a content provider of type  $y$ ;
  - \*  $by$  is the payment that the content provider needs to pay to the platform. Notice here a simplifying modeling maneuver: While we would like to make the CP’s payment a linear function of the traffic originating from it, which is roughly  $N_y$ , we make it instead a linear function of its quality  $y$ , which is proportional to  $N_y$ .

Thus, the utility function of a content provider with quality  $y$  is as follows:

$$\lambda \int_{x_t}^{\infty} \frac{T_0}{T} x^{1-\gamma} y dx - by - ay$$

– Revenue of the ISP, from charges imposed on consumers and CPs:

$$\mathcal{R} = c \int_{x_t}^{\infty} x^{-\gamma} dx + b \int_{y_t}^{\infty} y \times y^{-\beta} dy = c \int_{x_t}^{\infty} x^{-\gamma} dx + b \int_{y_t}^{\infty} y^{1-\beta} dy,$$

– Thus, the parameters of our model are these: power-law exponents  $\gamma$  and  $\beta$ ; the consumer concave function  $\phi$ ; and the CP utility parameters  $a$  (expected cost per unit of size) and  $\lambda$ . The decision variables are  $b$  and  $c$  (the prices charged).

### 3 Revenue Maximization

We calculate the optimum prices for the ISP to charge the two sides of the market. For technical reasons we start by finding the optimum  $b$  (CP fee) when  $c = 0$  (this is Theorem 1), and then proceed to the general case (Theorem 2). The proofs are in the Appendix A.

**Theorem 1.** *If  $c = 0$ , the optimal pricing strategy is*

$$b_{opt} = \begin{cases} a & a > \frac{\lambda}{2} \frac{x_0^{2-\gamma}}{\gamma-2} \\ \frac{\lambda}{\gamma-2} x_0^{2-\gamma} - a & 0 \leq a \leq \frac{\lambda}{2} \frac{x_0^{2-\gamma}}{\gamma-2} \end{cases}$$

**Theorem 2.** *If  $c > 0, a \geq 0$  and  $\phi(\cdot)$  is a positive increasing concave function, the optimal pricing strategy is*

– If  $a \leq \frac{1}{2}(\frac{\gamma-2}{\gamma-1}\phi'(\bar{Y}x_0) + \lambda)\bar{X}$ ,

$$\begin{cases} b_{opt} = \frac{\lambda}{\gamma-2}x_0^{2-\gamma} - a \\ c_{opt} = \phi(\frac{1}{\beta-2}y_0^{2-\beta}x_0) \end{cases}$$

– If  $a > \frac{1}{2}(\frac{\gamma-2}{\gamma-1}\phi'(\bar{Y}x_0) + \lambda)\bar{X}$ ,

$$\begin{cases} b_{opt} = (\frac{2\lambda}{\frac{\gamma-1}{\gamma-2}\phi'(\sqrt{\bar{Y}\frac{1}{\beta-2}y^{*2-\beta}x_0}) + \lambda} - 1)a \\ c_{opt} = \phi(\sqrt{\bar{Y}\frac{1}{\beta-2}y^{*2-\beta}x_0}) \end{cases}$$

where  $y^*$  is the solution of  $y_t$  which satisfies the following equation:

$$a - \frac{1}{2}\left(\frac{\gamma-2}{\gamma-1}\phi'\left(\frac{\sqrt{T_0}\sqrt{\int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_0}^{\infty} x^{1-\gamma} dx}}x_0\right) + \lambda\right)\frac{\sqrt{T_0}\sqrt{\int_{x_0}^{\infty} x^{1-\gamma} dx}}{\sqrt{\int_{y_t}^{\infty} y^{1-\beta} dy}} = 0$$

## 4 Socially Optimum Pricing

While excluding consumers is obviously inefficient, rather surprisingly including all CPs may not be socially optimal. The intuitive reason is that low quality CPs clutter the Internet and incur large costs without adding enough value. Again we must determine the optimal  $x_t$  and  $y_t$ , and the corresponding  $b$  and  $c$ . Let  $\mathcal{S}$  denote the social welfare. We have:

$$\mathcal{S} = \int_{x_t}^{\infty} v(x)x^{-\gamma} dx + \int_{y_t}^{\infty} v(y)y^{-\beta} dy - a \int_{y_t}^{\infty} y^{1-\beta} dy \quad (1)$$

where

$$v(x) = \phi\left(\int_{y_t}^{\infty} \frac{T_0}{T} y^{1-\beta} x dy\right) = \phi\left(\frac{\sqrt{T_0} \int_{y_t}^{\infty} y^{1-\beta} dy}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}}x\right) \quad (2)$$

and

$$v(y) = \lambda \int_{x_t}^{\infty} \frac{T_0}{T} x^{1-\gamma} y dx = \lambda \frac{\sqrt{T_0} \int_{x_t}^{\infty} x^{1-\gamma} dx}{\sqrt{\int_{y_t}^{\infty} y^{1-\beta} dy}}y \quad (3)$$

Plugging these two equation above into Eq. 1, we get

$$\mathcal{S} = \int_{x_t}^{\infty} \phi\left(\frac{\sqrt{T_0} \int_{y_t}^{\infty} y^{1-\beta} dy}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}}x\right)x^{-\gamma} dx + \lambda \sqrt{T_0} \sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx} \sqrt{\int_{y_t}^{\infty} y^{1-\beta} dy} - a \int_{y_t}^{\infty} y^{1-\beta} dy$$

We can prove the following.



**Theorem 3.** To maximize the social welfare, the optimal  $x_t = x_0$ , while the optimal  $y_t$  satisfies the following

$$y_t = \begin{cases} y_0 & \text{if } a \leq \frac{1}{2}(\int_{x_0}^{\infty} \phi'(x\bar{Y})x^{1-\gamma}dx + \lambda\bar{X}) \\ \hat{y} & \text{if } a > \frac{1}{2}(\int_{x_0}^{\infty} \phi'(x\bar{Y})x^{1-\gamma}dx + \lambda\bar{X}) \end{cases}$$

where  $\hat{y}$  is the solution of the following equation of  $y_t$ :

$$a - \frac{1}{2}(\int_{x_0}^{\infty} \phi'(x\sqrt{\bar{Y} \int_{y_t}^{\infty} y^{1-\beta}dy})x^{1-\gamma}dx + \lambda\bar{X}) \frac{\sqrt{\bar{Y}}}{\sqrt{\int_{y_t}^{\infty} y^{1-\beta}dy}} = 0$$

In terms of the pricing strategy,

- If  $a \leq \frac{1}{2}(\int_{x_0}^{\infty} \phi'(x\bar{Y})x^{1-\gamma}dx + \lambda\bar{X})$ ,

$$\begin{cases} \hat{b}_{opt} \leq \frac{\lambda}{\gamma-2}x_0^{2-\gamma} - a \\ \hat{c}_{opt} \leq \phi(\frac{1}{\beta-2}y_0^{2-\beta}x_0) \end{cases}$$

- If  $a > \frac{1}{2}(\int_{x_0}^{\infty} \phi'(x\bar{Y})x^{1-\gamma}dx + \lambda\bar{X})$

$$\begin{cases} \hat{b}_{opt} = (\frac{2\lambda}{\frac{1}{\bar{X}} \int_{x_0}^{\infty} \phi'(\sqrt{\bar{Y} \frac{1}{\beta-2} \hat{y}^{2-\beta} x} x^{1-\gamma} dx + \lambda)} - 1)a \\ \hat{c}_{opt} \leq \phi(\sqrt{\bar{Y} \frac{1}{\beta-2} \hat{y}^{2-\beta} x_0}) \end{cases}$$

*Proof.* Firstly, we consider the optimal  $x_t$  to maximize social welfare of the platform. As we know,  $\mathcal{S}$  is a function of  $x_t$  and  $y_t$ , which is denoted by  $S(x_t, y_t)$ .

$$\begin{aligned} \frac{\partial S(x_t, y_t)}{\partial x_t} &= \frac{1}{2} \frac{x_t^{1-\gamma}}{\int_{x_t}^{\infty} x^{1-\gamma} dx} \int_{x_t}^{\infty} \phi'(\frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x) \frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x^{1-\gamma} dx \\ &\quad - \phi(\frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\int_{x_t}^{\infty} x^{1-\gamma} dx} x_t) x_t^{-\gamma} - \frac{1}{2} \frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x_t^{1-\gamma} \\ &\leq \frac{1}{2} \frac{x_t^{-\gamma}}{\int_{x_t}^{\infty} x^{1-\gamma} dx} \int_{x_t}^{\infty} \phi'(\frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x_t) \frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x_t x^{1-\gamma} dx \\ &\quad - \phi(\frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\int_{x_t}^{\infty} x^{1-\gamma} dx} x_t) x_t^{-\gamma} - \frac{1}{2} \frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x_t^{1-\gamma} \\ &\leq \frac{1}{2} \frac{x_t^{-\gamma}}{\int_{x_t}^{\infty} x^{1-\gamma} dx} \phi(\frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\int_{x_t}^{\infty} x^{1-\gamma} dx} x_t) \int_{x_t}^{\infty} x^{1-\gamma} dx - \phi(\frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\int_{x_t}^{\infty} x^{1-\gamma} dx} x_t) x_t^{-\gamma} \\ &\quad - \frac{1}{2} \frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x_t^{1-\gamma} = -\frac{1}{2} \phi(\frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\int_{x_t}^{\infty} x^{1-\gamma} dx} x_t) x_t^{-\gamma} - \frac{1}{2} \frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x_t^{1-\gamma} \\ &\leq 0 \end{aligned} \tag{4}$$

The first inequality in above proof is based on the fact that  $\phi'$  is a decreasing function of  $x_t$ . The second inequality is because  $\forall x \geq 0, x\phi'(x) \leq \phi(x)$ . Therefore, the optimal  $x_t$  is  $x_0$ .

Next, we consider  $y_t$ .

$$\begin{aligned}
 \frac{\partial S(x_t, y_t)}{\partial y_t} &= ay_t^{1-\beta} - \frac{1}{2} \int_{x_t}^{\infty} \phi' \left( \frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x \right) \frac{\sqrt{T_0} y_t^{1-\beta}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx \int_{y_t}^{\infty} y^{1-\beta} dy}} x^{1-\gamma} dx \\
 &\quad - \frac{1}{2} \lambda \frac{\sqrt{T_0 \int_{x_t}^{\infty} x^{1-\gamma} dx}}{\sqrt{\int_{y_t}^{\infty} y^{1-\beta} dy}} y_t^{1-\beta} \\
 &= y_t^{1-\beta} \left( a - \frac{1}{2} \left( \int_{x_t}^{\infty} \phi' \left( \frac{\sqrt{T_0 \int_{y_t}^{\infty} y^{1-\beta} dy}}{\sqrt{\int_{x_t}^{\infty} x^{1-\gamma} dx}} x \right) x^{1-\gamma} dx + \lambda \int_{x_t}^{\infty} x^{1-\gamma} dx \right) \frac{\sqrt{\bar{Y}}}{\sqrt{\int_{y_t}^{\infty} y^{1-\beta} dy}} \right) \\
 &= y_t^{1-\beta} \left( a - \frac{1}{2} \left( \int_{x_0}^{\infty} \phi' \left( x \sqrt{\bar{Y} \int_{y_t}^{\infty} y^{1-\beta} dy}} \right) x^{1-\gamma} dx + \lambda \bar{X} \right) \frac{\sqrt{\bar{Y}}}{\sqrt{\int_{y_t}^{\infty} y^{1-\beta} dy}} \right)
 \end{aligned} \tag{5}$$

Based on the same discussion in the proof of Theorem 2,

$$h(y_t) = \frac{1}{2} \left( \int_{x_0}^{\infty} \phi' \left( x \sqrt{\bar{Y} \int_{y_t}^{\infty} y^{1-\beta} dy} \right) x^{1-\gamma} dx + \lambda \bar{X} \right) \frac{\sqrt{\bar{Y}}}{\sqrt{\int_{y_t}^{\infty} y^{1-\beta} dy}}$$

is an increasing function of  $y_t$ . Thus,

- If  $a \leq h(y_0)$ , then  $\frac{\partial S(x_t, y_t)}{\partial y_t} \leq 0$ . Thus the optimal  $y_t$  is  $y_0$ .  
In this case, the optimal pricing strategy is

$$\begin{cases} \hat{b}_{opt} \leq \frac{\lambda}{\gamma-2} x_0^{2-\gamma} - a \\ \hat{c}_{opt} \leq \phi \left( \frac{1}{\beta-2} y_0^{2-\beta} x_0 \right) \end{cases}$$

- If  $a > h(y_0)$ , then there exists a unique solution  $\hat{y}$  for  $h(y_t) - a = 0$ . Then the socially optimal pricing is

$$\begin{cases} \hat{b}_{opt} = \left( \frac{2\lambda}{\frac{1}{\bar{X}} \int_{x_0}^{\infty} \phi' \left( \sqrt{\bar{Y} \frac{1}{\beta-2} \hat{y}^{2-\beta} x} \right) x^{1-\gamma} dx + \lambda} - 1 \right) a \\ \hat{c}_{opt} \leq \phi \left( \sqrt{\bar{Y} \frac{1}{\beta-2} \hat{y}^{2-\beta} x_0} \right) \end{cases}$$

#### 4.1 Comparison of $\hat{y}$ and $y^*$

We would like to know the relationship between the socially optimum cut off point for CPs  $\hat{y}$  and its revenue maximizing counterpart  $y^*$ . This relationship depends on  $\gamma, \beta$ , and  $\phi$ . When  $\phi$  belongs to a natural class of concave functions — namely, fractional powers — such comparison is possible: Revenue maximization demands that more CPs be cut off than does efficiency, assuming CP costs are not very low.

Let us define two important constants  $\zeta = \max\{\frac{1}{2}(\frac{\gamma-2}{\gamma-1}\phi'(\bar{Y}x_0) + \lambda)\bar{X}, \frac{1}{2}(\int_{x_0}^{\infty} \phi'(x\bar{Y})x^{1-\gamma}dx + \lambda\bar{X})\}$  and  $\eta = \min\{\frac{1}{2}(\frac{\gamma-2}{\gamma-1}\phi'(\bar{Y}x_0) + \lambda)\bar{X}, \frac{1}{2}(\int_{x_0}^{\infty} \phi'(x\bar{Y})x^{1-\gamma}dx + \lambda\bar{X})\}$

**Theorem 4.** Suppose  $a > \zeta$  and  $\phi(x) = x^\theta$  where  $0 < \theta < 1$ . Then  $\hat{y} < y^*$ .

*Proof.* For  $y^*$ ,

$$\begin{aligned}
 a &= \frac{1}{2}\left(\frac{\gamma-2}{\gamma-1}\phi'\left(\frac{\sqrt{T_0}\sqrt{\int_{y^*}^{\infty} y^{1-\beta}dy}}{\sqrt{\int_{x_0}^{\infty} x^{1-\gamma}dx}}x_0\right) + \lambda\right)\frac{\sqrt{T_0}\sqrt{\int_{x_0}^{\infty} x^{1-\gamma}dx}}{\sqrt{\int_{y^*}^{\infty} y^{1-\beta}dy}} \\
 &= \frac{1}{2}\left(\frac{\gamma-2}{\gamma-1}\phi'\left(\sqrt{\bar{Y}\int_{y^*}^{\infty} y^{1-\beta}dy}x_0\right) + \lambda\right)\frac{\bar{X}\sqrt{\bar{Y}}}{\sqrt{\int_{y^*}^{\infty} y^{1-\beta}dy}} = g(y^*)
 \end{aligned} \tag{6}$$

For  $\hat{y}$ ,

$$a = \frac{1}{2}\left(\int_{x_0}^{\infty} \phi'(x\sqrt{\bar{Y}\int_{\hat{y}}^{\infty} y^{1-\beta}dy})x^{1-\gamma}dx + \lambda\bar{X}\right)\frac{\sqrt{\bar{Y}}}{\sqrt{\int_{\hat{y}}^{\infty} y^{1-\beta}dy}} = h(\hat{y}) \tag{7}$$

Suppose  $\hat{y} = y^* = y'$ , then

$$\begin{aligned}
 &h(y') - g(y') \\
 &= \frac{1}{2}\frac{\sqrt{\bar{Y}}}{\sqrt{\int_{y'}^{\infty} y^{1-\beta}dy}}y'\left(\int_{x_0}^{\infty} \phi'(x\sqrt{\bar{Y}\int_{y'}^{\infty} y^{1-\beta}dy})x^{1-\gamma}dx - \frac{\gamma-2}{\gamma-1}\phi'(x_0\sqrt{\bar{Y}\int_{y'}^{\infty} y^{1-\beta}dy})\bar{X}\right)
 \end{aligned} \tag{8}$$

Let  $\sqrt{\bar{Y}\int_{y'}^{\infty} y^{1-\beta}dy} = Z$ , then

$$\begin{aligned}
 &\int_{x_0}^{\infty} \phi'(Zx)x^{1-\gamma}dx - \frac{\gamma-2}{\gamma-1}\phi'(Zx_0)\bar{X} \\
 &= \int_{x_0}^{\infty} \theta Z^{\theta-1}x^{\theta-\gamma}dx - \frac{\gamma-2}{\gamma-1}\theta(Zx_0)^{\theta-1}\frac{1}{\gamma-2}x_0^{2-\gamma} \\
 &= \theta Z^{\theta-1}x_0^{1+\theta-\gamma}\left(\frac{1}{\gamma-\theta-1} - \frac{1}{\gamma-1}\right) > 0
 \end{aligned} \tag{9}$$

Thus, if  $\hat{y} = y^*$ ,  $h(\hat{y}) - g(y^*) > 0$ . Since  $g$  and  $h$  are both increasing functions, then  $\hat{y} < y^*$  if  $h(\hat{y}) = g(y^*) = a$ .

### 4.2 Welfare Comparison

To summarize our results so far:

- In both revenue and welfare maximization, no consumers are left outside the market.

- When CP costs are small ( $a \leq \eta$ ), then no CPs are cut off either.
- When  $\eta < a \leq \zeta$ , then no CPs are cut off for social optimality, however, some CPs will be cut off for revenue optimality.<sup>4</sup>
- But otherwise, some CPs must be cut off for efficiency (that is, net neutrality is socially suboptimal), while more will have to be cut off for revenue optimality.<sup>5</sup>

But the question now arises, how does the social welfare of net neutrality compare with that of revenue optimality? Simulations show that the answer depends on CP costs, that is to say,  $a$ . In the simulation  $\gamma = \beta = 2.5$  and  $\lambda = 0.1$  where  $\phi(x) = x^{1/2}$ . Figure 1 shows the social welfare curve for three different values of  $a$ :  $1.1\zeta, 1.5\zeta, 2\zeta$ . When  $a = 1.1\zeta$  (that is, close to the neutrality region) net neutrality has better social welfare than revenue, while when  $a = 1.5\zeta, 2\zeta$  the social welfare in revenue optimum case is quite a bit larger than the social welfare in net neutrality. In fact, we can show that there is a single transition in this regard (proof in the Appendix B):

**Theorem 5.** *If  $\phi(x) = x^\theta (0 < \theta < 1)$ , there exists a unique  $\bar{a}$  such that when  $a < \bar{a}$  net neutrality has better welfare than revenue maximization, while the opposite happens when  $a > \bar{a}$ .*

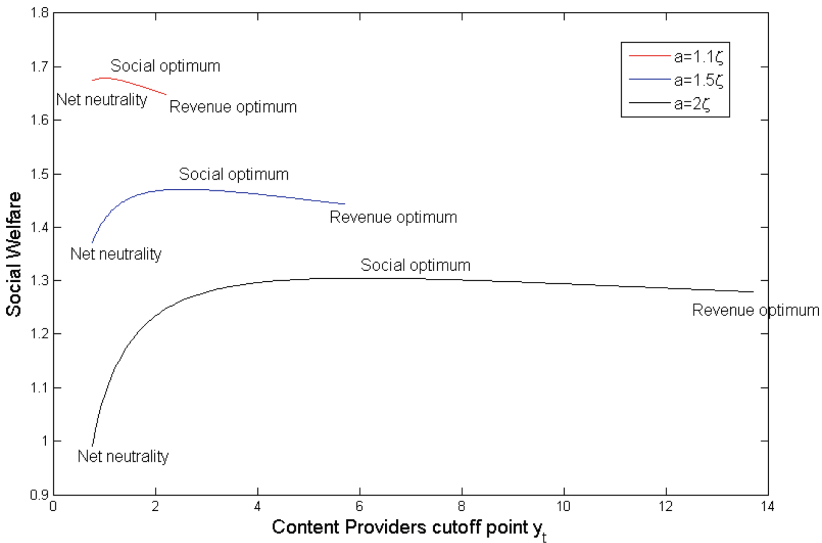


Fig. 1. Social Welfare Curve

<sup>4</sup> This is because  $\frac{1}{2}(\frac{\gamma-2}{\gamma-1}\phi'(\bar{Y}x_0) + \lambda)\bar{X} < \frac{1}{2}(\int_{x_0}^\infty \phi'(x\bar{Y})x^{1-\gamma}dx + \lambda\bar{X})$  when  $\phi$  is a fractional power function.

<sup>5</sup> Perhaps what is most striking in this figure (especially to somebody trained in approximation algorithms and the price of anarchy) is that, in all three cases and for these parameters and model, neither of the two extreme regimes (revenue maximization and net neutrality) is catastrophically suboptimal in social welfare.

## 5 Conclusion and Further Work

We have presented a parsimonious model of the Internet as a two-sided market with power-law distributed types from the two sides, with a simple cost structure for CPs, and utilities for the two sides based on simple and natural assumptions.

- Net neutrality is socially optimum only when CP productivity is very high. For lower levels of CP productivity (larger  $a$ ), net neutrality is better than ISP revenue maximization, but net neutrality is worse than ISP revenue maximization for even lower values. The preeminence of CP productivity as the determining factor of the optimum regulatory regime is one interesting insight from our model.

There are many possible extensions that seem very interesting, and some of them appear to be within reach.

1. Here we have adopted the “no payment” interpretation of net neutrality. What about the “non-differentiation” point of view? We have interesting preliminary results of this sort (see the Appendix C). Assume that part of the bandwidth is set aside for paying CPs. The point is that the payment counter-incentive will increase speed in this “channel” (this is the *Paris metro pricing* idea [16]). How large part of the total bandwidth should be so allocated, and how should it be priced? In the appendix we answer these questions, analytically and in more detail computationally, for the case  $a = 0$  — that is, zero costs for CPs. The general  $a$  case seems harder, but it would be interesting to crack it.
2. How could we make our model more realistic, without sacrificing much of its simplicity? We have tried other forms of CP costs and charges (for example, constant instead of linear in  $y$ ) without seeing qualitatively different results. But how about changing the utility model? One alternative model would weigh CP revenue by the type of the users it attracts. Another would use more elaborate and realistic speed functions, for example from queueing theory.
3. We have not considered *subsidies* of CPs by the ISP (negative  $b$ ; note that subsidies are common in two-way markets). Would they ever improve social welfare, or even ISP revenue?
4. A common argument against net neutrality is that it does not incentivize ISPs to invest in network technology. What can our model tell us about this? In our calculations we have used, for simplicity that the bandwidth  $B$  is one. We suspect that re-introducing  $B$  into our formulas might reveal interesting insights about incentives of the ISP to invest.
5. We have assumed a monopolist ISP; how would ISP competition affect the market? We suspect that many ISPs competing for consumers under revenue maximization would result in  $c = 0$ , and would charge CPs in near-identical ways, because each of them will be “selling” to the CPs a different lot (in expectation of the same size) of the same product: the consumers who (randomly) chose this ISP. Hence we suspect that our results summarized in Theorem 1 come close to obtaining yet another interesting comparison point, telling us how the ISP’s monopoly is affecting the efficiency of the market.

**Acknowledgement.** Research results reported in this work are partially supported by the National Natural Science Foundation of China (Grant No. 61632017) and a Project 985 grant of Shanghai Jiao Tong University. This work was done while the authors were visiting the Simons Institute for the Theory of Computing, Berkeley.

## Appendix

The Appendix is available through our ArXiv version.

## References

1. Armstrong, M.: Competition in two-sided markets. *Rand J. Econ.* **37**, 668–691 (2006)
2. Altman, E., Legout, A., Xu, Y.: Network non-neutrality debate: an economic analysis (2011). [arXiv:1012.5862v2](https://arxiv.org/abs/1012.5862v2)
3. Cheng, H., Bandyopadhyay, S., Guo, H.: The debate on net neutrality: a policy perspective. *Inf. Syst. Res.* **22**, 60–82 (2011)
4. Choi, J., Kim, B.: Net neutrality and investment incentives. *Rand J. Econ.* **41**, 446–471 (2010)
5. Economides, N., Hermalin, B.: The economics of network neutrality. *Rand J. Econ.* **43**, 602–629 (2012)
6. Economides, N., Tag, J.: Network neutrality in the Internet: a two-sided market analysis. *Inf. Econ. Policy* **19**, 215–248 (2007)
7. Faloutsos, M., Faloutsos, P., Faloutsos, C.: Proceedings of SOGMOD 1999, pp. 251–262 (1999)
8. Frieden, R.: What’s new in the network neutrality debate. *Mich. St. Law Rev.* (2015)
9. Gabaix, X.: Power laws in economics and finance. *Annu. Rev. Econ.* **1**, 255–293 (2009)
10. Kocsis, V., de Bijl, P.W.: Network neutrality and the nature of competition between network operators. *Int. Econ. Econ. Policy* **4**, 159–184 (2010)
11. Krämer, J., Lukas, J.: Network neutrality and congestion-sensitive content providers: implication for service innovation, broadband investment and regulation. In: Net Institute Working Paper, 10 September 2010
12. Lee, R., Wu, T.: Subsidiating creativity through network design: zero-pricing and net neutrality. *J. Econ. Perspect.* **23**, 61–76 (2009)
13. Mitzenmacher, M.: A brief history of generative models for power law and lognormal distributions. *Internet Math.* **1**(2), 226–251 (2004)
14. Musacchio, J., Schwartz, G., Walrand, J.: A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue. *Rev. Netw. Econ.* **8**(1), 22–39 (2009)
15. Njoroge, P., Ozdaglar, A., Stier-Moses, N., Weintraub, G.: Investment in two-sided markets and the net neutrality debate. *Rev. Netw. Econ.* **12**, 355–402 (2014)
16. Odlyzko, A.: Internet pricing and the history of communications. *Comput. Netw.* **36**(5/6), 493–517 (2001)
17. Peha, J.M., Lehr, W.H., Wilkie, S.: The state of the debate on network neutrality International. *J. Commun.* **1**, 709–716 (2007)

18. Saltzer, J., Reed, D., Clark, D.: End-to-end arguments in system design. *ACM Trans. Comput. Syst. (TOCS)* **2**(4), 277–288 (1984)
19. Rochet, J.C., Tirole, J.: Platform competition in two-sided markets. *J. Eur. Econ. Assoc.* **1**, 990–1029 (2003)
20. Tate, D.T.: Net neutrality 10 years later: a still unconvinced commissioner. *Fed. Commun. J.* **66**, 509–524 (2015)
21. [https://en.wikipedia.org/wiki/Net\\_neutrality](https://en.wikipedia.org/wiki/Net_neutrality)

# On-Demand or Spot? Selling the Cloud to Risk-Averse Customers

Darrell Hoy<sup>1</sup>(✉), Nicole Immorlica<sup>2</sup>, and Brendan Lucier<sup>2</sup>

<sup>1</sup> University of Maryland, College Park, USA  
darrell.hoy@gmail.com

<sup>2</sup> Microsoft Research, New York, USA

**Abstract.** In Amazon EC2, cloud resources are sold through a combination of an on-demand market, in which customers buy resources at a fixed price, and a spot market, in which customers bid for an uncertain supply of excess resources. Standard market environments suggest that an optimal design uses just one type of market. We show the prevalence of a dual market system can be explained by heterogeneous risk attitudes of customers. In our stylized model, we consider unit demand risk-averse bidders. We show the model admits a unique equilibrium, with higher revenue and higher welfare than using only spot markets. Furthermore, as risk aversion increases, the usage of the on-demand market increases. We conclude that risk attitudes are an important factor in cloud resource allocation and should be incorporated into models of cloud markets.

## 1 Introduction

Cloud computing allows clients to rent computing resources over the internet to perform a variety of computing tasks, from hosting simple web servers to computing complex financial models. By offloading these tasks to the cloud, clients avoid the necessity of procuring and maintaining expensive servers and infrastructure. The current market leader in this industry is Amazon who launched its cloud platform, Amazon Elastic Compute Unit (EC2), in 2006. Amazon uses its cloud internally for many of its own computations. Additionally, Amazon contracts with large clients who reserve instances of cloud resources for long usage periods. Due to natural variation in the nature of computing tasks from Amazon and its large clients, EC2 has a varying amount of leftover computing resources. Amazon sells these resources to small clients.

This leads to a natural question: how should a cloud provider price its resources to these small clients? The pricing model adopted by Amazon has two main components: an on-demand market and a spot market. In the on-demand market, clients may buy an instance of cloud resources at a fixed reservation price.<sup>1</sup> After resources have been allocated internally, to large clients, and to

---

Part of this work was completed while D. Hoy was an intern at Microsoft Research.

<sup>1</sup> This might more naturally be called a “reservation market” and we switch to this terminology in the remainder of the paper; however we stick to the term “on-demand” for the current discussion as this is the term used by Amazon.



clients in the on-demand market, extra supply might still remain. This supply is sold in the spot market. In a spot market, clients place bids for instances, and a price is set so that the available supply equals the total demand at that price.

Viewed through the lens of microeconomic theory, the persistence of this dual market is a curiosity at first glance. Indeed, in standard economic environments, a risk-neutral, expected-utility-maximizing client who desires a cloud resource should simply buy it in whichever market is expected to have the lower price – typically the spot market. This suggests all sales should happen in the spot market, leaving the on-demand market defunct.

That this is not reflective of reality stems from several factors. Most apparent is that clients are rarely risk-neutral. For example, it is easy to imagine that a company would have a soft budget set aside for computational costs. They would then spend freely within the confines of this budget, and extend the budget cautiously when necessary to meet their computing needs. This type of behavior suggests a tendency towards risk aversion on the part of the clients. As the budget is freely available, clients might prefer to “overspend” to guarantee the required resources at the on-demand price.

We show in a stylized setting that the presence of heterogenous risk attitudes can explain the prevalence of a combined on-demand and spot market. Specifically, this dual market induces a unique equilibrium in which more risk-averse customers (e.g., those with higher budgets) buy resources in the on-demand market and the others bid in the spot market. We show that this equilibrium outcome outperforms the outcomes achieved by running only one of the two types of markets on its own in many key objectives.

## 1.1 Results and Techniques

In order to highlight the impact of risk aversion on the market, we focus our analysis on a simple setting in which there is only one type of computational resource being sold (e.g., a server with one core for one hour), and each buyer demands only a single instance of this cloud resource at any given time<sup>2</sup>. Formally, we assume a continuum of buyers, where the type of a buyer consists of a value for an instance and a utility curve that maps outcomes (i.e., allocation and price paid) to payoffs. The utility curve describes a buyer’s attitude toward risk: for example, a buyer with a soft budget (as described above) would likely prefer to spend all of their budget all the time than to spend twice their budget half the time, and this preference is captured by a non-linear utility curve. The buyer types are described by a joint common prior. We assume the market is large; i.e., no single buyer has significant impact on the market outcome.

The market works as follows. First, the seller sets a price for on-demand instances. This price should be high enough to guarantee that supply exceeds

<sup>2</sup> Of course, this model abstracts away from many reasonable sources of risk aversion in the cloud, such as clients with diminishing marginal returns for multiple instances, the cost of prematurely terminating a long-running task. Even ignoring these factors, our model still generates heterogeneous preferences toward on-demand versus spot pricing.

demand, motivated by the fact that resources are always available for purchase in on-demand markets in practice. Buyers then realize their types (i.e., value/budget pairs) and choose whether to buy in the on-demand market. After these decisions have been made, the unsold supply receives an exogenous shock, modeling variation in the demand of large clients. Any remaining supply is then sold to the remaining buyers at a market-clearing price.<sup>3</sup>

We prove that this system has a unique (subgame-perfect) equilibrium for each choice of the on-demand market price. We do this by analyzing the relationship between the spot price distribution and the distributions of clients' types and corresponding supply and demand. It turns out that the distribution over spot prices up to a certain value  $v$  depends only on the distributions of clients' types in the range  $[0, v]$ , and hence one can explicitly solve for the price distribution recursively.

This equilibrium satisfies a monotonicity property: agents that are more risk-averse are more likely to purchase in the on-demand market, whereas agents that are less risk-averse are more likely to use the spot market. Furthermore, as the distribution shifts such that agents become more averse to risk (in the sense of first-order stochastic dominance), we show more clients end up buying instances in the on-demand market and the revenue correspondingly increases. This result is perhaps intuitive, but it is not obvious: as clients become more averse to risk, they shift towards the on-demand market and hence both decrease supply and demand in the spot market. This in turn could cause the spot price to shift either up or down, which would impact purchasing decisions of all clients. By further arguing about the equilibrium of the market, we show the shift towards the on-demand market in fact causes the spot price to increase thereby reinforcing the shift towards the on-demand market. Therefore, the equilibrium is monotone with such shifts in the value and budget distribution. This further illustrates the connection between the on-demand market and risk attitudes.

We leverage our equilibrium characterization to compare the dual market outcome to the outcome of a spot-only or on-demand-only market. We are interested in the welfare, efficiency, and revenue properties of these markets. The revenue of a market outcome is simply the sum of the payments, and is equal to the cloud provider's utility. The welfare of an outcome is the total utility of all market participants including the cloud provider. The efficiency is the total value of the cloud clients, ignoring payments. In risk neutral environments, the welfare and efficiency are equal, but with risk-averse clients the welfare can be less than the efficiency.

It is easy to see that a spot-only market is more efficient than a dual market, which in turn is more efficient than an on-demand-only market. This is because the spot market precisely generates the efficient outcome, even with exogenous

---

<sup>3</sup> Since our model abstracts away from inter-temporal effects, we do not explicitly model the impact of fluctuating spot prices and changes to on-demand prices over time. Investigating a repeated-game model of the market, and/or agents with time-dependent preferences (e.g., minimizing the cost of a large job subject to a deadline), is left as a direction for future work.

supply uncertainty: in a spot market, allocation is monotone in value and thus a lower-valued client is never served in place of the higher-valued one.

Surprisingly, these efficiency comparisons do not extend to welfare. As we show, the welfare of the dual market is better than the welfare of the spot market alone, regardless of the price set for the on-demand instances. In particular, this is true even when the on-demand market price is set to maximize the revenue of the cloud provider. This is not trivial: the on-demand market adds inefficiency, since clients with high value but low aversion to risk may not wish to purchase on-demand, whereas clients with lower value but higher risk-aversion might. This leads to circumstances where lower-valued clients win but higher-valued clients lose. However, since the clients that are winning in this scenario are actually more risk-averse, the transfer of payments to the cloud provider increases welfare. We show that the welfare increase due to additional transfers from risk-averse clients offset any inefficiencies in the allocation. Moreover, since this welfare comparison holds at every setting of the on-demand price, it applies in particular to the price that maximizes revenue. We show this price must also generate more revenue than a spot-only market, leading to increases in welfare and revenue.

In summary, a dual spot/on-demand market simultaneously improves both the revenue and welfare of a spot-only market. We also show by example that while an on-demand market is revenue-optimal for risk-neutral buyers, a dual market can generate strictly higher revenue when buyers are risk-averse. Furthermore, while a dual market may sometimes generate less revenue than an on-demand-only market, a dual market is always more efficient. This suggests that a cloud provider, especially one that holds a dominant position in the marketplace, might prefer a dual market system. This phenomenon is driven by heterogeneous risk attitudes that arise naturally in the context of cloud computation, leading us to posit that risk aversion is an important element to consider when one models the cloud marketplace.

## 1.2 Related Work

A number of papers explore cloud-computing market design. Zhang et al. [17] consider designing a truthful auction where uncertainty lies in the arrival of demand and value profiles of bidders, whether they have a large job with deadline or general demand over time. An et al. [2] design a negotiation-based mechanism for setting price contracts in the presence of demand uncertainty. Borgs et al. [5] consider the pricing problem faced by a seller setting on-demand prices over multiple time periods and uncertain supply, with agents who arrive and have different deadlines for their tasks. The paradigm of a dual spot+reserve mechanism has also received a lot of attention. Wang et al. [16] uses a Markov decision process to model the designer's choice of how to partition supply between the reserve and spot markets. Abhishek et al. [1] models the cloud market as a queuing model, in which a continuum of jobs arrive and have (private) waiting costs. They find that a fixed cost model provides greater expected revenue than a spot market. Additionally, recent works [8, 13] have focused on the problem faced by bidders in such a market: when to use the spot market and when to reserve.

Ben-Yehuda et al. [4] analyzes the expected spot prices in comparison to their reservation prices, and find that it is very likely that Amazon is intentionally manipulating the price or supply distribution so as to provide users with more uncertainty in the spot market. In all of the models described above, agents are risk-neutral and do not have budgets. As far as we are aware, our work is the first to use risk aversion to explain the prevalence of a spot+reserve market.

Auctions for cloud computation resources share similarities to electricity markets, where the split between a spot market and a so-called “futures” market is common. Indeed, the use of both markets has been advocated to account for risk-averse buyers and sellers (see e.g., [3]). One difference is that, unlike the on-demand market for cloud computation, futures markets for electricity are typically resolved years in advance.

While most work in auction theory assumes risk-neutral agents, some work has been done for auctions with risk-averse bidders. Optimal auctions have been characterized for simple settings [9,11], but the solutions are generally not expressible in closed form. It is therefore more common to study the simple auctions used in practice, with the general finding that second-price or spot-like auctions do poorly for revenue when compared with first-price auctions [6,7,15]. Matthews [12] has looked at the preferences of bidders in the auctions, and showed that first-price auctions not only can get more revenue than the second-price auction, but also can be preferred by bidders due to reduction in uncertainty around the payment. Our model of risk-aversion as the presence of a soft budget is closely related to the capacitated utilities model of Fu et al. [6], where the capacity in their model corresponds to value minus budget in our model. They show that with capacitated agents, a simple first-price auction with reserve has revenue that approximates the revenue of the optimal mechanism.

Our results are of a similar flavor to the eBay-style buy-it-now auction considered by Mathews and Katzman [10]. As in our model, they find that adding a buy-it-now option increases revenue, and as agents become more risk-averse, the optimal price increases. Their model differs from ours in that they consider an explicitly randomized allocation rule, rather than clearing the market at a spot price, in order to incentivize use of the buy-it-now (i.e., reservation) option.

## 2 Preliminaries

In our model, a single cloud provider (the seller) is selling compute resources to a continuum of clients (the bidders).

*Utility Structure.* Each bidder  $i$  has value  $v_i \in [0, 1]$  for a single compute instance. As is standard in the economics literature, we model the risk attitude of bidder  $i$  through a utility function  $u_i : \mathcal{R}_{\geq 0} \rightarrow \mathcal{R}_{\geq 0}$ . If the bidder obtains an instance and pays  $p$ , then her utility is  $u_i(v_i - p)$ . We will assume that  $u_i(0) = 0$ , that  $u_i$  is continuous and non-decreasing, and that  $u_i$  is not identically 0. Note that since  $u_i(0) = 0$ , we can think of  $v_i$  as the maximum price at which bidder  $i$  is willing to purchase an instance. A bidder that does not obtain an instance will pay nothing and have utility 0.

We focus our attention on agents that are *risk averse*. That is, we will assume that utility curves are weakly concave, as is standard when modeling risk aversion. We allow  $u_i$  to be linear, in which case bidder  $i$  is said to be risk-neutral.

Roughly speaking, an agent with a utility curve that is “more concave” will be more risk averse, in the sense that they are more likely to prefer guaranteed outcomes to uncertain lotteries. More formally, we say that utility function  $u$  is more risk averse than  $u^*$ , and write  $u \preceq u^*$ , if for every distribution  $L$  over non-negative real values and every fixed value  $d \geq 0$ , if  $\mathbf{E}_{x \sim L}[u(x)] \geq u(d)$ , then  $\mathbf{E}_{x \sim L}[u^*(x)] \geq u^*(d)$ . In other words, if an agent with utility curve  $u$  prefers a lottery  $L$  over a guaranteed payout of  $d$ , then an agent with utility curve  $u^*$  would prefer the lottery as well. This defines a partial order over utility curves. Note that, under this definition, all (weakly) concave utility curves are (weakly) more risk-averse than a risk-neutral (i.e., linear) curve. Note also that for twice-differentiable utility curves,  $u$  is more risk-averse than  $u^*$  if and only if the standard Arrow-Pratt measure of risk-aversion is nowhere lower for curve  $u$  than for curve  $u^*$  [14].<sup>4</sup>

*Demand Structure.* Types are distributed according to a joint distribution  $F$  on pairs  $(v, u)$ . For ease of exposition, we will assume throughout that  $F$  is supported on a finite collection of  $(v, u)$  pairs. Write  $V$  and  $U$  for the (finite) sets of values and utility curves that support  $F$ , and for  $(v, u) \in V \times U$  we will write  $f(v, u)$  for the probability that an agent has type  $(v, u)$ .

We will use  $F(v)$  to refer to the induced distribution over values; that is,  $F(v)$  is the probability that an agent’s value is at most  $v$ . We will assume a large-market condition, which is that the aggregate demand is distributed exactly according to the type distribution  $F$ .

*Supply Structure.* The supply of instances,  $q$ , is unknown to the bidders and seller until the instances are to be allocated. The supply is then drawn from a distribution,  $Q$ . We will normalize the supply so that  $q$  represents the fraction of the market that can be simultaneously served, hence  $q \in [0, 1]$ .

## 2.1 Auction Formats

We will consider three auction formats in this paper: spot auctions, reservation auctions (previously referred to as on-demand), and dual (or spot+reservation) auctions.

**Spot ( $M^s$ ).** One type of auction to run is a *market-clearing* auction, or a spot auction. In this auction, buyers submit bids. A market-clearing price  $p_s$  is chosen such that the quantity of bids exceeding  $p_s$  is equal to the supply. Under our unit-demand and large market assumptions, it is a dominant strategy for a bidder to bid her value; henceforth we assume the bids in the spot market equal the values. We observe that a market-clearing price

---

<sup>4</sup> We define risk aversion with respect to agent preferences directly, rather than via the Arrow-Pratt measure, to avoid requiring utility curves be twice differentiable.

always exists in our market, even in the presence of non-linear utilities: with available supply  $q$ , and distribution over values  $F$ , the market clearing price, written  $p_s(q)$ , is precisely<sup>5</sup> the value for which  $q = 1 - F(p_s(q))$ .

**Reservation ( $M^r$ ).** In a reservation-only (or “on-demand”) market, the auctioneer sets a fixed price  $p_r$  per instance, in advance of seeing the realization of supply. Price  $p_r$  need not be a market-clearing price. If there is not enough supply to satisfy the demand for instances at this price, the winning bidders are chosen uniformly at random from among those who wish to purchase.

**Spot+Reservation ( $M^{s+r}$ ).** In a spot and reservation market, the auctioneer first sets a fixed price  $p_r$  and runs a reservation auction. The remaining inventory of supply (if any) is then sold via a spot auction. The exact timeline of events in the spot and reservation auction  $M^{s+r}$  is as follows:

1. Auctioneer announces reservation price  $p_r$ .
2. Bidders realize types  $(v_i, u_i) \sim F(v, u)$ .
3. Each bidder decides whether to purchase an instance in the reservation auction, indicated by  $a^r \in \{0, 1\}$ . Let  $T = \sum_{v,u} a^r(v, u) f(v, u)$  be the total volume of reserved instances.
4. Auctioneer realizes supply  $q \sim Q$ , and reserved instances are allocated as in the reservation market described above.
5. If  $q > T$ , the auctioneer runs a market-clearing auction to clear the excess capacity. Let  $p_s(q)$  be the resulting market-clearing spot price.

Note that our specification does not ask bidders to decide whether or not to participate in the spot market. The fact that bidders are unit demand, and that the spot auction is truthful under our large market assumption, implies that in equilibrium bidders will bid (truthfully) in the spot auction if (and only if) they don’t buy an instance in the reservation auction.

For a given (implicit) strategy profile for mechanism  $M^{s+r}$ , we will write  $S(p_s)$  for the cumulative distribution function of the resulting spot prices.

*Solution Concept: Subgame-Perfect Equilibrium.* For each of these auctions, the solution concept we apply is *subgame-perfect equilibrium*. A strategy profile for a multi-stage game forms a subgame-perfect equilibrium (SPE) if, at every stage  $t$  of the game and every possible history of actions by players in previous stages, no agent can benefit by unilaterally deviating from her prescribed strategy from stage  $t$  onward.

For the spot auction and reservation auction, there is only one stage of the resulting game and hence equilibria are straightforward: each agent chooses to purchase her utility-maximizing quantity of instances given the specified price.

For mechanism  $M^{s+r}$ , we can characterize the SPE as follows. In the second (i.e., spot) stage of the mechanism, the equilibrium condition implies that agents always purchase instances if and only if their value is above the realized spot price. Thus, the only strategic choice to be made by agents is in the first (i.e., reservation) stage of the mechanism, where each agent must select whether to

---

<sup>5</sup> This price may not be unique if  $q = 0$  or  $q = 1$ . In these cases we define  $p_s(q)$  to be the supremum of prices satisfying the written condition, which will be  $\infty$  for  $q = 0$ .

purchase an instance in the reservation market. We will therefore define a strategy profile  $\mathbf{s}$  to be a mapping from a type  $(v, u)$  to an action  $\{0, 1\}$ , where  $\mathbf{s}(v, u)$  is interpreted as the number of instances to purchase in the reservation market. Note that the distribution over market-clearing prices in the second stage is completely determined by the actions of agents in the first stage, and hence is determined by  $\mathbf{s}$ . An equilibrium is then a strategy profile such that no agent can benefit by unilaterally deviating from strategy  $\mathbf{s}$  (i.e., by reserving more or fewer instances), given the distribution of spot prices implied by  $\mathbf{s}$ .

## 2.2 Objectives

We consider three objectives when evaluating mechanisms: revenue (*REV*), welfare (*WEL*) and efficiency (*EFF*). The revenue of a mechanism  $M$ ,  $REV(M)$ , is the sum of the payments made to the auctioneer. Note that for the spot+reservation mechanism,  $REV(M^{s+r}) = p_r T + \mathbf{E}_{q \sim Q}[p_s(q)(q - T)]$  where we used the fact that  $T \leq q$  with probability 1. The welfare of a mechanism  $WEL(M)$  is the sum of utilities of all agents, including the auctioneer (whose utility is precisely the revenue of the mechanism). The efficiency  $EFF(M)$  of a mechanism measures the value created, without considering the welfare lost due to the (non-linear) utility functions of agents. For the spot+reservation mechanism:  $EFF(M^{s+r}) = \mathbf{E}_{(v_i, u_i) \sim F}[v_i \cdot s(v_i, u_i) + v_i \cdot (1 - s(v_i, u_i)) \cdot S(v_i)]$

If an agent reserves, her value generated is  $v_i$ . If she does not reserve, her value generated is  $v_i \cdot S(v_i)$ , which is her value times the probability that the spot price is below her value. Note that if agents are risk-neutral (i.e., have the identity function as their utility functions), then  $EFF(M) = WEL(M)$ .

## 3 Equilibrium Behavior and Analysis

In this section, we analyze the choices of bidders and use this to characterize equilibrium of the spot+reservation market. We begin by noting the relationship between the spot price distribution and the distributions of supply, type, and reservation demand in equilibrium. Recall that  $Q$  denotes the CDF of the supply distribution.

**Lemma 1.** *Fix strategy profile  $\mathbf{s}$ , let  $S$  be the distribution of spot prices under  $\mathbf{s}$ , and let  $T(p)$  be the volume of reserved instances demanded from agents with value at most  $p$ , under  $\mathbf{s}$ . Then  $\mathbf{s}$  forms an equilibrium if and only if, for all  $p$ ,*

$$S(p) = 1 - Q(1 - F(p) + T(p)), \text{ and} \quad (1)$$

$$T(p) = \sum_{\substack{v \in V \\ v \leq p}} \sum_{u \in U} f(v, u) \cdot \mathbb{1}[u(v - p_r) \geq \mathbf{E}_{p \sim S}[u(\max\{v - p, 0\})]]. \quad (2)$$

*Proof.* The probability that the spot price is at most  $p$  is exactly the probability that the supply is greater than necessary to satisfy all of the demand for resources from bidders with higher marginal values than  $p$ , plus all reservation

demand for resources with lower marginal values than  $p$ . On the other hand, the volume of reserved resources demanded from agents with value at most  $p$ , at equilibrium, is precisely the probability that such an agent will prefer the deterministic reservation outcome to the lottery over outcomes determined by the distribution over spot market prices.

In light of Lemma 1, we will tend to equate equilibria with the resulting distributions  $S$  and  $T$ , rather than with an explicit strategy profile  $\mathbf{s}$ .

**Lemma 2.** *Purchasing in the reservation stage is monotone in the risk-aversion of  $u_i$ : if a  $(v_i, u_i)$  bidder (weakly) prefers to reserve an instance, then a  $(v_i, u_i^*)$  bidder with  $u_i^* \preceq u_i$  (weakly) prefers reserving.*

*Proof.* We begin by considering the special event in which the agent is not allocated an instance even if they reserve, due to the supply being insufficient to honor all reservations and the agent not being selected randomly as a winner. If this event occurs, the bidder’s utility will necessarily be 0, and this is independent of their utility curve and their chosen action (since, if  $q < T$ , no agent that enters the spot market will obtain an instance). It therefore suffices to consider the agent’s expected utility conditional on the event that the agent will be allocated an instance with certainty if they choose to reserve. With this in mind, the utilities of a unit demand agent from reserving or participating only in the spot market, respectively, are

$$\begin{aligned} u^r(v_i, u_i) &= u_i(v_i - p_r), \\ u^s(v_i, u_i) &= \mathbf{E}_{p_s \sim S}[u_i(v_i - p_s) \cdot \mathbf{1}_{p_s \geq v_i}]. \end{aligned}$$

We now want to show that  $u^r(v_i, u_i) \geq u^s(v_i, u_i)$  implies  $u^r(v_i, u_i^*) \geq u^s(v_i, u_i^*)$ .

Note that, fixing  $v_i$ , the spot market generates a certain lottery  $L$  over values  $(v_i - p_s)$ , and the reservation market generates a certain value  $v_i - p_r$ . Thus, from the definition of risk aversion, if an agent with utility curve  $u_i$  prefers the certain outcome to the lottery  $L$ , and  $u_i^* \succeq u_i$ , then an agent with utility curve  $u_i^*$  prefers the certain outcome as well.

### 3.1 Equilibrium Existence and Uniqueness

We are now ready to establish uniqueness of equilibrium. One subtlety about equilibrium uniqueness is the manner in which buyers break ties. If a positive mass of agents is indifferent between the spot and reservation markets, there may be multiple market outcomes consistent with those preferences. We will therefore fix some arbitrary manner in which bidders break ties, which could be randomized and heterogeneous across bidders. Our claim is that for any such tie-breaking rule, the resulting equilibrium will be unique.

**Lemma 3.** *There is a unique equilibrium of  $M^{s+r}$ . Moreover, this equilibrium is computable in time proportional to the size of the support of type distribution  $f$ .*



*Proof.* As shown in Lemma 1, the challenge of characterizing equilibrium essentially reduces to characterizing the fraction of bidders who reserve at a given price,  $T(p)$ . This is because  $T$  determines the distribution  $S$  over spot prices, and  $S$  (together with an arbitrary tie-breaking rule) uniquely determines the strategy profile  $s$ , since this can be inferred from the expected utility when choosing the spot market. Thus, to show uniqueness and existence of equilibrium, it suffices to show uniqueness of the functions  $S$  and  $T$ .

We will prove that, for all  $v \in V$ ,  $T(v)$  and  $S(v)$  are uniquely determined by the functions  $T$  and  $S$  restricted to values less than  $v$ . The result will then follow by induction on the elements of  $V$ .

Consider first an agent with value  $v = \min V$ . Recall that the spot price is always at least  $v$ . Thus, if  $p_r < v$  then the agent will always reserve, if  $p_r > v$  then the agent will always choose the spot market, and if  $p_r = v$  the agent will be indifferent and apply the fixed tie-breaking rule. In each case, the value of  $T(v)$  is uniquely determined, and thus  $S(v)$  is as well.

Now choose  $v > \min V$ , and suppose  $T$  and  $S$  are determined for all elements of  $V \cap [0, v)$ . We claim the distribution of the random variable  $\max\{v - p_s, 0\}$ , where  $p_s$  is distributed according to  $S$ , is then uniquely determined. This is because the non-zero values of this random variable are distributed according to  $S$  restricted to values in  $V \cap [0, v)$ . But, by Lemma 1, this random variable determines the value of  $T(v)$ , which in turn determines the value of  $S(v)$ . Thus  $T(v)$  and  $S(v)$  are uniquely determined by  $S$  and  $T$  on  $V \cap [0, v)$ , as required. Moreover, they can be explicitly computed by evaluating the summation in Lemma 1.


### 3.2 An Example: Soft Budgets

In this section we present a special case of risk-aversion, driven by soft budgets, and give an interpretation of our equilibrium characterization for this case.

Suppose that each buyer  $i$  is characterized by their value  $v_i$  for a compute instance and a soft budget  $b_i \in [0, v_i]$ . We think of  $b_i$  as a budget of funds that has been allocated to acquiring a compute instance. If the buyer obtains an instance but pays less than  $b_i$ , the residual budget is lost: it is as if they had paid  $b_i$ . On the other hand, if the buyer pays more than  $b_i$ , they suffer no additional penalty; they simply incur the cost of their payment.

This scenario is captured by a piecewise linear utility curve,  $u_i(z) = \min(z, v_i - b_i)$ . Note that if a buyer with budget  $b_i$  obtains an instance and pays  $p_i > b_i$  their utility is  $u_i(v_i - p_i) = v_i - p_i$ , whereas if they pay  $p_i < b_i$  their utility is  $u_i(v_i - p_i) = v_i - b_i$ . For a fixed value  $v_i$ , an agent with a higher budget is more risk-averse. To see this, note that if a buyer *strictly* prefers a lottery  $L$  to a deterministic outcome  $d$ , then it must be that  $d < v_i - b_i$  (since otherwise  $d$  must be utility-maximizing). In this case, decreasing the budget can only make the lottery more valuable, while not affecting the utility from the deterministic outcome. Thus, a decreased budget can only increase the propensity to select a lottery over a deterministic outcome.

The monotonicity result in Lemma 2 thus results in a partitioning of agents that prefer the spot market to the reservation market and vice versa by an



**Fig. 1.** With soft-budgets, a monotone-decreasing indifference curve partitions agents into those that reserve an instance and those who rely on the spot market.

indifference curve over budgets. See Fig. 1 for an illustration. Any agent with  $(v_i, b_i)$  below the curve prefers the spot market; any agent above the curve (such that  $v_i \geq b_i$ ), prefers the reservation market. Lemma 3 shows that this indifference curve is unique, given the distribution over agent types, and precisely specifies which agents choose to reserve and which enter the spot market.

## 4 Comparative Statics

In this section, we first consider the impact of changes to buyer risk attitudes. We show that as agents become more risk averse, more agents use the reservation market and revenue increases, for every setting of the reservation price. Second, we compare the reservation+spot mechanism to the spot market and the reservation market. We first show that the combination mechanism's outcomes are more efficient than running a reserve market on its own. We then show that it generates both more revenue and more welfare than running only a spot market. Formal proofs are deferred to the full version of the paper.

The results in this section hold under two assumptions on the reservation price set by the seller. First, we will make the assumption that  $p_r$  is set high enough so that, in the resulting equilibrium,  $\Pr_{q \sim Q}[q < T] = 0$ . That is, over the uncertainty in supply, the mechanism can serve the reserved instances with certainty. This assumption is motivated by the fact that these instances are typically viewed as guaranteed by the mechanism.

Another natural and practical property is that the reservation price  $p_r$  be set high enough that it will be greater than the expected spot price. That is,  $p_r$  is large enough that it is more costly, in expectation, to reserve a guaranteed instance than to bid for an instance in the spot market.

### 4.1 Effect of Increased Risk Aversion

We consider the impact of an increase in risk aversion. Consider type distribution  $F$  and a type distribution  $F^+$  induced by a pointwise transformation  $g^+(U, V) \rightarrow (U, V)$  applied to each point in  $F$  which weakly increases risk aversion and does not affect values. Specifically, for any  $(u^+, v^+) = g^+(u, v)$ ,  $v^+ = v$  and  $u^+ \preceq u$ . In the following lemma, we show that such a change can only increase the fraction of agents who choose to reserve, and can only increase revenue.

**Lemma 4.** *For mechanism  $M^{s+r}$ , and for any reserve price  $p_r$ , if risk aversion of agents increases then the fraction of agents who purchase in the reservation stage increases, as does the expected revenue of the mechanism.*

The intuition underlying Lemma 4 is as follows. The first order effect from a change in risk aversion is an increase in  $T$ , the fraction of users who choose to reserve at a given price. This increase in reservations translates into higher spot prices, since it reduces the quantity sold in the spot market. Higher spot prices in turn cause more users to prefer to reserve, which can only increase spot prices further. This can be shown by induction over agent values.

## 4.2 Comparing Mechanisms

We now compare welfare and revenue of  $M^{s+r}$  to  $M^s$  and  $M^r$ . Here we make use of the two assumptions discussed in Sect. 2.1: first, the reservation phase is not oversubscribed, i.e., the reservation price is set sufficiently high that there will be sufficient supply to fulfill the demand for reserved instances; and second, the reservation price is sufficiently high to be above the expected spot price.

We begin by comparing the revenue of the dual mechanism with the expected revenue of a spot-only market. Note that, trivially, the best revenue of the combined mechanism is at least the revenue of a spot market; this is because, in the combined mechanism, the reserve price can be set sufficiently high that all customers buy in the spot market. We show something stronger: for *every* choice of reservation price, the revenue of the combined mechanism is at least that of a spot market run in isolation.

**Lemma 5.** *For any choice of the reservation price satisfying our assumptions, the expected revenue of the reservation and spot mechanism is weakly greater than the revenue of the spot-only market.*

*Proof (sketch).* As in Lemma 4, as risk aversion increases, revenue increases for a fixed reservation price. Fix a reservation price and consider starting with a distribution of risk-neutral agents. These agents will all bid in the spot market and thus the outcome (and in particular, the revenue) will be identical to the spot-only mechanism. By deforming the utility curves of the agents in a manner that only increases risk aversion, until they match the correct distribution, and applying Lemma 4, we can conclude that the revenue of the dual market only increases while the revenue of the spot-only mechanism, which is unaffected by the utility curves, remains the same.

*Example 1.* This example illustrates the revenue of the dual mechanism can be strictly greater than both the spot and the reservation mechanisms. We consider an example in which agent utility curves are specified by soft budgets, as in Sect. 3.2. Recall that a budget of 0 corresponds to risk-neutrality. Take  $\epsilon > 0$  to be sufficiently small, and consider the following distribution over buyer types:

- with probability  $0.5 - \epsilon$ ,  $(v, b) = (5, 0)$
- with probability  $0.5$ ,  $(v, b) = (10, 10 - \epsilon)$
- with probability  $\epsilon$ ,  $(v, b) = (20, 0)$

The supply is distributed such that  $q = 1 - \epsilon$  with probability  $0.8$ , and otherwise  $q = 0.5 + \epsilon/2$ . In the spot-only auction, the spot price is  $5$  with supply  $q = 1 - \epsilon$  and  $10$  with supply  $q = 0.5 + \epsilon/2$ , giving revenue of  $5 - 3\epsilon$ . In the reservation market, the optimal reserve price is  $10$ , which generates a total revenue of  $5 + 9\epsilon$ .

Consider the dual mechanism with reservation price  $10 - \epsilon$ . At equilibrium, the buyers of type  $(10, 10 - \epsilon)$  strictly prefer to reserve (obtaining utility  $\epsilon$  with probability  $1$ , rather than utility  $\epsilon$  with probability  $0.8$ ), whereas the buyers of type  $(20, 0)$  strictly prefer to participate in the spot market (obtaining utility  $(20 - 5)$  with probability  $0.8$ , rather than utility  $10 + \epsilon$  with probability  $1$ ). The revenue is then  $7 - \frac{5}{2}\epsilon$ , greater than spot-only revenue  $5 - 3\epsilon$ , and reservation-only revenue  $5 + 9\epsilon$  for sufficiently small  $\epsilon$ .

**Lemma 6.** *For any choice of the reserve price satisfying the assumptions above, the expected efficiency of the reservation+spot mechanism is weakly greater than the efficiency of the reservation market with the same reservation price.*

*Proof.* Recall that the efficiency of a mechanism is the expected value generated by the agents, ignoring the welfare lost due to the nonlinear utility functions. For any realized supply  $q$  then, weakly more people are served in  $M^{s+r}$ , hence efficiency is greater.

**Theorem 1.** *In any equilibrium of the the spot and reservation mechanism where the reservation price is set above the expected spot price, the expected welfare of the reservation and spot mechanism is weakly greater than the expected welfare of the spot-only mechanism.*

*Proof (sketch).* Note that, relative to a spot market, introducing a reservation price adds inefficiency. This is because if a bidder is willing to reserve to get a guaranteed instance, any time they would not have one in the spot market, there is a higher valued bidder than them who could be allocated.

However, welfare is increased when a bidder chooses to reserve. Consider an agent with value  $v_i$  who is willing to reserve at price  $p_r$ . Reserving increases his utility and the auctioneer is receiving more revenue, because the reservation price is greater than the expected spot price (by assumption).

The full proof consists of three parts. First, we define a benchmark  $B^{s+r}$  that is just like  $M^{s+r}$  except agents who reserve pay the spot price instead of the reservation price. We then, we show that the welfare of  $M^{s+r}$  is greater than the welfare of  $B^{s+r}$ , which will follow largely because the expected reservation price is greater than the expected spot price. Finally, we show that spot prices increase when agents choose to reserve, which leads to the welfare of  $B^{s+r}$  being greater than the welfare of  $M^s$ , and hence the welfare of  $M^{s+r}$  is greater than  $M^s$ .

## References

1. Abhishek, V., Kash, I.A., Key, P.: Fixed and market pricing for cloud services. CoRR, abs/1201.5621(2012)
2. An, B., Lesser, V., Irwin, D., Zink, M.: Automated negotiation with decommitment for dynamic resource allocation in cloud computing. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, vol. 1, pp. 981–988. International Foundation for Autonomous Agents and Multiagent Systems (2011)
3. Ausubel, L.M., Cramton, P.: Using forward markets to improve electricity market design. *Util. Policy* **18**(4), 195–200 (2010)
4. Ben-Yehuda, O.A., Ben-Yehuda, M., Schuster, A., Tsafir, D.: Deconstructing amazon EC2 spot instance pricing. *ACM TEAC* **1**(3), 1–20 (2013)
5. Borgs, C., Candogan, O., Chayes, J., Lobel, I., Nazerzadeh, H.: Optimal multi-period pricing with service guarantees. *Manag. Sci.* **60**, 1792–1811 (2014)
6. Fu, H., Hartline, J., Hoy, D.: Prior-independent auctions for risk-averse agents. In: EC 2013(2013)
7. Hu, A., Matthews, S.A., Zou, L.: Risk aversion and optimal reserve prices in first- and second-price auctions. *J. Econ. Theory* **145**(3), 1188–1202 (2010)
8. Ma, D., Huang, J.: The pricing model of cloud computing services. In: EC 2012, pp. 263–269 (2012)
9. Maskin, E., Riley, J.: Optimal auctions with risk averse buyers. *Econometrica* **52**(6), 1473–1518 (1984)
10. Matthews, T., Katzman, B.: The role of varying risk attitudes in an auction with a buyout option. *Econ. Theory* **27**(3), 597–613 (2006)
11. Matthews, S.A.: Selling to risk averse buyers with unobservable tastes. *J. Econ. Theory* **30**(2), 370–400 (1983)
12. Matthews, S.A.: Comparing auctions for risk averse buyers: a buyer’s point of view. *Econometrica* **55**(3), 633–646 (1987)
13. Menache, I., Shamir, O., Jain, N.: Dynamic resource allocation for executing batch jobs in the cloud. In: ICAC 2014
14. Pratt, J.W.: Risk aversion in the small and in the large. *Econometrica* **32**(1), 122–136 (1964)
15. Riley, J., Samuelson, W.: Optimal auctions. *Am. Econ. Rev.* **71**(3), 381–392 (1981)
16. Wang, W., Li, B., Liang, B.: Towards optimal capacity segmentation with hybrid cloud pricing. In: ICDCS 2012, pp. 425–434. IEEE (2012)
17. Zhang, H., Li, B., Jiang, H., Liu, F., Vasilakos, A.V., Liu, J.: A framework for truthful online auctions in cloud computing with heterogeneous user demands. In: IEEE INFOCOM, pp. 1510–1518. IEEE (2013)

# Buying Data from Privacy-Aware Individuals: The Effect of Negative Payments

Weina Wang<sup>(✉)</sup>, Lei Ying, and Junshan Zhang

School of Electrical, Computer and Energy Engineering, Arizona State University,  
Tempe, AZ 85281, USA

{weina.wang, lei.ying.2, junshan.zhang}@asu.edu


**Abstract.** We study a market model where a data analyst uses monetary incentives to procure private data from strategic data subjects/individuals. To characterize individuals' privacy concerns, we consider a local model of differential privacy, where the individuals do not trust the analyst so privacy costs are incurred when the data is reported to the data analyst. We investigate a basic model where the private data are bits that represent the individuals' knowledge about an underlying state, and the analyst pays each individual according to all the reported data. The data analyst's goal is to design a payment mechanism such that at an equilibrium, she can learn the state with an accuracy goal met and the corresponding total expected payment minimized. What makes the mechanism design more challenging is that not only the data but also the privacy costs are unknown to the data analyst, where the costs reflect individuals' valuations of privacy and are modeled by "cost coefficients." To cope with the uncertainty in the cost coefficients and drive down the data analyst's cost, we utilize possible negative payments (which penalize individuals with "unacceptably" high valuations of privacy) and explore interim individual rationality. We design a family of payment mechanisms, each of which has a Bayesian Nash equilibrium where the individuals exhibit a threshold behavior: the individuals with cost coefficients above a threshold choose not to participate, and the individuals with cost coefficients below the threshold participate and report data with quality guarantee. By choosing appropriate parameters, we obtain a sequence of mechanisms, as the number of individuals grows large. Each mechanism in this sequence fulfills the accuracy goal at a Bayesian Nash equilibrium. The total expected payment at the equilibrium goes to zero; i.e., this sequence of mechanisms is asymptotically optimal.

## 1 Introduction

Exploiting human-related data such as medical records and financial data has created great value to the society. However, the ever-improving capability of data analysis in the advancing big data technology makes it possible to extract personal information undesirably, giving rise to technical barriers for data collection.

In short, big data analytics is a double edged sword. This in turn necessitates incentivizing data subjects/individuals for providing quality data while preserving privacy.

In this paper, we consider a market model where a data analyst uses monetary incentives to procure private data from strategic data subjects/individuals. Specifically, the data analyst elicits data from a population of  $N$  individuals. Each individual  $i$ 's private data is a binary *signal*  $S_i$  that reflects her knowledge about an underlying *state*, which is represented by a binary random variable  $W$ . Conditional on the state  $W$ , the signals are independently generated such that the probability for  $S_i$  to be the same as  $W$  is  $\theta$ , where  $0.5 < \theta < 1$ . The data analyst is interested in learning  $W$ . This structure is illustrated in Fig. 1.



**Fig. 1.** System model. The data analyst is interested in the state  $W$ , which is a binary random variable. Each individual  $i$  has a private binary signal  $S_i$  and a type  $C_i$  that characterizes her valuation of privacy.  $S_1, S_2, \dots, S_N$  are conditionally i.i.d. given  $W$ . Individual  $i$  reports randomized data  $X_i$ , which is generated based on  $S_i$  and  $C_i$ .

To provide monetary incentives, the data analyst announces a mechanism, which is a function that determines the amounts of payments to individuals according to their reported data. Since an individual's payment may depend on others' reports, this payment mechanism induces a game among the individuals. Due to privacy concerns, an individual experiences a cost when she releases data to the analyst. Her payoff is the difference between the payment and the privacy cost. The goal of the data analyst in this setting is to design a mechanism to achieve a desired learning accuracy at an equilibrium in a cost-effective manner.

**Privacy Cost Model.** To quantify the privacy costs, we consider a local model of differential privacy (an introduction of which can be found in [10]). In this local model, the individuals do not trust the data analyst with their data, so they have to evaluate their privacy costs carefully when reporting data to the analyst. To control the privacy cost, we assume in the paper that an individual adds random noise to her data and reports the resulting perturbed version. Intuitively, the more "noisy" the reported data is, the more privacy is retained, and thus the less privacy cost is incurred. An individual will weigh the privacy cost against the payment to choose the best way of perturbing her data. In contrast, in a centralized model of data privacy with a trustworthy data analyst (e.g., [15]), the action of providing data to the analyst itself, whether truthful or not, does not cause privacy loss. There privacy costs are incurred when the analyst releases the outcome of the mechanism, so the individuals cannot control their privacy costs except choosing to participate or not.

We remark that the different privacy cost models make the structures of the mechanism design problem fundamentally different. In a centralized model, the design goal is to have a mechanism that elicits truthful data reporting and its outcome satisfies the promised privacy guarantee. However, in the local model considered in this paper, truthfulness is no longer a focal design goal since it incurs high privacy costs to individuals that need to be compensated by payments. Instead, the data analyst seeks for mechanisms that cost-effectively elicit data with small enough perturbation, and consequently the analyst needs to manage equilibria consisting of noisy data reporting. Another major difference is that it is unnecessary to make the outcome of the mechanism guarantee privacy in the local model since the control of privacy remains in the hand of the individuals.

### **Unknown Privacy Valuations and the Impact of Negative Payments.**

We consider the natural setting where different individuals may have different valuations of privacy and their valuations are unknown to the data analyst. In this model, the analyst is not able to tune the payments to the privacy costs, which may result in overpayments when some individuals' costs are lower than expected. This uncertainty can also introduce unwanted noise in the reported data. To see this, consider a mechanism that always pays a nonnegative amount of payment to a participant. For an individual whose valuation of privacy is very high, participating and reporting only noise is a better strategy than opting out since she may still receive some nonnegative amount of payment without incurring any privacy cost. Then the payment does not buy the analyst any useful information from this individual, and moreover, the analyst has to work with these unusable reports during data analysis.

With these observations, we consider payment mechanisms that are interim individually rational; i.e., the expected payoff of each individual in an equilibrium is nonnegative, although the realizations of the payments can be negative. In practice, this can model the scenario where there are repeated data collection (e.g., to learn the ratings of different movies), and in some rounds the payments received by the individual may be negative, but in the long run, the average payoff will be nonnegative. Negative payments can be utilized to “filter out” individuals with high privacy costs; i.e., we design the mechanism such that their expected payoff is negative if they report only noise. This saves the data analyst's payments on poor quality data and simplifies the data analysis. We will see that we can actually drive the total cost to zero for the data analyst as the population size becomes large.

We remark that one possible approach to implement negative payments is to let the data analyst set up an online payment system using virtual currency or credits. Instead of paying real money to an individual every time she reports a data, virtual currency or credits can be added to or reduced from the user's account. An individual can be paid weekly or monthly with real dollars. Since the expected payment is nonnegative, the real-dollar payment over a long time period is nonnegative with a high probability. We remark that negative payments



may not be feasible in many scenarios. The focus of this paper is to reveal the fundamental benefit of negative payments to the data analyst when feasible.

**Main Results.** With the above formulation, we are interested in the following intriguing questions: (1) How will the individuals behave to reconcile the conflict between privacy and rewards? (2) How should the data analyst design the mechanism such that she can achieve her learning goal cost-effectively?

Let  $X_1, X_2, \dots, X_N$  denote the reported bits of the individuals. We model the privacy cost of an individual as a function of her privacy loss, which is measured by the level of (local) differential privacy [8, 9] of her data reporting strategy. This cost function of individual  $i$  is characterized by her *type*  $C_i$ : when individual  $i$  reports data with a (local) differential privacy level of  $\epsilon$  after observing her type  $C_i = c_i$ , her privacy loss is  $\epsilon$  and the corresponding privacy cost is  $c_i\epsilon$ . The type of an individual is also called her cost coefficient due to this linear form. We assume that the types are i.i.d. and an individual’s type is independent from her private data, which is applicable to the scenario where an individual’s valuation of privacy is intrinsic and thus is not affected by the specific private data she has. The reported data and cost coefficients are also illustrated in Fig. 1. We remark that both the settings where an individual’s valuation of privacy is independent (e.g., [14]) and correlated (e.g., [15]) with her private data have been studied in the literature. We further assume that it is possible for individuals to have valuations arbitrarily close to zero. In this paper, the prior distribution of the state, signals and types is public information. However, neither the private signals nor the types are known to the data analyst.

Our main result is centered around constructing a family of payment mechanisms indexed by parameters, which provide answers to the proposed questions from the following perspectives.

- **Behavior of Individuals with Privacy Concerns.** We show that the individuals exhibit a threshold behavior in a Bayesian Nash equilibrium of the proposed mechanism: the individuals with cost coefficients above a threshold  $c_{\text{th}}$  opt out, and the individuals with cost coefficients below  $c_{\text{th}}$  participate and report data with a privacy level no smaller than  $\epsilon$ , where  $c_{\text{th}}$  and  $\epsilon$  are parameters of the mechanism. Since a larger privacy level means that the data is less perturbed, the data analyst can incentivize the participants to limit the perturbation to a desired extent by choosing an appropriate  $\epsilon$ . It can be seen from this result that this mechanism resolves the otherwise nuisance that individuals with high privacy costs may participate and report only noise: they are “filtered out”, and the “remaining” participants all report data with quality guarantee.
- **Tradeoff Between Learning Accuracy and Cost.** We show that as the population size grows to infinity, the data analyst can learn the underlying state with arbitrarily small overall probability of error, with the total expected payment at the Bayesian Nash equilibrium going to zero. That is to say, if the data analyst can recruit a large number of individuals, she can choose appropriate parameters to fulfill her learning goal and in the meanwhile drive her cost to zero at a Bayesian Nash equilibrium. Since the total

equilibrium expected payment of any mechanism is nonnegative due to individual rationality, this implies that the designed mechanism with properly chosen parameters asymptotically minimizes the cost for achieving any accuracy goal.

**Related Work.** Market approaches for collecting data from privacy-aware individuals have led to a fruitful line of work [3, 5, 13–16, 20, 24, 26, 29, 32, 35]. These papers except [5, 29, 32] adopt the centralized model for privacy. The seminal work by Ghosh and Roth [16] and a line of subsequent work [13, 14, 20, 24, 26] considered the setting where the private data is verifiable so the individuals cannot misreport data, but they can strategically report their privacy costs. A recent work [5] considered a model where a data analyst procures possibly noisy estimates (data) from data providers. This can be thought of as a local privacy model, but still the data is verifiable. The setting of the work [3, 15, 29, 32, 35] is more similar to ours, where the individual have the option of misreporting data. The work [3, 15, 35] considered the centralized privacy model, where revealing data to the data analyst does not incur privacy costs. Then strategically reporting data can alter the individuals’ payments but does affect their privacy costs. The work [29, 32] considered the local privacy model but assumed the privacy cost functions are known to the data analyst. Our work studies this problem in a local privacy model, where neither the data nor the privacy cost functions are known. The mechanism thus needs to deal with the uncertainty in both sources and work with noisy reports.

The broader field of the interplay between differential privacy and mechanism design, first studied by McSherry and Talwar [22], is surveyed in [25]. The behavior of individuals with privacy concerns has been studied in [4], which investigates the types of games in which strategic individuals truthfully follow randomized response. The market approach for collecting private data also shares some structural similarity with the problem of information elicitation (e.g., [23]), especially the effort elicitation in the context of crowdsourcing (e.g., [2, 6, 21, 34]), where effort, instead of privacy concerns, affects the quality of the data and the cost of the individuals.

The local model of differential privacy, which generalizes the randomized response [33], has been studied in the literature [1, 4, 7–10, 17–19, 27, 28, 30]. The hypothesis testing formulation in our paper is similar to a setting in [18], where the authors find an optimal locally differentially private privatization mechanism that maximizes the statistical discrimination of the hypotheses. In practice, Google’s Chrome web browser has implemented the RAPPOR mechanism [11, 12] to collect users’ data using a locally differentially private protocol.

## 2 Model

We study the setting in which the data analyst is interested in learning an underlying state  $W$ , represented by a binary random variable. Consider a set  $[N] = \{1, 2, \dots, N\}$  of individuals. Each individual  $i$  possesses a binary signal

$S_i$ , which is her private data, and reports data  $X_i$ , which takes values in  $\mathcal{X} = \{0, 1, \perp\}$ , with  $\perp$  meaning “to opt out.” The data analyst announces a payment mechanism  $\mathbf{R}: \mathcal{X}^N \rightarrow \mathbb{R}^N$ , which takes the reported data  $\mathbf{X} = (X_1, \dots, X_N)$  as input and produces  $\mathbf{R}(\mathbf{X})$ , where  $R_i(\mathbf{X})$  is the payment to individual  $i$ . The model is illustrated in Fig. 1. The payment mechanism induces a game among the individuals. The elements of the game are as follows.

**Players.** The players in this game are the individuals, who are self-interested, rational and risk-neutral. Following conventional game theory notation, we let “ $-i$ ” denote all the individuals other than some given individual  $i$ .

**Prior Distributions.** The state  $W$  follows a probability distribution given by the PMF  $P_W$ . We assume that  $P_W(1) > 0$  and  $P_W(0) > 0$ . The individuals’ signals  $\mathbf{S} = (S_1, S_2, \dots, S_N)$  reflect their knowledge about the state  $W$ . Conditional on the state  $W$ , the signals  $S_1, S_2, \dots, S_N$  are independently generated according to  $\mathbb{P}(S_i = w \mid W = w) = \theta$  for  $w \in \{0, 1\}$ , where the parameter  $\theta$  with  $0.5 < \theta < 1$  is called the quality of signals. We refer to these conditional distributions as the signal structure of the model.

**Types and Strategies.** An individual  $i$ ’s type  $C_i$ , also called her cost coefficient, characterizes her valuation of privacy. We will elaborate on the assumptions on the types when we introduce the payoff functions below. Roughly, an individual with larger  $C_i$  experiences more privacy cost for the same privacy loss. A data reporting strategy for individual  $i$  is a plan on what to report according to her signal  $S_i$  and her type  $C_i$ . Thus it is a mapping  $\sigma_i: \{0, 1\} \times (0, +\infty) \rightarrow \mathcal{D}(\mathcal{X})$ , where  $\mathcal{D}(\mathcal{X})$  is the set of probability distributions on  $\mathcal{X} = \{0, 1, \perp\}$ , prescribing a distribution to the reported data  $X_i$  for each possible value pair of  $S_i$  and  $C_i$ . Therefore, the strategy corresponds to the set of conditional distributions of  $X_i$  given  $S_i$  and  $C_i$ . Since we will discuss different strategies of individual  $i$ , we denote these conditional probabilities by  $\mathbb{P}_{\sigma_i}(X_i = x_i \mid S_i = s_i, C_i = c_i)$  for  $x_i \in \{0, 1, \perp\}$ ,  $s_i \in \{0, 1\}$ , and  $c_i \in (0, +\infty)$ . Let  $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_N)$ , which is called a strategy profile. A strategy profile is said to be homogeneous if all the strategies in the profile are the same.

**Payoff Functions.** The payoff of each individual is the difference between the payment she receives and her privacy cost. An individual experiences a cost due to the privacy loss during data reporting. Recall that we model the privacy cost of an individual as consisting of two components: privacy loss and a privacy cost function, where the privacy loss depends on her data reporting strategy and the privacy cost function represents her valuation of privacy. For an individual  $i$ , conditional on her type  $C_i = c_i$ , we measure individual  $i$ ’s privacy loss for reporting data with strategy  $\sigma_i$  by the privacy level defined as follows:

$$\zeta(c_i, \sigma_i) = \max \left\{ \ln \frac{\mathbb{P}_{\sigma_i}(X_i \in \mathcal{E} \mid S_i = s_i, C_i = c_i)}{\mathbb{P}_{\sigma_i}(X_i \in \mathcal{E} \mid S_i = 1 - s_i, C_i = c_i)} : \mathcal{E} \subseteq \{0, 1, \perp\}, s_i \in \{0, 1\} \right\},$$

where we follow the convention that  $0/0 = 1$ . This measure of privacy loss is in the same vein as the local model of differential privacy [10, 19], which views each

individual's data as a database of size 1 and quantifies the privacy guarantee of her local randomizer by the differential privacy level. The difference here is that the strategy  $\sigma_i$  has another input  $C_i$ , since an individual can choose the way of perturbing her data according to her cost coefficient. Our measure of privacy loss is the differential privacy level of the strategy  $\sigma_i$  when  $C_i$  is given.

Then we model individual  $i$ 's cost incurred by this privacy loss as a linear function with  $C_i$  as the coefficient, i.e., the cost can be written as  $g(C_i, \sigma_i) = C_i \cdot \zeta(C_i, \sigma_i)$ . We call  $g$  the privacy cost function.

We assume that the coefficients  $C_1, C_2, \dots, C_N$  are i.i.d. positive random variables with CDF  $F_C$ , and they are independent of  $W$  and  $\mathcal{S}$ . The randomness of these coefficients captures the data analyst's uncertainty of individuals' valuations of privacy. The independence assumption is applicable to the scenario where individuals' valuations of privacy are intrinsic and thus are not affected by the specific private data they have. For ease of exposition, we further assume that  $F_C$  is a continuous function and  $F_C(c) > 0$  for any  $c > 0$ , which means that it is possible for individuals to have an arbitrarily low valuation of privacy. Similar analysis can be carried out for other models for the types (but the resulting accuracy-payment relation may be different).

**Mechanism Design.** The data analyst cannot force an individual to report data with a specific strategy. However, the data analyst can design the payment mechanism to impact individuals' strategies to drive the individuals to act in a desired way since the individuals are rational, i.e., they will choose the strategies that benefit them most. We consider the Bayesian Nash equilibria in a payment mechanism, viewing  $C_i$  as individual  $i$ 's type.

**Definition 1.** *A strategy profile  $\sigma$  is a Bayesian Nash equilibrium of a payment mechanism  $\mathbf{R}$  if for any individual  $i$ , any  $c_i > 0$  and any strategy  $\sigma'_i$ ,*

$$\mathbb{E}_{\sigma}[R_i(\mathbf{X}) - g(C_i, \sigma_i) \mid C_i = c_i] \geq \mathbb{E}_{(\sigma'_i, \sigma_{-i})}[R_i(\mathbf{X}) - g(C_i, \sigma'_i) \mid C_i = c_i],$$

where the subscript  $\sigma$  and  $(\sigma'_i, \sigma_{-i})$  indicate that the distribution of  $\mathbf{X}$  is determined by the strategy profile  $\sigma$  and  $(\sigma'_i, \sigma_{-i})$ , respectively.

The data analyst is interested in learning the state  $W$  from the reported data  $\mathbf{X}$ , so she performs hypothesis testing between the two hypotheses  $H_0: W = 0$  and  $H_1: W = 1$ . The learning accuracy is measured by the overall probability of error, denoted by  $p_e$ , which is  $P_W(0) \cdot (\text{Type I error}) + P_W(1) \cdot (\text{Type II error})$ . An accuracy goal can be written as  $p_e \leq p_e^{\max}$  for some  $p_e^{\max}$ .

Then the data analyst aims to design a payment mechanism such that her accuracy goal can be fulfilled at a Bayesian Nash equilibrium and the corresponding total expected payment is minimized. It is easy to see that the equilibrium total expected payment is nonnegative in any mechanism due to the nonnegativity of privacy cost functions and individual rationality. In this mechanism design problem, the joint distribution  $\mathcal{P}$  of the state  $W$ , the signal  $\mathcal{S}$  and the cost coefficients, which can be represented by  $(P_W, \theta, F_C)$ , is common knowledge. The data analyst announces the form of the payment mechanism and then the individuals report data simultaneously. The reported data  $\mathbf{X}$  is public.

Each individual  $i$ 's signal and type,  $S_i$  and  $C_i$ , are not observable to other individuals or the data analyst. No one has access to the state  $W$ .

### 3 Asymptotically Optimal Mechanisms

**Theorem 1.** *To achieve any accuracy goal of the data analyst, the total expected payment needed at an equilibrium is  $o(1)$ . Specifically, there exists a sequence of mechanisms, each of which is designed for a different population size  $N$ , such that the accuracy goal can be fulfilled at a Bayesian Nash equilibrium of every mechanism in the sequence, and the total expected payment goes to zero as the population size  $N$  goes to infinity; i.e., this sequence of mechanisms is asymptotically optimal.*

In the remainder of this section, we present the design of a family of payment mechanisms, parameterized by the population size  $N$ , the prior  $\mathcal{P}$ , a cost coefficient threshold parameter  $c_{\text{th}}$  and a data quality parameter  $\epsilon$ . The asymptotically optimal sequence of mechanisms in Theorem 1 is given by a sequence of mechanisms within this family with properly chosen parameters. In particular,  $c_{\text{th}}$  is a threshold on cost coefficients such that an individual is expected to participate if her coefficient does not exceed the threshold; and  $\epsilon$  is the target quality which is the level noise expected in the reported data. The formula for calculating  $c_{\text{th}}$  and  $\epsilon$  will be presented in Sect. 5. Theorem 1 is a high level description of Theorem 3, which will be derived in the remainder of this paper.

#### Payment Mechanism $\mathbf{R}^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}$

1. Each individual reports her data (which can also be “to opt out”).
2. Compute the number of participants  $n$ .
3. For non-participating individuals, the payment is zero.
4. If there is only one participant, the data analyst pays zero to this participant. Otherwise, for each participating individual  $i$ , compute the majority of other participants' reported data, denoted by  $M_{-i}$ . Then the data analyst pays individual  $i$  according to  $X_i$  and  $M_{-i}$  as follows:

$$R_i^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}(\mathbf{X}) = A_{X_i, M_{-i}} \frac{c_{\text{th}}(e^\epsilon + 1)^2}{2e^\epsilon} + B_{M_{-i}} \left( \frac{c_{\text{th}}(e^\epsilon + 1)}{e^\epsilon} + c_{\text{th}}\epsilon \right),$$

where  $A_{1,1}, A_{0,1}, A_{1,0}, A_{0,0}, B_1, B_0$  are given below.

Next we define the coefficients  $A_{1,1}, A_{0,1}, A_{1,0}, A_{0,0}, B_1, B_0$  used in the mechanism  $\mathbf{R}^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}$  through a series of calculations. In a nutshell,  $A_{1,1}$  and  $A_{0,0}$  determine the reward part of the payment to an individual when her reported data matches the majority of others; similarly,  $A_{0,1}$  and  $A_{1,0}$  determine the penalty part of the payment to an individual when her reported data does not match the majority of others. They incentivize the individuals to report data that reveals certain amount of information about their private signals. The coefficients  $B_1$  and  $B_0$  offset the payments for the cases that the majority of others'

reports is 1 and 0, respectively, to discourage the individuals with cost coefficients above threshold parameter  $c_{\text{th}}$  from participating. We remark that when an individual's reported data does not match with the majority of others, these coefficients make sure that the payment to this individual is negative.

The definition of the coefficients  $A_{1,1}$ ,  $A_{0,1}$ ,  $A_{1,0}$ ,  $A_{0,0}$ ,  $B_1$ ,  $B_0$  involves some intermediate quantities, the physical meanings of which will be given after we characterize a Bayesian Nash equilibrium of the mechanism in Sect. 4. Given a  $c_{\text{th}} \in (0, +\infty)$  and  $\epsilon \in (0, +\infty)$ , for each  $c_i \in (0, c_{\text{th}})$ , we consider the following equation with variable  $\xi$ :  $c_{\text{th}}(e^\epsilon + 1)^2 e^\xi = c_i e^\epsilon (e^\xi + 1)^2$ . It can be proved that this equation has a unique solution in  $(0, +\infty)$ . Let this solution define a function  $\xi(c_i)$ . Specifically,

$$\xi(c_i) = \ln \left( \frac{1}{\frac{1}{2} - \sqrt{\frac{1}{4} - \frac{c_i}{c_{\text{th}}} \frac{e^\epsilon}{(e^\epsilon + 1)^2}}} - 1 \right). \quad (1)$$

Let

$$\mu = \int_0^{c_{\text{th}}} \frac{e^{\xi(c_i)}}{e^{\xi(c_i)} + 1} dF_{C|C_i \leq c_{\text{th}}}(c_i), \quad \alpha = \theta\mu + (1 - \theta)(1 - \mu), \quad (2)$$

where  $F_{C|C_i \leq c_{\text{th}}}$  is the conditional distribution of  $C_i$  given  $C_i \leq c_{\text{th}}$ .

Given that the number of participants is  $n$  with  $n \geq 2$ , we define the following quantities. Consider a random variable that follows the binomial distribution with parameters  $n - 1$  and  $\alpha$ . Let  $\beta^{(n)}$  denote the probability that this random variable is greater than or equal to  $\lfloor \frac{n-1}{2} \rfloor + 1$ . For convenience, we define the following quantity to deal with technical details:

$$\gamma^{(n)} = \begin{cases} 1 - \left( \frac{n-1}{2} \right) \alpha^{\frac{n-1}{2}} (1 - \alpha)^{\frac{n-1}{2}} & \text{if } n - 1 \text{ is even,} \\ 1 & \text{if } n - 1 \text{ is odd.} \end{cases}$$

Let  $P_{\geq 1} = 1 - (1 - F_C(c_{\text{th}}))^{N-1}$ , where  $F_C$  is the CDF of  $C_i$ . We define

$$\begin{aligned} A_{1,1} &= \frac{P_W(1)\theta(1 - \beta^{(n)}) + P_W(0)(1 - \theta)(1 - (\gamma^{(n)} - \beta^{(n)}))}{P_{\geq 1}P_W(1)P_W(0)(2\theta - 1)(2\beta^{(n)} - \gamma^{(n)})}, \\ A_{0,1} &= -\frac{P_W(1)(1 - \theta)(1 - \beta^{(n)}) + P_W(0)\theta(1 - (\gamma^{(n)} - \beta^{(n)}))}{P_{\geq 1}P_W(1)P_W(0)(2\theta - 1)(2\beta^{(n)} - \gamma^{(n)})}, \\ A_{1,0} &= -\frac{P_W(1)\theta\beta^{(n)} + P_W(0)(1 - \theta)(\gamma^{(n)} - \beta^{(n)})}{P_{\geq 1}P_W(1)P_W(0)(2\theta - 1)(2\beta^{(n)} - \gamma^{(n)})}, \\ A_{0,0} &= \frac{P_W(1)(1 - \theta)\beta^{(n)} + P_W(0)\theta(\gamma^{(n)} - \beta^{(n)})}{P_{\geq 1}P_W(1)P_W(0)(2\theta - 1)(2\beta^{(n)} - \gamma^{(n)})}, \\ B_1 &= -\frac{P_W(1)(1 - \beta^{(n)}) - P_W(0)(1 - (\gamma^{(n)} - \beta^{(n)}))}{2P_{\geq 1}P_W(1)P_W(0)(2\beta^{(n)} - \gamma^{(n)})}, \\ B_0 &= \frac{P_W(1)\beta^{(n)} - P_W(0)(\gamma^{(n)} - \beta^{(n)})}{2P_{\geq 1}P_W(1)P_W(0)(2\beta^{(n)} - \gamma^{(n)})}. \end{aligned}$$

## 4 Bayesian Nash Equilibrium

In this section, we first characterize the individuals' behavior at a Bayesian Nash equilibrium of the designed mechanism. The equilibrium behavior affects the quality of the reported data and the payments. Then we leverage the properties of the Bayesian Nash equilibrium to explain the physical meanings of the quantities defined during the construction of the mechanism in Sect. 3.

**Theorem 2.** *The mechanism  $\mathbf{R}^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}$  yields a Bayesian Nash equilibrium  $\sigma$ , in which each individual  $i$ 's strategy  $\sigma_i$  is described as follows:*

- If  $c_i > c_{\text{th}}$ ,  $\mathbb{P}_{\sigma_i}(X_i = \perp \mid S_i = s_i, C_i = c_i) = 1$  for any  $s_i \in \{0, 1\}$ ; i.e., if individual  $i$ 's cost coefficient is larger than the parameter  $c_{\text{th}}$ , individual  $i$  declines to participate regardless of her signal.
- If  $c_i \leq c_{\text{th}}$ ,

$$\mathbb{P}_{\sigma_i}(X_i = 1 \mid S_i = 1, C_i = c_i) = \mathbb{P}_{\sigma_i}(X_i = 0 \mid S_i = 0, C_i = c_i) = \frac{e^{\xi(c_i)}}{e^{\xi(c_i)} + 1},$$

$$\mathbb{P}_{\sigma_i}(X_i = 0 \mid S_i = 1, C_i = c_i) = \mathbb{P}_{\sigma_i}(X_i = 1 \mid S_i = 0, C_i = c_i) = \frac{1}{e^{\xi(c_i)} + 1},$$

where  $\xi(c_i)$  is defined in (1); i.e., if individual  $i$ 's cost coefficient is no larger than the parameter  $c_{\text{th}}$ , individual  $i$  flips her signal with a probability depending on her cost coefficient to generate her reported data.

The following corollary describes the quality of the reported data and each participant's expected payment at the Bayesian Nash equilibrium in Theorem 2.

**Corollary 1.** *For the mechanism  $\mathbf{R}^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}$ , consider the Bayesian Nash equilibrium  $\sigma$  given in Theorem 2.*

- For each participating individual  $i$ ,

$$\mathbb{P}_{\sigma_i}(X_i = 1 \mid S_i = 1, i \text{ participates}) = \mathbb{P}_{\sigma_i}(X_i = 0 \mid S_i = 0, i \text{ participates}) = \mu,$$

where  $\mu$  is defined in (2) and  $\mu \geq \frac{e^\epsilon}{e^\epsilon + 1}$ .

- The expected payment to each participating individual  $i$  is bounded as

$$\mathbb{E}_\sigma[R_i^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}(\mathbf{X}) \mid i \text{ participates}] \leq c_{\text{th}}(1 + e^{-\epsilon} + \epsilon).$$

The proofs of Theorem 2 and Corollary 1 are presented in the full version [31]. Theorem 2 and Corollary 1 show how individuals with high privacy costs are “filtered out” in the equilibrium by negative payments. In other words, they will decide not to participate because the expected payment is negative, which is a result of the possible negative payments in the proposed mechanism. The “remaining” individuals, i.e., participants, all report data with quality guarantee. The roles of the parameters  $c_{\text{th}}$  and  $\epsilon$  in the designed mechanism  $\mathbf{R}^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}$  are as follows: The parameter  $c_{\text{th}}$  works as a threshold on the cost coefficients

for participation; The parameter  $\epsilon$  gives a guarantee on the probability that a participant's reported data is the same as the signal, which measures the quality of the reported data. We remark that in this equilibrium, each individual's exact cost coefficient is not revealed to other.

The physical meanings of the quantities  $\xi(c_i)$ ,  $\mu$ ,  $\alpha$ ,  $\beta^{(n)}$ ,  $\gamma^{(n)}$  and  $P_{\geq 1}$  defined during the construction of the mechanism in Sect. 3 can be well explained at the Bayesian Nash equilibrium given in Theorem 2. The quantity  $\xi(c_i)$  shows up in Theorem 2, characterizing the strategy  $\sigma_i$  of individual  $i$  when  $c_i \leq c_{\text{th}}$ . It is the differential privacy level of  $\sigma_i$  given  $C_i = c_i$  when  $c_i \leq c_{\text{th}}$ . Now let us condition on the event that individual  $i$  participates, which, by Theorem 2, is equivalent to the event  $C_i \leq c_{\text{th}}$ . The quantity  $\mu$  shows up in Corollary 1, and it is the probability that individual  $i$  truthfully reports her signal, given whatever the signal is. Then the quantity  $\alpha$  is the probability that the reported data  $X_i$  is consistent with the state  $W$ , given whatever the state is. Conditional on the event that there are  $n - 1$  participants among the individuals other than individual  $i$ , where  $n \geq 2$ , the quantities  $\beta_n$  and  $1 - (\gamma_n - \beta_n)$  are the probabilities that the majority of these participants' reported data agrees with the state, given that the state is 1 and 0, respectively. Finally, the quantity  $P_{\geq 1}$  is the probability that at least one individual among the individuals other than individual  $i$  participates.

## 5 Accuracy and Payment

In this section, we show that the data analyst can achieve any accuracy goal in the Bayesian Nash equilibrium with proper choice of parameters  $N$ ,  $c_{\text{th}}$  and  $\epsilon$ . The cost of the data analyst, which is the total expected payment at the equilibrium, goes to zero as the number of individuals goes to infinity. Since the privacy cost of an individual is always nonnegative, the total expected payment at an equilibrium of any mechanism is nonnegative due to individual rationality. Therefore, the designed mechanism asymptotically minimizes the cost for the data analyst to achieve any accuracy goal.

Recall that with the procured data  $\mathbf{X}$ , the data analyst learns the state  $W$  by performing hypothesis testing between the two hypotheses  $H_0: W = 0$  and  $H_1: W = 1$ . An accuracy goal can be written as  $p_e \leq p_e^{\max}$  for some  $p_e^{\max}$ , where  $p_e$  is the overall probability of error for hypothesis testing. We consider the maximum likelihood decision. The values for  $N$ ,  $c_{\text{th}}$ ,  $\epsilon$  are chosen using the procedure below. The intuition is that we first fix the quality that the analyst expects to obtain from each participant and the types of individuals the analyst would like to collect data from, and then the accuracy goal can be met when the population size is large enough to make sure that there are enough participants.

**Parameter Selection Procedure.** Pick any  $\epsilon$  such that  $\epsilon \in (0, +\infty)$ . Let

$$D(\epsilon) = \frac{1}{2} \ln \frac{(e^\epsilon + 1)^2}{4(\theta e^\epsilon + 1 - \theta)((1 - \theta)e^\epsilon + \theta)}, \quad n_e(\epsilon) = \frac{-\ln(\frac{1}{2}p_e^{\max})}{D(\epsilon)},$$

$$\rho(\epsilon) = \frac{1}{n_e(\epsilon)p_e^{\max}} + 2 + \sqrt{\frac{1}{(n_e(\epsilon))^2(p_e^{\max})^2} + \frac{2}{n_e(\epsilon)p_e^{\max}}}.$$



Then pick any integer  $N$  such that  $N > \rho(\epsilon)n_e(\epsilon)$ . For the selected  $N$ , let  $p_{\text{th}}(N, \epsilon) = \rho(\epsilon)n_e(\epsilon)/N$ , which is roughly the participation percentage, and then let  $c_{\text{th}}(N, \epsilon) = \inf\{c: F_C(c) = p_{\text{th}}(N, \epsilon)\}$ .

Recall that we assume  $F_C$  to be a continuous function, so the set  $\{c: F_C(c) = p_{\text{th}}(N, \epsilon)\}$  is nonempty and thus  $c_{\text{th}}(N, \epsilon) \geq 0$  is finite. An example of this parameter selection procedure (and the resulted upper bound on total expected payment) can be found in the full version [31].

**Theorem 3.** *For the mechanism  $\mathbf{R}^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}$ , consider the Bayesian Nash equilibrium  $\sigma$  given in Theorem 2. Given an accuracy goal  $p_e \leq p_e^{\max}$ , let  $(N, c_{\text{th}}, \epsilon)$  be chosen according to the parameter selection procedure above and the data analyst performs hypothesis testing using the maximum likelihood approach.*

– The decision function  $\psi$  has the following form:

$$\psi(\mathbf{X}) = \begin{cases} 1 & \text{if } \sum_i \mathbb{1}_{\{X_i=1\}} \geq \sum_i \mathbb{1}_{\{X_i=0\}}, \\ 0 & \text{otherwise;} \end{cases} \quad (3)$$

– The overall probability of error,  $p_e$ , meets the accuracy goal  $p_e \leq p_e^{\max}$ ;  
 – The total expected payment is bounded as

$$\mathbb{E}_{\sigma} \left[ \sum_{i=1}^N R_i^{(N, \mathcal{P}, c_{\text{th}}, \epsilon)}(\mathbf{X}) \right] \leq c_{\text{th}}(\epsilon, N) \rho(\epsilon) n_e(\epsilon) \cdot (1 + e^{-\epsilon} + \epsilon). \quad (4)$$

Since  $\rho(\epsilon)$  and  $n_e(\epsilon)$  are constants for given  $\epsilon$ , and  $c_{\text{th}}(\epsilon, N)$  goes to 0 as  $N \rightarrow \infty$ , this total expected payment goes to zero, with the accuracy goal met, as  $N \rightarrow \infty$ .

The proof of Theorem 3 is presented in the full version [31]. Theorem 3 shows that choosing parameters according to the parameter selection procedure for the designed family of mechanisms not only meets the accuracy goal of the data analyst but is also cost-effective. The intuition is that as  $N$  becomes large, the requirement on the participation percentage becomes lower, which allows the mechanism to collect data from individuals with lower privacy costs and thus drives down the data analyst’s cost. This suggests a way of constructing the asymptotically optimal sequence in Theorem 1: Fix an  $\epsilon \in (0, +\infty)$ , and then choose a sequence of mechanisms, each of which is designed for a different population size  $N$  and has parameter  $c_{\text{th}}$ , both of which are chosen according to the parameter selection procedure.

## 6 Conclusions

We considered incentive mechanisms for collecting private data from strategic, privacy-aware individuals, whose valuations of privacy are unknown. The data analyst is interested in learning an underlying state from the private data of individuals with minimum overall payment. We considered a local model of data

privacy, where the data analyst is not necessarily trusted, and data subjects are endowed with the ability to control their own privacy, which frees the data analyst from the responsibility of privacy protection. We designed a family of payment mechanisms for the data analyst, which utilize negative payments to prevent individuals with high privacy valuations from reporting only noise and cut down the cost of the data analyst. In each designed mechanism, the individuals exhibit a threshold behavior at a Bayesian Nash equilibrium: only those with cost coefficients below some threshold participate, and they report data with certain quality guarantee, where the threshold and the quality guarantee are both parameters of the mechanism. With appropriate choices of parameters, the data analyst can fulfill any accuracy goal with diminishing cost at the equilibrium as the number of individuals grows to infinity.

**Acknowledgments.** This work was supported partially by NSF grant CNS-1618768.

## References

1. Bassily, R., Smith, A.: Local, private, efficient protocols for succinct histograms. In: Proceedings of the Annual ACM on Symposium on Theory of Computing (STOC), Portland, OR, pp. 127–135 (2015)
2. Cai, Y., Daskalakis, C., Papadimitriou, C.: Optimum statistical estimation with strategic data sources. In: Proceedings of the Conference on Learning Theory (COLT), Paris, France. pp. 280–296. July 2015
3. Chen, Y., Chong, S., Kash, I.A., Moran, T., Vadhan, S.: Truthful mechanisms for agents that value privacy. In: Proceedings of the ACM Conference on Electronic commerce (EC), Philadelphia, PA. pp. 215–232 (2013)
4. Chen, Y., Sheffet, O., Vadhan, S.: Privacy games. In: Liu, T.-Y., Qi, Q., Ye, Y. (eds.) WINE 2014. LNCS, vol. 8877, pp. 371–385. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-13129-0\\_30](https://doi.org/10.1007/978-3-319-13129-0_30)
5. Cummings, R., Ligett, K., Roth, A., Wu, Z.S., Ziani, J.: Accuracy for sale: aggregating data with a variance constraint. In: Proceedings of the Conference on Innovations in Theoretical Computer Science (ITCS), Rehovot, Israel. pp. 317–324 (2015)
6. Dasgupta, A., Ghosh, A.: Crowdsourced judgement elicitation with endogenous proficiency. In: Proceedings of the International Conference on World Wide Web (WWW), Rio de Janeiro, Brazil, pp. 319–330, May 2013
7. Duchi, J.C., Jordan, M.I., Wainwright, M.J.: Local privacy and minimax bounds: sharp rates for probability estimation. In: Advances Neural Information Processing Systems (NIPS), Lake Tahoe, NV, pp. 1529–1537, December 2013
8. Dwork, C.: Differential privacy. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 1–12. Springer, Heidelberg (2006). doi:[10.1007/11787006\\_1](https://doi.org/10.1007/11787006_1)
9. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 265–284. Springer, Heidelberg (2006). doi:[10.1007/11681878\\_14](https://doi.org/10.1007/11681878_14)
10. Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* **9**(3–4), 211–407 (2014)

11. Erlingsson, Ú., Pihur, V., Korolova, A.: RAPPOR: randomized aggregatable privacy-preserving ordinal response. In: Proceedings of the ACM SIGSAC Conference on Computer and Communication Security (CCS), Scottsdale, AZ, pp. 1054–1067 (2014)
12. Fanti, G.C., Pihur, V., Erlingsson, Ú.: Building a RAPPOR with the unknown: privacy-preserving learning of associations and data dictionaries. [arXiv:1503.01214](https://arxiv.org/abs/1503.01214) [cs.CR] (2015)
13. Fleischer, L.K., Lyu, Y.: Approximately optimal auctions for selling privacy when costs are correlated with data. In: Proceedings of the ACM Conference on Electronic Commerce (EC), Valencia, Spain, pp. 568–585 (2012)
14. Ghosh, A., Ligett, K.: Privacy and coordination: computing on databases with endogenous participation. In: Proceedings of the ACM Conference on Electronic Commerce (EC), Philadelphia, PA, pp. 543–560 (2013)
15. Ghosh, A., Ligett, K., Roth, A., Schoenebeck, G.: Buying private data without verification. In: Proceedings of the ACM Conference on Economics and Computation (EC), Palo Alto, CA, pp. 931–948 (2014)
16. Ghosh, A., Roth, A.: Selling privacy at auction. In: Proceedings of the ACM Conference on Electronic Commerce (EC), San Jose, CA, pp. 199–208 (2011)
17. Hsu, J., Khanna, S., Roth, A.: Distributed private heavy hitters. In: Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.) ICALP 2012. LNCS, vol. 7391, pp. 461–472. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-31594-7\\_39](https://doi.org/10.1007/978-3-642-31594-7_39)
18. Kairouz, P., Oh, S., Viswanath, P.: Extremal mechanisms for local differential privacy. In: Advances Neural Information Processing Systems (NIPS), Montreal, Canada, pp. 2879–2887, December 2014
19. Kasiviswanathan, S.P., Lee, H.K., Nissim, K., Raskhodnikova, S., Smith, A.: What can we learn privately? *SIAM J. Comput.* **40**(3), 793–826 (2011)
20. Ligett, K., Roth, A.: Take it or leave it: running a survey when privacy comes at a cost. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 378–391. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-35311-6\\_28](https://doi.org/10.1007/978-3-642-35311-6_28)
21. Liu, Y., Chen, Y.: Learning to incentivize: eliciting effort via output agreement. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), July 2016, New York, NY
22. McSherry, F., Talwar, K.: Mechanism design via differential privacy. In: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science (FOCS), Providence, RI, pp. 94–103 (2007)
23. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting informative feedback: the peer-prediction method. In: Golbeck, J. (ed.) Computing with Social Trust. Human-Computer Interaction Series, pp. 185–212. Springer, London (2009)
24. Nissim, K., Vadhan, S., Xiao, D.: Redrawing the boundaries on purchasing data from privacy-sensitive individuals. In: Proceedings of the Conference Innovations in Theoretical Computer Science (ITCS), Princeton, NJ, pp. 411–422 (2014)
25. Pai, M.M., Roth, A.: Privacy and mechanism design. *SIGecom Exch.* **12**(1), 8–29 (2013)
26. Roth, A., Schoenebeck, G.: Conducting truthful surveys, cheaply. In: Proceedings of the ACM Conference Electronic Commerce (EC), Valencia, Spain, pp. 826–843 (2012)
27. Shokri, R.: Privacy games: optimal user-centric data obfuscation. In: Proceedings of the Privacy Enhancing Technologies (PETS), Philadelphia, PA, pp. 299–315 (2015)

28. Wang, W., Ying, L., Zhang, J.: On the relation between identifiability, differential privacy, and mutual-information privacy. In: Proceedings of the Annual Allerton Conference on Communication, Control and Computing, Monticello, IL, pp. 1086–1092, September 2014
29. Wang, W., Ying, L., Zhang, J.: A game-theoretic approach to quality control for collecting privacy-preserving data. In: Proceedings of the Annual Allerton Conference on Communication, Control and Computing, Monticello, IL, pp. 474–479, September 2015
30. Wang, W., Ying, L., Zhang, J.: A minimax distortion view of differentially private query release. In: Proceedings of the Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, pp. 1046–1050, November 2015
31. Wang, W., Ying, L., Zhang, J.: Buying data from privacy-aware individuals: the effect of negative payments. Technical report, Arizona State University, Tempe, AZ (2016)
32. Wang, W., Ying, L., Zhang, J.: The value of privacy: strategic data subjects, incentive mechanisms and fundamental limits. In: Proceedings of the Annual ACM SIGMETRICS Conference on Antibes Juan-les-Pins, France, June 2016
33. Warner, S.L.: Randomized response: a survey technique for eliminating evasive answer bias. *J. Amer. Stat. Assoc.* **60**(309), 63–69 (1965)
34. Witkowski, J., Bachrach, Y., Key, P., Parkes, D.C.: Dwelling on the negative: incentivizing effort in peer prediction. In: Proceedings of the AAAI Conference on Human Computation and Crowdsourcing (HCOMP), Palm Springs, CA, November 2013
35. Xiao, D.: Is privacy compatible with truthfulness? In: Proceedings of the Conference on Innovations in Theoretical Computer Science (ITCS), Berkeley, CA, pp. 67–86 (2013)

# Bidding Strategies for Fantasy-Sports Auctions

Aris Anagnostopoulos<sup>1</sup>, Ruggiero Cavallo<sup>2</sup>(✉), Stefano Leonardi<sup>1</sup>,  
and Maxim Sviridenko<sup>2</sup>

<sup>1</sup> Sapienza University of Rome, Rome, Italy

{[aris](mailto:aris@uniroma1.it),[leonardi](mailto:leonardi@uniroma1.it)}@dis.uniroma1.it

<sup>2</sup> Yahoo Research, New York, USA

[cavallo@yahoo-inc.com](mailto:cavallo@yahoo-inc.com), [sviri@yahoo-inc.com](mailto:sviri@yahoo-inc.com)

**Abstract.** Fantasy sports is a fast-growing, multi-billion dollar industry [10] in which competitors assemble virtual teams of athletes from real professional sports leagues and obtain points based on the statistical performance of those athletes in actual games. Users (team *managers*) can add, drop, and trade players throughout the season, but the pivotal event is the player draft that initiates the competition. One common drafting mechanism is the so-called *auction draft*: managers bid on athletes in rounds until all positions on each roster have been filled. Managers start with the same initial virtual budget and take turns successively nominating athletes to be auctioned, with the winner of each round making a virtual payment that diminishes his budget for future rounds. Each manager tries to obtain players that maximize the expected performance of his own team. In this paper we initiate the study of bidding strategies for fantasy sports auction drafts, focusing on the design and analysis of simple strategies that achieve good worst-case performance, obtaining a constant fraction of the best value possible, regardless of competing managers' bids. Our findings may be useful in guiding bidding behavior of fantasy sports participants, and perhaps more importantly may provide the basis for a competitive *auto-draft* mechanism to be used as a bidding proxy for participants who are absent from their league's draft.

## 1 Introduction

In fantasy sports, individuals compete against each other by becoming virtual managers of a team of professional athletes, choosing players and modifying rosters over the course of a season, competing based on the statistical performance of the athletes composing their respective teams. Athlete statistics from real-life games are converted into “fantasy points”, which are compiled and aggregated. Points may be manually calculated by a participant designated as “league commissioner” who coordinates and manages the overall league, or they may be compiled and calculated by online platforms tracking game results. Managers draft, trade, and cut athletes over the course of the season in response to changing evaluations of athlete potentials, analogously to real sports.

---

A. Anagnostopoulos and S. Leonardi—This author's work was partly supported by the EU FET project MULTIPLEX no. 317532.

Fantasy sports are a multi-billion dollar industry [10]. According to the Fantasy Sports Trade Association (FSTA), in 2015 there were 56.8 million people playing fantasy sports in the USA and Canada. On average, fantasy-sports participants (age 18+) spend \$465 on league-related costs, single-player challenge games, and league-related materials over a 12-month period [1]. Each fantasy-sports league can have up to 20 teams, although most managers prefer to have leagues of at most 12, presumably because large leagues require drafting lower-performance and lesser-known athletes.

There are two standard ways to pick (draft) a fantasy team: snake drafts and auctions. Even though the majority of fantasy leagues use the standard snake draft (taking turns choosing players), anecdotal evidence suggests that auction-style drafts—wherein managers take turns *nominating* players, who are ultimately allocated based on competitive bidding—are popular among the most experienced and engaged users. If one of the managers is not present during the draft, there is typically an *auto-draft* algorithm, which makes decisions for the absentee manager. Furthermore, in some cases the auto-draft system is employed as a practice tool for inexperienced users. The real-time fantasy auction draft involves sophisticated decision making and strategies [2, 7], and being able to practice such an auction in advance is important for one’s success during the actual draft.

This paper is motivated by the goal of designing an auto-draft proxy bidder good enough to keep a manager competitive—in some broad sense—with the rest of the league. Because this is a competitive and strategic environment, it is natural to take a game-theoretic approach and consider strategy profiles that form Nash equilibria. However, we find that plausible pure strategy equilibria often do not even exist, and in the cases where they do exist their computation would be completely impractical for non-computer-aided participants, thus removing any predictive value. We therefore instead focus on strategies with good *guarantees*, creating teams with a certain competitive value even when faced with optimal adversarial opponent bidders.

## 1.1 Description of the Real Auction

A fantasy sports league is composed of a set of  $k$  team *managers* (or *users*)  $u_1, \dots, u_k$ —where  $k$  usually ranges from 3 to 20—who form teams from a pool of  $n$  *athletes* (or *players*)  $P_1, \dots, P_n$ .<sup>1</sup> Each fantasy team must be composed of  $m$  athletes. The number of athletes  $m$  depends on the sport and fantasy games provider; for example, Yahoo Daily Fantasy NFL team rosters have 9 slots: 1 quarterback, 2 running backs, 3 wide receivers, 1 tight end, 1 flexible position, and 1 slot for a defensive team.

As mentioned above, leagues are formed via snake draft or auction draft, with this choice determined by the initiator of the league. In snake drafts users

---

<sup>1</sup> For instance, in the NFL each actual team consists of a 53-man roster (plus a 5-athlete practice squad). There are 32 NFL teams, making for a total of  $n \geq 1,696$  athletes, all of whom are eligible to be drafted.

successively take athletes, with no bidding or competition. But auction drafts, which are the focus of the current paper, are strategically complex. Each fantasy manager is given a fixed budget  $B$  to draft a team with a set number of athletes. For instance, Yahoo provides each team manager with a \$200 budget, although the precise value is not extremely relevant. Managers take turns successively (in some pre-determined order) nominating athletes for bidding, which then proceeds via an *English auction* as follows. The default bid ascribed to the nominating manager is \$1, and it can be raised to any dollar value within the nominating manager's budget. Managers are given a fixed amount of time (say 10s) to place a higher bid; if any manager does so, the clock resets and users have another 10s to place a yet higher bid, and so on. Managers are free to bid any amount as long as they have available funds. When bidding stops, the manager who submitted the high bid wins the athlete and his bid amount is subtracted from the funds available to him for future bidding.

There is no incentive to retain budget after the auction, because leftover money cannot be used for anything else. Managers are allowed to bid only to a dollar amount that leaves at least one dollar for each currently unfilled spot on the team. Further, each athlete comes with a fixed position (there are quarterbacks, running backs, wide receivers, tight ends, kickers, and defensemen for NFL fantasy athletes). Each team must meet a fixed distribution of positions that depends on the sport and fantasy-game provider (for example, an NFL manager has to acquire exactly 2 running backs). The goal of a manager is to fill her positions in a way that maximizes the overall value of the team according to the stats her players accrue during real play throughout the season.

## 1.2 The Auction that We Analyze

In this work we initiate the study of fantasy-sports auctions. We make some simplifying assumptions to make the problem amenable to a rigorous theoretical investigation. Our first simplifying assumption is that team managers agree on the *value* of every athlete. This assumption, clearly somewhat of a departure from reality, is made primarily to simplify presentation; in fact, if value estimates vary across managers, our main results continue to apply with respect to the manager's subjective value estimates for the players in the league (see the Conclusion section for further discussion of this). Moreover, a rough correspondence in value estimates can be expected because previous season statistics (average number of fantasy points per game, etc.) and *next-season projections* are widely available on specialized services on the web (e.g., at [rotowire.com](http://rotowire.com)), which many managers consult. Formally, we assume that each athlete  $P_i$  has an associated value  $v_i$  denoting the expected number of fantasy points he will earn throughout the season;  $v_i$  is a shared belief, common knowledge to all managers.

A second simplifying assumption is that each phase of the auction draft is a second-price sealed-bid auction in which managers are allowed to place arbitrary fractional bids (limited only by their remaining budget). Even though this also deviates from reality, the English auction described in the previous section yields similar results, in theory. If more than one manager announces the highest bid

in our sealed bid auction, we allocate the athlete to each high-bidding manager with equal probability. The exception to this rule is when all managers place the minimal bid. In this case we assume that the athlete is won by the nominating manager, deterministically. That is, no manager can impose an athlete on others simply by nominating him. This reflects the real-life policy: an athlete always goes to a nominating manager if there is no competing bid from other managers.

We assume that managers are allowed to bid anywhere between \$0 and their entire remaining budget for each athlete (recall that in reality bids are capped to ensure at least \$1 remains for every open roster position). If a manager expends the entire budget before filling her roster, she can pick up lesser-preferred athletes for free after all other managers have filled their rosters. If there are multiple managers with zero budget left at the end of the auction, the remaining athletes are chosen by managers according to the nomination order.

Having fixed the price of each athlete, we can assume that for each position type the player pool has exactly the number of athletes required to complete each team. For example, if there are 8 NFL managers, and each manager must hire 2 running backs, then there exist exactly 16 running backs overall; these will be the most valuable, as no manager will prefer a lower-ranked running back (and it is easy to show that no advantage can be gained by putting one up for bid). We can therefore assume that  $n = km$ .

Finally, it will be useful to define the value

$$V = \frac{1}{k} \sum_{i=1}^n v_i$$

This can be considered the *fair share* of total value each manager would get if all managers were equally competitive. Against optimally bidding adversaries, the best a manager can reasonably hope for is to draft a team of value not much worse than  $V$ .

### 1.3 Summary of Results

After describing related work in Sect. 2, we begin in Sect. 3 by showing that pure strategy subgame perfect Nash equilibria do not generally exist in the game representing the fantasy auction draft. In fact, even for what is virtually the simplest example one can conjure (2 managers, 4 athletes to be drafted), unless the athletes are nominated in a particular pre-fixed order, there is no pure strategy Nash equilibrium. That entails—allowing the nomination choices to be in fact part of the strategy space—that there can be no pure-strategy subgame perfect equilibrium of the auction draft.

This, along with other practical critiques of an equilibrium-based analysis, motivates us to focus instead on worst-case analysis. In other words, can we describe a bidding algorithm that performs well no matter what opponent managers do? In Sect. 4 we provide an analysis of the simplified case where athletes are automatically nominated in non-increasing order of their values ( $v_1 \geq v_2 \geq \dots \geq v_n$ ). For each athlete  $P_i$  we define his *fair price* as  $c_i = \frac{v_i}{V} B$ ,



and we use this fair price as the basis of the following simple (and nonadaptive) bidding strategy: letting  $b_i$  be the manager’s current leftover budget before the  $i^{\text{th}}$  athlete is nominated, the manager always places bid  $\min\{\alpha c_i, b_i\}$ , with  $\alpha = 3/2$ . We show that a manager playing this strategy is guaranteed to obtain value at least  $V/3$ , regardless of the other managers’ bids. Moreover, we show that this is the best one can do for a strategy of this class, in that any alternative choice of  $\alpha$  will yield a worst-case value no greater than  $V/3$ .

In Sect. 5 we analyze a more complex bidding strategy for the general case in which nominations are made in a general adaptive fashion according to manager strategies. We prove that the proposed bidding strategy is able to guarantee a final team with value at least  $V/16$ .

## 2 Related Work

For a general reference to auction theory we refer to [15]. Dynamic auctions have been extensively studied in a literature that has recently been very active [3]. Note that there is no mechanism design component in the current paper, because the draft mechanism is determined by the platform and we are rather studying how to best engage with it. We focus on the strategic implications for users who in principle can play very sophisticated bidding strategies. Users can strategize on the order of nomination of the players and on the bids posted in each round of the auction. Existence of equilibria in dynamic auctions—subgame perfect equilibria [17]—indeed requires very strong rationality assumptions on agents [14].

In this work we do not attempt to analyze the structure of complex equilibria, which, when they even exist, are implausible to reach. We instead propose bidding strategies that achieve the formation of teams with a guaranteed share of the total value of the player pool, even in the presence of irrational opponents. Another difference from classical dynamic auctions is that users cannot participate in all rounds of the auction. Users pass if their budget is expired or if the role of the nominated player has already been filled on the team.

The imposition of financial constraints such as budgets is also known to alter the properties of even simple standard auctions [6]. For example, the VCG mechanism [19] does not retain all its glorious properties if payments are restricted by a budget [9]. Dynamic auctions with budgets have recently been investigated in the setting of ad auctions for sponsored search [12].

Although bidding and player nomination strategies are at the crux of our work, the reader may note some commonalities with the vast literature on the fair division of goods [4, 16]. Given the symmetric, full-information setting we consider, the best outcome that a strategy played by all team managers can achieve is the proportional share of the total value of the players. Proportionality is also one of the main goals in the fair-division literature. A second commonality is the allotment of an equal budget to all users to play in all rounds of the auction. Competitive equilibrium from equal incomes (CEEI) [5, 13] also assumes the same budget given to each agent to acquire a set of indivisible goods. However,

differently from CEEI, the prices of the goods in fantasy auctions are decided by iterative auction rounds rather than by a centralized pricing mechanism.

We also note the connection to an array of prior work analyzing bidding strategies and autonomous bidding across a diverse spectrum of competitive environments (see, e.g., [8, 11, 18]).

### 3 On the Absence of Equilibria

Given that our setting is one of strategic interaction amongst self-interested agents, at first blush game theory seems the most natural way to approach an analysis and evaluation of bidding strategies. However, there are issues in this approach. For any realistic version of the fantasy auction draft game, computing equilibria is virtually guaranteed to be computationally intractable. Thus, in a single-shot world where the participants (save for the one being served by our bidding proxy) are human rather than computational, it is highly implausible to ascribe predictive power to any given strategy profile merely because it constitutes a Nash equilibrium.

But there is an even more fundamental issue: generally there will *not even exist* any pure strategy subgame perfect equilibria in fantasy draft auctions. Consider the following simple example:

**Example 1.** There are two users with equal budgets; each team roster has two slots; there are four athletes eligible for nomination, two of unequal positive value, and two of value 0.

*Claim 1.* In Example 1, if the lower (nonzero) value athlete is nominated first, there exists no pure strategy Nash equilibrium forward from that point.

*Proof.* First note that in any equilibrium, regardless of what transpires in the first round, in the round in which the other nonzero athlete is nominated each agent must bid as much of his remaining budget as necessary to win (if possible). Then, if someone wins the first athlete for a nonzero price, he will lose the second (high-value) athlete; if he can bid instead to lose the first athlete in order to win the second, this is beneficial.

If bids for the first athlete are 0 and  $\epsilon > 0$ , the first manager has a beneficial deviation in instead bidding  $\delta \in (0, \epsilon)$ , losing the low-value athlete and taking the opposing manager out of contention for the high-value athlete.

If bids for the first athlete are  $\delta > 0$  and  $\epsilon > 0$ , with  $\epsilon > \delta$ , the manager bidding  $\epsilon$  has a beneficial deviation in instead bidding  $\gamma \in (0, \delta)$ , taking the opposing athlete out of contention for the high-value athlete, as above. Similarly, if the two managers bid the same value  $\epsilon > 0$ , they each have positive probability of winning, and thus each has a beneficial deviation in instead bidding  $\gamma \in (0, \epsilon)$ .

If bids for the first athlete are 0 and 0, then the nominating manager gets that athlete, and then can bid  $B$  in the second round and also win the high-value athlete with probability  $1/2$ . The other manager has a beneficial deviation in bidding any  $\epsilon > 0$  first, since he will then take the low-value athlete for free (with certainty), and still win the high-value athlete with probability  $1/2$ .

This exhausts the space of possible bids in round 1, and none are consistent with an equilibrium strategy profile.

Since the low-value athlete being nominated first is a possible path of the full game, and since there is no equilibrium strategy profile forward from that subgame, we have the following corollary:

**Corollary 1.** *In Example 1 there exists no pure-strategy subgame perfect Nash equilibrium of the game as a whole.*

However, if the *high-value* athlete is nominated first, there is a straightforward pure strategy Nash equilibrium of that subgame: both users bid their entire budgets for each nonzero value athlete. Each user will win one nonzero value athlete: with probability 1/2 it will be the high-value athlete; deviating by underbidding in the first round will only ensure that the user gets the low-value athlete (or possibly nothing if he bids 0 in the first round).

Therefore, if we look at the game as a whole, recognizing that athlete nomination is a part of the strategizing, is there an equilibrium? There is, but it is not subgame perfect (which we knew must be the case from Corollary 1), and thus not very satisfying in any kind of predictive sense. Any equilibrium strategy profile must specify what happens off the equilibrium path, and there is no way to do so for the “lower-value athlete nominated first” path that is consistent with rational users (since there is no equilibrium there). Thus, the existence of pure strategy equilibria inherently depends on a model of user behavior that subverts the main rationale for considering equilibrium in the first place, that is, the idea that agents are rational. Other examples may have no equilibria whatsoever.

## 4 Analysis of Fair-Price Bidding

In this section we start our adversarial analysis by considering the case where athlete nomination does not form part of the strategy space. Instead, players are nominated in non-increasing value order, and the strategy space of each manager consists in bidding for these players. Recall that we have a second-price auction, and in the case of (positive bid) ties the athlete is allocated uniformly at random among the managers that placed the highest bid.

Recall that there are  $k$  team managers and  $m$  slots per team, that  $v_i$  is the value of the athlete with  $i^{\text{th}}$  highest value, and that  $V = \frac{1}{k} \sum_{i=1}^{km} v_i$ .

For each athlete  $P_i$  we define his *fair price* as:

$$c_i = v_i \frac{B}{V}.$$

We first consider the following natural nonadaptive strategy, which we call *simple fair-price bidding*. We will show that it may lead to a bad team, but can be made robust via a slight modification, motivating in this way our variant solution. Let  $b_i$  be the current leftover budget before the  $i^{\text{th}}$  athlete is nominated. In simple fair-price bidding the manager always places a bid equal to  $\min\{c_i, b_i\}$ . That is,

$$\text{SIMPLE-FAIR-PRICE-BID}(b_i, v_i) = \min \left\{ \frac{v_i}{V} B, b_i \right\}.$$

We will show that a manager following the simple fair-price bidding strategy could end up with a team of value arbitrarily close to  $V/k$  (which is bad). Consider the case where  $v_1 = \dots = v_{k-1} = V(1 - \varepsilon)$  and  $v_k = \dots = v_{km} = \frac{V(1+k\varepsilon-\varepsilon)}{km-k+1}$ . Our manager bids  $(1-\varepsilon)B$  for the first  $k-1$  athletes; imagine that the other managers bid  $(1-\varepsilon/2)B$ . This means that after  $k-1$  rounds our manager did not win any athletes and has his full budget  $B$  available for bidding, whereas the remaining  $k-1$  managers have each won one athlete of value  $V(1-\varepsilon)$  and have remaining budget  $\varepsilon B/2$ .

We assume that  $\varepsilon$  is chosen to be small enough that

$$\frac{V(1+k\varepsilon-\varepsilon)}{km-k+1} \cdot \frac{B}{V} > \varepsilon \frac{B}{2}.$$

In this case our manager bids the fair value of  $\frac{(1+k\varepsilon-\varepsilon)B}{km-k+1}$  for each of the next  $k$  athletes and wins each of these bids. Therefore, the value of the final team for our manager, who uses the *simple fair-price bidding* strategy in the auction, is

$$\frac{(1+k\varepsilon-\varepsilon)Vm}{km-k+1},$$

which is arbitrarily close to  $V/k$  for large enough  $m$  and small enough  $\varepsilon$ .

This example motivates us to modify the simple fair-price bidding strategy. In the modified fair-price bidding strategy, with parameter  $\alpha \geq 1$ , the manager always places bid  $\min\{\alpha c_i, b_i\}$ . That is,

$$\text{FAIR-PRICE-BID}(b_i, v_i) = \min\left\{\frac{\alpha v_i}{V}B, b_i\right\}.$$

**Theorem 1.** *The expected value of the team generated by FAIR-PRICE-BIDDING with parameter  $\alpha = 3/2$  is at least  $V/3$ , regardless of the other managers' bidding strategies.*

*Proof.* Assume we have  $r \geq 0$  athletes with values  $v_1 \geq \dots \geq v_r \geq V/\alpha$ . Our manager bids her whole budget  $B$  for these athletes. After  $r$  rounds of the auction our manager gets an athlete of value  $v_i \geq V/\alpha$  with probability  $p \geq \min\{1, r/k\}$ . If  $r \geq k$  then our manager always gets one of these high-value athletes and there is nothing more to prove. So now assume that  $r < k$ .

Let  $W = \sum_{i=1}^r v_i$ . We claim that the expected value  $E'$  of the final team selected by our manager conditioned on the fact that she wins one of the high-value athletes is at least  $W/r$ . Let  $p_i \geq \frac{1}{r-i+1}$  be the probability that our manager wins the  $i^{\text{th}}$  round of the auction conditioned on the fact that she did not win any of the previous rounds. Then we can estimate:

$$E' \geq v_1 p_1 + v_2 (1 - p_1) p_2 + \dots + v_r \prod_{i=1}^{r-1} (1 - p_i) p_r \geq \sum_{i=1}^r v_i / r$$

by using the standard majorization inequality (it can be proven easily by induction on  $t$ ) and the fact that:

$$\begin{aligned}
 p_1 + (1 - p_1)p_2 + \dots + \prod_{i=1}^{t-1} (1 - p_i)p_t &= 1 - \prod_{i=1}^t (1 - p_i) \\
 &\geq 1 - \prod_{i=1}^t \left(1 - \frac{1}{r - i + 1}\right) = 1 - \frac{r - t}{r} = \frac{t}{r}.
 \end{aligned}$$

If our manager does not get a high-value athlete (it happens with probability  $1 - p$ ), then the remaining  $k - r$  managers with budgets  $B$  participate in the auction to acquire athletes with total value  $kV - W$ . Also, all the athletes auctioned in this case have values  $v_i < V/\alpha$ . When all the managers have spent their entire budgets or filled all the spots on their rosters, the remaining athletes are distributed at random in a fair way. We consider two cases:

1. There is a moment during the auction when our manager cannot bid the target price of  $\alpha c_j$  for the current athlete  $P_j$ , that is,  $b_j < \alpha v_j \cdot B/V$ . We claim that the value of the athletes on our manager's roster at this moment is already at least  $\frac{V}{2\alpha}$ .

Assume to the contrary that the value of the athletes on our manager's roster at this moment is strictly less than  $\frac{V}{2\alpha}$ . Then  $b_j > B/2$  and we derive that  $v_j > \frac{V}{2\alpha}$ . Therefore, all athletes nominated so far have values higher than  $\frac{V}{2\alpha}$  and our manager did not win any of them (or we are done). Therefore  $b_j = B$ . But recall that  $v_j < V/\alpha$  (the athlete is not one of the  $r$  high-value athletes), which means that our manager can bid the target price of  $\alpha c_j$ , arriving at a contradiction. Therefore, the value of the athletes on our manager's roster is at least  $\frac{V}{2\alpha}$ . Let  $p'$  be the probability that the auction ends in this case.

2. Now take the case where we never reach a situation in which we cannot bid the target price, and our manager fills all the spots on her roster. Assume that  $p''$  is the probability of this case, and note that  $p + p' + p'' = 1$ .

At any moment during the auction our manager either bids the target price of  $\alpha c_j$  for the athlete or does not bid anything because she already filled positions on her roster of the same type as the position type of the athlete who is nominated now. Let  $X$  be the value of the team of our manager at the end of the auction. Consider manager  $u$  who is one of the remaining  $k - r - 1$  managers. We divide the team of manager  $u$  into two sets of athletes  $S^+(u)$  and  $S^-(u)$ , where  $S^+(u)$  is the set of athletes that  $u$  won for price greater than or equal to the target price of  $\alpha c_j$  and  $S^-(u)$  is the set of athletes that  $u$  won for price strictly smaller than the target price.

We claim that  $\sum_{i \in S^+(u)} v_i \leq V/\alpha$  and  $\sum_{i \in S^-(u)} v_i \leq X$ . To show the former, assume that  $\sum_{i \in S^+(u)} v_i > V/\alpha$ . Then  $u$  has paid at least

$$\sum_{i \in S^+(u)} \frac{\alpha v_i}{V} B = \alpha \frac{B}{V} \sum_{i \in S^+(u)} v_i > \alpha \frac{B}{V} \cdot \frac{V}{\alpha} = B,$$

leading to a contradiction, because the initial budget of manager  $u$  is  $B$ . The only way for  $u$  to win a bid for an athlete in  $S^-(u)$  is if our manager (who bids the target value for each athlete she has a remaining slot for on her team) does not place a bid. That could only happen when our manager already filled the spots on its roster corresponding to the position of the athlete who is currently nominated. That means all athletes in  $S^-(u)$  have smaller value than all athletes of the same type on our manager roster, that is,  $\sum_{i \in S^-(u)} v_i \leq X$ .

Analogously, for each manager  $u$  who got one of the first  $r$  athletes, the value of the remaining athletes on his team is at most  $X$ . Therefore,

$$kV \leq kX + (k - r - 1) \frac{V}{\alpha} + W,$$

which entails that

$$X \geq V - \frac{W}{k} - \frac{k - r - 1}{k} \cdot \frac{V}{\alpha}.$$

Overall, the expected value of the final team for our manager is lower-bounded by:

$$\begin{aligned} \frac{W}{r}p + \frac{V}{2\alpha}p' + Xp'' &\geq \frac{W}{r}p + \frac{V}{2\alpha}p' + \left( V - \frac{W}{k} - \frac{k - r - 1}{k} \frac{V}{\alpha} \right) p'' \\ &\geq \frac{W}{r} \cdot \left( p - \frac{r}{k} p'' \right) + \frac{V}{2\alpha}p' + \left( 1 - \frac{k - r}{\alpha k} \right) Vp'' \\ &\geq \frac{V}{\alpha} \cdot \left( p - \frac{r}{k} p'' \right) + \frac{V}{2\alpha}p' + \left( 1 - \frac{k - r}{\alpha k} \right) Vp'' \\ &= \frac{V}{\alpha} \cdot \left( p - \frac{r}{k} p'' \right) + \frac{V}{2\alpha}p' + \left( 1 - \frac{k - r}{k} \right) \frac{V}{\alpha} p'' + \left( 1 - \frac{1}{\alpha} \right) Vp'' \\ &= \frac{V}{\alpha} p + \frac{V}{2\alpha} p' + \left( 1 - \frac{1}{\alpha} \right) Vp'' \geq \frac{V}{3}. \end{aligned}$$

We now give examples that show that our analysis of the algorithm is tight, in that for any choice of  $\alpha$ , given  $k > 2$  managers, there exists an example where our manager collects no more than a third of the fair value.

First off, for any  $\alpha \leq 1$ , the example given at the beginning of this section serves to demonstrate the claim. Now consider any  $\alpha > 3/2$  and the following example: there are two types of athletes,  $k$  managers, and team composition constraints where we need to choose exactly one athlete of type one and  $m - 1$  athletes of type two. The athletes of type one have value  $v_1 = \dots = v_k = \frac{1}{3}V(1 + \varepsilon)$  ( $k$  athletes). There are  $2k - 2$  athletes of type two with value  $v_{k+1} = \dots = v_{3k-2} = \frac{1}{3}V$ . The remaining athletes are also of type two but have value

$$v_{3k-1} = \dots = v_{km} = \frac{kV - kv_1 - (2k - 2)v_{k+1}}{km - k - 2(k - 1)} \approx \frac{2V}{3km},$$

where the approximation is close for large enough  $m$  and small enough  $\varepsilon$ .

During the first bidding round our manager bids  $\min\{\alpha B(1 + \varepsilon)/3, B\} > B(1 + \varepsilon)/2$  for the first athlete. The other managers let our manager win with this bid and force her to pay  $B(1 + \varepsilon)/2$  (by bidding this amount), and they get the remaining  $k - 1$  athletes of the first type for free (one athlete per each manager). During the next  $2k - 2$  bidding rounds our manager bids her whole remaining budget of  $B(1 - \varepsilon)/2$  whereas the other managers bid  $B/2$  and win each of these rounds. After that our manager wins the next  $m - 1$  bidding rounds and obtains a team of value  $v_1 + (m - 1)v_{3k-1} \approx V/3$ .

We now consider the case where  $\alpha \leq 3/2$ . Consider the example where there is only one type of athlete, there are  $k$  managers, and athletes have the following values:  $v_1 = \dots = v_{k-1} = \frac{2}{3}V(1 - \varepsilon)$  and  $v_k = \dots = v_{km} = \frac{kV - (k-1)v_1}{km - k + 1}$ . Our manager bids  $2\alpha B(1 - \varepsilon)/3 < B$  for each of the first  $k - 1$  athletes and loses all of those rounds to other managers who bid their whole budgets  $B$ . After that our manager wins  $m$  bids and ends up with a team of value

$$m \cdot \frac{kV - (k - 1)v_1}{km - k + 1},$$

which is arbitrarily close to  $V/3$  for large enough  $m$  and small enough  $\varepsilon$ .

## 5 Arbitrary Nomination Order

We now consider the more general (and realistic) setting in which each manager is allowed to nominate an arbitrary athlete when his turn comes up, with the order of nominators pre-determined arbitrarily.

### 5.1 Algorithm Description

Our algorithm for this expanded version of the problem, which we call SELECTIVE-FAIR-BIDDING, will depend on three parameters:  $\alpha$ ,  $\beta$ , and  $\gamma$ , with  $\beta > \alpha \geq 1$  and  $\gamma \geq 1$ . We define two groups of athletes:  $L = \{i : v_i \geq V/\beta\}$  and  $S = \{i : v_i < V/\beta\}$ .

If  $\sum_{i \in L} v_i \geq \sum_{i \in S} v_i$ , our manager ignores (bids 0 for) all athletes with indices in the set  $S$  and only bids for athletes with indices in  $L$ . The bid value is always the whole budget  $B$ .

Now we describe the strategy in the case that  $\sum_{i \in L} v_i < \sum_{i \in S} v_i$ . Letting  $t$  be the number of distinct types of athletes that each team must have, let  $d_j$  be the number of athletes of type  $j = 1, \dots, t$  the team must have, with  $\sum_{j=1}^t d_j = m$ . All athletes in the set  $S$  are partitioned into  $t$  sets  $S_1, \dots, S_t$  by type. Let

$$A_j = \frac{\sum_{i \in S_j} v_i}{kd_j}$$

be the average value for the athlete of type  $j = 1, \dots, t$  (among the athletes in group  $S$ ) per available spot on the team roster.

Our manager only bids for an athlete  $P_i$  if either  $i \in L$ , or  $i \in S_j$  and  $v_i \geq A_j/\gamma$  for some  $j \in \{1, \dots, t\}$ . We will call such athletes *desirable*. The bid value for desirable athletes is  $\min\{\alpha c_i, b_t\}$  where  $b_t$  is the remaining budget in the current period  $t$  of the bidding procedure and  $c_i$  is the fair value of athlete  $i$ . The athletes from the set  $S$  of value less than  $A_j/\gamma$  will be called *undesirable*, and our manager bids 0 for all of them.

Our manager always nominates the highest-value available athlete that can fill one of the positions on her roster.

### 5.2 Analysis

**Theorem 2.** *The expected value of the team generated by SELECTIVE-FAIR-BIDDING with  $\alpha = 16/3$ ,  $\beta = 8$ , and  $\gamma = 2$  is at least  $V/16$ .*

*Proof Sketch.* The proof can be divided into two cases, the second of which we only sketch here due to limitations on space (the proof is otherwise complete).

**Case 1.** Assume that  $\sum_{i \in L} v_i \geq \sum_{i \in S} v_i$ . Recall that our manager bids her whole budget  $B$  for each athlete in  $L$  until she either wins one or there are no more athletes in  $L$  left. There are  $|L|$  rounds of the auction that are relevant to our manager (as she bids zero in the others); let  $p_j \geq 1/k$  be the probability that our manager wins an athlete during the  $j$ th round our manager bids in, for  $j = 1, \dots, |L|$ , conditioned on the fact that she did not win an athlete during the previous rounds. And let  $v_{L(j)}$  be the value of the athlete nominated in the  $j$ th round our manager bids in. Then the expected value of the athletes won by our manager is at least

$$E' = v_{L(1)}p_1 + v_{L(2)}(1 - p_1)p_2 + \dots + v_{L(|L|)} \prod_{i=1}^{|L|-1} (1 - p_i)p_{|L|}.$$

If  $\prod_{i=1}^{|L|} (1 - p_i) \leq 1/2$ , then

$$E' \geq \left(1 - \prod_{i=1}^{|L|} (1 - p_i)\right) \min_i v_{L(i)} \geq \frac{V}{2\beta}.$$

Otherwise, that is, if  $\prod_{i=1}^{|L|} (1 - p_i) > 1/2$ , then

$$E' > \sum_{i \in L} \frac{v_{L(i)}p_i}{2} \geq \sum_{i \in L} \frac{v_{L(i)}}{2k} \geq \sum_{i=1}^n \frac{v_{L(i)}}{4k} = \frac{V}{4}.$$

**Case 2.** Assume that  $\sum_{i \in L} v_i < \sum_{i \in S} v_i$ . This is the more difficult of the two cases, and due to space constraints we must omit much of the analysis. However, it can be established that in this case:

$$E' \geq V \left( \frac{1}{2} \min \left\{ \frac{1}{\gamma}, 1 - \frac{1}{\gamma} \right\} - \frac{1}{\alpha} \right),$$



If we choose  $\gamma = 2$ , combining bounds for the various cases, the expected value of our manager's team at the end of the auction is at least:

$$V \cdot \min \left\{ \frac{1}{4} - \frac{1}{\alpha}, \frac{1}{\alpha} - \frac{1}{\beta}, \frac{1}{4}, \frac{1}{2\beta} \right\}.$$

Now choosing  $\alpha = 16/3$  and  $\beta = 8$ , we derive the lower bound of  $V/16$  on the expected value of the final team.

## 6 Conclusion

In this paper we initiated the study of fantasy auction drafts, which play an important role in the large and growing market of fantasy sports. We abstracted the problem by defining a simple but realistic auction, which models the real-life process. We studied pure-strategy Nash equilibria and showed that even for two players, there are no pure-strategy subgame perfect Nash equilibria. We thus turned our attention to worst-case outcomes and designed a deterministic algorithm for bidding in fantasy auction drafts, which guarantees the creation of a team with a total value that in expectation (over the random choices of the allocation mechanism in case of ties in the highest bids) is at least a constant ( $1/16$ ) approximation of the optimal possible.

Throughout the paper we assumed that for each athlete  $P_i$  there exists a universal value  $v_i$ ; yet our results are more general. Our worst-case bounds hold also for the setting where each manager  $u_j$  can have an idiosyncratic value  $v_{i,j}$  for each athlete  $P_i$ . Our algorithm can guarantee the creation of a team that has a value of at least  $V_j/16$ , where

$$V_j = \frac{1}{k} \sum_{i=1}^{km} v_{i,j}.$$

Consideration of equilibrium outcomes was a nonstarter for our purposes quite apart from concerns about collusion and the like. However, one can show rather easily that if other managers do collude, it can have an adverse effect on the quality of the team our manager ends up with. Yet our algorithm is *collusion resistant* with respect to the worst case: even if other managers collude, the guarantees remain unchanged.

The main open theoretical question regards closing the gap between what we guarantee with our strategy and what is *possible* to guarantee. We have given an algorithm that provides a constant approximation ( $1/16$ ) to the best one can hope for. Can the constant be improved? Can one show any upper bound?

Finally, there are empirical angles that we hope will be pursued. The fact that the worst-case guarantee of the algorithm we propose is  $V/16$  does not at all indicate that it will not be competitive in practice. If such bidding strategies are ultimately implemented via auto-draft proxies in actual fantasy auctions, this will yield an intriguing dataset that can be studied and perhaps used as the basis for an empirically grounded iteration on the work we initiated here.

## References

1. FSTA industry demographics. <http://fsta.org/research/industry-demographics/>
2. Beller, M., Fantasy football 2015 draft preview: auction strategies. Sports Illustrated (2015)
3. Bergemann, D., Said, M.: Dynamic Auctions: A Survey. Cowles Foundation Discussion Papers 1757R, Cowles Foundation for Research in Economics, Yale University, March 2010
4. Brams, S.J., Taylor, A.D., Division, F.: From Cake-Cutting to Dispute Resolution. Cambridge University Press, Cambridge (1996)
5. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium. *J. Polit. Econ.* **119**, 1061–1103 (2011)
6. Che, Y.-K., Gale, I.: Standard auctions with financially constrained bidders. *Rev. Econ. Stud.* **65**, 1–21 (1998)
7. Cockcroft, T.H.: Fantasy football auction strategy. [http://espn.go.com/fantasy/football/story/\\_/id/13084199/fantasy-football-auction-strategy-fantasy-football](http://espn.go.com/fantasy/football/story/_/id/13084199/fantasy-football-auction-strategy-fantasy-football). Accessed 18 August 2015
8. Csirik, J.A., Littman, M.L., Singh, S., Stone, P.: FAucS: an FCC spectrum auction simulator for autonomous bidding agents. In: Proceedings of the Second International Workshop on Electronic Commerce, pp. 139–151 (2001)
9. Dobzinski, S., Lavi, R., Nisan, N.: Multi-unit auctions with budget limits. In: FOCS, pp. 260–269 (2008)
10. Goff, B.: The \$70 billion fantasy football market. *Forbes* (2013)
11. Greenwald, A., Boyan, J.: Bidding under uncertainty: theory and experiments. In: Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence, pp. 209–216 (2004)
12. Gummadi, R., Key, P.B., Proutiere, A.: Optimal bidding strategies in dynamic auctions with budget constraints. In: 49th Annual Allerton Conference on Communication, Control, and Computing, Allerton Park and Retreat Center, Monticello, IL, USA, 28–30 September 2011, p. 588 (2011)
13. Hylland, A., Zeckhauser, R.: Efficient allocation of individuals to positions. *J. Polit. Econ.* **87**, 293–314 (1979)
14. Iyer, K., Johari, R., Sundararajan, M.: Mean field equilibria of dynamic auctions with learning. In: Proceedings of the 12th ACM Conference on Electronic Commerce, EC 2011. ACM, New York, pp. 339–340 (2011)
15. Milgrom, P.R.: Putting Auction Theory to Work. Cambridge University Press, Cambridge (2004)
16. Moulin, H.: Fair Division and Collective Welfare. MIT Press, Cambridge (2004)
17. Osborne, M.J.: An Introduction to Game Theory. Oxford University Press, New York (2003)
18. Reeves, D.M., Wellman, M.P., MacKie-Mason, J.K., Osepayshvili, A.: Exploring bidding strategies for market-based scheduling. *Decis. Support Syst.* **39**, 67–85 (2005)
19. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Financ.* **16**, 8–37 (1961)

# Competitive Equilibria for Non-quasilinear Bidders in Combinatorial Auctions

Rad Niazadeh<sup>1</sup>(✉) and Christopher A. Wilkens<sup>2</sup>

<sup>1</sup> Department of Computer Science, Cornell University, Ithaca, USA  
rad@cs.cornell.edu

<sup>2</sup> Yahoo Research, Santa Clara, USA  
cwilkens@yahoo-inc.com

**Abstract.** quasilinearity is a ubiquitous and questionable assumption in the standard study of Walrasian equilibria. Quasilinearity implies that a buyer’s value for goods purchased in a Walrasian equilibrium is always additive with goods purchased with unspent money. It is a particularly suspect assumption in combinatorial auctions, where buyers’ complex preferences over goods would naturally extend beyond the items obtained in the Walrasian equilibrium.

We study Walrasian equilibria in combinatorial auctions when quasilinearity is not assumed. We show that existence can be reduced to an Arrow-Debreu style market with one divisible good and many indivisible goods, and that a “fractional” Walrasian equilibrium always exists. We also show that standard integral Walrasian equilibria are related to integral solutions of an induced configuration LP associated with a fractional Walrasian equilibrium, generalizing known results for both quasilinear and non-quasilinear settings.

## 1 Introduction

Money is inherently useless; it only holds value because of the promise that it can be used to buy something useful. Thus, an agent’s utility for money will depend substantially on what she already has and the alternative ways that it can be spent. A student who saves money on her habitual cup of coffee can spend it on many things. On the other hand, a corporate event planner who is given a dedicated budget for drinks may not receive any benefit for discounted coffee; on the contrary, she cannot spend the money on herself, and her budget may get cut next time if she doesn’t spend enough on the current event.

General market equilibria capture the ephemerality of money. Arrow and Debreu’s exchange model is simple: agents have goods; they sell those goods, then buy what they want most. In this setting, money has no inherent value and is simply a lubricant facilitating exchange. This works because Arrow and Debreu capture the entire economy, so there is nothing outside the market on which to spend money.

---

R. Niazadeh—Supported by Google PhD Fellowship. This research was mostly done when the first author was doing an internship at Yahoo Research.

In contrast, Walrasian (competitive) equilibria only capture a slice of the overall market. To do so, they must attribute utility to unspent money, since it can be spent elsewhere. In a Walrasian equilibrium model, agents arrive with money and only spend it if the received value outweighs the cost.

The default way to capture the implicit value of money is through a quasilinear utility function, i.e. a bidder's utility  $u$  is the difference between her value  $v$  and the price she pays  $p$ . Quasilinearity assumes that an agent's total value is additive in what she gets now and what she gets from an outside option, and that the amount of utility she gets from an outside option scales linearly with the amount of money she applies to it. Both are plausible modeling assumptions, but assuming that they are always true is only somewhat more defensible than assuming that a bidder always has an additive valuation over all items in the market.

We can illustrate one simple violation with our coffee example. Suppose our student will buy exactly one cup of coffee each day. If she doesn't buy coffee now, she will spend \$3 on coffee elsewhere, so her value for coffee now is \$3; however, if we give her *both* a cup of coffee *and* \$3, she will spend the extra \$3 on something completely different, like a movie ticket (she already has her cup of coffee). Whether this movie ticket gives her the same utility as a cup of coffee, or half the utility, or a quarter of the utility — there is nothing in our market model to imply that her utility from an extra \$3 is in any way tied to her utility for a cup of coffee, except that it is plausibly *at most* her utility for a cup of coffee.<sup>1</sup> In effect, the student's marginal utility for \$3 is completely different depending on whether or not she gets coffee now.

Relaxing quasilinearity for Walrasian equilibria is thus a fundamental question, particularly in combinatorial auction domains where bidders are assumed to have complex preferences over sets of items. This is the topic of our paper. Existing relaxations of quasilinearity focus on unit demand agents, and the literature is quite limited. The existence of Walrasian equilibria was first shown by Quinzii (1984). Later, Alaei et al. (2011) show that Walrasian equilibria exist using a direct approach and have the same kind of lattice structure as standard Walrasian equilibria for unit demand bidders. Maskin (1987) studies a superficially different problem and adds a single divisible good (i.e. money) to a standard general market model with indivisible goods; his result is that market equilibria always exist. We study Walrasian equilibria with general utilities in an even more obvious setting: combinatorial auctions. In a combinatorial auction, bidders may have complex preferences over the multiple goods being sold. Thus it is natural that one should extend these complex preferences to items outside the Walrasian micromarket by allowing non-quasilinear preferences for money. Our main results establish conditions under which a Walrasian equilibrium exists.

---

<sup>1</sup> We know that she chose a cup of coffee over a movie ticket initially, so that implies her value for a cup of coffee is less than her value for a movie ticket. On the other hand, there might also be complementarities here if the student is unable to enjoy the movie without first having a cup of coffee...

## Fractional Walrasian Equilibria and Market Equilibria

Walrasian equilibria capture a microcosm of a larger market. To do so, they capture both an agent’s utility for unspent money and goods’ inherent indivisibility at small scales. As our prior discussion suggests, money is simply a proxy for the portion of the general market outside the goods available in the Walrasian equilibrium. It is therefore not surprising that the first step in our work constructs a reduction from a Walrasian equilibrium problem to a general Arrow-Debreu market with an extra good (money), and an extra agent (the seller). Together, the extra good and agent capture the market outside the Walrasian equilibrium. This market is special because all goods except money have supply 1 and are indivisible. Maskin (1987) studies a similar market without adding the seller as an agent. Our first result shows that the equilibria are the same:

**Lemma 1 (Informal).** *A set of prices and allocations for a combinatorial auction are a Walrasian equilibrium if and only if they correspond to a market equilibrium of the associated special Arrow-Debreu market.*

The relationship with the special Arrow-Debreu market also lays the foundation for our first main result — “fractional” Walrasian equilibria always exist:

**Theorem 1 (Informal).** *In a combinatorial auction setting, a fractional Walrasian equilibrium always exists, that is there exists a set of prices and, for each player, a distribution over sets of goods whose support is the collection of all demanded sets under the equilibrium prices, and market clears in expectation.*

This follows by proving that an auction’s associated general market model always has a fixed point, and that demand of a bidder at the fixed point can always be decomposed into a distribution over goods.

## Configuration LPs and True Walrasian Equilibria

To understand true (non-fractional) Walrasian equilibria, we must understand the combinatorial auction’s *configuration linear program*. In a quasilinear setting, the configuration LP captures the way goods can be (fractionally) assigned to bidders; the LP’s objective is the total value generated by the assignment. Walrasian equilibria here are known to be equivalent to integral optima of an auction’s configuration LP.

Unfortunately, the configuration LP is not available without quasilinearity because a bidder’s value is not well-defined. Instead, we introduce an *induced configuration LP* that is associated with a particular price vector  $p^*$ . This LP is constructed by fixing bidders’ utilities at  $p^*$  and assuming they are otherwise quasilinear. Integral optima are again related to Walrasian equilibria, but only if  $p^*$  already supported a fractional Walrasian equilibrium:

**Theorem 2 (Informal).** *A Walrasian equilibrium for a combinatorial auction exists if and only if there is a price vector  $p^*$  supporting a fractional Walrasian equilibrium for which the induced configuration LP has an integral optimum.*

This theorem has a couple of interesting corollaries. First, we can see how it relates to the results of Maskin (1987), Quinzii (1984), and Alaei et al. (2011):

**Corollary 1.** *If the induced configuration LP is always integral, a combinatorial auction always has a Walrasian equilibrium.*

In unit-demand settings, the induced configuration LP is a matching LP at any fixed price  $p^*$ . Thus, it is integral and always has an integral optimum, and Walrasian equilibria will always exist. This generalizes the results of Maskin, Quinzii and Alaei et al. — Maskin studied a sibling of our general market model and effectively showed that when bidders are unit demand, there is always an integral solution for every fixed point, while Alaei et al. directly show that Walrasian equilibria exist for unit demand settings.

Another simple but useful corollary happens when the configuration LP is independent of  $p^*$ :

**Corollary 2.** *If the induced configuration LP is independent of  $p^*$ , then a combinatorial auction has a Walrasian equilibrium if and only if the configuration LP has an integral optimum, regardless of any properties of  $p^*$ .*

The dependence on  $p^*$  cancels when utilities are quasilinear, and it explains why we can simply talk about the configuration LP without talking about a specific set of prices  $p^*$ .

Together, these results build a picture of the equilibrium landscape outside quasilinearity — Walrasian equilibria still exist at least in a fractional form, but testing existence in general is substantially more complicated. Existence is still related to a configuration LP, but that LP can only be defined once prices  $p^*$  are in hand.

**Related Work.** The objective of this paper is establishing existence characterization for competitive equilibria in combinatorial auction. The problem is closely related to the existing literature in economics and theoretical computer science from different directions. First, there is a prominent literature on competitive equilibrium in combinatorial auctions for the quasilinear setting. These papers characterize existence conditions for a Walrasian equilibrium and provide practical necessary conditions for a competitive equilibrium to exist in the quasilinear setting such as gross substitute, e.g. Gul and Stacchetti (1999), Bikhchandani and Mamer (1997), Bevia et al. (1999), Murota and Tamura (2001). People have also thought about using the properties of competitive equilibria in quasilinear settings when valuations are satisfying the gross substitute, such as lattice structure (Gul and Stacchetti 1999), in order to design ascending auctions and identifying the connections between the well-known VCG mechanism (Vickrey 1961; Groves 1973; Clarke 1971) and Walrasian equilibria (Kelso Jr and Crawford 1982; Cramton et al. 2006; Nisan et al. 2007). Closely related is the literature on assignment games and core allocations, in which they tried to generalize the stable matching concept (either one-to-one matching, one-to-many matching, or many-to-many matching) to two-sided markets with indivisible goods and quasilinear utilities, e.g. Shapley and Shubik (1971) and Echenique et al. (2004).

The second direction that connects our work to the literature is the existing work on non-quasilinear utilities (or non-transferable utilities) and two-sided matching markets or general Arrow-Debreu market (Arrow and Debreu, 1954)

with only one divisible good and unit-demand agents. The problem formulation is as introduced by Demange and Gale (1985). The existence of competitive equilibria for non-quasilinear utilities was first proved by Quinzii (1984), Gale (1984), Svensson (1984), and later by Kaneko and Yamamoto (1986). They showed if there is a single divisible good (say money) in an economy and if agents are unit-demand then there still exists a competitive equilibrium under certain reasonable (monotonicity) conditions. There is also the work of Maskin (1987) on fair allocation of indivisible goods with money, that provides a simpler proof for the existence of the equilibrium with indivisible goods and only one divisible item in the unit-demand case. More recently, there has been a work on two-sided matching markets with non-transferable utilities by Alaei et al. (2011) that followed a different combinatorial approach. They showed the existence of competitive equilibrium when utility functions are monotone, and generalized the lattice structure and properties associated with the minimum lattice point to a general non-quasilinear setting.

## 2 Settings and Notations

We are looking at a *combinatorial auction*, in which we have a set  $\mathcal{I}$  of  $m$  items and a set  $\mathcal{B}$  of  $n$  buyers interested in these items. For every  $x \in \{0, 1\}^m$  and price  $p \in \mathbb{R}_+$ , let  $u_i(x, p)$  be the utility of bidder  $i$  if she gets bundle  $X = \{j \in \mathcal{I} : x_j = 1\}$  of items under the price  $p$ . We assume utilities are strictly decreasing and continuous with respect to the price  $p$ , and increasing with respect to  $x_j$  for every item  $j$ . The competitive equilibrium (also known as *Walrasian equilibrium*) can be defined in this setting as follows.

**Definition 1.** A *Walrasian Equilibrium (WE)* is a pair of allocation and prices  $(\{x^{(i)}\}_{i=1}^n, \{p_j\}_{j=1}^m)$  that satisfies the following conditions:

- $\forall i \in \mathcal{B}$  and  $j \in \mathcal{I} : x^{(i)} \in \{0, 1\}^m$  and  $p_j \in \mathbb{R}_+$ .
- **[Feasibility]**  $\forall j \in \mathcal{I} : \sum_{i=1}^n x_j^{(i)} \leq 1$ .
- **[Satisfaction]**  $\forall i \in \mathcal{B} : x^{(i)} \in \operatorname{argmax}_{x' \in \{0, 1\}^m} u_i(x', \sum_{j=1}^m p_j x'_j)$ ,
- **[Market clearance]**  $\forall j \in \mathcal{I}$ , if  $p_j > 0$  then  $\sum_{i=1}^n x_j^{(i)} = 1$ .

Besides WE, we also need to define a fractional equilibrium, in which each buyer has a distribution over bundles of items. However, such an allocation is only feasible in expectation, meaning that each item gets allocated with probability less than or equal to 1. Note that such an equilibrium cannot be realized in reality and it is just a solution concept that will shed insight on the structure of WE, as we show later in this paper. More precisely, we have the following definition:

**Definition 2.** In a combinatorial auction  $(\mathcal{I}, \mathcal{B}, \{u_i(\cdot)\}_{i \in \mathcal{B}})$ , a fractional WE is defined to be a pair of allocation and prices  $(\{x_{i,S}\}, \{p_j\})$  such that

- $\forall (i, S) \in \mathcal{B} \times 2^{\mathcal{I}}$  and  $j \in \mathcal{I}$ :  $x_{i,S} \in \mathbb{R}_+$  and  $p_j \in \mathbb{R}_+$ .
- **[Feasibility]**  $\forall j \in \mathcal{I}$ :  $\sum_{i \in \mathcal{B}, S: j \in S} x_{i,S} \leq 1$ ,  $\forall i \in \mathcal{B}$ :  $\sum_{S \subseteq \mathcal{I}} x_{i,S} = 1$ .
- **[Satisfaction]**  $\forall i \in \mathcal{B}$ , if  $x_{i,S} > 0$  then  $\mathbf{1}_S \in \operatorname{argmax}_{x' \in \{0,1\}^m} u_i(x', \sum_{j=1}^m p_j x'_j)$ .
- **[Market clearance]**  $\forall j \in \mathcal{I}$ , if  $p_j > 0$  then  $\sum_{i \in \mathcal{B}, S: j \in S} x_{i,S} = 1$ .

In this paper, we also talk about a special case of *Arrow-Debreu markets* in which all commodities except one are indivisible. In such markets, there is a set  $\mathcal{A}$  of  $N$  of agents in the market who are interested in trading a set  $\mathcal{C}$  of  $M$  commodities. We assume all commodities are indivisible except the last commodity  $j = M$  (we sometimes call this commodity ‘money’). Each agent  $i$  will bring an endowment  $w^{(i)} \in \mathbb{R}_+^M$  of commodities to the market to trade. Agent  $i$  gets a utility of  $\tilde{u}_i(x)$  for an allocation  $x \in \{0, 1\}^{M-1} \times \mathbb{R}_+$  of commodities. Moreover, we assume  $\tilde{u}_i(x)$  is strictly increasing with respect to  $x_j$  for all  $j \in \mathcal{C}$  and continuous with respect to allocation of money, i.e.  $x_M$ . We next define the *general market equilibrium* for such a market.

**Definition 3.** A *General Market Equilibrium (GME)* is a pair of allocation and prices  $(\{x^{(i)}\}_{i=1}^N, \{p_j\}_{j=1}^M)$  that satisfies the following conditions:

- $\forall i \in \mathcal{A}$  and  $j \in \mathcal{C} \setminus \{M\}$ :  $x_j^{(i)} \in \{0, 1\}$  and  $p_j \in \mathbb{R}_+$ .
- $\forall i \in \mathcal{A}$ :  $x_M^{(i)} \in \mathbb{R}_+$  and  $P_M \in \mathbb{R}_+$ .
- **[Satisfaction]**  $\forall i \in \mathcal{A}$ :  $x^{(i)} \in \operatorname{argmax}_{j \neq M: x'_j \in \{0,1\}, x'_M \in \mathbb{R}_+} \tilde{u}_i(x') \text{ s.t. } \sum_{j=1}^M x'_j p_j \leq \sum_{j=1}^M w_j^{(i)} p_j$ .
- **[Market clearance]**  $\forall j \in \mathcal{C}$ , if  $p_j > 0$  then  $\sum_{i=1}^N x_j^{(i)} = \sum_{i=1}^N w_j^{(i)}$ .

To define more notations for an Arrow-Debreu market with only one divisible good, let  $D_i(\{p_j\}_{j \in \mathcal{C}})$  be the collection of feasible allocation of commodities to agent  $i$ , such that each maximizes utility of agent  $i$  under prices  $\{p_j\}_{j \in \mathcal{C}}$  and they satisfy the budget constraint of agent  $i$ . In other words:

$$D_i(\{p_j\}_{j \in \mathcal{C}}) \triangleq \operatorname{argmax}_{j \neq M: x'_j \in \{0,1\}, x'_M \in \mathbb{R}_+} \tilde{u}_i(x') \text{ s.t. } \sum_{j=1}^M x'_j p_j \leq \sum_{j=1}^M w_j^{(i)} p_j \quad (1)$$

Let  $\bar{D}_i(\{p_j\}_{j \in \mathcal{C}})$  be all vectors in  $D_i(\{p_j\}_{j \in \mathcal{C}})$  when we delete the allocation of the divisible good, i.e. last coordinate, from all vectors. Define total demand to be  $D(\{p_j\}_{j \in \mathcal{C}}) \triangleq \sum_{i \in \mathcal{A}} D_i(\{p_j\}_{j \in \mathcal{C}})$  and the total demand for items to be  $\bar{D}(\{p_j\}_{j \in \mathcal{C}}) \triangleq \sum_{i \in \mathcal{A}} \bar{D}_i(\{p_j\}_{j \in \mathcal{C}})$ , where summations are Minkowski summations of sets. Moreover, let  $\tilde{D}_i = \operatorname{Conv}(D_i)$  and  $\tilde{D} = \operatorname{Conv}(D)$  where  $\operatorname{Conv}(\cdot)$  is the convex hull of its argument. Clearly, all these sets are finite (because utility



is strictly increasing in money and hence given an allocation of indivisible items the allocation of money will be the unique number that fills the budget slack) and hence convex hulls are well defined. Also, from the definition of convex hull and Minkowski summation,  $\tilde{D} = \sum_{i \in \mathcal{A}} \tilde{D}_i$ .

### 3 Reduction from Combinatorial Auction to Arrow-Debreu Market

We start by defining a general Arrow-Debreu market with one divisible good.

**Definition 4.** *Given a combinatorial auction  $(\mathcal{I}, \mathcal{B}, \{u_i(\cdot)\}_{i \in \mathcal{B}})$ , its corresponding Arrow-Debreu market  $(\mathcal{C}, \mathcal{A}, \{\tilde{u}_i(\cdot)\}_{i \in \mathcal{A}}, \{w^{(i)}\}_{i \in \mathcal{A}})$  is the following:*

- There are  $M = m + 1$  commodities, where the first  $m$  indivisible commodities in  $\mathcal{C}$  are items in  $\mathcal{I}$ , and the last divisible commodity is a special commodity called ‘money’.
- There are  $N = n + 1$  agents:
  - For every  $i \in [n]$  we have  $\tilde{u}_i(x, y) = u_i(x, \frac{Z}{n} - y)$  for every  $x \in \{0, 1\}^m, y \in \mathbb{R}_+$ , where  $Z$  is large enough such that  $Z > \sum_{i=1}^n u_i(\mathbf{1}, 0)$ ,
  - The last agent is a special agent called the ‘seller’ and her utility is computed as  $\tilde{u}_{n+1}(x, y) = y$  for every  $x \in \{0, 1\}^m, y \in \mathbb{R}_+$ .
- For every  $i \in [n]$ , endowment of agent  $i$  is  $w^{(i)} = (0, \dots, 0, \frac{Z}{n})$ . For the seller,  $w_{n+1} = (1, \dots, 1, 0)$ .

We now have the following lemma, which basically shows that the correspondence described in Definition 4 preserves the equilibrium. The proof is provided in the online version of this paper (Niazadeh and Wilkens 2016).

**Lemma 2.** *A pair  $(\{x^{(i)}\}_{i=1}^n, \{p_j\}_{j=1}^m)$  is a WE for the combinatorial auction  $(\mathcal{I}, \mathcal{B}, \{u_i(\cdot)\}_{i \in \mathcal{B}})$  if and only if there exists a GME  $(\{\tilde{x}^{(i)}\}_{i=1}^N, \{\tilde{p}_j\}_{j=1}^M)$  for its corresponding Arrow-Debreu market denoted by  $(\mathcal{C}, \mathcal{A}, \{\tilde{u}_i(\cdot)\}_{i \in \mathcal{A}}, \{w^{(i)}\}_{i \in \mathcal{A}})$  as in Definition 4 such that  $\forall (i, j) \in [N - 1] \times [M - 1] : \tilde{x}_j^{(i)} = x_j^{(i)}, \tilde{x}_M^{(N)} = \sum_{j=1}^m p_j, \forall j \in [M - 1] : \frac{\tilde{p}_j}{\tilde{p}_M} = p_j$  and  $\tilde{p}_M > 0$ .*

### 4 A Generalization of Configuration LP

In this section, we start exploring the connections between welfare maximization and WE for non-quasilinear utilities in combinatorial auctions. In the quasilinear world, where  $u_i(x, p) = v_i(x) - p$ , the following connections are known:

- The combinatorial auction configuration LP, i.e. the following linear program

$$\begin{aligned}
 & \text{maximize} && \sum_{i \in \mathcal{B}, S \subseteq \mathcal{I}} x_{i,S} v_i(\mathbf{1}_S) \\
 & \text{subject to} && \sum_{S \subseteq \mathcal{I}} x_{i,S} \leq 1, && i \in \mathcal{B}. \\
 & && \sum_{i \in \mathcal{B}} \sum_{S \subseteq \mathcal{I}: j \in S} x_{i,S} \leq 1, && j \in \mathcal{I}. \\
 & && x_{i,S} \geq 0, && i \in \mathcal{B}, S \subseteq \mathcal{I}.
 \end{aligned}$$

that characterizes maximum welfare allocations, has an integral optimal solution if and only if WE exists.

- A vector of prices is a WE price vector if it forms an optimal solution to the dual of the configuration LP. Moreover, if a dual solution can be supported by an integral feasible primal, then it is a WE price vector.

The question we address here is how can one generalize these concepts to the case of non-quasilinear utilities, in the hope that they shed some insights on existence and structural properties of WE for non-quasilinear utilities. To this end, we first define an *equivalent quasilinear value function* for each bidder, which will act similar to the value function in the quasilinear configuration LP.

**Definition 5.** Fix a vector of prices  $p^*$ . For bidder  $i$  with utility function  $u_i(x, p)$ , the equivalent quasilinear value function at price vector  $p^*$  is defined as

$$v_i^{(p^*)}(x) \triangleq u_i(x, \sum_{j=1}^m x_j p_j^*) + \sum_{j=1}^m p_j^* \quad (2)$$

As it can be seen from the definition, the equivalent quasi-linear value function, together with prices  $p^*$ , will generate the same utility as the original utility function, if we assume quasi-linearity. Given the definition of an equivalent quasilinear value function for each bidder at a fixed price vector  $p^*$ , here is a natural generalization to the configuration LP. The program maximizes welfare with respect to the equivalent quasilinear value function.

**Definition 6 (Induced configuration LP at price  $p^*$ ).** Fixing a price vector  $p^*$ , the induced configuration LP at price  $p^*$  is defined as the following linear program with variables  $\{x_{i,S}\}_{i \in \mathcal{B}, S \subseteq \mathcal{I}}$  (allocation):

$$\begin{aligned} & \text{maximize} && \sum_{i \in \mathcal{B}, S \subseteq \mathcal{I}} x_{i,S} v_i^{(p^*)}(\mathbf{1}_S) \\ & \text{subject to} && \sum_{S \subseteq \mathcal{I}} x_{i,S} \leq 1, && i \in \mathcal{B}. \\ & && \sum_{i \in \mathcal{B}} \sum_{S \subseteq \mathcal{I}: j \in S} x_{i,S} \leq 1, && j \in \mathcal{I}. \\ & && x_{i,S} \geq 0, && i \in \mathcal{B}, S \subseteq \mathcal{I}. \end{aligned}$$

Similar to the quasilinear utilities, one can look at the dual program of the linear program in Definition 6 which sheds more insights on the structure of the WE, as we show later in this paper.

**Definition 7 (Dual induced configuration LP at price  $p^*$ ).** Fixing a price vector  $p^*$ , the dual of the induced configuration LP in Definition 6 is the following linear program with variable  $\{u_i\}_{i \in \mathcal{B}}$  (utilities) and  $\{p_j\}_{j \in \mathcal{I}}$  (prices).

$$\begin{aligned} & \text{minimize} && \sum_{i \in \mathcal{B}} u_i + \sum_{j \in \mathcal{I}} p_j \\ & \text{subject to} && \sum_{j \in S} p_j + u_i \geq v_i^{(p^*)}(\mathbf{1}_S), && i \in \mathcal{B}, S \subseteq \mathcal{I}. \\ & && u_i \geq 0, p_j \geq 0, && i \in \mathcal{B}, j \in \mathcal{I}. \end{aligned}$$

In the next section we show how the linear programs in Definitions 6 and 7 are related to the existence of WE in non-quasilinear settings.

## 5 Main Results and Their Applications

Our main result is proving the existence of WE under necessary and sufficient structural conditions, and bridging the gap between the concept of WE and configuration LP for non-quasilinear utilities. More accurately, we show the induced configuration LP in Definition 6 is strong enough to provide us with necessary and sufficient conditions for the existence of equilibrium, however we have to look at this LP when  $p^*$  is also an equilibrium price vector. Using the reduction in Sect. 4 to general markets, we show such item prices always exist. Moreover, it turned out that by using the primal-dual LP machinery one can show such prices will get supported by an integral allocation to form a WE if and only if the corresponding induced configuration LP has an integral optimal solution.

**Definition 8.** Fix a combinatorial auction  $(\mathcal{I}, \mathcal{B}, \{u_i(\cdot)\}_{i \in \mathcal{B}})$  and consider its corresponding Arrow-Debreu market  $(\mathcal{C}, \mathcal{A}, \{\tilde{u}_i(\cdot)\}_{i \in \mathcal{A}}, \{w^{(i)}\}_{i \in \mathcal{A}})$ , as defined in Sect. 4. The market correspondence  $\phi$  is defined as follows.

$$\forall (p, d) \in \mathbb{R}_+^M \times \mathbb{R}_+^M : \phi(p, d) = (\tilde{D}(p), \mathcal{F}(d)), \quad (3)$$

where  $\mathcal{F}(d) \triangleq \underset{\hat{p} \in \mathbb{R}_+^M}{\operatorname{argmax}} \hat{p} \cdot (d - (1, 1, \dots, 1, Z))$  and  $\tilde{D}(p)$  is the convex hull of total demand set  $D(p)$ . We say a point  $(p, d)$  is a fixed point of the market correspondence if

$$(p, d) \in \phi(p, d) = (\tilde{D}(p), \mathcal{F}(d)). \quad (4)$$

As we will show later, the fixed point of market correspondence  $\phi$  defined in Definition 8 always exists, under monotonicity assumptions on utility functions. This fixed point is essentially giving us an equilibrium price vector that can be supported by a *fractional* allocation of items to buyers in a way that it produces an envy-free market clearing outcome (i.e. an outcome that everyone gets an optimal allocation under prices and market clears). However, we expect integral allocations in a WE of the combinatorial auction. To address this we utilize the induced configuration LP defined in Definition 6 and its dual in Definition 7 at the fixed point price to see if the supporting fractional allocation can basically be decomposed into integral allocations. This helps us to find a structural characterization for WE. Putting all the pieces together, we get two main results.

**Theorem 3.** Given a combinatorial auction  $(\mathcal{I}, \mathcal{B}, \{u_i(\cdot)\}_{i \in \mathcal{B}})$  in which for every buyer  $i$  the utility function  $u_i(x, p)$  is increasing with respect to allocation of items, and strictly decreasing and continuous with respect to the money, a fractional WE (Definition 2) always exists.

**Theorem 4.** *Given a combinatorial auction  $(\mathcal{I}, \mathcal{B}, \{u_i(\cdot)\}_{i \in \mathcal{B}})$  that satisfies conditions in Theorem 3, and its corresponding Arrow-Debreu market  $(\mathcal{C}, \mathcal{A}, \{\tilde{u}_i(\cdot)\}_{i \in \mathcal{A}}, \{w^{(i)}\}_{i \in \mathcal{A}})$  as in Definition 4, a pair of prices and allocation  $(\{p_j\}_{j \in \mathcal{I}}, \{x^{(i)}\}_{i \in \mathcal{B}})$  is a WE if and only if:*

- $\exists \tilde{p} \in \mathbb{R}_+^M$  and  $\tilde{d} \in \mathbb{R}_+^M$  s.t.  $\tilde{p}_M > 0, \tilde{d}_M = Z, j \in [m] : p_j = \frac{\tilde{p}_j}{\tilde{p}_M}$  and  $(\tilde{p}, \tilde{d})$  is a fixed point of  $\phi$ .
- $\{x_{i,S}\}$  is an optimal integral solution for the induced configuration LP at prices  $p^* = p$ , where  $\forall i \in \mathcal{B}, S \subseteq \mathcal{I} : x_{i,S} = 1 \iff j \in S : x_j^{(i)} = 1$ .

While our main results in Theorems 3 and 4 characterize structural necessary and sufficient conditions for the existence of competitive equilibria, there are simple corollaries of these theorems that are of interest. The first corollary, whose proof can be seen from the proofs of Theorems 3 and 4 in Sect. 6, states the relationship between the dual of induced configuration LP at some price  $p^*$  and prices in a competitive equilibrium (either fractional or integral).

**Corollary 3.** *For a combinatorial auction  $(\mathcal{I}, \mathcal{B}, \{u_i(\cdot)\}_{i \in \mathcal{B}})$  and its corresponding Arrow-Debreu market  $(\mathcal{C}, \mathcal{A}, \{\tilde{u}_i(\cdot)\}_{i \in \mathcal{A}}, \{w^{(i)}\}_{i \in \mathcal{A}})$ , these statements are equivalent:*

- Price vector  $p$  is a fractional WE price vector, as in Definition 2.
- There exists a price vector  $\tilde{p} \in \mathbb{R}_+^M$  s.t.  $(\tilde{p}, \tilde{d})$  is a fixed point of the market correspondence  $\phi$  and  $p_j = \frac{\tilde{p}_j}{\tilde{p}_M}$ .
- Let  $u_i = \max_{S \subseteq \mathcal{I}} v_i^{(p)}(\mathbf{1}_S) - \sum_{j \in S} p_j$ . Then  $(u, p)$  is an optimal solution to the dual of induced configuration LP at price  $p$ .

Another corollary of our main result is the connection between fractional WE, as in Definition 2, and true (integral) WE, as in Definition 1. In fact, by directly applying Corollary 3 to Theorem 4 one can restate Theorem 4 through the following corollary, which bypasses the relationship to markets and reveals the relationship between fractional and integral WE.

**Corollary 4.** *Given a combinatorial auction  $(\mathcal{I}, \mathcal{B}, \{u_i(\cdot)\}_{i \in \mathcal{B}})$  that satisfies conditions in Theorem 3, a pair of prices and allocation  $(\{p_j\}_{j \in \mathcal{I}}, \{x^{(i)}\}_{i \in \mathcal{B}})$  is a WE if and only if:*

- $p$  is a price vector of a fractional WE, as in Definition 2.
- $\{x_{i,S}\}$  is an optimal integral solution for the induced configuration LP at prices  $p^* = p$ , where  $\forall i \in \mathcal{B}, S \subseteq \mathcal{I} : x_{i,S} = 1 \iff j \in S : x_j^{(i)} = 1$ .

The next corollary of our results is the existence of competitive equilibrium for the special case of unit-demand bidders (or matching markets). In this case, each buyer is interested in at most one item and a feasible allocation is an integral matching. Our result, somehow surprisingly, will give a simple proof for the existence of WE in this setting, which has been observed and proved in the literature first by Quinzii (1984), and later by Alaei et al. (2011) and

Maskin (1987) - Maskin and Quinzii studied the matching markets with one divisible goods and showed the existence of an envy-free outcome in such a market, while Alaei et al. directly showed that competitive equilibrium exists by a combinatorial proof.

**Corollary 5.** (Alaei et al. 2011; Quinzii 1984; Maskin 1987) *In a special case of unit-demand bidders, if the utilities are increasing with respect to the allocation and strictly decreasing and continuous with respect to money, competitive equilibrium always exists.*

*Proof.* Pick any price vector  $p^*$  and look at the induced configuration LP at this price. Interestingly, the feasible polytope of such LP is the matching polytope. We know matching polytope is integral (Schrijver 1983), and hence there always exists an integral optimal solution to the induced configuration LP at any price vector  $p^*$ . As we show later in the proof of Theorem 3, there always exists a fractional WE price vector  $p$  under the conditions in the statement of the corollary. Now, induced configuration LP at price  $p$  has an optimal integral primal solution  $x$  that supports any optimal solution  $(\hat{p}, \hat{u})$  of the dual program (meaning that for each bidders  $i$  the item she gets in  $x$  is her preferred item under prices  $\hat{p}$ , she gets a utility  $\hat{u}_i$  and also market clears). According to Corollary 3,  $p$  is also an optimal dual solution for the induced configuration LP at price  $p$ , and hence  $x$  supports  $p$ . So  $(x, p)$  forms a WE for the matching market.

The last corollary of our result is a simple proof for the classic result of Gul and Stacchetti (1999), in which they demonstrate the relationship between competitive equilibria and configuration LP in the case of quasilinear utilities (which is the case when  $u_i(x, p) = v_i(x) - p$  for all  $i \in \mathcal{B}$ , where  $v_i(\mathbf{1}_S)$  denotes the value of bidder  $i$  for bundle  $S$ ). In fact, in the special case of quasilinear, the induced configuration LP at any price  $p^*$  will not be a function of  $p^*$ , as  $v_i^{(p^*)}(\mathbf{1}_S) = u_i(\mathbf{1}_S, \sum_{j \in S} p_j^*) + \sum_{j \in S} p_j^* = v_i(\mathbf{1}_S)$ . We therefore have the following corollary.

**Corollary 6.** (Gul and Stacchetti 1999) *Given a combinatorial auction with quasilinear utilities, a WE exists if and only if the configuration LP has an integral optimal solution.*

*Proof.* Mixing Theorem 3 and Corollary 4, we conclude there always exist a price vector  $p$  such that it is a fractional WE price vector and together with  $\{u_i\}$ , where  $u_i = \max_{S \subseteq \mathcal{I}} v_i(\mathbf{1}_S) - \sum_{j \in S} p_j$ , forms an optimal solution to the dual of configuration LP. So, using Theorem 4 a WE exists if and only if there exists an integral optimal solution for the configuration LP.

## 6 Proofs of the Main Results

### 6.1 Proof of Theorem 3

We begin by looking at the market correspondence  $\phi$  (Definition 8). We have the following lemma, whose proof is basically by Kakutani's fixed point theorem (Kakutani et al. 1941) and is provided in the online version of this

paper (Niazadeh and Wilkens 2016). Checking the conditions of this theorem is technical and we omit the details for the sake of brevity. We assert that the proof is similar to the fixed-point proof in (Arrow and Debreu 1954) or (Maskin 1987) with some minor modifications.

**Lemma 3.** *If for all  $i \in \mathcal{B}$ ,  $u_i(x, p)$  satisfies the following conditions:*

- *Continuous with respect to  $p$ ,*
- *Increasing with respect to  $x_j : j \in \mathcal{I}$ ,*
- *Strictly decreasing with respect to  $p$ ,*

*then the market correspondence  $\phi$  will have a fixed point.*

## 6.2 Proof of Theorem 4.

**[Part 1, ‘if’ Direction].** Suppose  $(\tilde{p}, \tilde{d})$  is a fixed point of the correspondence  $\phi$  (this fixed point always exists, due to Lemma 3, and  $\tilde{p}_M > 0$ ). Now, following the proof of Theorem 3, there exists a fractional WE  $(\{\tilde{x}_{i,S}\}_{i \in \mathcal{B}, S \subseteq \mathcal{I}}, \{p_j\}_{j \in \mathcal{I}})$ , as in Definition 2, such that  $j \in [m] : p_j = \frac{\tilde{p}_j}{\tilde{p}_M}$ . Now fix  $p^* = p$  and consider the induced configuration LP at  $p^*$ . Note that  $\{\tilde{x}_{i,S}\}$  is a feasible solution for this LP by the definition of fractional WE. For every  $i \in \mathcal{B}$  let  $u_i = \max_{S \subseteq \mathcal{I}} u_i(\mathbf{1}_S, \sum_{j \in S} p_j)$ . Then  $(u, p)$  will form a feasible solution for the dual of induced configuration LP at price  $p^* = p$ , simply because  $\forall i \in \mathcal{B} : u_i \geq 0, \forall j \in \mathcal{I} : p_j \geq 0$  and we have:

$$\forall i \in \mathcal{B}, S \subseteq \mathcal{I} : u_i + \sum_{j \in S} p_j \geq u_i(\mathbf{1}_S, \sum_{j \in S} p_j) + \sum_{j \in S} p_j = v_i^{(p^*)}(\mathbf{1}_S) \quad (5)$$

We next prove that  $(u, p)$  will form an optimal solution for the dual of induced configuration LP at price  $p^* = p$ , by supporting this feasible dual solution with the feasible primal solution  $\{\tilde{x}_{i,S}\}$ . This can be done by the method of complementary slackness as following:

- if  $\tilde{x}_{i,S} > 0$ , then by Definition 2 we have  $\mathbf{1}_S \in \operatorname{argmax}_{x' \in \{0,1\}^m} u_i(x', \sum_{j=1}^m p_j x'_j)$ .

Therefore,

$$u_i + \sum_{j \in S} p_j = u_i(\mathbf{1}_S, \sum_{j \in S} p_j) + \sum_{j \in S} p_j = v_i^{(p^*)}(\mathbf{1}_S). \quad (6)$$

- if  $p_j > 0$ , then by Definition 2 we have  $\sum_{i \in \mathcal{B}, S: j \in S} \tilde{x}_{i,S} = \tilde{d}_j = 1$ .
- Due to the proof of Theorem 3, there always exists at least one point in the convex combination of demanded vectors of buyer  $i$ , and hence if  $u_i > 0$  then  $\sum_{S \subseteq \mathcal{I}} x_{i,S} = 1$ .

So  $(u, p)$  is an optimal dual solution. Now suppose  $\{x_{i,S}\}$  be an optimal integral solution to configuration LP at price  $p$ . Accordingly,  $\{x_{i,S}\}$  and  $(u, p)$  should satisfy complementary slackness conditions. These conditions show why  $(\{x_{i,S}\}, p)$  will form a WE:

- Proof of satisfaction: due to complementary slackness, if buyer  $i$  gets bundles  $S$ , then  $x_{i,S} = 1 > 0$ , and therefore:

$$u_i + \sum_{j \in S} p_j = v_i^{(p^*)}(\mathbf{1}_S) \Rightarrow u_i(\mathbf{1}_S, \sum_{j \in S} p_j) = u_i = \max_{S' \subseteq \mathcal{I}} u_i(\mathbf{1}_{S'}, \sum_{j \in S'} p_j). \quad (7)$$

- Proof of market clearance: due to complementary slackness, if  $p_j > 0$ , then

$$\sum_{i \in \mathcal{B}, S: j \in S} x_{i,S} = 1$$

as desired.

So, putting all pieces together, we conclude that  $(x, p)$  is a WE for the combinatorial auction, where  $p_j = \frac{\tilde{p}_j}{\tilde{p}_M}$ ,  $(\tilde{p}, \tilde{d})$  is a fixed point of the market correspondence  $\phi$ , as stated in the Theorem 4, and  $x_j^{(i)} = 1 \iff j \in S$ ,  $x_{i,S} = 1$ .  $\square$

**[Part 2, ‘only If’ Direction].** The proof is provided in the online version of this paper (Niazadeh and Wilkens 2016).

## 7 Conclusion

In the study of Walrasian equilibria, it is standard to assume that bidders have utilities that are quasilinear in money. Unfortunately, this is a strong assumption that has attracted little attention. We strive to study Walrasian equilibria in general combinatorial auction settings without assuming utilities are quasilinear, and our main results shed light on when they exist. Unsurprisingly, we find that some of the strong results for quasilinear bidders break when we relax our utility model. We show structure that does exist, and how it connects to a few key results for general quasilinear and unit demand non-quasilinear settings; however, we have only touched a small fraction of what is known about the quasilinear setting, and that is one source of interesting open questions. For example:

- *What natural properties of utility functions guarantee the existence of Walrasian equilibria?* In quasilinear settings, it is known that gross substitutes is sufficient in a combinatorial auction.
- *When do equilibria have a lattice structure?* It is known that in quasilinear settings and in unit-demand non-quasilinear ones, Walrasian equilibria have a lattice structure. Does this exist more generally?

Another direction for research surrounds generalizations of Walrasian equilibria:

- *When do combinatorial Walrasian equilibria exist in general?* In quasilinear settings where Walrasian equilibria fail to exist, one line of research shows that a kind of combinatorial equilibrium does exist (Feldman *et al.* 2016).

A third direction for research is to ask when quasilinearity is justified:

- *What kinds of micromarkets naturally lead to quasilinear relationships with the global market?* Taking the view that Walrasian equilibria capture a small slice of a market, one should be able to identify conditions under which a micromarket naturally has a quasilinear relationship with the other options available to an agent.

## References

- Alaei, S., Jain, K., Malekian, A.: Competitive equilibrium in two sided matching markets with general utility functions. *ACM SIGecom Exchanges* **10**(2), 34–36 (2011)
- Arrow, K.J., Debreu, G.: Existence of an equilibrium for a competitive economy. *Econometrica: J. Econom. Soc.* **22**, 265–290 (1954)
- Bevia, C., Quinzii, M., Silva, J.A.: Buying several indivisible goods. *Math. Soc. Sci.* **37**(1), 1–23 (1999)
- Bikhchandani, S., Mamer, J.W.: Competitive equilibrium in an exchange economy with indivisibilities. *J. Econ. Theor.* **74**(2), 385–413 (1997)
- Clarke, E.H.: Multipart pricing of public goods. *Public Choice* **11**(1), 17–33 (1971)
- Cramton, P.C., Shoham, Y., Steinberg, R., et al.: *Combinatorial Auctions*, vol. 475. MIT press, Cambridge (2006)
- Demange, G., Gale, D.: The strategy structure of two-sided matching markets. *Econometrica: J. Econom. Soc.* **53**, 873–888 (1985)
- Echenique, F., Oviedo, J.: A theory of stability in many-to-many matching markets (2004)
- Feldman, M., Gravin, N., Lucier, B.: Combinatorial walrasian equilibrium. *SIAM J. Comput.* **45**(1), 29–48 (2016)
- Gale, D.: Equilibrium in a discrete exchange economy with money. *Int. J. Game Theor.* **13**(1), 61–64 (1984)
- Groves, T.: Incentives in teams. *Econometrica: J. Econom. Soc.* **41**, 617–631 (1973)
- Gul, F., Stacchetti, E.: Walrasian equilibrium with gross substitutes. *J. Econ. Theor.* **87**(1), 95–124 (1999)
- Kakutani, S., et al.: *A Generalization of Brouwer’s Fixed Point Theorem*. Duke University Press, Durham (1941)
- Kaneko, M., Yamamoto, Y.: The existence and computation of competitive equilibria in markets with an indivisible commodity. *J. Econ. Theor.* **38**(1), 118–136 (1986)
- Kelso Jr., A.S., Crawford, V.P.: Job matching, coalition formation, gross substitutes. *Econometrica: J. Econom. Soc.* **50**, 1483–1504 (1982)
- Maskin, E.S.: On the fair allocation of indivisible goods. In: Feiwel, G.R. (ed.) *Arrow and the Foundations of the Theory of Economic Policy*, pp. 341–349. Springer, Heidelberg (1987)
- Murota, K., Tamura, A.: *Computation of Competitive Equilibria of Indivisible Commodities Via M-convex Submodular Flow Problem*. Kyoto University, Research Institute for Mathematical Sciences, Kyoto (2001)
- Niazadeh, R., Wilkens, C.A.: Competitive equilibria for non-quasilinear bidders in combinatorial auctions. *CoRR*, abs/1606.06846 (2016)
- Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: *Algorithmic Game Theory*, vol. 1. Cambridge University Press, Cambridge (2007)



- Quinzii, M.: Core and competitive equilibria with indivisibilities. *Int. J. Game Theor.* **13**(1), 41–60 (1984)
- Schrijver, A.: Short proofs on the matching polyhedron. *J. Comb. Theor. Ser. B* **34**(1), 104–108 (1983)
- Shapley, L.S., Shubik, M.: The assignment game I: the core. *Int. J. Game Theor.* **1**(1), 111–130 (1971)
- Svensson, L.-G.: Competitive equilibria with indivisible goods. *J. Econ.* **44**(4), 373–386 (1984)
- Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Finan.* **16**(1), 8–37 (1961)

# Correlated and Coarse Equilibria of Single-Item Auctions

Michal Feldman<sup>1,2</sup>, Brendan Lucier<sup>2(✉)</sup>, and Noam Nisan<sup>2,3</sup>

<sup>1</sup> Tel Aviv University, Tel Aviv, Israel  
`michal.feldman@cs.tau.ac.il`

<sup>2</sup> Microsoft Research, Redmond, USA  
`brlucier@microsoft.com`

<sup>3</sup> Hebrew University, Jerusalem, Israel  
`noam@cs.huji.ac.il`

**Abstract.** We study correlated equilibria and coarse equilibria of simple first-price single-item auctions in the simplest auction model of full information. Nash equilibria are known to always yield full efficiency and a revenue that is at least the second-highest value. We prove that the same is true for all *correlated equilibria*, even those in which agents overbid – i.e., bid above their values.

*Coarse equilibria*, in contrast, may yield lower efficiency and revenue. We show that the revenue can be as low as 26% of the second-highest value in a coarse equilibrium, even if agents are assumed not to overbid, and this is tight. We also show that when players do not overbid, the worst-case bound on social welfare at coarse equilibrium improves from 63% of the highest value to 81%, and this bound is tight as well.

## 1 Introduction

A very basic tenet of economic theory is to analyze strategic situations such as games or markets *in equilibrium*. The logic being that systems will typically reach an equilibrium point, following some dynamic, a dynamic that may be difficult to understand or analyze. Of course, in order for the equilibrium concept to be predictive, it must correspond to outcomes of the types of dynamics we consider possible. In Game Theory, the leading equilibrium concept is a Nash equilibrium.

In Algorithmic Game Theory, Nash equilibrium is not the only notion of equilibrium that is considered. On the one hand, it is typically computationally-hard to find a Nash equilibrium, and so it is questionable whether a Nash equilibrium can be viewed as a reasonable prediction of an outcome of a game. In contrast, there are a host of natural “learning-like” dynamics that converge to

---

M. Feldman—Supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement number 337122.

N. Nisan—Supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences and by Israel-USA Bi-national Science Foundation (BSF) grant number 2014389.

more general notions of equilibria, specifically to *correlated equilibria* or to the even more general *coarse equilibria*<sup>1</sup> which often seem to be more natural predications than Nash equilibria. It is thus common to consider also these more general notions of equilibrium in scenarios studied in Algorithmic Mechanism Design.

This extension of our concept of the class of possible equilibria has a bright side and a dark side. On the negative side, if we accept that one of these generalized equilibria notions is a possible outcome, then we need to ensure that all such equilibria produce whatever result is desired by us (in terms of “Price of Anarchy”, we can get worse bounds as we need to take the worst case performance over a wider set of equilibria.) On the positive side, in cases where we can control the equilibrium reached (e.g. by coding specific dynamics into software), we may take advantage of the extra flexibility to obtain better equilibrium points (in terms of “Price of Stability”, we may get better bounds, as we can choose an equilibrium within a wider class).

In this paper we study correlated equilibria and coarse equilibria in the simplest auction model: a full-information first-price auction of a single item. This is the simplest instance in the class of simultaneous auctions, which has received much attention lately [1, 3, 9, 12, 18]. Here too the literature is concerned both about the difficulty of reaching equilibria [2, 5–7] and about the additional loss of efficiency or revenue in these generalized types of equilibria.<sup>2</sup> A loss of efficiency here corresponds to misallocation of the item (i.e., the winner not being the bidder with the highest value), while a loss of revenue may also be viewed as a type of implicit self-stabilizing collusion between the bidders [14].

For concreteness, consider the case where Alice has value 1 for the item and Bob has value 2, where the values are common knowledge and they are participating in a first price auction. In this game the strategy space of each bidder is the set of possible bids (non-negative numbers), and the outcome from a pair of bids  $a$  by Alice and  $b$  by Bob is that Alice wins whenever  $a > b$  and pays  $a$  (so her utility is  $1 - a$  and Bob’s utility is 0) and Bob wins whenever  $b > a$  and pays  $b$  (so his utility is  $2 - b$  and Alice’s utility is 0). For simplicity, let us assume that ties are broken in favor of Bob, i.e. that he wins whenever  $b \geq a$ .

It is quite easy to analyze the pure Nash equilibria of this game: for every value of  $1 \leq v \leq 2$  there exists an equilibrium where both Alice and Bob bid the same value  $v$ , and Bob wins the tie. In the general case of first price auctions, the price  $v$  may be anything between the first price and the second price in the auction.<sup>3</sup> While non-trivial, it is also not difficult to analyze the mixed Nash equilibria of this game, where it turns out that every mixed Nash equilibrium is outcome-equivalent to a pure Nash equilibrium [4]. “Outcome equivalent” means that we get the same distribution of the identity of the winner and of his payment. In particular, all mixed (or pure) Nash equilibria of a full-information first price auction attain perfect social welfare (i.e., the player with highest value always

<sup>1</sup> Sometimes called “coarse correlated equilibria” [19] or “Hannan consistent” [10].

<sup>2</sup> For example, while pure Nash equilibria of simultaneous first-price auctions are known to be fully efficient, mixed Nash equilibria may not [12].

<sup>3</sup> This requires that the player with the highest value – Bob in our example – wins the tie; otherwise no pure equilibrium exists but arbitrarily close  $\epsilon$ -equilibria do.

wins) and have revenue that is bounded below by the second highest value in the auction (and from above by the highest value).

What about correlated equilibria and coarse equilibria? Correlated equilibria give a richer class of outcomes, since certainly a single correlated equilibrium can mix between several pure equilibria. Our first result shows that this is all that can be obtained, so in particular, correlated equilibria also yield perfect social welfare and a revenue that is at least the second highest value.

**Theorem:** Every correlated equilibrium of a first-price auction is outcome-equivalent to a mixture of pure Nash equilibria.

In [14] a similar theorem was proved for the special case of symmetric bidders, even in Bayesian settings. Whether or not correlated equilibria can be richer in non-symmetric Bayesian settings remains open.<sup>4</sup> There are several other known cases of games where correlated equilibria cannot improve upon Nash equilibria (see [15] and references therein). In [8] it was shown, in a more general setting than the one described here, that there is a *unique* correlated equilibrium if one eliminates *weakly dominated strategies*. That is, if no player ever bids above their value. Indeed, the only correlated equilibrium satisfying this constraint is the pure Nash equilibrium in which both agents bid the second-highest value.

We then turn our attention to coarse equilibria, and it turns out that a wider set of outcomes becomes possible. In [17] a coarse equilibrium is exhibited in a two-player single-item auction that is not outcome-equivalent to a mixture of pure auctions. In fact, its welfare is only  $1 - 1/e \approx 63\%$  of the optimum. This matches the general Price of Anarchy upper bound given in [18] (established via the smoothness technique [16, 18]), which applies even to general multi-item simultaneous auctions with XOS bidders. For multi-item simultaneous auctions, this bound of  $1 - 1/e$  is tight even with respect to Nash equilibria [4], but for single-item auctions, as we have seen, it is only tight for coarse equilibria and not for correlated or for Nash equilibria.

The example that attains this low welfare has the undesirable property that it uses weakly dominated strategies. That is, in this example, the support of the coarse equilibrium contains strategies where one of the players bids above his value. This use of dominated bidding strategies seems highly unnatural, so it is natural to ask whether there exist other inefficient coarse equilibria that do not use overbidding. Consider the example given in Table 1, of a coarse equilibrium, where  $\epsilon$  is some small enough constant (e.g.  $\epsilon = 10^{-4}$ ).

One may directly verify that this is indeed a coarse equilibrium.<sup>5</sup> This finite equilibrium allocates the item to Alice sometimes, and so the social welfare that

<sup>4</sup> Note, however, that in asymmetric Bayesian settings, even in (Bayesian) Nash equilibria, the winner is not necessarily the bidder with the highest valuation [13].

<sup>5</sup> Ignoring  $O(\epsilon)$  terms, at equilibrium we have: For Alice:  $u_A = 0.02 * 1 + 0.02 * 0.9 = 0.038$  while deviating to 0 would yield utility 0.02, deviating to 0.1 yield utility 0.036, deviating to 0.5 yield utility 0.035, deviating to 0.8 yield utility 0.036, deviating to 0.9 yield utility 0.37 and deviating to 1 yield utility 0. For Bob we have  $u_B = 0.03 * 1.5 + 0.11 * 1.2 + 0.19 * 1.1 + 0.63 * 1 = 1.016$ , but deviating to 1 would give utility 1, and deviating to anything below 1 would lose with probability of at least 63% leading to utility that is certainly less than 1.

**Table 1.** A coarse equilibrium in an auction with  $v_{Alice} = 1$ ,  $v_{Bob} = 2$ 

| Probability | Alice's Bid      | Bob's Bid |
|-------------|------------------|-----------|
| 2%          | $\epsilon$       | 0         |
| 2%          | $0.1 + \epsilon$ | 0.1       |
| 3%          | $0.5 - \epsilon$ | 0.5       |
| 11%         | $0.8 - \epsilon$ | 0.8       |
| 19%         | $0.9 - \epsilon$ | 0.9       |
| 63%         | $1 - \epsilon$   | 1         |

it reaches is not perfect! Also notice that the winner pays at most 1 for the item, but sometimes pays strictly less than 1, and thus the revenue is strictly smaller than the second price!

This leads us to a natural question: what is the lowest welfare possible in a coarse equilibrium where no player overbids? We might term this ratio “PoUA” – “the price of *undominated* anarchy”. We show that indeed insisting that players never overbid ensures a significantly higher share of welfare, and provide tight bounds for it. To the best of our knowledge, this is the first indication that a no-overbidding restriction improves worst case guarantees in first-price auctions.<sup>6</sup>

**Theorem:** In every coarse equilibrium of a single-item first-price auction where players never bid above their value, the social welfare is at least a  $c$  fraction of the optimal, where  $c \approx 0.813$ .

**Theorem:** There exists a single-item first-price auction with two players that has a coarse equilibrium where players never bid above their value, whose social welfare is only a  $c \approx 0.813$  fraction of the optimal welfare.

We then focus our attention on the revenue of the auction. While Nash equilibria and correlated equilibria always yield revenue that is at least the second highest value, our example above has shown that coarse equilibrium may yield lower revenue. We ask how low may this revenue be, and provide a tight bound:

**Theorem:** In every coarse equilibrium of any single-item first-price auction, the revenue is at least  $1 - 2/e \approx 26\%$  of the second highest value.

**Theorem:** There exists a single-item first-price auction with two players that has a coarse equilibrium where players never bid above their value whose revenue is only  $1 - 2/e \approx 26\%$  of the second highest value.

Notice that here we get the same bound whether or not players may bid above their value. This lower bound is obtained in a symmetric instance (i.e., where the two players have the same value). We remark that in large symmetric instances the revenue approaches the value of the players. We also show that as the gap between the highest value and the second highest value increases,

---

<sup>6</sup> In contrast, for multi-item simultaneous auctions, the  $1 - 1/e$  bound is tight for XOS valuations even without overbidding [3].

the revenue approaches the second highest value. This is in contrast to social welfare, where noted above the inefficiency may persist even when the gap in the players' values is arbitrarily large.

## 2 Preliminaries

We will focus on an auction with  $n$  players and a single item for sale. Player  $i$  has value  $v_i$  for the item, and we index the players so that  $v_1 \geq v_2 \geq \dots \geq v_n$ .

The auction proceeds as follows: the players simultaneously submit real-valued bids,  $\mathbf{x} = (x_1, \dots, x_n)$ . Ties are broken according to a fixed tie-breaking function, which maps the (maximal) bids to a winner. Player  $i$  wins when  $x_i \geq x_j$  for all  $j$  and the tie at value  $x_i$  (if any) is broken in favor of player  $i$ , which we denote by  $x_i \succsim \mathbf{x}_{-i}$ . The winner pays his or her bid.

Given a joint distribution  $D$  over the bids of the players, the expected payment of player  $i$  is  $E_{\mathbf{x} \sim D}[x_i \cdot 1_{x_i \succsim \mathbf{x}_{-i}}]$ , so his expected utility is  $u_i = v_i \cdot Pr_{\mathbf{x} \sim D}[x_i \succsim \mathbf{x}_{-i}] - E_{\mathbf{x} \sim D}[x_i \cdot 1_{x_i \succsim \mathbf{x}_{-i}}]$ .

We will study correlated and coarse correlated equilibria. The following definitions are tailored to our auction setting; a more general definition of correlated equilibria for infinite games can be found in Hart and Schmeidler [11].

**Definition 1.** *A joint distribution  $D$  over bids is a correlated equilibrium if, for every player  $i$  and every (measurable) deviation function  $b_i: \mathbb{R} \rightarrow \mathbb{R}$  of player  $i$ , it holds that*

$$\begin{aligned} v_i \cdot Pr_{\mathbf{x} \sim D}[x_i \succsim \mathbf{x}_{-i}] - E_{\mathbf{x} \sim D}[x_i \cdot 1_{x_i \succsim \mathbf{x}_{-i}}] \\ \geq v_i \cdot Pr_{\mathbf{x} \sim D}[b_i(x_i) \succsim \mathbf{x}_{-i}] - E_{\mathbf{x} \sim D}[b_i(x_i) \cdot 1_{b_i(x_i) \succsim \mathbf{x}_{-i}}]. \end{aligned}$$

**Definition 2.** *A joint distribution  $D$  over bids is a coarse correlated equilibrium (or coarse equilibrium for short) if, for every player  $i$  and for every unilateral deviation  $x'_i \in \mathbb{R}$  of player  $i$ , it holds that*

$$v_i \cdot Pr_{\mathbf{x} \sim D}[x_i \succsim \mathbf{x}_{-i}] - E_{\mathbf{x} \sim D}[x_i \cdot 1_{x_i \succsim \mathbf{x}_{-i}}] \geq v_i \cdot Pr_{\mathbf{x} \sim D}[x'_i \succsim \mathbf{x}_{-i}] - E_{\mathbf{x} \sim D}[x'_i \cdot 1_{x'_i \succsim \mathbf{x}_{-i}}].$$

In each of these definitions, we can interpret  $x_i$  as a bid that is recommended to agent  $i$  by a coordinator of the equilibrium. We will sometimes refer to agent  $i$  as being “told to bid  $x_i$ ” when this interpretation is convenient. Under this interpretation, correlated equilibria are immune to deviations that can condition on the recommended bid, whereas coarse equilibria need only be immune to unconditional deviations (i.e., constant bidding functions).

### 2.1 Tie Breaking

Before we continue, a word about tie-breaking is in order. All of our theorems that claim something for all equilibria will hold for *every* tie breaking rule. In our constructions, we will allow ourself to choose a tie breaking rule to our liking. Note however that if the tie breaking rule is not to the reader's liking, in a joint

distribution over  $\mathbf{x}$  we can always avoid any dependence on it by increasing one of the maximal bids by  $\epsilon$ , which would give us an  $\epsilon$ -equilibrium (rather than an exact one). Thus every construction of an equilibrium that we provide using a particular tie-breaking rule immediately implies also an  $\epsilon$ -equilibrium for any tie breaking rule and any  $\epsilon > 0$ .

## 2.2 The Distribution on the Winning Price

When analysing revenue and welfare in an equilibrium, it will be most convenient to consider the single-dimensional distribution on the winning price, i.e., on  $\max_i\{x_i\}$ . We will denote the cumulative distribution on the winning price by  $F$ . The revenue can be easily expressed in terms of  $F$  as  $Revenue = E_{x \sim F}[x] = \int (1 - F(x))dx$  (where the integration is over the support of the distribution).

The starting point for our analysis of *coarse equilibria* is the following. Suppose there are  $n = 2$  players, say Alice and Bob, with values 1 and  $v \leq 1$  respectively. If Alice chooses to deviate from the equilibrium to some fixed bid  $x$ , then her utility will be  $(1 - x) \cdot F_{Bob}(x) \geq (1 - x) \cdot F(x)$  (this expression ignores the possibility of a tie), where  $F_{Bob}$  is the cumulative distribution of Bob's bid, which is certainly stochastically dominated by the cumulative distribution on the winning bid. In our constructions we will typically have both Alice and Bob always bidding the same value  $x = y$  (where this joint value is distributed according to  $F$ ), and thus will have  $F_{Bob} = F$ .

Denoting Alice's utility at equilibrium by  $\alpha$ , a necessary condition that this deviation is not profitable is thus  $\alpha \geq (1 - x) \cdot F(x)$ , i.e. that  $F(x) \leq \alpha/(1 - x)$ . Similarly for Bob we must have  $F(x) \leq \beta/(v - x)$ , where  $\beta$  is Bob's utility at equilibrium. Thus if  $F$  corresponds to a coarse equilibrium then it must be stochastically dominated by the minimum of these two expressions.

The following simple calculation states the closed form expression for the revenue of a distribution of this form.

**Lemma 1.** *Let the cumulative distribution function  $G = G_{a,b}$  be defined by  $G(x) = a/(b - x)$  for  $0 \leq x \leq b - a$ . Then,  $E_G[x] = b - a + a \ln(a/b)$ .*

## 3 Correlated Equilibrium

The goal of this section is to establish that every correlated equilibrium of a first-price auction is outcome-equivalent to a mixture of pure Nash equilibria. This characterization implies that the revenue of the auctioneer is always at least the second-highest value,  $v_2$ . We begin by showing that the winning bid is never lower than the second-highest of the players' values.

**Lemma 2.** *For every correlated equilibrium  $D$ ,  $Pr_{\mathbf{x} \sim D}[\max_i\{x_i\} < v_2] = 0$*

*Proof.* Assume otherwise. We will derive a contradiction by finding a utility-improving deviation for one of the two highest-valued bidders.

Let  $S = \{p | Pr_{\mathbf{x} \sim D}[\max_i \{x_i\} < p] > 0\}$ , and let  $p^* = \inf(S) < v_2$ . That is,  $p^*$  is the infimum of the support of winning bids, which by assumption is less than  $v_2$ . Fix some  $p^* < p < (p^* + v_2)/2$ , and define  $\delta = Pr_{\mathbf{x} \sim D}[\max_i \{x_i\} < p] > 0$ .

Consider players 1 and 2. (Recall that players are indexed from largest value to smallest.) One of the two players must be winning with a bid less than  $p$  with probability at most  $\delta/2$ , say player  $j$ . That is,  $Pr_{\mathbf{x} \sim D}[p > x_j \succsim \mathbf{x}_{-j}] \leq \delta/2$ , recalling that  $x_j \succsim \mathbf{x}_{-j}$  means that either  $x_j$  is strictly larger than the other bids, or that it is weakly larger and the tie is broken in favor of player  $j$ .

Player  $j$  is the bidder for which we will construct a deviation. As expected, we will exploit the non-coarse nature of the equilibrium, constructing a deviation for bidder  $j$  that depends on the bid suggested by the (supposed) correlated equilibrium. Choose  $\epsilon > 0$  so that  $p + 2\epsilon < (p^* + v_2)/2$  (which is possible by the choice of  $p$  above). Then define a deviation function  $b_j$  by  $b_j(x_j) = p + \epsilon$  for all  $x_j \leq p$ , and  $b_j(x_j) = x_j$  for all  $x_j > p$ . That is, when being told any value  $x_j \leq p$ , player  $j$  bids instead  $p + \epsilon$ . On the up side, this will certainly win all the cases where  $\max_i \{x_i\} < p$ , increasing the probability of winning by at least  $\delta/2$  and thus increasing his utility by at least  $\delta \cdot (v_2 - p - \epsilon)/2$ . On the down side, player  $j$  now pays  $p + \epsilon$  when he wins rather than his original bid  $x_j$ . Since player  $j$  never won when  $x_j < p^*$  (as  $Pr_{\mathbf{x} \sim D}[\max_i \{x_i\} < p^*] = 0$ , by definition of  $p^*$ ) he pays at most  $(p - p^*) + \epsilon$  more whenever he wins, which happened with probability of at most  $\delta/2$ . Thus the down side of player  $j$ 's utility from deviation is at most  $\delta \cdot (p - p^* + \epsilon)/2$ . So the deviation is profitable whenever  $\delta \cdot (v_2 - p - \epsilon)/2 > \delta \cdot (p - p^* + \epsilon)/2$  which is the case due to our choice of  $\epsilon$ .

We next show that only the players with the highest value can win in a correlated equilibrium.

**Lemma 3.** *For every correlated equilibrium  $D$ , and any player  $i$  such that  $v_i < v_1$ , we have  $Pr_{\mathbf{x} \sim D}[x_i \succsim \mathbf{x}_{-i}] = 0$ .*

*Proof.* By Lemma 2, no bidder ever wins with a bid (and hence price) strictly less than  $v_2$ . On the other hand, if any player  $i > 1$  ever wins with a price that is strictly more than  $v_i$ , their utility will be negative, making a deviation to a bid 0 profitable. This immediately implies the desired result if  $v_2 = v_1$ , so from this point onward we will assume  $v_2 < v_1$ .

The only case we further need to consider is if some player  $i \geq 2$  with  $v_i = v_2$  wins at price exactly  $v_2$ , say with some probability  $\delta > 0$ . But then player 1 would prefer to deviate from any  $x_1 \leq v_2$  to  $v_2 + \epsilon$ , gaining utility of at least  $\delta(v_1 - v_2 - \epsilon)$  due to winning all cases in which  $\max_i \{x_i\} \leq v_2$ , and losing at most  $\epsilon$  due to the additional payment (since, by Lemma 2 and the first-price nature of the auction, player 1 never pays less than  $v_2$  when she wins). Choosing  $\epsilon$  small enough, this deviation becomes profitable.

In conclusion we have a complete characterization of correlated equilibria in terms of their outcomes. Clearly every mixture of pure equilibria is a correlated equilibrium, and this turns out to be all that is possible:



**Theorem 1.** *Every correlated equilibrium of the single-item first-price auction is equivalent (in terms of winning probabilities and payments) to a mixture of pure equilibria (where Alice always wins the ties).*

*Proof.* First suppose  $v_1 > v_2$ . By Lemma 3, player 1 always wins and never pays less than  $v_2$ , so she must always bid at least  $v_2$ . Clearly player 1 can never bid more than  $v_1$  since that will give her negative utility (as she does always win). Thus player 1's bid  $x_1$  is supported on the interval  $[v_2, v_1]$  and she always wins. The outcome is thus equivalent to that of a similar distribution on the pure equilibria in which all players bid  $x \in [v_2, v_1]$  (with player 1 winning the ties).

Next suppose  $v_1 = v_2$ . By Lemma 3, only the maximum-valued players ever win, and the winner always pays at least  $v_1$ . The utility of every player is therefore exactly 0. The outcome is thus equivalent to a similar distribution on the pure equilibria, in which all players bid  $v_1$  and ties are broken in favor of the appropriate maximum-value player.

## 4 Price of Undominated Anarchy

The following theorem shows that if players do not overbid, the welfare guarantee in any coarse correlated equilibrium improves from 63% to 81%.

**Theorem 2.** *In every coarse equilibrium of the single-item first-price auction where players never bid above their value, the social welfare is at least a 0.813559... fraction of the optimal.*

*Proof.* Our approach to the proof will be to consider the distribution of the price paid by the winner of the auction. We will bound the CDF of this distribution, using the coarse equilibrium condition that no bidder wishes to unilaterally deviate to any constant bid  $x$  that is at most their value. Since the social welfare is the sum of the expected revenue and the expected buyer utilities, we can then translate these bounds on the prices directly into a bound on welfare.

Let us start by normalizing the values of the players in the auction: let us call the player with highest value Alice, and normalize this value to 1, and let us call the player with second highest value Bob, so his value is  $v \leq 1$ . There could be other players in the auction but our analysis will ignore them. Fix a coarse equilibrium of that auction. Let us further denote Alice's utility in the equilibrium by  $\alpha$  and Bob's utility by  $\beta$ . Since Bob never uses dominated strategies, he always bids at most  $v$  and thus Alice can always deviate to  $v + \epsilon$  obtaining a utility of  $1 - v - \epsilon$ , for any positive  $\epsilon$ . We must therefore have  $\alpha \geq 1 - v$ .

Denote by  $F_{CE}$  the cumulative distribution on the price paid by the winner of the auction. The fact that Alice does not want to deviate implies that  $F_{CE}(x) \leq \alpha/(1 - x)$  for all  $0 \leq x \leq 1 - \alpha$ . The fact that Bob does not want to deviate implies that  $F_{CE}(x) \leq \beta/(v - x)$  for all  $0 \leq x \leq v - \beta$ . Thus, the distribution  $F_{CE}$  stochastically dominates the following distribution whose cumulative distribution function is:

$$F(x) = \begin{cases} \min\left\{\frac{\alpha}{1-x}, \frac{\beta}{v-x}\right\} & 0 \leq x \leq \max(v - \beta, 1 - \alpha) \\ 1 & x > \max(v - \beta, 1 - \alpha) \end{cases}$$

The revenue raised by the auction is simply the expected value of the winning price, which is bounded from below by the expected value of  $x$  that is drawn according to  $F$ . Thus,  $Revenue \geq \int_0^1 (1 - F(x))dx$ , and a lower bound on the welfare is obtained by adding this revenue to the sum of utilities; i.e., to  $\alpha + \beta$ . We will calculate such a lower bound, over all possible values of  $\alpha \geq 1 - v$ ,  $\beta$ , and  $v \leq 1$ . That is, we will show that for all possible values of  $\alpha, \beta$ , and  $v$ , we have that  $\alpha + \beta + \int_0^1 (1 - F(x))dx \geq 0.813559\dots$

In calculating  $\int_0^1 (1 - F(x))dx$  we will split into two cases.

**Case 1:**  $\beta \geq v\alpha$ . This is the easy case since here  $\beta/(v - x) \geq \alpha/(1 - x)$  for all  $0 \leq x \leq v$  and thus  $F$  simplifies to  $F(x) = \alpha/(1 - x)$  for all  $0 \leq x \leq 1 - \alpha$ , and so our integral simplifies to

$$Revenue = \int_0^{1-\alpha} \left(1 - \frac{\alpha}{1-x}\right) dx = 1 - \alpha + \alpha \log \alpha.$$

Thus a lower bound on the welfare is  $\alpha + \beta + 1 - \alpha + \alpha \log \alpha$ . In this case we had that  $\beta \geq \alpha v \geq \alpha(1 - \alpha)$  so our lower bound on welfare, over all  $\beta$  and  $v$  is

$$Welfare \geq 1 + \alpha(1 - \alpha) + \alpha \log \alpha.$$

The last expression attains its minimum of 0.838... over all  $0 \leq \alpha \leq 1$  at  $\alpha = 0.203\dots$  (where  $\alpha$  is the solution to the equation  $2x - \log x - 2 = 0$ ) and so we have that for the case  $\beta \geq v\alpha$  the welfare is at least  $0.838\dots > 0.813559\dots$ <sup>7</sup>

**Case 2:**  $\beta < v\alpha$ . This is the more complex case. In this case we have that  $\beta/(v - x) < \alpha/(1 - x)$  exactly when  $x < \theta = (\alpha v - \beta)/(\alpha - \beta)$ , and thus the revenue is obtained as

$$\begin{aligned} Revenue &= \int_0^\theta (1 - \beta/(v - x))dx + \int_\theta^{1-\alpha} (1 - \alpha/(1 - x))dx \\ &= \alpha \log \left(\frac{\alpha - \beta}{1 - v}\right) + \beta \log \left(\frac{\beta(1 - v)}{v(\alpha - \beta)}\right) + 1 - \alpha. \end{aligned}$$

Our lower bound for the welfare is thus

$$Welfare \geq \beta + \alpha \log \left(\frac{\alpha - \beta}{1 - v}\right) + \beta \log \left(\frac{\beta(1 - v)}{v(\alpha - \beta)}\right) + 1.$$

Taking the derivative with respect to  $v$ , we get the expression  $(\alpha v - \beta)/((1 - v)v)$  which is always positive in our range and thus for every  $\alpha$  and  $\beta$ , the minimum is obtained at the lowest possible value  $v = 1 - \alpha$ .

<sup>7</sup> To get an auction with these parameters we need to specify when each of the players wins in a way that will achieve these values of  $\alpha$  and  $\beta$ . The following parameters yield these utilities: Alice and Bob bid the same value of  $x$  distributed according to the same  $F$  that provided the lower bound:  $\alpha$  that is the solution of the equation  $2x - \log x - 2 = 0$ ,  $v = 1 - \alpha$  and  $\beta = v\alpha$ . Alice wins whenever  $x = 0$  and Bob wins otherwise. Thus the probability that Alice wins is  $\alpha = F(0)$  and she pays nothing, indeed obtaining utility of  $\alpha$ . Bob wins probability  $p = 1 - \alpha$  and pays the entire revenue obtaining net utility of  $pv - Revenue = (1 - \alpha)(1 - \alpha) - (1 - \alpha + \alpha \log \alpha)$  which for our  $\alpha$  is indeed  $(1 - \alpha)\alpha = \beta$ .

Substituting this value of  $v$ , we get that the minimum possible welfare is the minimum of the function

$$\beta + \alpha \log \left( \frac{\alpha - \beta}{\alpha} \right) + \beta \log \left( \frac{\beta \alpha}{(1 - \alpha)(\alpha - \beta)} \right) + 1. \quad (1)$$

The following claim shows that the minimum of this function is 0.813559..., as promised (proof deferred to the full version), completing the proof of Theorem 2.

*Claim.* The minimum of the function in Eq. (1) is 0.813559....

We show this bound is tight by exhibiting an auction with matching welfare.

**Theorem 3.** *There exists a single-item two-player auction with player values 1 and  $v \leq 1$ , and a coarse equilibrium of that auction where players never bid above their values, whose social welfare matches the bound from Theorem 2 (0.813559...).*

*Proof.* Our approach is to construct an equilibrium in which the distribution over prices paid precisely matches the “bounding” distribution  $F$  from the proof of Theorem 2, and the agent utilities precisely match the values for which the welfare expression attained its minimum in that proof. Call the player with value 1 Alice, and the player with value  $v \leq 1$  Bob.

Guided by the proof of Theorem 2, we will choose a parameter  $\alpha$ , then set

$$\beta = \frac{\alpha - \alpha^2}{e\alpha - \alpha + 1} \quad \text{and} \quad v = 1 - \alpha. \quad (2)$$

We will arrange the parameters so that  $\alpha$  and  $\beta$  are Alice’s and Bob’s utilities at equilibrium, respectively.

Define  $F(x) = \min\{\frac{\alpha}{1-x}, \frac{\beta}{v-x}\}$ , for  $x \in [0, v]$ . In the equilibrium we construct, a value will be drawn from the distribution with CDF  $F$  and both players will bid that value. Note that neither Alice nor Bob has a profitable deviation in such an equilibrium, as long as their utilities are  $\alpha$  and  $\beta$ , respectively. Thus, to show that an equilibrium exists for a certain choice of  $\alpha$ , we must specify when each of the players wins so that they achieve the utilities  $\alpha$  and  $\beta$ .

We will show that an equilibrium exists for all  $\alpha \in [0.27, 0.28]$ . This will imply the desired result, since in particular this includes the value of  $\alpha$  for which the welfare bound from Theorem 2 is achieved. Recall from the proof of Theorem 2 that, if an equilibrium exists, its welfare will be

$$W = \alpha \log \left( \frac{e\alpha}{(e - 1)\alpha + 1} \right) + 1.$$

Write  $q$  for the solution to  $W = v(1 - q) + q$ , so that

$$q = \frac{W - v}{1 - v}. \quad (3)$$

We first claim that if we are able to specify when Alice wins, so that she wins with probability  $q$  and her utility is  $\alpha$ , then it necessarily follows that Bob will

have utility  $\beta$ . This is because, writing  $p_A$  and  $p_B$  for the expected payment of Alice and Bob respectively,

$$q + (1 - q)v = W = p_A + p_B + \alpha + \beta.$$

So if indeed  $q - p_A = \alpha$ , we can conclude that  $(1 - q)v - p_B = \beta$  and hence Bob's utility is precisely  $\beta$ . We will therefore focus on Alice's utility for the remainder of the proof. We can substitute the expressions for  $\beta$  and  $v$  (Eq. (2)) into our expression for  $q$  to yield

$$q = 2 + \log \left( \frac{\alpha}{1 + (e - 1)\alpha} \right). \tag{4}$$

This expression is non-decreasing on the interval  $[0.27, 0.28]$ , so we can conclude (by evaluating the expression on the endpoints) that  $q \in [0.3, 0.4]$  for  $\alpha \in [0.27, 0.28]$ .

The minimum total utility that can be achieved by Alice, while winning with probability  $q$ , is if she wins when prices are highest. That is, whenever the price is at or above  $F^{-1}(1 - q)$ . Under this specification, the utility of Alice would be

$$u_{min} = q - \int_{F^{-1}(1 - q)}^v xF'(x)dx.$$

Similarly, the maximum possible utility achievable by Alice is if she wins when prices are lowest; that is, when prices are at or below  $F^{-1}(q)$ . Under this choice, the utility of Alice would be

$$u_{max} = q - \int_0^{F^{-1}(q)} xF'(x)dx.$$

Since  $F$  is continuous on the range  $(0, v)$ , it is enough to show that  $\alpha \in [u_{min}, u_{max}]$ , since this implies the existence of an interval upon which Alice could win so that her utility is exactly  $\alpha$ .

**Proposition 1.** *For any  $\alpha \in [0.27, 0.28]$  it holds that  $\alpha \in [u_{min}, u_{max}]$ .*

The proof of Proposition 1 appears in the full version of the paper. The high-level idea behind the proof is to first show that  $F(x) = \frac{\beta}{v - x}$  for  $x \in [0, F^{-1}(q)]$  and  $F(x) = \frac{\alpha}{1 - x}$  for  $x \in [F^{-1}(1 - q), 1]$ . With this we can derive closed-form formulas for  $u_{min}$  and  $u_{max}$ . The desired inequalities of Proposition 1 then follow from standard functional analysis, concluding the proof of Theorem 3.

## 5 Revenue in Coarse Equilibria

We start with a construction of a two-bidder first-price auction that admits a coarse equilibrium whose revenue is  $1 - 2/e$  fraction of the second highest bid.

**Lemma 4.** *There exists a coarse equilibrium of a single-item two-player auction with player values 1 and 1 whose revenue is  $1 - 2/e \leq 0.27$ .*

*Proof.* Here is a coarse equilibrium: the two players bid  $(x, x)$  where  $x$  is distributed according to the cumulative distribution function  $F(x) = e^{-1}/(1-x)$  (for all  $0 \leq x \leq 1 - 1/e$ ), and each of them wins exactly half the time (at each price). Applying the calculation in the previous lemma, the total revenue of this auction is  $1 - e^{-1} + e^{-1} \ln e^{-1} = 1 - 2e^{-1}$ , and each player's utility is thus  $e^{-1}$ . A possible deviation of one of the players to  $x$  will yield utility  $F(x)(1-x) = e^{-1}$  and is thus not strictly profitable. Thus we are indeed in a coarse equilibrium.

We show that the construction above is essentially the worst possible case across all first price auctions. We first establish this bound for the two-bidder case, then prove the general theorem by reducing an auction with an arbitrary number of bidders and arbitrary values to the two-bidder case. To state this cleanly, we will fix the value of the second highest bidder, Bob, to 1 and let Alice's value  $v$  be any quantity that is at least 1.

**Lemma 5.** *Consider a coarse equilibrium of the single-item 2-player first price auction where Bob has value 1 and Alice has value  $v \geq 1$ . Then, the revenue of the seller is at least  $1 - 2/e \geq 0.26$ .*

The proof of Lemma 5 appears in the full version of the paper. The main idea is to consider the distribution of prices paid at equilibrium, and use the equilibrium conditions to bound its cumulative distribution function. Subject to these conditions, one can show that the expected price paid is maximized when the distribution is  $F$  from the proof of Lemma 4. We can now easily conclude the main theorem of this section.

**Theorem 4.** *In every first price auction, with any number of bidders, the revenue in every coarse equilibrium is at least a  $1 - 2/e \geq 0.26$  fraction of the second highest value.*

*Proof.* Take an equilibrium of an auction with  $k$  bidders with values  $v_1 \geq v_2 \geq \dots \geq v_k$ . We will now construct an equilibrium of the two-player auction with values  $v_1 \geq v_2$  that has the same revenue as does the original auction. After scaling, the main lemma bounds the revenue of the two-player auction to be at least  $(1 - 2/e)v_2$  and so this is also the bound on the original one.

To get the coarse equilibrium for the two player auction, simply take the same distribution on bids as in the original auction, but assigning the winning bids of players  $i \geq 3$  to one of the first two bidders (arbitrarily). Notice that since none of the players  $i \geq 3$  had a negative utility in the original auction (otherwise they would deviate to 0), and furthermore, each of the first two players gets at least as much utility from winning as do any of the players  $i \geq 3$ , thus we are only increasing the utilities of the first and second player in the new equilibrium. On the other hand, notice that we have not changed the utilities from deviations at all since these utilities depend only on the distribution of the winning price and not on the identity of the winner. It follows that the first two players still do not want to deviate and so we have a coarse equilibrium in the two-player game.

We remark that as the competition increases, the auctioneer’s revenue grows. For the case of two symmetric bidders (with value 1), Lemma 4 shows a coarse equilibrium with revenue  $1 - 2/e$ . For the case of  $n$  symmetric bidders we show the following.

**Theorem 5.** *In every first price auction, with any number of symmetric bidders with value  $v$ , the revenue in every coarse equilibrium is at least  $(1 - \frac{n}{e^{n-1}})v$ . This is tight.*

*Proof.* We first show that the revenue is always at least  $(1 - \frac{n}{e^{n-1}})v$ . Let  $F$  be the distribution of the price. The sum of the bidders’ utilities is  $v - E[x]$  (where  $x$  is distributed according to  $F$ ). Clearly, one of them has utility at most  $\frac{1}{n}(v - E[x])$ ; denote this value by  $\alpha$ . Since no deviation to any  $x$  is profitable for that player, it holds that  $F(x)(v - x) \leq \alpha$  for all  $x$ , that is  $F(x) \leq \frac{\alpha}{v-x}$ . It follows that the expected value of  $x$  according to  $F$  is at least the expected value of  $x$  according to the distribution  $\alpha/(v - x)$  which is  $v - \alpha + \alpha \ln(\alpha/v)$  (by Lemma 1). Substitute  $E[x] = v - \alpha n$  (by the definition of  $\alpha$ ) to get  $\alpha \leq v/e^{n-1}$ . It follows that  $E[x] = v - \alpha n \geq v(1 - \frac{n}{e^{n-1}})$ .

We now construct a coarse equilibrium with revenue at most  $(1 - \frac{n}{e^{n-1}})v$ . Consider a profile where bidders bid  $x$  according to the distribution  $F(x) = \alpha/(v - x)$ , where  $\alpha = v/e^{n-1}$ ; and each bidder wins with probability  $1/n$ . The expected payment is  $E[x] = v - \alpha + \alpha \ln(\alpha/v)$ . The expected utility of a bidder is  $1/n(v - E[x])$  and this should be at least  $\alpha$  (the deviation utility). Solving for  $\alpha$ , we get  $\alpha \leq v/e^{n-1}$ . So this is an equilibrium, and the revenue is  $E[x] = v - \alpha + \alpha \ln(\alpha/v) = v(1 - \frac{n}{e^{n-1}})$ .

We can also show that as the gap between the highest value and the second highest value increases, the revenue must get close to the second highest value. To state this in the cleanset way, we will fix the value of the second highest bidder, Bob, to 1 and let Alice’s value  $v$  approach infinity.

**Theorem 6.** *For every  $\epsilon > 0$  there exists  $v_0 = O(\epsilon^{-4})$  such that in any auction where Alice has value  $v \geq v_0$  and Bob has value 1 (and perhaps other players with other values), the revenue is at least  $1 - \epsilon$ .*

*Proof.* Assume by way of contradiction that the total revenue is less than  $1 - \epsilon$ . It follows that with probability of at least  $\epsilon/3$  the price paid by the winner is at most  $1 - \epsilon/3$  (otherwise the revenue would be bounded below by  $(1 - \epsilon/3)^2 \geq 1 - \epsilon$ , for small enough  $\epsilon$ ). It follows that Bob must win the item with probability of at least  $\epsilon^2/9$  as otherwise his utility would be less than that while deviating to  $1 - \epsilon/3$  would ensure utility of at least that. Now the bound on the revenue implies that the probability that the winning price is very high, greater than  $18/\epsilon^2$  can be at most  $\epsilon^2/18$ . Now consider a deviation of Alice to  $18/\epsilon^2$ : her probability of winning goes up by at least  $\epsilon^2/9 - \epsilon^2/18$  (the probability that Bob wins minus the probability of any bids above  $18/\epsilon^2$ ). Her utility changes as follows: on the up side it increases by at least  $\epsilon^2 v/18$  due to the increased winning probability, and on the down side it decreases by at most  $18/\epsilon^2$  due to the increased price. The deviation must be beneficial whenever  $v > 18^2/\epsilon^4$ .

## References

1. Bhawalkar, K., Roughgarden, T.: Welfare guarantees for combinatorial auctions with item bidding. In: Proceedings of the 22nd annual ACM-SIAM Symposium on Discrete Algorithms (2011)
2. Cai, Y., Papadimitriou, C.H.: Simultaneous Bayesian auctions and computational complexity. In: ACM Conference on Economics and Computation, EC 2014, Stanford, CA, USA, 8–12 June 2014, pp. 895–910 (2014)
3. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian combinatorial auctions. In: Proceedings of the 35th International Colloquium on Automata, Languages and Programming (2008)
4. Christodoulou, G., Kovács, A., Sgouritsa, A., Tang, B.: Tight bounds for the price of anarchy of simultaneous first price auctions. CoRR abs/1312.2371 (2013)
5. Daskalakis, C., Syrgkanis, V.: Learning in auctions: regret is hard, envy is easy. CoRR abs/1511.01411 (2015)
6. Devanur, N.R., Morgenstern, J., Syrgkanis, V., Weinberg, S.M.: Simple auctions with simple strategies. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC 2015, Portland, OR, USA, 15–19 June 2015, pp. 305–322 (2015)
7. Dobzinski, S., Fu, H., Kleinberg, R.D.: On the complexity of computing an equilibrium in combinatorial auctions. In: Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, 4–6 January 2015, pp. 110–122 (2015)
8. Dütting, P., Kesselheim, T., Tardos, É.: Mechanism with unique learnable equilibria. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC 2014, pp. 877–894 (2014)
9. Feldman, M., Fu, H., Gravin, N., Lucier, B.: Simultaneous auctions are (almost) efficient. In: Proceedings of the 45th annual ACM Symposium on Theory of Computing (2013)
10. Hart, S., Mas-Colell, A.: A general class of adaptive strategies. *J. Econo. Theory* **98**(1), 26–54 (2001)
11. Hart, S., Schmeidler, D.: Existence of correlated equilibria. *Math. Oper. Res.* **14**(1), 18–25 (1989)
12. Hassidim, A., Kaplan, H., Mansour, Y., Nisan, N.: Non-price equilibria in markets of discrete goods. In: Proceedings of the 12th ACM Conference on Electronic Commerce (2011)
13. Kaplan, T., Zamir, S.: Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case. *Econ. Theor.* **50**(2), 269–302 (2012)
14. Lopomo, G., Marx, L., Sun, P.: Bidder collusion at first-price auctions. *Rev. Econ. Design* **15**(3), 177–211 (2011)
15. Moulin, H., Gupta, S.S., Ray, I.: Coarse correlated equilibria in an abatement game. Working paper (2013)
16. Roughgarden, T.: Intrinsic robustness of the price of anarchy. *J. ACM* **62**(5), 32 (2015)
17. Syrgkanis, V.: Efficiency of mechanisms in complex markets. Ph.D. thesis, Cornell University, 8 (2014)
18. Syrgkanis, V., Tardos, E.: Composable and efficient mechanisms. In: Proceedings of the 45th Annual ACM Symposium on Theory of Computing (2013)
19. Young, H.P.: Strategic Learning and Its Limits. Oxford University Press, Oxford (2004)

# Pricing to Maximize Revenue and Welfare Simultaneously in Large Markets

Elliot Anshelevich, Koushik Kar, and Shreyas Sekar<sup>(✉)</sup>

Rensselaer Polytechnic Institute, Troy, NY, USA  
eanshel@cs.rpi.edu, koushik@ecse.rpi.edu, sekars@rpi.edu

**Abstract.** We study large markets with a single seller who can produce many types of goods, and many multi-minded buyers. The seller chooses posted prices for its many items, and the buyers purchase bundles to maximize their utility. For this setting, we consider the following questions: what fraction of the optimum social welfare does a revenue maximizing solution achieve? Are there pricing mechanisms which achieve both good revenue and good welfare simultaneously? To address these questions, we give envy-free pricing schemes which are guaranteed to result in both good revenue and welfare, as long as the buyer valuations for the goods they desire have a nice (although reasonable) structure, e.g., the aggregate buyer demand has a monotone hazard rate or is not too convex. We also show that our pricing schemes have implications for any solution which achieves high revenue: specifically that in many settings, prices that maximize (approximately) profit also result in high social welfare. Our results holds for general multi-minded buyers in large markets with production costs; we also provide improved guarantees for the important special case of unit-demand buyers.

## 1 Introduction

Social Welfare and Profit<sup>1</sup> are the two canonical objectives in the extensive literature dealing with *envy-free* algorithmic pricing. The study of these two objectives, in isolation from each other, has inspired the design of novel pricing mechanisms for revenue maximization [4, 14] in a variety of interesting markets, and an equally voluminous body of work on welfare maximization [12, 16]. While the significance of profit and social welfare is clear, it is easy to overlook the fact that the two objectives do not exist in a vacuum. For instance, although a monopolistic seller may only be interested in profits, myopically increasing prices while compromising on buyer welfare can lead to poor long-term revenue. This is distinctly true for large markets with repeated engagement where singularly optimizing for one objective while ignoring the other (as in the existing literature) could adversely affect the health of the marketplace [3]. Therefore, not only is it desirable to promote the design of holistic pricing solutions that optimize on both counts simultaneously, it is also crucial to gain a better understanding

---

<sup>1</sup> For convenience, we will use revenue and profit interchangeably in this work.



of how existing algorithms perform in a *bicriteria* sense. Against this backdrop, we seek to address the following questions.

*What fraction of the optimum social welfare does a revenue maximizing solution achieve? Are there pricing mechanisms which achieve both good revenue and good welfare simultaneously?*

Both in economics and in computer science [18], it is well understood that the goals of maximizing revenue and social welfare are often at odds with each other. Bearing this in mind, we seek to quantify the exact amount of friction between these two objectives in large markets. In particular, we are interested in understanding the surplus achieved by a profit maximizing solution, a problem that has received considerable attention in Auction theory [1, 18]. The fact that we restrict our attention to the revenue end of the spectrum is motivated partly by the observation that welfare maximizing prices can result in negligible profits (see Example 3) even for trivial instances. However, unlike most analogous work in the theory of auctions, we are interested in understanding these trade-offs as well as designing bicriteria approximation algorithms in multi-item markets where the seller’s modus operandi involves posting prices on the individual goods. In that sense, this work is a high-level extension of the recent body of work on envy-free revenue-maximization [2, 14, 17] towards additional ambitious objectives.

### 1.1 Market Model: Item Pricing for Multi-minded Buyers

In this work, we adopt a simple posted-pricing mechanism that captures the operation of most real-life large markets: the seller posts a single price per good, and each buyer purchases a bundle of goods that maximizes their utility. The seller controls a set  $T$  of available goods, and can produce any desired quantity  $x_t$  of a good  $t \in T$ , for which he incurs a cost of  $C_t(x_t)$ . The market consists of a large number of buyers who are *multi-minded*, meaning that each buyer  $i$  has a ‘desired set of bundles of goods’: the buyer has the same value  $v_i$  for each of these bundles and under a given set of prices, purchases the bundle that maximizes her utility.

Multi-minded buyers represent a class of computationally attractive yet combinatorially non-trivial buyer valuations that have recently been featured in a number of papers [8, 19]. Perhaps, more importantly, the class strictly generalizes highly popular models such as *unit-demand* and *single-minded* valuations. Secondly, the convex production costs considered in our framework strictly generalize models with limited (or unlimited) supply, which are usually the norm in the pricing literature. As [2, 6] point out, limited supply is often too rigid for realistic, large markets where the seller may be able to increase production, albeit at a higher cost. Bicriteria algorithms notwithstanding, our work actually presents the first profit-maximization algorithms for general multi-minded buyers even with limited supply.

Our model captures several scenarios of interest wherein a typically profit maximizing seller may be driven to ensure good overall social utility. For

instance, consider a market for plug-in electric vehicle (PEV) charging stations: each good represents a time slot, and each buyer may desire specific (sets of) slots based on time constraints and charging capacity. Varying demand and electricity generation costs necessitate differential pricing across time slots [2, 5]. In such a large market application, it is clearly in the seller’s long-term interest to not drive away a significant population of its customer base.

**Circumventing Computational Complexity via Oblivious Guarantees.**

One of the challenges in essentially any non-trivial setting (including *all* the settings which we consider), is that computing profit-maximizing prices is NP-Hard. This is largely due to the fact that the seller is not allowed to price-discriminate, i.e., it must charge the same price for each good to all the buyers, instead of having different prices for each buyer. In view of the computational barriers surrounding the optimal profit solution, a seller which cares about profit may use a variety of strategies, from approximation algorithms to heuristics. The uncertainty regarding the actual strategy adopted by the seller in turn casts aspersions on the practical significance of our goal of characterizing the social welfare at optimal-revenue solutions. One of the contributions of this work is a framework that allows us to completely circumvent the complexity question: *our guarantees on the social welfare do not depend on the exact details of the pricing mechanism used by the seller, and instead would hold for a wide variety of pricing mechanisms, as long as these prices achieve decent revenue guarantees.*

**Inverse Demand Functions and  $\alpha$ -Strong Regularity.** In order to concisely represent the large number of buyers in the market, we classify the buyers into a finite set of buyer types  $B$  such that all of the multi-minded buyers belonging to a certain type desire the same set of bundles. Then, each buyer type can be fully captured by a subset of  $2^T$  along with an *inverse demand distribution*  $\lambda_i(x)$  describing the valuations of buyers having this type. Formally, for any buyer type  $i \in B$ ,  $\lambda_i(x) = p$  implies exactly  $x$  amount of buyers of type  $i$  have a valuation of  $p$  or more for each bundle in their common desired set. Although different buyer types may have different demand functions, it is natural to assume that the valuations of all buyers are often sampled (albeit differently) from some global distribution. Because of this, we will make the assumption that the buyer valuations for every type have the same support  $[\lambda^{min}, \lambda^{max}]$ .

A first stab at the problem reveals that the above framework is too coarse to obtain meaningful trade-offs between welfare and profit. Indeed, it is not hard to reason that a precise characterization of the revenue-welfare trade-offs would depend heavily on the distributions of the buyer valuations. To better understand this dependence, we study a class of inverse demand functions parameterized by a single parameter  $\alpha \in [0, 1]$  known as  $\alpha$ -strongly regular distributions.

**Definition 1 ( $\alpha$ -Strong Regularity [10]).** *A buyer type  $i$  is said to have an  $\alpha$ -Strongly regular demand function ( $\alpha$ -SR) for  $\alpha \in [0, 1]$  if for any  $x_1 < x_2$ , we have  $\frac{\lambda_i(x_2)}{|\lambda_i(x_2)|} - \frac{\lambda_i(x_1)}{|\lambda_i(x_1)|} \leq \alpha(x_2 - x_1)$ .*

$\alpha$ -Strongly regular distributions were introduced in [10] as a strict generalization of *monotone hazard rate* (MHR) distributions that smoothly interpolate between

MHR ( $\alpha = 0$ ) and regular distributions ( $\alpha = 1$ ). Our main contribution is the design of mechanisms that simultaneously obtain good revenue and welfare for small  $\alpha$ , and degrade gracefully as  $\alpha$  increases. Note that even the set of  $\alpha$ -SR functions with  $\alpha = 0$ , for which we obtain the strongest results, contains a large class of important distributions, including exponential (e.g.,  $e^{-x}$ ), polynomial (e.g.,  $1 - x^2$ ), and all log-concave functions. The reader is asked to refer to Sect. 2 for a more detailed discussion regarding this class of distributions.

## 1.2 Our Contributions

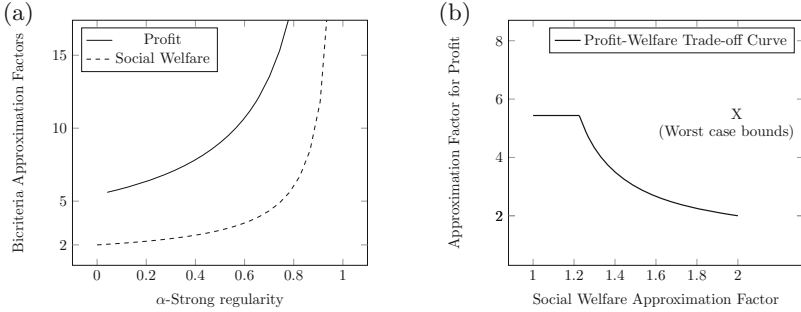
The primary algorithmic contribution of this work is a new (profit, welfare)-bicriteria approximation for general markets with multi-minded buyers and production costs, stated below.

**(Informal Theorem).** *We can compute in poly-time a set of item prices that guarantee a  $\Theta(\frac{\log \Delta}{1-\alpha})$ -approximation to the optimum profit, and a  $\Theta(\frac{1}{1-\alpha})$ -approximation for welfare, where  $\Delta$  is the ratio of the size of the largest bundle desired by any buyer to the smallest one.*

There are several exciting aspects to this result: (i) Ours is also the first non-trivial profit-maximization algorithm for multi-minded buyers in the envy-free literature. (ii) When buyers desire small bundles (e.g., for unit-demand valuations where all bundles are unit-size), our solution extracts a constant portion of the social welfare as revenue, as illustrated in Fig. 1(a). Moreover, for the important special case of unit-demand valuations ( $\Delta = 1$ ), we provide a much simpler pricing mechanism with slightly better constant approximation factors than in the theorem statement. (iv) Finally, even when buyers desire large bundles, it is reasonable to expect that in markets with similar types of goods, the various bundles are of approximately the same size, i.e.,  $\Delta$  is small. In the PEV example, one expects different electric vehicles to have similar charging capacity.

**Profit-Welfare Trade-Offs.** The approximation guarantees in the theorem statement are only the worst-case bounds derived independently for each objective. In fact, as illustrated in Fig. 1(b), we prove that the two worst-case factors never occur simultaneously and the actual bicriteria bound lies on a trade-off curve, resulting in improved approximations for at least one objective, i.e., if the actual welfare is close to the worst-case guarantee, then the profit is much better than in the theorem and vice-versa.

**Social Welfare of Other Revenue-Maximizing Solutions.** All of the revenue guarantees in this paper are shown by comparing the profit of our solution to the optimum welfare, an approach that has strong implications towards bounding the social welfare of other profit-maximizing solutions. Specifically, we design a simple framework using our bicriteria bounds as a black-box result and show that *any* pricing mechanism which achieves revenue better than our (efficiently computable) mechanism is guaranteed to deliver at least a  $\Theta(\frac{\log \Delta}{1-\alpha})$ -approximation to the optimum welfare. This holds whether the seller computes



**Fig. 1.** (a) Exact bounds for profit and welfare as a function of  $\alpha$  for unit-demand buyers. We obtain constant-factor bicriteria approximations when the demand is close to *MHR* ( $\alpha = 0$ ), and good guarantees for larger  $\alpha$  even when the demand is close to the *equal-revenue distribution* ( $\alpha = 1$ ). (b) Actual Revenue-Welfare Curve for unit-demand buyers,  $\alpha = 0$ : The exact bicriteria guarantees lie on the trade-off curve, so that one of the two approximation factors is significantly better than the worst-case bound (point X). For e.g., when welfare is only half-optimal, the revenue factor improves from  $2e$  to 2.

revenue-maximizing prices (an NP-Hard problem even for unit-demand with  $\alpha = 0$ ), or uses a more efficient mechanism. Thus, one of the main messages of this paper is that even a seller interested solely in maximizing profits can guarantee a good social welfare without sacrificing any revenue. For example, in unit-demand markets with MHR valuations, there is no reason for such a seller to not also achieve at least a  $2e$ -approximation to the optimum social welfare irrespective of their preferred pricing mechanism.

**Technical Difficulties.** Although our large market model falls in the realm of settings where it is possible to efficiently compute social welfare maximizing prices, exploiting this for profit-maximization as in [4, 14] leads to poor approximation guarantees, for e.g.,  $O(\frac{\lambda^{max}}{\lambda^{min}})$ -bounds even for unit-demand instances. Instead, our techniques rely crucially on exploiting the structure of  $\alpha$ -strongly regular functions to efficiently compute prices that compromise neither on revenue nor welfare. Finally, in this work, we will focus solely on deterministic pricing mechanisms. While randomized mechanisms that mix between welfare and profit maximizing solutions are theoretically interesting, the ensuing price fluctuations render them unsuitable for many settings of interest [13, 20].

**Related Work: Existing Bicriteria Algorithms.** The primary barrier towards designing envy-free revenue-maximizing prices — a lack of insight regarding the optimum solution — is also the chief architect behind the existence of many (implicit) bi-criteria approximation algorithms in the algorithmic pricing literature. More concretely, a majority of the revenue-maximization algorithms in the literature [4, 7, 14, 17] achieve their approximation factors for revenue by comparing it to the optimum social welfare. Exploiting these revenue-

welfare ties further, it is not hard to see that such a  $\beta$ -approximation algorithm for revenue trivially results in a  $(\beta, \beta)$ -bicriteria approximation. For instance, the results from [4, 14] immediately imply  $(\Theta(\log |B|), \Theta(\log |B|))$ -bicriteria approximation algorithms for unit-demand and unlimited supply markets respectively.

In contrast to the trivial  $(\beta, \beta)$  type bounds in previous work, our specific focus on bicriteria approximations leads to significant improvements in social welfare without sacrificing much profit. We also remark that while specific bicriteria bounds were also provided in [2], their results only apply to the easiest version of our setting ( $\alpha = 0, \Delta = 1$ ). Moreover, the setting studied in this work, that of multi-minded buyers and production costs is significantly more general than most previous work on envy-free pricing, which looked at unit-demand or single-minded valuations in markets with (un)limited supply.

### 1.3 Other Related Work

While (revenue,welfare)-bicriteria approximations have not specifically been studied beyond the single good case, the broader understanding of trade-offs between the two objectives has been a prominent motif in the pricing literature. We first highlight two overarching differences between our results and other work on profit-welfare trade-offs, specifically in auctions: (i) we study reasonably general combinatorial markets with multi-minded buyers and production cost functions, and not just limited-supply settings with unit-demand buyers as in other work, and (ii) unlike similar (types of) results in Bayesian auctions, our pricing mechanisms are non-discriminatory, and therefore, envy-free.

Characterizing the efficiency of revenue-optimal mechanisms is a fundamental question that has spurred multiple avenues of research; most pertinent to the questions posed in this work are the tight bounds on the (in)efficiency of the Myerson revenue-maximizing mechanism for single good settings appearing in [1, 18]: in particular, [18] provides welfare bounds for general single-parameter auctions as a function of the distribution of buyer valuations, as we do in this work.

Moving beyond lower bounds, other researchers have adopted a more constructive approach by explicitly taking into account both the objectives either via bicriteria mechanisms [11, 22], or by optimizing linear combinations of revenue and welfare [20], or even characterizing the revenue-welfare Pareto curve [13]. We reiterate that all of the above papers consider simple single good settings, where the revenue-optimal mechanism is well understood. Moreover, in comparison to the revenue-welfare Pareto curves in [13], the implicit trade-off curves in our work are of a different nature as they are obtained for a single instance on top of the worst-case bounds. Finally, multi-objective trade-offs are quite popular in the Sponsored Search literature [3, 12, 21], which are repeated engagement markets with tight competition. Such settings can essentially be viewed as a special case of unit-demand markets.

## 2 Model and Preliminaries

Our market model comprises of a single seller controlling a set  $T$  of goods and a large number of infinitesimal buyers. The buyers can be concisely represented using a finite set of multi-minded buyer types  $B$ : for a given type  $i \in B$ , all the buyers having this type desire the same set of item bundles  $B_i \subseteq 2^T$ , and each buyer is indifferent between the bundles in  $B_i$ . Notice that when all of the desired bundles are of unit cardinality, our model reduces to the *unit-demand* case; when each buyer type desires only a single bundle ( $|B_i| = 1$ ), we get *single-minded* valuations. Finally, buyers belonging to the same type may hold different valuations for the same bundles, this is modeled by way of an *inverse demand function*  $\lambda_i(x)$  for every  $i \in B$ ;  $\lambda_i(x) = p$  implies that exactly  $x$  amount of buyers of type  $i$  value the bundles in  $B_i$  at valuation  $p$  or more. Given  $\lambda_i(x) = p$ , it is not hard to see that the total utility derived when  $x$  amount of buyers purchase some bundle at price  $p$  is  $u_i(x) = \int_{z=0}^x \lambda_i(z) dz$ .

The market operates according to a natural pricing mechanism with the seller posting a price  $p_t$  for each good  $t \in T$ . Buyers purchase one of the utility-maximizing bundles available to them, i.e., a buyer belonging to type  $i$  will purchase the cheapest bundle in  $B_i$  as long as its price is no larger than her valuation for the same. Therefore, if  $\bar{p}_i$  denotes the bundle in  $B_i$  with the smallest price and  $\bar{x}_i$  is the population of buyers of this type who purchased some bundle, then  $\lambda_i(\bar{x}_i) = \bar{p}_i$ .

**Pricing Solutions, Social Welfare, and Revenue.** We use  $(\mathbf{p}, \mathbf{x}, \mathbf{y})$  to represent the outcome of the market mechanism. Here  $\mathbf{p}$  is the vector of prices,  $\mathbf{x}$  denotes the allocation to the buyers with  $x_i$  being the total amount of good purchased by buyers of type  $i$ , and finally  $y_t$  is the total amount of good  $t \in T$  sold to the buyers. We now define the two main metrics that form the crux of this paper.

- The social welfare of a solution  $(\mathbf{p}, \mathbf{x}, \mathbf{y})$  is defined to be the total utility of all the buyers and the seller and therefore, is equal to the utility of the buyers minus the production cost incurred by the seller, i.e.,

$$SW(\mathbf{p}, \mathbf{x}, \mathbf{y}) = \sum_{i \in B} \int_{x=0}^{x_i} \lambda_i(x) dx - \sum_{t \in T} C_t(y_t).$$

- The (seller's) profit at the solution  $(\mathbf{p}, \mathbf{x}, \mathbf{y})$  is the total income due to each good in the market minus the total production cost incurred, i.e.,

$$\pi(\mathbf{p}, \mathbf{x}, \mathbf{y}) = \sum_{t \in T} [p_t y_t - C_t(y_t)].$$

Notice that the social welfare is independent of the prices, and depends only on  $(\mathbf{x}, \mathbf{y})$ . One of the main goals of this paper is to obtain a lower bound on the social welfare of the profit maximizing solution. We will use  $(\mathbf{p}^{opt}, \mathbf{x}^{opt}, \mathbf{y}^{opt})$

to denote the maximum profit solution, and  $(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*)$  to denote the solution maximizing welfare. Sometimes we will also use  $SW^* = SW(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*)$  and  $\pi^{opt} = \pi(\mathbf{p}^{opt}, \mathbf{x}^{opt}, \mathbf{y}^{opt})$ . Thus the quantity we are interested in is  $\frac{SW(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*)}{SW(\mathbf{p}^{opt}, \mathbf{x}^{opt}, \mathbf{y}^{opt})}$ .

We remark that the maximum social welfare solution should ideally be represented as  $(\mathbf{x}^*, \mathbf{y}^*)$ . However, it is not particularly hard to see that there always exist prices  $\mathbf{p}^*$ , defined as  $p_t^* = c_t(y_t^*)$  for all  $t \in T$ , such that  $(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*)$  is a valid (envy-free) solution of our pricing problem. More importantly, such a welfare maximizing solution can be computed efficiently via a simple convex program.

**Structure of the Demand and Cost Functions.** In this work, we take both the inverse demand and production cost functions to be continuously differentiable. In addition, we also make the standard assumption that the utility function  $u_i(x)$  is concave for all  $i \in B$  and therefore its derivative  $\lambda_i(x)$  is non-increasing with  $x$ . Finally, we consider production cost functions that are *doubly convex* i.e., both  $C_t$  and its derivative  $c_t(x) = \frac{d}{dx}C_t(x)$  are convex and non-decreasing for all  $t \in T$  and further,  $C_t(0) = c_t(0) = 0$ . A number of well studied cost functions fall within our framework [6].

As mentioned previously, it is often natural to assume that the inverse demand distributions for different buyer types have the same support  $[\lambda^{min}, \lambda^{max}]$ . In fact, *all* of our results hold under the more general *uniform peak* assumption, which will be assumed for the rest of this work.

**Definition 2.** (*Uniform Peak Assumption*) For every  $i \in B$ ,  $\lambda_i(0) = \lambda^{max}$ .

**$\alpha$ -SR inverse demand** (Definition 1) Recently,  $\alpha$ -strong regularity has gained some popularity [9, 10] as an elegant characterization of the class of regular functions, which encompasses most well-studied demand distributions including polynomial ( $\lambda(x) = 1 - x^2$ ;  $\alpha = 0$ ), exponential ( $\lambda(x) = e^{-x}$ ;  $\alpha = 0$ ), power law ( $\lambda(x) = \frac{1}{\sqrt{x}}$ ;  $\alpha = \frac{1}{2}$ ), and the equal-revenue distribution ( $\lambda(x) = \frac{1}{x}$ ;  $\alpha = 1$ ). For such functions, one can interpret  $\alpha$  as a measure of the convexity of the function as larger values of  $\alpha$  imply greater convexity or alternatively as a bound on the volatility of the inverse demand  $\lambda_i(x)$  as every  $\alpha$ -SR demand function satisfies  $\frac{d}{dx} \left( \frac{\lambda(x)}{|\lambda'(x)|} \right) \leq \alpha$ . As expected, the equal-revenue distribution ( $\alpha = 1$ ) leads to the worst-case bounds for all of our results: in fact even in single good, single buyer type markets, the revenue-optimal solution for  $\alpha = 1$  only extracts a negligible fraction of the optimum social welfare. However, what is surprising (as evidenced by Fig. 1(a)) is that we obtain reasonably good performance guarantees even when  $\alpha$  is larger than  $\frac{1}{2}$ .

### 3 Warm-Up: Profit and Welfare for Unit-Demand Markets

As a first step towards stating our general results in Sect. 5, we consider the important special case of unit-demand markets [7, 14]. Recall that unit-demand

valuations are a simple sub-class of multi-minded functions, wherein for each buyer (type)  $i \in B$ , the bundles desired by  $i$  are singleton sets. Our main result in this section is a simple pricing rule that achieves a  $(\Theta(\frac{1}{1-\alpha}), \frac{2-\alpha}{1-\alpha})$ -bicriteria approximation algorithm for revenue and welfare respectively when the buyers have  $\alpha$ -SR inverse demand functions. While we present a generalized version of this theorem in Sect. 5, the algorithm in that case is much more involved.

That said, our reasons for dedicating an entire section to unit-demand buyers is two-fold: (i) the algorithm presented in this section is extremely simple and the constant factors hidden by the asymptotic bound are smaller, and (ii) the unit-demand case provides a platform for us to discuss the various implications of our results including the profit-welfare trade-off and the ability to derive welfare bounds for the profit-maximizing solution. Before stating our main theorem, we give a simple example to highlight the poor revenue guarantees obtained by the welfare maximizing prices even for single-good markets.

*Example 3.* Consider a single good market with a negligible production cost function, say  $C(x) = \epsilon x$ . Obviously, there is only one buyer type, and suppose that its inverse demand is  $\lambda(x) = 1 - x$  ( $\alpha$ -SR for  $\alpha = 0$ ). It is easy to observe that at the social welfare maximizing solution, the good is priced at  $p^* = \epsilon$ , resulting in near-zero revenue. On the contrary, one can price at  $p^{opt} = \frac{1}{2}$  to obtain a constant fraction of the optimum welfare as profit.

**Theorem 4.** *For any unit-demand instance where buyers have  $\alpha$ -strongly regular inverse demand functions, there is a poly-time  $(\zeta, \frac{2-\alpha}{1-\alpha})$ -bicriteria approximation algorithm for revenue and social welfare respectively, where*

$$\zeta = 2\left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} + \frac{\alpha}{1-\alpha} = \Theta\left(\frac{1}{1-\alpha}\right).$$

The exact guarantees for profit and welfare are illustrated in Fig.1(a). Detailed proofs of all of our results can be found in the full version of this paper. (Algorithm) The bicriteria approximation factor is achieved by the following simple pricing mechanism

- Compute the max-welfare solution  $(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*)$ .
- For every good  $t$ , set its price  $\tilde{p}_t = \max(p_t^*, \lambda^{max}(1 - \alpha)^{\frac{1}{\alpha}})$ .

**Proof Sketch:** We now give some intuition for why this pricing solution produces a good approximation to both profit and welfare. Let  $\tilde{p} = \lambda^{max}(1 - \alpha)^{\frac{1}{\alpha}}$ . We begin by analyzing these types of *thresholded* pricing schemes, in which the price is simply the maximum of the optimum price  $p_t^*$  and a constant; our algorithm uses such a scheme with constant  $\tilde{p}$ . In the following lemma, we show that such solutions have nice structure; essentially we can think of buyers who purchase goods priced at  $\tilde{p}$  and those who purchase goods priced at  $p_t^*$  as separate systems.

**Lemma 5.** *Suppose that  $((\tilde{p})_{t \in T}, (\tilde{x})_{i \in B}, (\tilde{y})_{t \in T})$  is a pricing solution resulting from a thresholded pricing vector. Then,*



1. The market can be clustered into two mutually disjoint sets of buyers and goods  $(B^H, T^H)$  and  $(B^L, T^L)$  so that the buyers in each cluster only purchase the goods in the same cluster and (a) for  $(i, t) \in (B^H, T^H)$ ,  $\tilde{p}_t = p_t^*$ ,  $\tilde{x}_i = x_i^*$ , and  $\tilde{y}_t = y_t^*$ ; (b) for  $(i, t) \in (B^L, T^L)$ ,  $\tilde{p}_t \geq p_t^*$ , and  $c_t(\tilde{y}_t) \leq c_t(y_t^*)$ .
2.  $\tilde{\mathbf{y}}$  is a welfare-maximizing allocation with respect to the demand vector  $\tilde{\mathbf{x}}$ .

We then prove that due to our choice of  $\tilde{p}$  the following two bounds hold: (i)  $SW(\tilde{\mathbf{p}}, \tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \leq O(\frac{1}{1-\alpha})\pi(\tilde{\mathbf{p}}, \tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ , and (ii)  $SW(\mathbf{p}^*, \mathbf{x}^*, \mathbf{y}^*) - SW(\tilde{\mathbf{p}}, \tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \leq O(\frac{1}{1-\alpha})\pi(\tilde{\mathbf{p}}, \tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ . Once we show these lower bounds on the profit of our pricing scheme, the approximation guarantee follows trivially.  $\square$

Henceforth, we will unambiguously use  $\zeta$  to denote the exact approximation factor appearing in the statement of the above theorem. Interestingly, in the proof of Theorem 4, the revenue guarantee of  $\zeta$  is shown with respect to the optimum social welfare, which in turn has important consequences for other profit-maximizing solutions (see Sect. 4).

**Profit-Welfare Trade-offs.** We further exploit the close ties between revenue and social welfare, and present a revenue-welfare trade-off that improves upon the bicriteria bound in Theorem 4 by showing that at least one of revenue or welfare is better than the factor guaranteed by the theorem. The bounds in Theorem 4 are actually somewhat misleading as they represent the worst-case bound for each objective, which is derived independently from the other objective by simply bounding the worst-case revenue (or welfare) over all instances. However, the worst-case performance for revenue ( $\zeta$ ) and the worst-case performance for social welfare ( $\frac{2-\alpha}{1-\alpha}$ ) need not and as we show, *do not* occur for the same instance. As a matter of fact, for a given instance, if the actual welfare obtained is close to the guarantee provided in Theorem 4, the large gap between welfare and revenue as in Fig. 1(a) completely vanishes and the approximation factors coincide. We first state the main structural claim that enables this trade-off.

*Claim.* Suppose that a pricing algorithm  $Alg$  satisfies  $SW(Alg) \leq c_1\pi(Alg)$ ,  $SW^* - SW(Alg) \leq c_2\pi(Alg)$ . Then for every instance there exists some  $1 \leq c \leq c_2 + 1$  such that  $Alg$  is a bicriteria approximation  $(\min(cc_1, \frac{cc_2}{c-1}), c)$  for revenue and welfare respectively.

Unfortunately, the designer has no control over the factor  $c$  as its exact value depends on the particular instance. Applying this claim to Theorem 4 yields the following corollary.

**Corollary 6.** *For every given instance, there exists a constant  $1 \leq c \leq \frac{2-\alpha}{1-\alpha}$  so that the solution returned by the prices of Theorem 4 has a social welfare that is within a factor  $c$  of the optimum welfare and such that*

$$\pi^{opt} \leq SW^* \leq \min\left(\frac{c}{c-1}, \frac{1}{1-\alpha}, \zeta\right)\pi(Alg).$$

For example, for MHR demand functions the statement of Theorem 4 makes it seem that this pricing scheme may return a solution which is a 2e-approximation for revenue and a 2-approximation for welfare. Corollary 6 points out that the actual results are far *better*.

## 4 Consequences for Solutions with High Revenue

In Sects. 3 and 5, we give efficient pricing mechanisms which simultaneously achieve good approximations for both revenue and welfare. Consider, however, a seller whose main priority is to simply maximize profits. This seller may choose to use a different pricing mechanism with better revenue guarantees than the ones offered in this paper. For example, the seller may choose prices which are guaranteed to come closer to achieving optimum revenue (these are efficiently computable for unit-demand settings [2, 15] under certain additional assumptions), or even use a large amount of resources to solve the intractable problem of actually computing the prices  $\mathbf{p}^{opt}$  which yield the highest possible revenue. One of the main messages of our paper is as follows:

*No matter what pricing mechanism the seller uses to optimize revenue, they can instead use a pricing mechanism which guarantees at least  $1/\zeta$  fraction of the optimum welfare, without sacrificing any revenue.*

In this section, we use a simple albeit highly general framework to derive results of this form. Consider an arbitrary profit maximization algorithm  $Alg$  that achieves a good approximation with respect to  $\pi^{opt}$ . How do we go about characterizing the social welfare at these solutions? The following theorem uses an existing profit-maximization algorithm whose guarantees hold with respect to the optimum welfare as a black-box to bound the welfare due to  $Alg$ .

**Theorem 7.** *Consider a benchmark profit maximization algorithm  $Alg^b$  whose profit  $\pi(Alg^b)$  is always within a factor  $c$  of  $SW^*$  for some  $c \geq 1$ . Consider any other pricing algorithm  $Alg$  for the same class of valuations that obtains at least as much profit as guaranteed by  $Alg^b$  on all instances. Then the social welfare obtained by  $Alg$  is at least a factor  $\frac{1}{c}$  times that of the optimum welfare.*

**Implications for Unit-Demand Markets.** We now apply the framework provided by Theorem 7 for unit-demand using our result from Theorem 4 as a black-box to obtain our first oblivious guarantee: namely *that any algorithm for unit-demand markets that obtains at least as much profit as guaranteed by the algorithm of Theorem 4 on all instances guarantees a social welfare that is within a factor  $\zeta$  of the optimum welfare.*

Thus, consider the case when a seller is using any arbitrary pricing mechanism  $Alg$ , with the main goal being to maximize profit. By simply computing the revenue given by our pricing schemes from Sect. 3, and then choosing the one which guarantees better revenue (i.e., choosing between  $Alg$  and our pricing scheme), we form a new pricing algorithm which does not sacrifice any revenue compared to  $Alg$ , and due to the above theorem, is *also* guaranteed to have good social welfare. Moreover, we also get the following trivial consequence:

**Corollary 8.** *The ratio of the optimum social welfare  $SW^*$  to the social welfare at the maximum profit solution  $SW(\mathbf{p}^{opt}, \mathbf{x}^{opt}, \mathbf{y}^{opt})$  is at most  $\zeta = \Theta\left(\frac{1}{1-\alpha}\right)$ .*

For MHR demand functions ( $\alpha = 0$ ), this implies that even for sellers who only care about profits, there is essentially no excuse not to also guarantee at least  $1/2e$  of the optimum social welfare. Thus, for the settings we consider, one can strive for truly high revenue, without sacrificing much in welfare.

## 5 Multi-minded Buyers

We now move on to our most general case with multi-minded buyers, wherein every buyer wishes to purchase one bundle from a desired set (of bundles). We use  $\ell^{max}$  to denote the cardinality of the maximum sized bundle desired by any buyer type, and  $\ell^{min}$  for the minimum sized bundle. The main result in this section is a bicriteria approximation algorithm that extends our results for the unit-demand case. Our algorithm still achieves a  $\Theta(\frac{1}{1-\alpha})$ -approximation to the optimum welfare; as for profit, we obtain a  $\Theta(\frac{1}{1-\alpha})$  bound further discounted by a  $\log(\Delta)$  factor, where  $\Delta = \frac{\ell^{max}}{\ell^{min}}$ . Moreover, as in the previous section, our profit bound is obtained in terms of the optimum social welfare, which allows the consequences mentioned in Sect. 4 to hold, i.e., we are able to show that the true lower bounds are better than the theorem states.

**Theorem 9.** *For any given instance with multi-minded buyers, there exists a poly-time  $(\Theta(\frac{\log(\Delta)}{1-\alpha}), \Theta(\frac{1}{1-\alpha}))$ -bicriteria approximation algorithm for profit and welfare respectively.*

**Proof Sketch:** The proof is quite involved, so here we provide a high level overview of the various ingredients that combine to form the proof.

*Step 1: Benchmark Solution.* The first step involves defining a benchmark solution  $(\mathbf{x}^b, \mathbf{y}^b)$  such that its social welfare  $SW^b$  and the *non envy-free profit* ( $\pi^b := \sum_{i \in B} \lambda_i(x_i^b)x_i^b - C(\mathbf{y}^b)$ ) obtained via discriminatory payments both approximate  $SW^*$  up to a  $\Theta(\frac{1}{1-\alpha})$  factor. However, this does not immediately lead to a bicriteria approximation as there may not exist any pricing vector that implements such a solution. Therefore, we will attempt to compute item prices which approximate the benchmark solution with respect to both the objectives.

*Step 2: Augmented Walrasian Equilibrium.* As a first step towards approximating the benchmark solution, we extend the notion of a *Walrasian equilibrium with reserve prices* that was introduced in [14] to settings where buyers purchase bundles of arbitrary sizes. Such a solution does not admit any infinitesimal buyer whose valuation is smaller than the reserve price and can be easily implemented using a pricing solution. Our plan then, is to form a sequence of equilibrium solutions resulting from different reserve prices that together behave like the benchmark solution; then, we will consider the ‘best’ among these solutions.

*Step 3: Sequence of Solutions.* We identify a starting point for our sequence of solutions via a crucial claim showing that there exists a carefully chosen reserve  $\tilde{r}$  such that the Walrasian equilibrium at this reserve extracts a good fraction of  $SW^b$  and, when buyers desire bundles of the same cardinality, also results in good profit compared to  $\pi^b$ . When there is a large disparity in the bundle sizes,

however, the profit returned by the solution does not meet our approximation guarantee as buyers may gravitate towards the smaller sized bundles. To fix this, we consider reserve prices that are scaled versions of the original reserve  $\tilde{r}$ .

In conclusion, our algorithm works by computing a series of pricing solutions corresponding to the Walrasian equilibria at reserve prices  $2^j \tilde{r}$  for  $j = 0$  to  $1 + \log(\Delta)$  (with welfare  $SW(j)$ , profit  $\pi(j)$ ) and returning the smallest index of  $j$  at which the profit equals the term guaranteed by the theorem. The rest of the proof follows from the following chain of lemmas that we prove: (i)  $SW(0)$  is comparable to  $SW^b$ , (ii)  $SW(j) - SW(j + 1) \leq 3\pi(j) + 3\pi(j + 1)$ , and (iii)  $SW(1 + \log(\Delta)) \leq O(\frac{1}{1-\alpha})\pi(1 + \log(\Delta))$ .  $\square$

**Consequences for Other High-Revenue Solutions.** A direct application of Theorem 7 using our newly obtained bounds on multi-minded buyers as an intermediate yields the following claim.

*Claim.* Let  $Alg$  be any algorithm that obtains at least as much profit as guaranteed by the algorithm of Theorem 9 on all instances. Then the social welfare obtained by  $Alg$  is at most a factor  $\Theta(\frac{\log(\Delta)}{1-\alpha})$  away from the optimum welfare.

## 6 Conclusions and Future Directions

In this work, we were able to provide envy-free posted pricing algorithms that simultaneously approximate both profit and social welfare for markets with quite general buyer valuations and production costs. We used our profit-maximization guarantees as a black-box and showed that any solution with reasonable profit guarantees (including the maximum profit solution) generates good welfare. In the process, we provide a partial characterization of the exact friction between these two objectives. The multi-objective approach to pricing taken in this work is motivated by the fact that different types of agents in the system care about different objectives, i.e., sellers care about maximizing profit whereas larger welfare benefits the buyers. Given that social welfare is a combination of both profit and buyer surplus, an equally natural bicriteria approach would involve the maximization of profit and surplus, as opposed to welfare. Unfortunately, one can design simple instances where the only possible non-trivial envy-free pricing solutions result in zero surplus. This raises a more fundamental open question: can we identify markers in repeated engagement markets other than welfare and profit that correspond to the needs of the various agents, and that admit non-trivial multi-objective approximation guarantees?

All the omitted proofs can be found in a full version of this paper available publicly on arXiv: <https://arxiv.org/abs/1610.04071>.

**Acknowledgements.** This work was supported in part by NSF awards CCF-1527497 and CNS-1218374.

## References

1. Abhishek, V., Hajek, B.E.: Efficiency loss in revenue optimal auctions. In: Proceedings of IEEE CDC (2010)
2. Anshelevich, E., Kar, K., Sekar, S.: Envy-free pricing in large markets: approximating revenue and welfare. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015. LNCS, vol. 9134, pp. 52–64. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-47672-7\\_5](https://doi.org/10.1007/978-3-662-47672-7_5)
3. Bachrach, Y., Ceppi, S., Kash, I.A., Key, P., Kurokawa, D.: Optimising trade-offs among stakeholders in ad auctions. In: Proceedings of ACM EC (2014)
4. Balcan, M.-F., Blum, A., Mansour, Y.: Item pricing for revenue maximization. In: Proceedings of ACM EC (2008)
5. Bhattacharya, S., Kar, K., Chow, J.H., Gupta, A.: Extended second price auctions for plug-in electric vehicle (PEV) charging in smart distribution grids. In: Proceedings of IEEE American Control Conference (ACC) (2014)
6. Blum, A., Gupta, A., Mansour, Y., Sharma, A.: Welfare and profit maximization with production costs. In: Proceedings of FOCS (2011)
7. Briest, P., Krysta, P.: Buying cheap is expensive: approximability of combinatorial pricing problems. *SIAM J. Comput.* **40**(6), 1554–1586 (2011)
8. Buchbinder, N., Gonen, R.: Incentive compatible multi-unit combinatorial auctions: a primal dual approach. *Algorithmica* **72**(1), 167–190 (2015)
9. Cole, R., Rao, S.: Applications of  $\alpha$ -strongly regular distributions to bayesian auctions. In: Markakis, E., Schäfer, G. (eds.) WINE 2015. LNCS, vol. 9470, pp. 244–257. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-48995-6\\_18](https://doi.org/10.1007/978-3-662-48995-6_18)
10. Cole, R., Roughgarden, T.: The sample complexity of revenue maximization. In: Proceedings of STOC (2014)
11. Daskalakis, C., Pierrakos, G.: Simple, optimal and efficient auctions. In: Chen, N., Elkind, E., Koutsoupias, E. (eds.) WINE 2011. LNCS, vol. 7090, pp. 109–121. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25510-6\\_10](https://doi.org/10.1007/978-3-642-25510-6_10)
12. Deng, X., Goldberg, P., Sun, Y., Tang, B., Zhang, J.: Pricing ad slots with consecutive multi-unit demand. In: Vöcking, B. (ed.) SAGT 2013. LNCS, vol. 8146, pp. 255–266. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-41392-6\\_22](https://doi.org/10.1007/978-3-642-41392-6_22)
13. Diakonikolas, I., Papadimitriou, C., Pierrakos, G., Singer, Y.: Efficiency-revenue trade-offs in auctions. In: Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.) ICALP 2012. LNCS, vol. 7392, pp. 488–499. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-31585-5\\_44](https://doi.org/10.1007/978-3-642-31585-5_44)
14. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: Proceedings of SODA (2005)
15. Hartline, J.D., Koltun, V.: Near-optimal pricing in near-linear time. In: Dehne, F., López-Ortiz, A., Sack, J.-R. (eds.) WADS 2005. LNCS, vol. 3608, pp. 422–431. Springer, Heidelberg (2005). doi:[10.1007/11534273\\_37](https://doi.org/10.1007/11534273_37)
16. Hsu, J., Morgenstern, J., Rogers, R.M., Roth, A., Vohra, R.: Do prices coordinate markets? In: Proceedings of STOC (2016)
17. Im, S., Pinyan, L., Wang, Y.: Envy-free pricing with general supply constraints for unit demand consumers. *J. Comput. Sci. Technol.* **27**(4), 702–709 (2012)
18. Kleinberg, R., Yuan, Y.: On the ratio of revenue to welfare in single-parameter mechanism design. In: Proceedings of ACM EC (2013)
19. Krysta, P., Telelis, O., Ventre, C.: Mechanisms for multi-unit combinatorial auctions with a few distinct goods. *J. Artif. Intell. Res. (JAIR)* **53**, 721–744 (2015)

20. Likhodedov, A., Sandholm, T.: Mechanism for optimally trading off revenue and efficiency in multi-unit auctions. In: Proceedings of ACM EC (2004)
21. Lucier, B., Paes Leme, R., Tardos, E.: On revenue in the generalized second price auction. In: Proceedings of WWW (2012)
22. Sivan, B., Syrgkanis, V., Tamuz, O.: Lower bounds on revenue of approximately optimal auctions. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 526–531. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-35311-6\\_42](https://doi.org/10.1007/978-3-642-35311-6_42)

# A Prior-Independent Revenue-Maximizing Auction for Multiple Additive Bidders

Kira Goldner<sup>(✉)</sup> and Anna R. Karlin

Department of Computer Science and Engineering, University of Washington,  
Box 352350, Seattle, WA 98195, USA  
{kgoldner, karlin}@cs.washington.edu

**Abstract.** Recent work by Babaioff et al. [1], Yao [30], and Cai et al. [7] shows how to construct an approximately optimal auction for additive bidders, given access to the priors from which the bidders' values are drawn. In this paper, building on the single sample approach of Dhangwatnotai et al. [15], we show how the auctioneer can obtain approximately optimal expected revenue in this setting *without* knowing the priors, as long as the item distributions are regular.

**Keywords:** Mechanism design · Approximately optimal auctions · Prior-independence · Additive bidders

## 1 Introduction

In a *multiple additive bidders setting*, there are  $n$  agents and a seller selling a set of  $m$  distinct items. Each agent  $i$  has a private value  $v_{ij}$  for item  $j$ , and value  $v_i(S) = \sum_{j \in S} v_{ij}$  for the set of items  $S$ . The seller runs an auction to determine who (if anyone) to sell each item to and at what price. The auction (or mechanism) takes as input the collection of bids, and determines a feasible allocation and a price to charge each agent. The seller knows ahead of time the distribution from which each  $v_{ij}$  is drawn.<sup>1</sup> A key question is how to design a truthful and optimal<sup>2</sup> (or approximately optimal) auction.

This is a notoriously difficult problem, but in the past decade, several breakthrough results have been obtained. There are three main lines of work related to optimal auctions for additive bidders. For the case of finite type spaces, [3–6, 29] are able to use linear and convex programming techniques to formulate and solve the optimal auction problem. This gives a black-box reduction from mechanism design to algorithm design that yields a polynomial time algorithm for revenue maximization in additive settings. A second strand of work [1, 7, 20, 22, 30]

---

This research was done in part while the authors were visiting the Simons Institute for Theoretical Computer Science. The authors are funded by the National Science Foundation under CCF grant 1420381.

<sup>1</sup> Thus, from the seller's perspective this value is a random variable  $V_{ij}$ .

<sup>2</sup> i.e., revenue-maximizing, in expectation.

handles arbitrary distributions and develops approximately optimal auctions. Finally, [10–12, 16, 18, 19] use duality frameworks to optimally solve the problem for certain settings with a small number of items, and to provide necessary and sufficient conditions under which grand bundle selling is optimal.

In this paper, we consider the question of *prior-independent* optimal mechanism design in the multiple additive bidders setting. By prior-independent we mean two things: first, that there exist prior distributions from which the agents’ values are drawn (as in all the work discussed above), and, second, that the mechanism designer has no knowledge of these priors. Thus, without any knowledge of the priors, we seek to construct a mechanism that guarantees a constant fraction of the expected profit achieved by the optimal mechanism tailored to the particular prior distributions. This guarantee should hold no matter what the distributions happen to be, as long as they satisfy the fairly standard condition of regularity. A growing body of work obtains prior-independent mechanisms in a number of settings [13, 15, 17, 25, 26].

The main result of this paper is an auction that achieves this goal for the additive bidder setting when the  $V_{ij}$ ’s are all independent and drawn from regular distributions. We give a mechanism that requires only a single sample from the distribution of each  $V_{ij}$ , and when there are at least two bidders from any prior distribution, we can implement a sample mechanism as a prior-independent mechanism. Thus, we add to the short list of prior-independent results in multi-parameter settings [13, 26].

Our work builds on the breakthrough results of Babaioff, Imorlica, Lucier, and Weinberg [1] and Yao [30], on the one hand, and Dhangwatnotai, Roughgarden, and Yan [15] on prior-independent mechanism design on the other hand. A crucial lemma in [15]<sup>3</sup> is that, for a single-item single-bidder problem, access to a *single sample* from a regular distribution is sufficient to approximate the optimal revenue, which in this case is the revenue that results from pricing at the reserve price for the distribution.

Amazingly, the approximately optimal auctions of Babaioff et al., Yao, and Cai et al. essentially only use second-price auctions and reserve pricing for either single items or bundles of items, and therefore, the single sample paradigm nearly suffices to construct a prior-independent version of these auctions. There is only one detail to resolve and that relates to the issue of pricing bundles: the sum of regular random variables is not necessarily regular. However, delving into the proof from Babaioff et al., we find that the solution to this problem essentially writes itself: in the “bad” case, when bundle pricing is necessary for approximating the optimal revenue, it happens to be that the relevant random variable concentrates so that, in fact, a sample bundle price is sufficient.

## 1.1 Other Related Work

An important line of recent research [9, 14, 21, 23, 24] has explored the sample complexity of auctions. For example, [15] shows that with a single sample, one

<sup>3</sup> This is a reinterpretation of the Bulow-Klemperer Theorem [2].



can design an auction that gets a constant factor approximation to the optimal single item auction. How much better can you do with more samples? This question has been explored in a number of auction settings; e.g. by Morgenstern and Roughgarden [23] in the additive bidder setting we study in this paper.

### 1.2 Organization

After preliminaries, we show in Sect. 2 how to approximate the optimal single additive bidder revenue when given access to a sample from each item distribution. Then, in Sect. 3, we use the latter result to give two approximately optimal auctions for the multiple additive bidders setting: one that is given access to a sample from every item distribution from each bidder and one is that is fully prior-independent. In Sect. 4 we discuss an improved analysis for bidders with finite support distributions. We conclude with open problems in Sect. 5.

### 1.3 Preliminaries

In this paper, we consider the setting of a revenue-maximizing monopolist seller with  $m$  items to sell to  $n$  additive bidders. Each bidder  $i$  has his value  $V_{ij}$  for item  $j$  drawn from an unknown prior distribution  $F_{ij}$ . All bidders are additive: for any set of items  $S$ , bidder  $i$ 's value for the set is

$$V_i(S) = \sum_{j \in S} V_{ij}.$$

We will assume that each of the distributions  $F_{ij}$  is regular. That is,  $\varphi_{ij}(v) = v - \frac{1-F_{ij}(v)}{f_{ij}(v)}$  is non-decreasing.

We also use the following notation:

- The *revenue curve*  $R(\cdot)$  gives the expected revenue for selling an item at a price  $x$  to a bidder with value  $V$  drawn from distribution  $F$ . That is,  $R(x) := x \cdot \Pr[V \geq x]$ .
- The *monopoly price*  $r^*$  is the price that maximizes revenue:  $r^* := \operatorname{argmax}_x R(x)$ .
- Consider a single additive bidder with value  $V_j \sim G_j$  for item  $j$ . Then  $\text{SREV}(V_1, \dots, V_m)$  denotes the optimal expected revenue that can be obtained by posting a price for each item  $j$  individually. That is,

$$\text{SREV}(V_1, \dots, V_m) := \sum_j \max_{x_j} R_j(x_j),$$

where  $R_j$  is the revenue curve associated with the distribution  $G_j$ .

- $\text{BREV}(V_1, \dots, V_m)$  denotes the optimal expected revenue for posting a price on the “grand bundle” of all of the items to this same additive bidder with  $V_j \sim G_j$ . That is,

$$\text{BREV}(V_1, \dots, V_m) := \max_x x \cdot \Pr\left[\sum_j V_j \geq x\right].$$

- For any number  $x$ , let  $(x)^+$  denote  $\max\{x, 0\}$ .

## 2 A Prior Independent Mechanism for a Single Additive Bidder, Given Samples

Our mechanism draws heavily on two prior results. The first demonstrates that access to a single sample from a bidder’s distribution can be used to obtain a  $\frac{1}{2}$ -approximation of optimal revenue in the single-item setting.

**Theorem 1 (Dhangwatnotai, Roughgarden, and Yan 2010).** *Consider a bidder whose value for a particular item is drawn from  $F$ , a regular distribution with monopoly price  $r^*$  and revenue function  $R(\cdot)$ . Let  $S \sim F$  be a random sample from the distribution  $F$ . Then, for every nonnegative number  $t$ ,*

$$\mathbb{E}(R(\max\{t, S\})) \geq \frac{1}{2}R(\max\{t, r^*\}).$$

*Therefore, in particular, for  $t = 0$ , the expected revenue from posting a price of  $S$  yields at least half of the optimal posted price revenue, which is  $R(r^*)$ .*

The second result we use demonstrates that a combination of two very simple mechanisms can be used to obtain a constant factor of the optimal revenue in the single additive bidder setting.

**Theorem 2 (Babaioff, Immorlica, Lucier, and Weinberg 2014).** *Consider a single additive bidder with value  $V_j$  for item  $j$  drawn independently from distribution  $G_j$ . Denote by  $\text{OPT}(V_1, \dots, V_m)$  the revenue of the optimal mechanism. Let  $t = \text{SREV}(V_1, \dots, V_m)$  denote the optimal expected revenue from selling the items separately, and define  $V := \sum_{j=1}^m V_j$ , the bidder’s value for the grand bundle. Then*

- *If  $\mathbb{E}[V \mid V_j \leq t \ \forall j] \leq 4\text{SREV}(V_1, \dots, V_m)$ , then*

$$\mathbb{E}[\text{OPT}(V_1, \dots, V_m)] \leq 6 \text{SREV}(V_1, \dots, V_m).$$

- *Otherwise, if  $\mathbb{E}[V \mid V_j \leq t \ \forall j] > 4\text{SREV}(V_1, \dots, V_m)$ , then*

$$\Pr \left[ V \geq \frac{2}{5} \cdot \mathbb{E}[V \mid V_j \leq t \ \forall j] \right] \geq \frac{47}{72}$$

and

$$\mathbb{E}[\text{OPT}(V_1, \dots, V_m)] \leq 2 \text{SREV}(V_1, \dots, V_m) + \mathbb{E}[V \mid V_j \leq t \ \forall j].$$

From this, Babaioff et al. obtain the following corollary:

**Corollary 1 (Babaioff, Immorlica, Lucier, and Weinberg 2014).** *Consider a single additive bidder with value  $V_j$  for item  $j$  drawn independently from distribution  $G_j$ . Let  $\text{SREV}(V_1, \dots, V_m)$  denote the optimal expected revenue from selling the items separately and let  $\text{BREV}(V_1, \dots, V_m)$  denote the optimal expected revenue from selling the grand bundle. Then*

$$\mathbb{E}[\text{OPT}(V_1, \dots, V_m)] \leq 6 \cdot \max\{\text{SREV}(V_1, \dots, V_m) + \text{BREV}(V_1, \dots, V_m)\}.$$

We now combine the single additive bidder analysis with samples from the distributions to give an approximately optimal mechanism for a single additive bidder that does not rely on knowledge of the priors, but rather uses a single sample from each distribution.

The multi-bidder analogue of the better of selling separately or selling the grand bundle is a two-part tariff mechanism as used in [7, 8, 30]. Here, each bidder is offered a list of item prices and an entry fee. Typically in these mechanisms, some item prices are determined first. Then, the buyer’s surplus values above each item’s price are analyzed to understand either (a) whether to increase the item prices or to use an entry-fee or (b) how to compute the entry fee. This is equivalent to the buyer’s prior distribution for each item shifted down by the item’s price. We define the mechanism in a slightly more general way than is necessary here with a parameter  $\Delta$  in order to easily extend to the case where we want to analyze the shifted distributions. In the single bidder setting, we do not shift the distributions, so we will set the shift  $\Delta_j = 0$  for all  $j$ . However, this parameter will allow us to use this mechanism as a black box in the multiple bidders setting.

**Definition 1.** *Define the Sample Mechanism as follows. Given a set  $(S_1, \dots, S_m)$  of samples from an additive bidder’s distribution, and a set of non-negative values  $\Delta_1, \dots, \Delta_m$ ,*

- (a) *with probability  $\frac{1}{2}$ : Offer a price of  $\max\{\Delta_j, S_j\}$  for each item  $j$  separately.*
- (b) *with probability  $\frac{1}{2}$ : Offer the bidder a price of  $S^+ = \sum_{j=1}^m (S_j - \Delta_j)^+$  to enter the auction. If he pays the entrance fee, he can take any item  $j$  he wants at price  $\Delta_j$ . (When  $\Delta = \mathbf{0}$ , this is simply pricing the grand bundle.)*

Denote the revenue from this mechanism as  $\text{SAMP}(V_1, \dots, V_m; \Delta)$ .

**Theorem 3.** *Consider a single additive bidder with value  $V_j$  for item  $j$  drawn independently from regular distribution  $G_j$ . Let  $\Delta_1, \dots, \Delta_j \geq 0$ , and define  $V_j^+ = (V_j - \Delta_j)^+$ . The Sample Mechanism has expected revenue which is a constant fraction of the optimal expected revenue for  $(V_1^+, \dots, V_m^+)$ .*

*Proof.* The first step of the Sample Mechanism obtains expected revenue which is a constant fraction of  $\text{SREV}(V_1^+, \dots, V_m^+)$ . To see this, note that

$$\begin{aligned} \text{SREV}(V_1^+, \dots, V_m^+) &= \sum_j \max_{x \geq \Delta_j} [(x - \Delta_j)^+ (1 - G_j(x))] \\ &\leq \sum_j \max_{x \geq \Delta_j} R_j(x) \leq \sum_j R_j(\max\{\Delta_j, r_j^*\}), \end{aligned}$$

where  $R_j(\cdot)$  is the revenue curve for  $G_j$  and  $r_j^*$  is the optimal reserve price for  $G_j$ . The last inequality follows from regularity of  $G_j$ , and thus concavity of  $R_j(\cdot)$ . An application of Theorem 1 shows that offering the bidder a price of  $\max\{\Delta_j, S_j\}$  for item  $j$  yields expected revenue at least half of  $R_j(\max\{\Delta_j, r_j^*\})$ , hence step (a) is a  $\frac{1}{4}$ -approximation to  $\text{SREV}(V_1^+, \dots, V_m^+)$ .

To complete the proof, we show how the Sample Mechanism approximates  $\mathbb{E}[\text{OPT}(V_1^+, \dots, V_m^+)]$  within a constant factor using the two cases of Theorem 2. To this end, let  $t = \text{SREV}(V_1^+, \dots, V_m^+)$  and define  $C := \mathbb{E}[V^+ \mid V_j^+ \leq t \ \forall j]$ , where  $V^+ = \sum_j V_j^+$ . If  $C \leq 4t$ , we are in the first case of Theorem 2 applied to the random variables  $(V_1^+, \dots, V_m^+)$ , hence

$$\mathbb{E}[\text{OPT}(V_1^+, \dots, V_m^+)] \leq 6 \text{SREV}(V_1^+, \dots, V_m^+) \leq 24 \text{SAMP}(V_1, \dots, V_m; \mathbf{\Delta}).$$

Otherwise,  $C > 4t$ , and we also need to consider the revenue from the second step of the Sample Mechanism. In this case, from Theorem 2, we have

$$\mathbb{E}[\text{OPT}(V_1^+, \dots, V_m^+)] \leq 2 \text{SREV}(V_1^+, \dots, V_m^+) + C \quad \text{and} \quad \Pr \left[ V^+ \geq \frac{2}{5}C \right] \geq \frac{47}{72}.$$

Next, recall that  $S^+ = \sum_{j=1}^m (S_j - \Delta_j)^+$  is entry fee that is offered. Observe that if  $V^+ > S^+$ , then the bidder will enter the auction, since his utility will then be  $\sum_{j|V_j \geq \Delta_j} (V_j - \Delta_j) - S^+ = V^+ - S^+$ . Therefore, the expected revenue from step (b), the entry fee portion of the auction, can be bounded as follows:

$$\begin{aligned} & \mathbb{E} [S^+ \mid V^+ \geq S^+] \Pr[V^+ \geq S^+] \\ & \geq \mathbb{E} \left[ S^+ \mid V^+ \geq S^+, V^+, S^+ \geq \frac{2}{5}C \right] \cdot \Pr \left[ V^+, S^+ \geq \frac{2}{5}C, V^+ \geq S^+ \right] \\ & \geq \frac{2}{5}C \cdot \Pr \left[ V^+ \geq S^+ \mid V^+, S^+ \geq \frac{2}{5}C \right] \cdot \left( \frac{47}{72} \right)^2 \\ & > \frac{1}{2} \cdot \frac{1}{6}C \\ & = \frac{1}{12}C. \end{aligned} \tag{1}$$

The third line follows from the second part of Theorem 2, and the independence of  $V^+$  and  $S^+$ . The fourth line follows from the fact that  $V^+$  and  $S^+$  are identically distributed.

In this case, it is now clear that the Sample Mechanism obtains a constant factor approximation:

$$\mathbb{E}[\text{OPT}(V_1^+, \dots, V_m^+)] \leq 2 \text{SREV}(V_1^+, \dots, V_m^+) + C \leq 24 \text{SAMP}(V_1, \dots, V_m; \mathbf{\Delta}).$$

Thus in either case,  $\mathbb{E}[\text{OPT}(V_1^+, \dots, V_m^+)] \leq 24 \text{SAMP}(V_1, \dots, V_m; \mathbf{\Delta})$ . This mechanism loses a factor of 4 compared to the prior-dependent  $\max\{\text{SREV}, \text{BREV}\}$  mechanism in [1].

**Corollary 2.** *Consider a single additive bidder with value  $V_j$  for item  $j$  drawn independently from regular distribution  $G_j$ . The Sample Mechanism with  $\Delta_j = 0$  for all  $j$  has expected revenue which is a constant fraction of the optimal expected revenue  $\mathbb{E}[\text{OPT}(V_1, \dots, V_m)]$ .*

### 3 Multiple Additive Bidders

Our mechanism builds on the following breakthrough result from which Yao constructs a simple, approximately optimal mechanism for the multiple additive bidders setting.

**Theorem 4 (Yao 2015).** *Consider  $n$  additive bidders, where  $V_{ij}$  is the value bidder  $i$  has for item  $j$ , where  $1 \leq i \leq n$  and  $1 \leq j \leq m$ . Assume the set of random variables  $\{V_{ij}\}_{j=1}^m$  are independent for each  $i$ . Define the following auxiliary random variables:*

$$X_{ij} := \max_{k \neq i} V_{kj} \quad \text{and} \quad A_{ij} := (V_{ij} - X_{ij})^+.$$

Then for  $\mathbf{V}_j = (V_{1j}, \dots, V_{nj})$ , the expected revenue of the optimal mechanism for the multiple additive bidders setting satisfies

$$\mathbb{E}[\text{OPT}(\mathbf{V}_1, \dots, \mathbf{V}_m)] \leq 8 \sum_i \mathbb{E}[\text{OPT}(A_{i1}, \dots, A_{im})] + 9 \mathbb{E}[\text{SPA}(\mathbf{V}_1, \dots, \mathbf{V}_m)],$$

where  $\text{SPA}(\cdot)$  is the revenue from running a separate second-price auction for each item and  $\text{OPT}(A_{i1}, \dots, A_{im})$  denotes the revenue obtained by the optimal single additive bidder auction, when that bidder's value for item  $j$  is  $A_{ij}$ .

#### 3.1 A Sample Auction for Multiple Bidders

By randomly choosing to either (a) run a second-price auction separately on each item or (b) run sample mechanisms on each bidder with  $\Delta_{ij} = X_{ij}$ , we can achieve a constant-fraction of the optimal revenue from only samples in the multiple-additive-bidder setting as well.

**Definition 2.** *Define the Multiple-Additive-Bidders Sample Mechanism parameterized by  $p$  as follows, given a sample  $S_{ij}$  from each bidder  $i$ 's distribution for item  $j$ :*

- (a) *with probability  $p$ : Run a Second-Price Auction on each item  $j$ . That is, offer each bidder  $i$  the option to take item  $j$  at a price equal to  $X_{ij}$ .*
- (b) *with probability  $1 - p$ : Offer each bidder  $i$  an entry fee of  $\sum_j (S_{ij} - X_{ij})^+$ . Any bidder willing to pay the entry fee can then take<sup>4</sup> item  $j$  at price  $X_{ij}$ .*

Let  $\text{MAB-SAMP}(\mathbf{V}_1, \dots, \mathbf{V}_m; p)$  denote the revenue from the Multiple-Additive-Bidders Sample Mechanism with parameter  $p$ .

**Theorem 5.** *In the setting of Theorem 4, when the random variables  $V_{ij}$  are all independent and, for each  $j$ , the random variables  $V_{ij}$  is drawn from regular distribution  $F_{ij}$  for each bidder  $i$  and item  $j$ , with access to a sample  $S_{ij}$  from each  $F_{ij}$ , the Multiple-Additive-Bidders Sample Mechanism with parameter  $p = \frac{9}{201}$  obtains at least a constant fraction of the optimal expected revenue.*

<sup>4</sup> This guarantees that each item is taken by at most one bidder.

*Proof.* Recall that  $X_{ij} := \max_{k \neq i} V_{kj}$  is the highest bid for item  $j$  excluding bidder  $i$ 's bid, and  $A_{ij} := (V_{ij} - X_{ij})^+$  is the surplus from buyer  $i$ 's value for item  $j$  over this price.

Given sample  $S_{ij}$  for each item  $j$ , an application of Theorem 3 on each bidder  $i$  where  $\Delta_{ij} = X_{ij}$  gives that  $\text{OPT}(A_{i1}, \dots, A_{im}) \leq 24 \text{SAMP}(V_{i1}, \dots, V_{im}; \mathbf{X}_i)$ , where  $\mathbf{X}_i = (X_{i1}, \dots, X_{im})$ . Using Theorem 4, this gives that

$$\begin{aligned} \mathbb{E}[\text{OPT}(\mathbf{V}_1, \dots, \mathbf{V}_m)] &\leq 8 \sum_i \mathbb{E}[\text{OPT}(A_{i1}, \dots, A_{im})] + 9 \mathbb{E}[\text{SPA}(\mathbf{V}_1, \dots, \mathbf{V}_m)] \\ &\leq 8 \cdot 24 \sum_i \mathbb{E}_{\mathbf{X}_i}[\text{SAMP}(V_{i1}, \dots, V_{im}; \mathbf{X}_i)] \\ &\quad + 9 \mathbb{E}[\text{SPA}(\mathbf{V}_1, \dots, \mathbf{V}_m)] \\ &\leq 201 \text{MAB-SAMP}(\mathbf{V}_1, \dots, \mathbf{V}_m; \frac{9}{201}) \end{aligned}$$

Running the sample mechanisms with probability  $\frac{192}{201}$  and a second-price auction separately on each item with probability  $\frac{9}{201}$  gives a 201-approximation. Similarly to the single bidder case, this loses less than a factor of 4 compared to the Bundling Mechanism in [30] which requires full knowledge of all of the prior distributions and achieves a 57-approximation.

### 3.2 A Prior-Independent Auction

We can also use sample mechanisms to sell to multiple additive bidders without extra samples. Analogously to [15], if the seller can identify which bidders come from the same distribution, she can take a sample bidder from each group  $a$  of identically distributed bidders and use it to set the prices for the rest of the group. This requires at least two bidders from each distribution group  $a$ . The mechanism is the same as the Multiple-Additive-Bidders Sample Mechanism, but with randomly excluded bidders used as samples.

**Definition 3.** *Define the Multiple-Additive-Bidders Prior-Independent Mechanism parameterized by  $p$  as follows:*

- (a) *with probability  $p$ : Run a Second-Price Auction on each item  $j$ . That is, offer each bidder  $i$  the option to take item  $j$  at a price equal to  $X_{ij}$ .*
- (b) *with probability  $1-p$ : Remove a random bidder  $i_a$  from each group of bidders  $a$  and let  $S_{aj}$  be his bid for item  $j$  (i.e.,  $S_{aj} := V_{i_a j}$ ). Let  $\mathcal{S}$  be the set of bidders sampled from each group  $a$ . Also, let  $\beta_{ij} = \max_{k \notin \mathcal{S}, k \neq i} V_{kj}$ . Offer each remaining bidder  $i$  from group  $a$  an entry fee of  $\sum_j (S_{aj} - \beta_{ij})^+$ . Any bidder willing to pay the entry fee can then take item  $j$  at price  $\beta_{ij}$ .*

Let  $\text{MAB-PI}(\mathbf{V}_1, \dots, \mathbf{V}_m; p)$  denote the revenue from the Multiple-Additive-Bidders Prior-Independent Mechanism with parameter  $p$ .

**Theorem 6.** *In the setting of Theorem 4, when the random variables  $V_{ij}$  are all independent and, for each  $j$ , the random variables  $V_{ij}$  is drawn from regular distribution  $F_{ij}$  for each bidder  $i$  and item  $j$ , with at least 2 bidders from every distribution group  $a$ , the Multiple-Additive-Bidders Prior-Independent Mechanism with parameter  $p = \frac{9}{1161}$  obtains at least a constant fraction of the optimal expected revenue.*

*Proof.* If  $n_a$  is the number of bidders from distribution group  $a$  and  $n_a^{\min}$  is the number of bidders in the smallest such group, then

$$\begin{aligned} \sum_i \mathbb{E}[\text{OPT}(A_{i1}, \dots, A_{im})] &\leq \sum_a \frac{n_a}{n_a - 1} \sum_{i \in a, i \notin S} \mathbb{E}[\text{OPT}(A_{i1}, \dots, A_{im})] \\ &\leq \frac{n_a^{\min}}{n_a^{\min} - 1} \sum_{i \notin S} \mathbb{E}[\text{OPT}(A_{i1}, \dots, A_{im})]. \end{aligned}$$

Also notice that

$$A_{ij} := \left( V_{ij} - \max\{\max\{S_{aj}\}, \beta_{ij}\} \right)^+ \quad \text{and define} \quad V_{ij}^+ := (V_{ij} - \beta_{ij})^+.$$

Clearly, the random variable  $V_{ij}^+$  dominates the random variable  $A_{ij}$  (i.e.,  $\Pr(V_{ij}^+ \geq x) \geq \Pr(A_{ij} \geq x)$  for all  $x$ ). Therefore,

$$\text{SREV}(A_{i1}, \dots, A_{im}) \leq \text{SREV}(V_{i1}^+, \dots, V_{im}^+)$$

and

$$\text{BREV}(A_{i1}, \dots, A_{im}) \leq \text{BREV}(V_{i1}^+, \dots, V_{im}^+)$$

Thus, by Corollary 1, it suffices to obtain a constant fraction of  $\mathbb{E}[\text{OPT}(V_{i1}^+, \dots, V_{im}^+)]$  for each  $i$ .

Using the analysis from Theorem 5, we put it all together to see that

$$\begin{aligned} \mathbb{E}[\text{OPT}(\mathbf{V}_1, \dots, \mathbf{V}_m)] &\leq 8 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \sum_{i \notin S} \mathbb{E}[\text{OPT}(A_{i1}, \dots, A_{im})] \\ &\quad + 9 \mathbb{E}[\text{SPA}(\mathbf{V}_1, \dots, \mathbf{V}_m)] \\ &\leq 8 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \cdot 6 \sum_{i \notin S} \mathbb{E}[\text{OPT}(V_{i1}^+, \dots, V_{im}^+)] \\ &\quad + 9 \mathbb{E}[\text{SPA}(\mathbf{V}_1, \dots, \mathbf{V}_m)] \\ &\leq 8 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \cdot 6 \cdot 24 \sum_{i \notin S} \mathbb{E}_{\beta_i}[\text{SAMP}(V_{i1}, \dots, V_{im}; \beta_i)] \\ &\quad + 9 \mathbb{E}[\text{SPA}(\mathbf{V}_1, \dots, \mathbf{V}_m)] \\ &\leq 1161 \cdot \frac{n_a^{\min}}{n_a^{\min} - 1} \text{MAB-PI}(\mathbf{V}_1, \dots, \mathbf{V}_m; \frac{9}{1161}) \end{aligned}$$

where of course, since  $n_a^{\min} \geq 2$ , then  $\frac{n_a^{\min}}{n_a^{\min} - 1} \leq 2$ .

Note that the loss due to excluding bidders to use as samples is a factor of  $6 \frac{n_a^{\min}}{n_a^{\min} - 1}$ .

## 4 Finite Support Distributions

Cai, Devanur, and Weinberg [7] present a new framework that analyzes revenue from multiple additive bidders with finite support distributions (over discrete value spaces) via a similar core-tail decomposition. These results also hold for discretizing a continuous value space and losing at most a factor of  $1 + \varepsilon$  in the revenue due to the discretization. Utilizing this analysis improves the constant of our approximation.

Precisely, they show that

$$E[\text{OPT}(\mathbf{V}_1, \dots, \mathbf{V}_m)] \leq 4 \text{SREV}(\mathbf{V}_1, \dots, \mathbf{V}_m) + \text{CORE}$$

where, for the highest other bid  $X_{ij} := \max_{k \neq i} V_{kj}$ ,

$$\text{CORE} = \mathbb{E} \left[ \sum_{i=1}^n \sum_{j=1}^m (V_{ij} - X_{ij})^+ \mathbb{1}_{V_{ij} \in [X_{ij}, X_{ij} + t_i]} \right]$$

and

$$t_i = \text{SREV}((V_{i1} - X_{i1})^+, \dots, (V_{im} - X_{im})^+).$$

Then if  $A_{ij} = (V_{ij} - X_{ij})^+ \cdot \mathbb{1}_{V_{ij} \in [X_{ij}, X_{ij} + t_i]}$ , we have that

$$\text{CORE} = \sum_{i=1}^n \mathbb{E}_{\mathbf{V}_{-i}} [\mathbb{E}_{V_i} [\sum_{j=1}^m A_{ij}]].$$

Then in a proof nearly identical to that of Theorem 3, we can show that  $\text{CORE} \leq 24 \sum_i \mathbb{E}_{\mathbf{V}_{-i}} [\text{SAMP}(V_{i1}, \dots, V_{im}; \mathbf{X}_i)]$ . For each bidder, we set  $\Delta_j = X_{ij}$  and bound  $\mathbb{E}[\sum_{j=1}^m A_{ij}] \leq 24 \text{SAMP}(V_{i1}, \dots, V_{im}; \mathbf{\Delta})$ .

In one case,  $\mathbb{E}[\sum_{j=1}^m A_{ij}] \leq 4t = 4 \text{SREV}((V_{i1} - X_{i1})^+, \dots, (V_{im} - X_{im})^+) \leq 16 \text{SAMP}(V_{i1}, \dots, V_{im}; \mathbf{X}_i)$ . In the other case,  $\mathbb{E}[\sum_{j=1}^m A_{ij}] > 4t$ . In this case, similarly to the proof of Theorem 2, we get that

$$\begin{aligned} \Pr \left[ \left| \sum_j A_{ij} - \mathbb{E} \left[ \sum_{j=1}^m A_{ij} \right] \right| \geq \frac{3}{5} \mathbb{E} \left[ \sum_{j=1}^m A_{ij} \right] \right] &\leq \frac{\text{var}(\sum_{j=1}^m A_{ij})}{\frac{3^2}{5} \mathbb{E}[\sum_{j=1}^m A_{ij}]^2} && \text{by Chebyshev's inequality} \\ &< \frac{\text{var}(\sum_{j=1}^m A_{ij})}{\frac{9}{25} \cdot 16t^2} && \text{since } \mathbb{E} \left[ \sum_j A_{ij} \right] > 4t \\ &\leq \frac{2t^2}{\frac{9}{25} \cdot 16t^2} = \frac{25}{72}. \end{aligned}$$

The final inequality follows from the fact that  $\text{var}(\sum_{j=1}^m A_{ij}) \leq t^2$  by Lemma 9 of [7].



Identically to the proof of Theorem 3, we get that if we offer the bidder a price of  $S^+ = \sum_{j=1}^m (S_{ij} - X_{ij})^+$  that

$$\mathbb{E}\left[\sum_{j=1}^m A_{ij}\right] \leq 12 \mathbb{E}[S^+ \cdot \Pr[\sum_{j=1}^m A_{ij} \geq S^+]] \leq 24 \text{SAMP}(V_{i1}, \dots, V_{im}; \mathbf{X}_i).$$

Hence  $\text{CORE} \leq 24 \sum_i \mathbb{E}_{\mathbf{V}_{-i}}[\text{SAMP}(V_{i1}, \dots, V_{im}; \mathbf{X}_i)]$ . Moreover, since  $(V_{ij} - \beta_{ij})^+$  stochastically dominates  $(V_{ij} - X_{ij})^+$ , then we only lose a factor of  $\frac{n_a^{\min}}{n_a^{\min}-1}$  for excluding bidders to use as samples to make a prior-independent auction.

As in the proof of Theorem 3, an application of Theorem 1 gives that a second-price auction on each item is a constant factor of the revenue from selling each item separately. Then again, since  $\beta_{ij} = \max_{k \notin \mathcal{S}, k \neq i} V_{kj}$ ,

$$\begin{aligned} \mathbb{E}[\text{OPT}(\mathbf{V}_1, \dots, \mathbf{V}_m)] &\leq 4 \text{SREV}(\mathbf{V}_1, \dots, \mathbf{V}_m) + \text{CORE} \\ &\leq 8 \text{SPA}(\mathbf{V}_1, \dots, \mathbf{V}_m) + \\ &\quad 24 \cdot \frac{n_a^{\min}}{n_a^{\min}-1} \sum_{i \notin \mathcal{S}} \mathbb{E}_{\beta_i}[\text{SAMP}(V_{i1}, \dots, V_{im}; \beta_i)] \\ &\leq 32 \cdot \frac{n_a^{\min}}{n_a^{\min}-1} \text{MAB-PI}(\mathbf{V}_1, \dots, \mathbf{V}_m; \frac{1}{4}) \end{aligned}$$

We lose a factor of 4 compared to the mechanism of [7] when given samples, and a factor of  $4 \cdot \frac{n_a^{\min}}{n_a^{\min}-1}$  without samples.

## 5 Open Problems

**Beyond Additive Bidders.** One interesting problem for future work is to design prior-independent mechanisms for more general valuations. Recent work in revenue maximization for more general multi-item settings gives mechanisms that have constant-factor approximation guarantees for a single subadditive buyer [28] and for multiple matroid-constrained buyers [8]. Both of these results rely on an analysis that chooses prices in the bidders’ distributions that would sell with a constrained ex-ante probability. As these probabilities are aimed at segmenting off the tails of the distributions and samples are unlikely to come from the tail, it is unclear how to design a prior independent mechanism for these settings.

**Lower Bounds.** Another interesting open problem is to obtain a lower bound on the gap in revenue between the optimal mechanism and the Sample Mechanism, and for the Multiple-Additive-Bidders Sample Mechanism as well. However, stronger lower bounds are still open problems for the mechanisms from [1, 7, 30] as well.

A lower bound from [15] shows that the factor of 2 in Theorem 1 is tight when a bidder’s distribution for a single item is the distribution where the revenue curve is a triangle, that is, where  $F(v) = \frac{v}{v+1}$  on  $[0, H)$  as  $H \rightarrow \infty$ .

The best known lower bound on the approximation of the  $\max\{\text{SREV}, \text{BREV}\}$  is a factor of 2. The example, given by Rubinstein [27], has  $n$  items from the equal revenue distribution and  $n$  rare but expensive items. The optimal revenue gets an equal fraction of revenue from each group; however, selling the grand bundle does well for the first set and poorly for the second while selling separately captures the revenue of the second set but not the first. Of course, this gap gives the Sample Mechanism a lower bound of 2 as well.

## References

1. Babaioff, M., Immorlica, N., Lucier, B., Weinberg, S.M.: A simple and approximately optimal mechanism for an additive buyer. In: Proceedings of the 55th Annual Symposium on Foundations of Computer Science, FOCS 2014 (2014). <http://arxiv.org/abs/1405.6146>
2. Bulow, J., Klemperer, P.: Auctions vs. negotiations. Technical report, National Bureau of Economic Research (1994)
3. Cai, Y., Daskalakis, C., Weinberg, S.M.: An algorithmic characterization of multi-dimensional mechanisms. In: Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing, pp. 459–478. ACM (2012)
4. Cai, Y., Daskalakis, C., Weinberg, S.M.: Optimal multi-dimensional mechanism design: reducing revenue to welfare maximization. In: 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS), pp. 130–139. IEEE (2012)
5. Cai, Y., Daskalakis, C., Weinberg, S.M.: Reducing revenue to welfare maximization: approximation algorithms and other generalizations. In: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 578–595. SIAM (2013)
6. Cai, Y., Daskalakis, C., Weinberg, S.M.: Understanding incentives: mechanism design becomes algorithm design. In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS), pp. 618–627. IEEE (2013)
7. Cai, Y., Devanur, N.R., Weinberg, S.M.: A duality based unified approach to Bayesian mechanism design. In: Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing. ACM (2016)
8. Chawla, S., Miller, J.B.: Mechanism design for subadditive agents via an ex ante relaxation. In: Proceedings of the 2016 ACM Conference on Economics and Computation, EC 2016, pp. 579–596. ACM, New York (2016). <http://doi.acm.org/10.1145/2940716.2940756>
9. Cole, R., Roughgarden, T.: The sample complexity of revenue maximization. In: Proceedings of the 46th Annual ACM Symposium on Theory of Computing, pp. 243–252. ACM (2014)
10. Daskalakis, C.: Multi-item auctions defying intuition? ACM SIGecom Exch. **14**(1), 41–75 (2015)
11. Daskalakis, C., Deckelbaum, A., Tzamos, C.: Mechanism design via optimal transport. In: Proceedings of the 14th ACM Conference on Electronic Commerce, EC 2013, pp. 269–286 (2013). <http://doi.acm.org/10.1145/2482540.2482593>
12. Daskalakis, C., Deckelbaum, A., Tzamos, C.: Strong duality for a multiple-good monopolist. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation, pp. 449–450. ACM (2015)

13. Devanur, N., Hartline, J., Karlin, A., Nguyen, T.: Prior-independent multi-parameter mechanism design. In: Chen, N., Elkind, E., Koutsoupias, E. (eds.) WINE 2011. LNCS, vol. 7090, pp. 122–133. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25510-6\\_11](https://doi.org/10.1007/978-3-642-25510-6_11). <http://users.eecs.northwestern.edu/~hartline/papers/auctions-WINE-11.pdf>
14. Devanur, N.R., Huang, Z., Psomas, C.A.: The sample complexity of auctions with side information. arXiv preprint [arXiv:1511.02296](https://arxiv.org/abs/1511.02296) (2015)
15. Dhangwatnotai, P., Roughgarden, T., Yan, Q.: Revenue maximization with a single sample. In: Proceedings of the 11th ACM Conference on Electronic Commerce, EC 2010, pp. 129–138. ACM, New York (2010). <http://doi.acm.org/10.1145/1807342.1807364>
16. Fiat, A., Goldner, K., Karlin, A.R., Koutsoupias, E.: The FedEx problem. In: Proceedings of the 2016 ACM Conference on Economics and Computation, EC 2016, pp. 21–22. ACM, New York (2016). <http://doi.acm.org/10.1145/2940716.2940752>
17. Fu, H., Hartline, J., Hoy, D.: Prior-independent auctions for risk-averse agents. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, pp. 471–488. ACM (2013)
18. Giannakopoulos, Y., Koutsoupias, E.: Duality and optimality of auctions for uniform distributions. In: Proceedings of the 15th ACM Conference on Economics and Computation, EC 2014, pp. 259–276 (2014). <http://doi.acm.org/10.1145/2600057.2602883>
19. Giannakopoulos, Y., Koutsoupias, E.: Selling two goods optimally. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015. LNCS, vol. 9135, pp. 650–662. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-47666-6\\_52](https://doi.org/10.1007/978-3-662-47666-6_52)
20. Hart, S., Nisan, N.: Approximate revenue maximization with multiple items. arXiv preprint [arXiv:1204.1846](https://arxiv.org/abs/1204.1846) (2012)
21. Huang, Z., Mansour, Y., Roughgarden, T.: Making the most of your samples. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation, pp. 45–60. ACM (2015)
22. Li, X., Yao, A.C.C.: On revenue maximization for selling multiple independently distributed items. Proc. Natl. Acad. Sci. **110**(28), 11232–11237 (2013). <http://dx.doi.org/10.1073/pnas.1309533110>
23. Morgenstern, J., Roughgarden, T.: Learning simple auctions. CoRR abs/1604.03171 (2016). <http://arxiv.org/abs/1604.03171>
24. Morgenstern, J.H., Roughgarden, T.: On the pseudo-dimension of nearly optimal auctions. In: Advances in Neural Information Processing Systems, pp. 136–144 (2015)
25. Roughgarden, T., Talgam-Cohen, I.: Optimal and near-optimal mechanism design with interdependent values. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC 2013, pp. 767–784. ACM, New York (2013). <http://doi.acm.org/10.1145/2482540.2482606>
26. Roughgarden, T., Talgam-Cohen, I., Yan, Q.: Supply-limiting mechanisms. In: Proceedings of the 13th ACM Conference on Electronic Commerce, pp. 844–861. ACM (2012). [http://theory.stanford.edu/~tim/papers/limit\\_conf.pdf](http://theory.stanford.edu/~tim/papers/limit_conf.pdf)
27. Rubinstein, A.: On the computational complexity of optimal simple mechanisms. In: Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, ITCS 2016, pp. 21–28. ACM, New York (2016). <http://doi.acm.org/10.1145/2840728.2840736>

28. Rubinstein, A., Weinberg, S.M.: Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation, pp. 377–394. ACM (2015). <http://arxiv.org/abs/1501.07637>
29. Weinberg, S.M.: Algorithms for strategic agents. Ph.D. thesis, Massachusetts Institute of Technology (2014). <http://people.csail.mit.edu/smweinberg/thesis.pdf>
30. Yao, A.C.C.: An  $n$ -to-1 bidder reduction for multi-item auctions and its applications. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 92–109 (2015). <http://arxiv.org/abs/1406.3278>

# Optimal Mechanism for Selling Two Items to a Single Buyer Having Uniformly Distributed Valuations

Thirumulanathan D.<sup>(✉)</sup>, Rajesh Sundaresan, and Y. Narahari

Indian Institute of Science, Bengaluru, India  
thirumulanathan@gmail.com

**Abstract.** We consider the design of an optimal mechanism for a seller selling two items to a single buyer so that the expected revenue to the seller is maximized. The buyer's valuation of the two items is assumed to be the uniform distribution over an arbitrary rectangle  $[c_1, c_1 + b_1] \times [c_2, c_2 + b_2]$  in the positive quadrant. The solution to the case when  $(c_1, c_2) = (0, 0)$  was already known. We provide an explicit solution for arbitrary nonnegative values of  $(c_1, c_2, b_1, b_2)$ . We prove that the optimal mechanism is to sell the two items according to one of eight simple menus. We also prove that the solution is deterministic when either  $c_1$  or  $c_2$  is beyond a threshold. Finally, we conjecture that our methodology can be extended to a wider class of distributions. We also provide some preliminary results to support the conjecture.

## 1 Introduction

*Optimal mechanism design* is the problem of finding a mechanism that generates the highest expected revenue to the seller. While the solution to the problem of selling a single item is well-known (Myerson [8]), optimal mechanism design for selling multiple items is a harder problem. Even the simplest setting with two items and one buyer remains as yet unsolved. Partial characterizations and solutions to some special cases are known. For example, Rochet and Choné [10] provided a partial characterization of the optimal mechanism as one whose parameters solve a system of partial differential equations. Manelli and Vincent [6, 7] derived the optimal mechanism when the number of items  $m = 2$ , and the valuation  $z = (z_1, z_2)$  of those items is such that  $z \sim \text{Unif}[0, 1]^2$ . Giannakopoulos and Koutsoupias [4] provided a solution when  $m \leq 6$  and  $z = (z_1, \dots, z_m) \sim \text{Unif}[0, 1]^m$ . In another paper [5], the same authors provided closed form solutions when  $m = 2$  and the distribution of  $z$  satisfies some sufficient conditions. Wang and Tang [12] proved that the optimal mechanisms have simple menus, when the distribution of  $z$  satisfies a certain power rate condition (that is satisfied by the uniform distribution). Each menu has at most four regions of constant allocations. However, the exact menus and associated allocations were left open.

---

Supported by the Defence Research and Development Organisation, Govt. of India.

Daskalakis et al. [1–3] considered the most general class of distributions till date, and gave an optimal solution when the distributions give rise to a so-called “well-formed” canonical partition (to be described in Sect. 4) of the support set.

The papers [1–5] rely on a result of Rochet [11] that transforms the search for an optimal mechanism into a search for a positive function. This function represents the valuation of the buyer minus the payment to the seller, and maximizes the expected payment subject to the function being positive, increasing, convex, and 1-Lipschitz. The above papers identify a dual problem, solve it, and exploit this solution to identify a primal solution. The assumption that the support set  $D$  of the distribution is  $[0, b_i]^2$  is crucially used in finding the dual solution. We are aware of only two examples where the support sets  $[c, c + 1]^2$ ,  $c > 0$ , and  $[4, 16] \times [4, 7]$ , for which the solutions are known, are not bordered by the coordinate axes. These were considered by Pavlov [9] and Daskalakis et al. [2], respectively. Daskalakis et al. [1, 2] do consider other distributions but they must satisfy  $f(z)(z \cdot n(z)) = 0$  on the boundaries of  $D$ , where  $n(z)$  is the normal to the boundary at  $z$ . The uniform distribution on arbitrary rectangles (which we consider in this paper) has  $f(z)(z \cdot n(z)) < 0$  in general on the left and bottom boundaries, and this requires additional nontrivial care in its handling.

## 1.1 Contribution, Method, and Outline

We solve the two-item single-buyer optimal mechanism design problem when  $z \sim \text{Unif}[c_1, c_1 + b_1] \times [c_2, c_2 + b_2]$ , for arbitrary nonnegative values of  $(c_1, c_2, b_1, b_2)$ . We prove that the structure of the optimal solution falls within a class of eight simple menus. Each has at most four constant allocation regions. In each region, goods are allocated with fixed probabilities. Our proof is constructive and provides a method to compute the exact solution. We also establish an interesting property of the optimal mechanism: given any value of  $c_1$ , we find a threshold value of  $c_2$  beyond which the mechanism becomes deterministic.

Our method is as follows. From [2, 3], we know that the dual problem is an optimal transport problem that transfers mass from the support set  $D$  to itself, subject to the constraint that the difference between the mass densities before and after the transfer *convex dominates* a signed measure defined by the distribution of the valuations. When  $(c_1, c_2) = (0, 0)$ , Daskalakis et al. [1] provided a solution where the difference between the densities not just convex dominates the signed measure, but equals it. In another work, Daskalakis et al. [2] provided an example ( $z \sim \text{Unif}[4, 16] \times [4, 7]$ ) where the difference between the densities strictly convex dominates the signed measure. They construct a line measure that convex dominates 0, add it to the signed measure, and then solve the example using the same method that they used to solve the  $(c_1, c_2) = (0, 0)$  case. They call this line measure as the “shuffling measure”.

Their method can be used to find the optimal solution for more general distributions, provided we know what shuffling measure must be used to arrive at the solution. It is not clear, a priori, what shuffling measure must be used, even for the restricted setting of uniform distributions. We prove that for the setting of uniform distributions, the optimal solution is always arrived at by using a shuffling (line) measure of certain forms. We do the following.

- We start with the shuffling (line) measure in [2]. We parametrize this measure using its depth ( $p_{a_1}$ ), slope ( $a_1$ ), and the length ( $m_1$ ), and find the relation between these parameters so that the line measure convex dominates 0.
- We then derive the conditions that these parameters must satisfy in order to solve the optimal transport problem. The conditions turn out to be polynomial equations of degree at most 4.
- We identify conditions on the parameters ( $c_1, c_2, b_1, b_2$ ) so that the solutions to the polynomials yield a valid menu (allocation probability is at most 1 and the canonical partition is within  $D$ ). We thus arrive at eight different menus. We prove that the optimal menu is one of the eight, for any  $(c_1, c_2, b_1, b_2) \geq 0$ .
- We then generalize the shuffling measure, and use it to solve an example linear density function. We thus conjecture that this class of shuffling measures will help identify the optimal mechanism for distributions with negative constant power rate.

Our work thus provides a method to construct an appropriate shuffling measure, and hence to arrive at the optimal solution, for various distributions whose support sets are not bordered by the coordinate axes. In our view, this is a non-trivial step towards understanding optimal mechanisms in multi-item settings.

The rest of the paper is organized as follows. In Sect. 2, we formulate an optimization problem that describes the two-item single-buyer optimal mechanism. In Sect. 3, we discuss the space of solutions and highlight a few interesting outcomes. In Sect. 4, we define the shuffling measures required to find the optimal mechanism for the uniform distribution on arbitrary rectangles. We prove that the optimal solution is one of the eight simple menus. In Sect. 5, we discuss the conjecture and possible extension to other classes of distributions.

## 2 Preliminaries

Consider a two-item one-buyer setting. The buyer's valuation is  $z = (z_1, z_2)$  for the two items, sampled according to the joint density  $f(z) = f_1(z_1)f_2(z_2)$ , where  $f_1(z_1)$  and  $f_2(z_2)$  are marginal densities. The support set of  $f$  is defined as  $D := \{z : f(z) > 0\}$ . Throughout this paper, we restrict attention to an arbitrary rectangle  $D = [c_1, c_1 + b_1] \times [c_2, c_2 + b_2]$ , where  $c_1, c_2, b_1, b_2$  are all nonnegative. A quasilinear mechanism comprises an allocation function  $q : D \rightarrow [0, 1]^2$  and a payment function  $t : D \rightarrow \mathbb{R}_+$  that represent, respectively, the probabilities of allocation of the items to the buyer and the amount that the buyer pays. In other words, for a reported valuation vector  $\hat{z} = (\hat{z}_1, \hat{z}_2)$ , item  $i$  is allocated with probability  $q_i(\hat{z})$ , where  $(q_1(\hat{z}), q_2(\hat{z})) = q(\hat{z})$ , and the seller collects a revenue of  $t(\hat{z})$  from the buyer. If the buyer's true valuation is  $z$ , and he reports  $\hat{z}$ , then his utility is  $\hat{u}(z, \hat{z}) := z \cdot q(\hat{z}) - t(\hat{z})$ , which is the valuation minus the payment.

A quasilinear mechanism is *incentive compatible* when truth telling is a weakly dominant strategy for the buyer, i.e.,  $\hat{u}(z, z) \geq \hat{u}(z, \hat{z})$  for every  $z, \hat{z} \in D$ . In this case the buyer's realized utility is

$$u(z) := \hat{u}(z, z) = z \cdot q(z) - t(z). \quad (1)$$

The following result is well known:

**Theorem 1** [11]. *A quasilinear mechanism  $(q, t)$ , with  $u(z) = z \cdot q(z) - t(z)$ , is incentive compatible iff  $u$  is convex and  $\nabla u(z) = q(z)$  for a.e.  $z \in D$ .*

An incentive compatible mechanism is *individually rational* if the buyer is not worse off by participating in the mechanism, i.e.,  $u(z) \geq 0$  for every  $z \in D$ .

An optimal mechanism is one that maximizes the expected revenue to the seller subject to incentive compatibility and individual rationality. By virtue of Theorem 1 and (1), an optimal mechanism solves the problem

$$\max_u \mathbb{E}_{z \sim f}[t(z)] = \int_D (z \cdot \nabla u(z) - u(z))f(z) dz \tag{2}$$

subject to  $\{(a) u$  convex,  $(b) \nabla u(z) \in [0, 1]^2$  a.e.  $z \in D$ ,  $(c) u(z) \geq 0 \forall z \in D.\}$

The condition  $\nabla u \in [0, 1]^2$  holds because, by Theorem 1,  $\nabla u(z) = q(z) \in [0, 1]^2$ . Since the components of  $\nabla u$  are nonnegative, the nonnegativity condition in (2) is equivalent to having  $u(c_1, c_2) \geq 0$ . We further set  $u(c_1, c_2) = 0$ , since the objective function can be only lower when  $u(c_1, c_2)$  is larger. Using integration by parts, the objective function of (2) can be written as  $\int_D u(z)\mu(z) dz + \int_{\partial D} u(z)\mu_s(z) d\sigma(z)$ , where  $\mu(z) := -z \cdot \nabla f(z) - 3f(z)$  for all  $z \in D$ , and  $\mu_s(z) := (z \cdot n(z))f(z)$  for all  $z \in \partial D$ .

The vector  $n(z)$  is the normal to the surface  $\partial D$  at  $z$ . We regard  $\mu$  and  $\mu_s$  as the densities of signed measures on  $D$  and  $\partial D$  that are absolutely continuous with respect to the two-dimensional and one-dimensional Lebesgue measures  $dz$  and  $d\sigma(z)$ , respectively. Often, by an abuse of notation, we use  $\mu$  and  $\mu_s$  to represent the measures instead of just the densities. By taking  $u(z) = 1 \forall z \in D$ , we observe that

$$\begin{aligned} \int_D \mu(z) dz + \int_{\partial D} \mu_s(z) d\sigma(z) &= \int_D u(z)\mu(z) dz + \int_{\partial D} u(z)\mu_s(z) d\sigma(z) \\ &= \int_D (z \cdot \nabla u(z) - u(z))f(z) dz = -1. \end{aligned}$$

By defining a point measure  $\mu_p := \delta_{\{(c_1, c_2)\}}$ , we have that  $\bar{\mu} := \mu + \mu_s + \mu_p$  has  $\bar{\mu}(D) = 0$ . (By  $\bar{\mu} = \mu + \mu_s + \mu_p$ , we mean  $\bar{\mu}(A) = \mu(A) + \mu_s(A \cap \partial D) + \mu_p(A \cap \{(c_1, c_2)\})$ .) The optimization problem (2) can now be written as

$$\max_u \int_D u d\bar{\mu} \tag{3}$$

subject to  $\{(a) u$  convex,  $(b) \nabla u(z) \in [0, 1]^2$  a.e.  $z \in D$ ,  $(c) u(c_1, c_2) = 0.\}$

The objective function remains unchanged upon the addition of the point measure  $\mu_p$  since  $u(c_1, c_2) = 0$ . We now recall the definition of the convex ordering relation. A function  $f$  is increasing if  $z \geq z'$  component-wise implies  $f(z) \geq f(z')$ .

**Definition 2** [2]. *Let  $\alpha$  and  $\beta$  be measures defined on a set  $D$ . We say  $\alpha$  convex dominates  $\beta$  ( $\alpha \succeq_{cvx} \beta$ ) if  $\int_D f d\alpha \geq \int_D f d\beta$  for all convex and increasing  $f$ .*



The dual of problem (3) is found to be [2, Theorem 2]

$$\inf_{\gamma: D \times D \rightarrow \mathbb{R}_+} \int_{D \times D} c(x, y) d\gamma(x, y) \tag{4}$$

subject to  $\{\gamma_1 - \gamma_2 \succeq_{c_v x} \bar{\mu}, \text{ where } \gamma(\cdot, D) = \gamma_1(\cdot) \text{ and } \gamma(D, \cdot) = \gamma_2(\cdot)\}$

where  $c(x, y) = \sum_{i=1}^2 (x_i - y_i)_+$ . The next lemma gives a sufficient condition for strong duality.

**Lemma 3** [2, Corollary 1]. *Let  $u^*$  and  $\gamma^*$  be feasible for the aforementioned primal (3) and dual (4) problems, respectively. Then  $\int_D u^* d\bar{\mu} = \int_{D \times D} c d\gamma^*$  iff (i)  $\int_D u^* d(\gamma_1 - \gamma_2) = \int_D u^* d\bar{\mu}$ , and (ii)  $u^*(x) - u^*(y) = c(x, y), \gamma^* - a.e.$*

Our problem now reduces to that of finding a  $\gamma$  such that  $\gamma_1 - \gamma_2$  convex dominates  $\bar{\mu}$ , and satisfies the conditions stated in Lemma 3. The key, nontrivial technical contribution of our paper is to identify such a  $\gamma$  when  $z \sim \text{Unif}[c_1, c_1 + b_1] \times [c_2, c_2 + b_2]$ , for all  $(c_1, c_2, b_1, b_2) \geq 0$ .

### 3 A Discussion of Optimal Solutions for the General Case

We now discuss the space of solutions for any nonnegative  $c_1, c_2, b_1, b_2$ . The main result of this paper is that the optimal mechanism is given as follows.

When the values of  $c_1$  and  $c_2$  are small, in the sense that

$$\{c_1 \leq b_1, c_2 \leq 2b_2(b_1 + c_1)/(b_1 + 3c_1)\} \cup \{c_2 \leq b_2, c_1 \leq 2b_1(b_2 + c_2)/(b_2 + 3c_2)\}, \tag{5}$$

the optimal menu is one of the four menus in Fig. 1. We discuss this in detail in Sect. 4.1. We describe the exact menu and associated allocations in Theorem 6.

When the value of  $c_2$  is large but  $c_1$  is small, in the sense that

$$\{c_1 \leq b_1, c_2 > 2b_2(b_1 + c_1)/(b_1 + 3c_1)\},$$

the optimal menu is one of the three menus in Fig. 2. We discuss this in detail in Sect. 4.2. We give the exact menu and associated allocations in Theorem 8.

When the value of  $c_1$  is large but  $c_2$  is small, in the sense that

$$\{c_2 \leq b_2, c_1 > 2b_1(b_2 + c_2)/(b_2 + 3c_2)\},$$


the optimal menu is one of the three menus in Fig. 3. Notice that this is symmetric to the case where  $c_2$  is large but  $c_1$  is small.

In the remaining cases when  $c_1$  and  $c_2$  are large, in the sense that


$$\{c_1 \geq b_1, c_2 \geq b_2\},$$

the optimal mechanism is given by pure bundling as in, for example, Fig. 1d.


Figure 4 provides a self-explanatory ‘phase diagram’ of optimal menus as a function of  $\frac{c_1}{b_1}$  and  $\frac{c_2}{b_2}$ . Interesting cases occur when either  $c_1 \leq b_1$  or  $c_2 \leq b_2$ .



**Fig. 1.** The structure of optimal mechanism when both  $c_1$  and  $c_2$  are small. The  $(q_1, q_2)$  in each region indicates the corresponding allocation probabilities.



**Fig. 2.** The structure of optimal mechanism when  $c_1$  is small,  $c_2$  is large.



**Fig. 3.** The structure of optimal mechanism when  $c_2$  is small,  $c_1$  is large.

Defining the power rate  $\Delta : D \rightarrow \mathbb{R}$  as  $\Delta(z_1, z_2) := -3 - z_1 f_1'(z_1)/f_1(z_1) - z_2 f_2'(z_2)/f_2(z_2)$ , Wang and Tang [12] proved that the optimal mechanism for a distribution with  $\Delta$  equaling a negative constant, has at most four constant allocation regions. The uniform distribution has power rate  $-3$  for all  $z \in D$ . Observe that in each of the menus in Figs. 1, 2, and 3, the number of regions is at most four, in agreement with the result in [12].

When  $c_1$  is small and  $c_2$  is large, in the sense that  $\{c_1 \leq b_1, c_2 \geq 2b_2(b_1/(b_1 - c_1))^2\}$ , we show (in Theorem 8) that the optimal mechanism is to sell the second item with probability 1 for the least valuation  $c_2$ , and sell item 1 at a reserve price as indicated by Myerson’s revenue maximizing mechanism. This is of course intuitive. Similar is the case when  $c_2$  is small and  $c_1$  is sufficiently large.

When both  $c_1$  and  $c_2$  are large,  $\{c_1 \geq b_1, c_2 \geq b_2\}$ , the optimal mechanism is to bundle the two goods and sell the bundle at the reserve price. The reserve price  $c_1 + c_2 + p$  is such that  $p = \left(\sqrt{(c_1 + c_2)^2 + 6b_1b_2} - c_1 - c_2\right)/3$ . So as  $c_1 + c_2 \rightarrow \infty$ , the reserve price is  $c_1 + c_2 + O(b_1b_2/(c_1 + c_2))$ . When  $c_1 < b_1$ , the mechanism is deterministic for any  $c_2 \geq 2b_2(b_1/(b_1 - c_1))^2$ , and when  $c_1 \geq b_1$ , it is deterministic for any  $c_2 \geq b_2$ .

### 4 The Solution for the Uniform Density on a Rectangle

In this section, we determine the optimal mechanism when  $z \sim \text{Unif}[c_1, c_1 + b_1] \times [c_2, c_2 + b_2]$ . The solution proceeds similar to the general characterization in [1, Sect. 10] that uses the optimal transport method. We compute the components of  $\bar{\mu}$  (i.e.,  $\mu, \mu_s, \mu_p$ ), with  $f(z) = \frac{1}{b_1b_2}$  for  $z \in D = [c_1, c_1 + b_1] \times [c_2, c_2 + b_2]$ , as

$$\begin{aligned} \text{(area density)} \quad \mu(z) &= -3/(b_1b_2), \quad z \in D \\ \text{(line density)} \quad \mu_s(z) &= \sum_{i=1}^2 (-c_i \mathbf{1}(z_i = c_i) + (c_i + b_i) \mathbf{1}(z_i = c_i + b_i))/(b_1b_2), \\ & \hspace{25em} z \in \partial D \\ \text{(point measure)} \quad \mu_p(\{c_1, c_2\}) &= 1. \end{aligned}$$

We define the functions  $s_i : [c_i, c_i + b_i] \rightarrow [c_{-i}, c_{-i} + b_{-i}]$ ,  $i = 1, 2$ , as follows.


$$s_i(z_i) := \left\{ z_{-i}^* \in [c_{-i}, c_{-i} + b_{-i}] : \int_{z_{-i}^*}^{c_{-i} + b_{-i}} \bar{\mu}(z_i, z_{-i}) dz_{-i} + \bar{\mu}(z_i, c_{-i} + b_{-i}) = 0 \right\}.$$

The integral denotes the integration of the densities of  $\bar{\mu}$ , and the equation an equation of appropriate densities. Now we define the *zero set*  $Z$  as the set formed by the intersection of  $\{(z_1, z_2) : z_2 \leq s_1(z_1)\}$ ,  $\{(z_1, z_2) : z_1 \leq s_2(z_2)\}$ , and  $\{(z_1, z_2) : z_1 + z_2 \leq c_1 + c_2 + p\}$ , with  $p$  chosen to satisfy  $\bar{\mu}(Z) = 0$ . We denote the point of intersection between the curves  $z_2 = s_1(z_1)$  and  $z_1 + z_2 = c_1 + c_2 + p$  by  $P$ , and that between  $z_1 = s_2(z_2)$  and  $z_1 + z_2 = c_1 + c_2 + p$  by  $Q$ . Let  $P = (P_1, P_2)$  and  $Q = (Q_1, Q_2)$  denote the respective co-ordinates of  $P$  and  $Q$ . The regions  $A, B$ , and  $W$ , see Fig. 5, are defined as


$$\begin{aligned} A &:= ([c_1, P_1] \times [P_2, c_2 + b_2]) \setminus Z; \quad B := ([Q_1, c_1 + b_1] \times [c_2, Q_2]) \setminus Z; \\ & \hspace{15em} W := D \setminus (Z \cup A \cup B). \end{aligned}$$

The partition of  $D$  into  $A, B, Z$ , and  $W$  is termed a *canonical partition*.

Suppose  $(c_1, c_2) = (0, 0)$ . Then we have  $s_i(z_i) = 2b_{-i}/3$  for all  $z_i \in [0, b_i]$ . So the  $s_i$  functions turn out to be constants when  $(c_1, c_2) = (0, 0)$ . When  $c_i > 0$ , the function  $s_i$  is not defined for  $z_i = c_i$ , because the density of  $\bar{\mu}(c_i, z_{-i})$  is negative



**Fig. 4.** A phase diagram of the optimal mechanism.



**Fig. 5.** The canonical partition of the set  $D$ .

for all  $z_{-i}$ . Thus when  $(c_1, c_2) \neq (0, 0)$ , we need to shuffle  $\bar{\mu}$  to get around this issue. We add to  $\bar{\mu}$  a “shuffling measure”  $\bar{\alpha}$  that convex dominates 0. We then find  $s_i(z_i)$  associated with  $\bar{\mu} + \bar{\alpha}$ , instead of  $\bar{\mu}$ .

We now investigate the possible structure of such a shuffling measure. We know (from [9]) that the optimal menu is as in Fig. 1a when  $z \sim \text{Unif}[c, c + 1]^2$ ,  $c \leq 0.077$ , and also that  $a_1 = a_2 = 0$  when  $z \sim \text{Unif}[0, 1]^2$ . In other words, the functions  $s_i(z_i)$  are constant when  $c = 0$ , and linear when  $c \in (0, 0.077]$ . Thus we anticipate that adding a linear shuffling measure on the top-left and bottom-right boundaries, with point measures at the top-left and bottom-right corners to neutralize the negative line densities at  $z_i = c_i$ , would yield the solution.

Daskalakis et al. [2] identify such a shuffling measure when they solve for  $D = [4, 16] \times [4, 7]$ . In this paper, we identify shuffling measures for all rectangles  $D$  on the positive quadrant. We first suitably parametrize the shuffling measure  $\bar{\alpha}$ . For each candidate  $\bar{\alpha}$ , we identify the partition associated with  $\bar{\mu} + \bar{\alpha}$ . Then for each rectangle  $D$ , we identify the parameters of  $\bar{\alpha}$  so that the partition becomes a canonical partition. We begin by describing a simple shuffling measure  $\bar{\alpha}$ .


### 4.1 The Menus When $c_1$ and $c_2$ Are Small

For  $p_{a_1}, m_1 > 0$ , define a linear function  $\alpha_s^{(1)} : [c_1, c_1 + m_1] \rightarrow \mathbb{R}$  given by

$$\alpha_s^{(1)}(x) := (2b_2 - c_2 - 3p_{a_1} + 3a_1(x - c_1))/(b_1b_2), \quad x \in [c_1, c_1 + m_1]. \quad (6)$$

As before, we will reuse  $\alpha_s^{(1)}$  to denote the measure with density given by (6). Define  $\alpha_p^{(1)} := c_1(b_2 - p_{a_1})/(b_1b_2)\delta_{c_1}$ , a point measure of mass  $c_1(b_2 - p_{a_1})/(b_1b_2)$  at location  $c_1$ . Finally, define  $\bar{\alpha}^{(1)} := \alpha_s^{(1)} + \alpha_p^{(1)}$ . See Fig. 6a. We assert the following. All proofs are found in the archival version of this paper.

**Proposition 4.** Suppose  $m_1 = \frac{4c_1(b_2 - p_{a_1})}{c_2 - 2b_2 + 3p_{a_1}}$  and  $a_1 = \frac{(c_2 - 2b_2 + 3p_{a_1})^2}{8c_1(b_2 - p_{a_1})}$ . Then,  $\bar{\alpha}^{(1)}([c_1, c_1 + m_1]) = \int_{[c_1, c_1 + m_1]} x \bar{\alpha}^{(1)}(dx) = 0$ , and hence  $\int_{[c_1, c_1 + m_1]} f d\bar{\alpha}^{(1)} = 0$  for any affine function  $f$  on  $[c_1, c_1 + m_1]$ . Furthermore,  $\bar{\alpha}^{(1)} \succeq_{cvx} 0$ .



**Fig. 6.** (a) The measure  $\bar{\alpha}$ . (b) The measure  $\bar{\beta}$ .

For  $p_{a_2}, m_2 > 0$ , we define a similar line measure  $\alpha_s^{(2)}$  at the interval  $[c_2, c_2 + m_2]$ , a point measure  $\alpha_p^{(2)}$  at  $c_2$ , and  $\bar{\alpha}^{(2)} := \alpha_s^{(2)} + \alpha_p^{(2)}$ . Define  $\tilde{D}^{(1)} := [c_1, c_1 + m_1] \times \{c_2 + b_2\}$ , an interval on the top boundary of  $D$  starting from the top-left corner, and  $\tilde{D}^{(2)} := \{c_1 + b_1\} \times [c_2, c_2 + m_2]$ , an interval on the right boundary of  $D$  starting from bottom-right corner. We add  $\bar{\alpha}^{(i)}$  measure at  $\tilde{D}^{(i)}$ . Define  $\bar{\alpha} := \bar{\alpha}^{(1)} + \bar{\alpha}^{(2)}$  on  $D$ . The functions  $s_i(\cdot)$  associated with  $\bar{\mu} + \bar{\alpha}$  are given by

$$s_i(z_i) = \begin{cases} c_{-i} + p_{a_i} - a_i(z_i - c_i), & c_i \leq z_i \leq c_i + m_i \\ \frac{2}{3}(c_{-i} + b_{-i}), & c_i + m_i < z_i \leq c_i + b_i. \end{cases}$$

We have thus found a shuffling measure  $\bar{\alpha}$  such that  $s_i(\cdot)$  associated with  $\bar{\mu} + \bar{\alpha}$  are defined everywhere in  $[c_i, c_i + b_i)$ , and also such that  $\bar{\alpha} \succeq_{cv,x} 0$ . The functions  $s_i(z_i)$  are indicated in Fig. 5. Observe that fixing  $p_{a_i}$  fixes the slopes of segments  $PP'$  and  $QQ'$  to be  $-a_1$  and  $-1/a_2$ , where

$$a_1 = \frac{(c_2 - 2b_2 + 3p_{a_1})^2}{8c_1(b_2 - p_{a_1})}, \quad a_2 = \frac{(c_1 - 2b_1 + 3p_{a_2})^2}{8c_2(b_1 - p_{a_2})}. \tag{7}$$

The point  $P = (c_1 + m_1, c_2 + p_{a_1} - m_1 a_1)$ , and the point  $Q = (c_1 + p_{a_2} - m_2 a_2, c_2 + m_2)$ , are also fixed to be

$$\begin{aligned} P &= (c_1 + 4c_1(b_2 - p_{a_1}) / (c_2 - 2b_2 + 3p_{a_1}), c_2 + (2b_2 - c_2 - p_{a_1}) / 2), \\ Q &= (c_1 + (2b_1 - c_1 - p_{a_2}) / 2, c_2 + 4c_2(b_1 - p_{a_2}) / (c_1 - 2b_1 + 3p_{a_2})). \end{aligned} \tag{8}$$

To complete the canonical partition, it remains to construct the zero set  $Z$ . Recall that the zero set is the intersection of the sets  $\{z_1 \leq s_2(z_2)\}$ ,  $\{z_2 \leq s_1(z_1)\}$ , and  $\{z_1 + z_2 \leq c_1 + c_2 + p\}$ , where  $p$  is fixed so that  $\bar{\mu}(Z) = 0$ . The construction of  $Z$  in Fig. 1a thus imposes two constraints: (1) The points  $P$  and  $Q$  lie on the line  $z_1 + z_2 = c_1 + c_2 + p$ , and (2)  $\bar{\mu}(Z) = 0$ .

The points  $P$  and  $Q$  lie on the line  $z_1 + z_2 = c_1 + c_2 + p$  for some  $p$ , when  $P_1 + P_2 = Q_1 + Q_2$  holds. Substituting the expressions for  $P$  and  $Q$  (from (8)) and simplifying, we have

$$\begin{aligned} &(c_1 - c_2 - 2(b_1 - b_2) - (p_{a_1} - p_{a_2}))(c_1 - 2b_1 + 3p_{a_2})(c_2 - 2b_2 + 3p_{a_1}) \\ &+ 8c_1(b_2 - p_{a_1})(c_1 - 2b_1 + 3p_{a_2}) - 8c_2(b_1 - p_{a_2})(c_2 - 2b_2 + 3p_{a_1}) = 0. \end{aligned} \tag{9}$$

We now have  $P_1 + P_2 = Q_1 + Q_2 = c_1 + c_2 + p$ , from which we derive  $p$  as

$$p = \frac{4c_1(b_2 - p_{a_1})}{c_2 - 2b_2 + 3p_{a_1}} + \frac{2b_2 - c_2 - p_{a_1}}{2} = \frac{2b_1 - c_1 - p_{a_2}}{2} + \frac{4c_2(b_1 - p_{a_2})}{c_1 - 2b_1 + 3p_{a_2}}. \quad (10)$$

The values of  $p_{a_1}$  and  $p_{a_2}$  are thus constrained to satisfy (9). Now we have only one free variable, say  $p_{a_1}$ , which when fixed identifies the canonical partition. This is done via the constraint  $\bar{\mu}(Z) = 0$ . Note that the regions  $A$  and  $B$  have been constructed to satisfy  $\bar{\mu}(A) = \bar{\mu}(B) = 0$ , and we have already established in Sect. 2 that  $\bar{\mu}(D) = 0$ . So asking for  $\bar{\mu}(Z) = 0$  is the same as asking  $\bar{\mu}(W) = 0$ . From an examination of Fig. 1a and simple geometry considerations, it is immediate that  $-\bar{\mu}(W) = 0$  iff

$$3(b_1 - (P_1 - c_1))(b_2 - (Q_2 - c_2)) - (c_2 + b_2)(b_1 - (P_1 - c_1)) - (c_1 + b_1)(b_2 - (Q_2 - c_2)) - \frac{3}{2}(P_2 - Q_2)(Q_1 - P_1) = 0. \quad (11)$$

This upon substitution of the expressions for  $P$  and  $Q$  yields

$$\begin{aligned} & 3\Pi_{i=1}^2 (b_i(c_{-i} - 2b_{-i} + 3p_{a_i}) - 4c_i(b_{-i} - p_{a_i})) \\ & - \sum_{i=1}^2 ((b_i(c_{-i} - 2b_{-i} + 3p_{a_i}) - 4c_i(b_{-i} - p_{a_i}))(c_i - 2b_i + 3p_{a_{-i}})(c_{-i} + b_{-i})) \\ & - 3/8\Pi_{i=1}^2 ((2b_{-i} - c_{-i} - p_{a_i})(c_i - 2b_i + 3p_{a_{-i}}) - 4c_{-i}(b_i - p_{a_{-i}})) = 0. \quad (12) \end{aligned}$$

We now proceed to find the values of  $c_i, b_i$  for which  $\exists (p_{a_1}, p_{a_2}) \in [0, b_2] \times [0, b_1]$  solving (9) and (12) simultaneously. We observe from (7) and (8) that the slopes  $a_i$  and the points  $P$  and  $Q$  change as a function of  $p_{a_i}$ , as follows.

**Observation 5.** *Assume  $p_{a_i} \in [(2b_{-i} - c_{-i})/3, b_{-i}]$ . Then, an increase in  $p_{a_1}$  increases  $a_1$ , and moves  $P$  towards south-west (i.e., decreases both  $P_1$  and  $P_2$ ). Similarly, an increase in  $p_{a_2}$  increases  $a_2$ , and moves  $Q$  towards south-west.*

We now define  $p_{a_i}^* := \frac{2b_{-i} - c_{-i}}{3} - \frac{4c_i}{9} + \frac{2}{9}\sqrt{2c_i(2c_i + 3(b_{-i} + c_{-i}))}$  and  $r_i := \frac{2b_{-i}(2b_i + 5c_i) - c_{-i}(2b_i - 3c_i)}{3(2b_i + 3c_i)}$ . Observe that  $a_i = 1$  when  $p_{a_i} = p_{a_i}^*$ , and that  $p_{a_i}^* \in [(2b_{-i} - c_{-i})/3, b_{-i}]$ . We consider four cases based on the values of  $r_i$  and  $p_{a_i}^*$ .

Consider case 1, when  $r_i \leq p_{a_i}^*$ ,  $i = 1, 2$ . Then we have  $a_i \leq 1$  at  $p_{a_i} = r_i$ . We now make a series of claims, proofs of which are in the archival version.

1. When  $c_i, b_i$  satisfy (5), then  $r_i \in [(2b_{-i} - c_{-i})/3, b_{-i}]$ ,  $i = 1, 2$ . When  $p_{a_i} = r_i$ , we have  $P = Q$ , i.e., the points  $P$  and  $Q$  coincide. Moreover,  $\bar{\mu}(W) \geq 0$ .
2. Fixing  $p_{a_i} = r_i$ , we now increase  $p_{a_1}$ , and adjust  $p_{a_2}$  so that  $P_1 + P_2 = Q_1 + Q_2$  holds (i.e., (9) is satisfied). Then, we claim that  $p_{a_2}$  increases as well. We also claim that  $P$  lies to the north-west of  $Q$ , if  $p_{a_i} \in [r_i, p_{a_i}^*]$ .
3. If  $P$  is to the north-west of  $Q$ , then  $\bar{\mu}(W)$  decreases with increase in  $p_{a_i}$ .
4. If  $\bar{\mu}(W) = 0$  for some  $p_{a_i} \in [r_i, p_{a_i}^*]$ , then the optimal menu is as in Fig. 1a.

5. Suppose now that  $p_{a_2}$  equals  $p_{a_2}^*$ , but still  $\bar{\mu}(W) > 0$  and  $p_{a_1} < p_{a_1}^*$ . Then,  $a_2$  equals 1, and we stop further increase in  $p_{a_2}$ . The menu in Fig. 1a coincides with the menu in Fig. 1b, and the regions  $B$  and  $W$  of the canonical partition merge. The picture is as in Fig. 7a. As we increase  $p_{a_1}$  to  $p'_{a_1} > p_{a_1}$  further, we hold  $Q$ ,  $p_{a_2}$ , and  $s_2(z_2)$  corresponding to the  $p_{a_1}$  that led to  $a_2 = 1$ . So the partition moves from Fig. 7a to b, when we increase  $p_{a_1}$ . We now claim that increasing  $p_{a_1}$  moves  $P$  towards south-west, and also decreases  $\bar{\mu}(W)$ .
6. Equating  $-\bar{\mu}(W) = 0$  for Fig. 1b, and substituting for  $a_1$  and  $p$ , we have

$$\begin{aligned}
 & -8c_1b_2^2 + (c_2b_1 - b_2b_1 - b_2c_1)(c_2 - 2b_2) + (c_2/2 - b_1)(c_2 - 2b_2)^2 - 3/8(c_2 - 2b_2)^3 \\
 & + (c_1(4c_2 - 3b_2) + 3b_1b_2 + 2(c_2 - 2c_1)(c_2 - 2b_2) - 15/8(c_2 - 2b_2)^2)p_{a_1} \\
 & + (3c_2/2 - 21/8(c_2 - 2b_2))p_{a_1}^2 - (9/8)p_{a_1}^3 = 0.
 \end{aligned} \tag{13}$$

If  $\bar{\mu}(W) = 0$  for some  $p_{a_1} \leq p_{a_1}^*$ , then the optimal menu is as in Fig. 1b.


7. Suppose now that  $p_{a_1}$  equals  $p_{a_1}^*$ , but still  $\bar{\mu}(W) > 0$ . Then  $a_1$  equals 1, and the picture is as in Fig. 8a. The regions  $A$  and  $W$  in the canonical partition merge, and the menu in Fig. 1b coincides with the menu in Fig. 1d. At this point, we stop further increase in  $p_{a_1}$ , and decrease  $p$ . We hold  $P$  and  $s_1(z_1)$  corresponding to  $p_{a_1} = p_{a_1}^*$ . The partition moves from Fig. 8a to b. We claim that a decrease in  $p$  decreases  $\bar{\mu}(W)$ . Furthermore, it is easy to see that  $\bar{\mu}(W) < 0$  when  $p = 0^+$ . Hence, by the continuity of  $\bar{\mu}(W)$  in the  $p$  parameter, there must be a  $p$  when  $\bar{\mu}(W) = 0$ . This happens when  $p = p^* := \left( \sqrt{(c_1 + c_2)^2 + 6b_1b_2} - c_1 - c_2 \right) / 3$ . The resulting partition is optimal, and the menu is as in Fig. 1d.
8. If in point 5, we have  $p_{a_1} = p_{a_1}^*$ , but still  $\bar{\mu}(W) > 0$  and  $p_{a_2} < p_{a_2}^*$ , then all arguments hold by symmetry, but the menu is as in Fig. 1c instead of 1b.

Consider case 2, when  $r_1 \leq p_{a_1}^*$  but  $r_2 > p_{a_2}^*$ . Then we have  $a_1 \leq 1$  and  $a_2 > 1$  at  $p_{a_i} = r_i$ . In this case, we hold  $Q$  and  $s_2(z_2)$  corresponding to  $p_{a_2} = r_2$ , but change the line joining  $Q$  and  $p_{a_2}$  to a line of slope  $-1$ , as illustrated in Fig. 9. Now the menu is the same as the menu in Fig. 1b, with  $p_{a_1} = r_1$ . We claim that the case is now similar to case 1, from point 5 onwards. The results are symmetric for case 3, when  $r_2 \leq p_{a_2}^*$  but  $r_1 > p_{a_1}^*$ . Finally, we claim that the optimal menu is as in Fig. 1d for case 4, when  $r_i \geq p_{a_i}^*$ ,  $i = 1, 2$ .


The decision tree in Fig. 10 summarizes the procedure described in cases 1–4. We now have the following theorem.

**Theorem 6.** *If  $c_i, b_i$  satisfy (5), then the optimal menu is one of the four menus in Fig. 1. The exact menu, and the values of  $(p_{a_i}, a_i, p)$ , are given as follows.*


- (a) *The optimal menu is as in Fig. 1a, if  $\exists p_{a_i} \in [r_i, p_{a_i}^*]$ ,  $i = 1, 2$ , solving (9) and (12) simultaneously. Values of  $(a_i, p)$  are found from (7) and (10).*
- (b) *The optimal menu is as in Fig. 1b, if (i) the condition in (a) fails, and (ii)  $\exists p_{a_1} \in [r_1, p_{a_1}^*]$  that solves (13). Values of  $(a_1, p)$  are found from (7) and (10). The optimal menu is as in Fig. 1c, when an analogous statement holds.*
- (c) *The optimal menu is as in Fig. 1d with  $p = p^*$ , if conditions in (a),(b) fail.*




**Fig. 7.** Illustration of the transition from Fig. 1a to b.



**Fig. 8.** Illustration of the transition from Fig. 1b to d.



**Fig. 9.** An illustration of change of slope in case 2, when  $r_1 \leq p_{a_1}^*$  but  $r_2 > p_{a_2}^*$ .



**Fig. 10.** The decision tree illustrating how to choose optimal menu in Fig. 1. Case  $i$  starts from State  $i$ ,  $i = 1, 2, 3, 4$ .

## 4.2 The Menus for Other Values of $c_1$ and $c_2$

Consider the case when  $c_1 \leq b_1$ , but  $c_2 > 2b_2(b_1 + c_1)/(b_1 + 3c_1)$ . Notice that  $c_2 > b_2$  for any value of  $c_1 \leq b_1$ , and thus this case violates (5). When  $c_2$  crosses  $2b_2(b_1 + c_1)/(b_1 + 3c_1)$ , we claim that the point  $P$  moves outside  $D$  (proof in the archival version). At  $c_2 = 2b_2(b_1 + c_1)/(b_1 + 3c_1)$ , we have  $P_2 = c_2$ , i.e.,  $P$  touches the bottom boundary of  $D$ . The menu in Fig. 1b coincides with the menu in Fig. 2b. We anticipate that adding on the top-left boundary a shuffling measure that is linear in  $[c_1, c_1 + p_{a_1}/a_1]$  and is a constant in  $[c_1 + p_{a_1}/a_1, c_1 + p]$ , would yield the menu in Fig. 2b. We now construct such a shuffling measure  $\tilde{\beta}$  (see Fig. 6b). Define the density  $\beta_s$  in the interval  $[c_1, c_1 + p]$  as

$$\beta_s(x) := (2b_2 + (3a_1(x - c_1) - c_2 - 3p_{a_1})\mathbf{1}(x \leq c_1 + p_{a_1}/a_1))/(b_1b_2). \quad (14)$$

As before, we will reuse  $\beta_s$  to denote the measure with density given by (14). Define  $\beta_p := c_1(b_2 - p_{a_1})/(b_1b_2)\delta_{c_1}$ , a point measure at  $c_1$  with mass  $c_1(b_2 - p_{a_1})/(b_1b_2)$ . Define  $\tilde{\beta} := \beta_s + \beta_p$ . This shuffling measure is added at the interval  $\tilde{D}^{(1)} = [c_1, c_1 + p] \times \{c_2 + b_2\}$ . Observe that the structures of  $\tilde{\beta}$  and  $\bar{\alpha}$  differ only in the interval  $[c_1 + p_{a_1}/a_1, c_1 + p]$ . A jump occurs when  $\beta_s$  reaches  $2b_2 - c_2$ .

The function  $s_1 : [c_1, c_1 + p] \rightarrow [c_2, c_2 + b_2]$ , associated with  $\tilde{\mu} + \tilde{\beta}$ , is given by  $s_1(z_1) = c_2 + (p_{a_1} - a_1(z_1 - c_1))_+$ . We now derive the parameters  $(p_{a_1}, a_1, p)$ , by imposing (1)  $\tilde{\beta} \succeq_{cvx} 0$ , (2)  $\int_{\tilde{D}^{(1)}} f d\tilde{\beta} = 0$  for any affine  $f$ , and (3)  $\tilde{\mu}(W) = 0$ .

**Proposition 7.** *Consider the menu in Fig. 2b. Then,  $\tilde{\mu}(W)$  equals 0 when  $p = (b_1 - c_1)/2$ . Furthermore, the conditions (1)  $\tilde{\beta} \succeq_{cvx} 0$ , and (2)  $\int_{\tilde{D}^{(1)}} f d\tilde{\beta} = 0$  for any affine  $f$ , are satisfied, if  $p_{a_1}$  is a solution to*



$$2b_1^2b_2^2c_2 - c_2^2b_2(b_1 - c_1)^2 + p_{a_1}(2b_1^2b_2^2 - 4b_1b_2c_1c_2 - 3c_2b_2(b_1 - c_1)^2) + p_{a_1}^2(2c_1^2c_2 - 4b_1b_2c_1 - 9b_2(b_1 - c_1)^2/4) + p_{a_1}^3(2c_1^2) = 0, \quad (15)$$

and if  $a_1$  satisfies

$$a_1 = p_{a_1}((3/2)p_{a_1} + c_2)/(b_1b_2 - c_1p_{a_1}) = p_{a_1} \left( \sqrt{2(p_{a_1} + c_2)/b_2} \right) / (b_1 - c_1). \quad (16)$$

We now have the following theorem.

**Theorem 8.** *Let  $\{c_1 \leq b_1, c_2 \in [2b_2(b_1 + c_1)/(b_1 + 3c_1), 2b_2(b_1/(b_1 - c_1))^2]\}$ . Then,  $\exists p_{a_1} \in [(2b_2 - c_2)_+, b_2]$  that solves (15). If  $a_1$  found from (16) is at most 1, then the optimal menu is as in Fig. 2b, with  $(p_{a_1}, a_1)$  as found above, and  $p = (b_1 - c_1)/2$ . If instead,  $a_1 > 1$ , then the optimal menu is as in Fig. 2a, with  $p = p^*$ . The optimal menu is as in Fig. 2c, when  $\{c_1 \leq b_1, c_2 \geq 2b_2(b_1/(b_1 - c_1))^2\}$ .*

In the proof, we analyze how  $(a_1, p_{a_1}/a_1, p)$  change, when  $p_{a_1}$  changes. We then decrease  $p_{a_1}$  starting from  $b_2$ , modify  $(a_1, p)$  according to certain equations, and arrive at some  $p_{a_1} \geq (2b_2 - c_2)_+$  that satisfies all the required conditions. The dynamics of the proof is similar to how we arrive at optimal menus in Sect. 4.1.

For the case when  $c_2 < b_2$  but  $c_1 > 2b_1(b_2 + c_2)/(b_2 + 3c_2)$ , the arguments are symmetric. An assertion similar to Theorem 8 can be made. When  $c_i \geq b_i$ , we have the following theorem.

**Theorem 9.** *The optimal menu is as in Fig. 1d, when  $\{c_1 \geq b_1, c_2 \geq b_2\}$ .*

## 5 Extension to Other Distributions

Optimal mechanism for uniform distribution over any rectangle was found only using  $\bar{\alpha}$  and its variants as shuffling measures. We can now ask if there is a generalization of  $\bar{\alpha}$  for other distributions. For distributions with constant negative power rate, we anticipate that the optimal menus would be of the same form as in uniform distributions (based on the result of Wang and Tang [12]), and that they can be arrived using similar  $\bar{\alpha}$ . We now consider an example of such a distribution. Let  $f_1(z) = f_2(z) = 2z/(2c + 1)$  when  $z \in [c, c + 1]$ . The power rate  $\Delta = -5$ , a negative constant. We use a “generalized”  $\bar{\alpha}$ , whose components are

$$\begin{aligned} \alpha_s(z_1, c_2 + b_2) &= f_1(z_1)(-\Delta(1 - F_2(p_{a_1} - a_1(z_1 - c_1))) - (c_2 + b_2)f_2(c_2 + b_2)), \\ & \quad z \in [c_1, P_1], \\ \alpha_p(c_1, c_2 + b_2) &= c_1f_1(c_1)(1 - F_2(p_{a_1})). \end{aligned} \quad (17)$$

Observe that this  $\bar{\alpha}$  reduces to the  $\bar{\alpha}$  used in the case of uniform distributions, when we substitute  $f_i(z_i) = 1$ , and  $\Delta = -3$ . We now have the following theorem.

**Theorem 10.** *Let  $f_i(z) = 2z/(2c + 1)$ ,  $z \in [c, c + 1]$ . For  $c = 0.1$ , the optimal menu is as in Fig. 1a, with  $p_{a_1} = 0.79615$ ,  $a_1 = 0.23198$ , and  $P_1 = 0.364655$ .*

The proof traces the same steps as in Sect. 4.1. This suggests that our method of constructing  $\bar{\alpha}$  is useful for a wider class of distributions. We conjecture that our method works for all distributions with constant negative power rate.

## References

1. Daskalakis, C., Deckelbaum, A., Tzamos, C.: Mechanism design via optimal transport. In: The 14th ACM Conference on Electronic Commerce, EC 2013, pp. 269–286. ACM, New York (2013)
2. Daskalakis, C., Deckelbaum, A., Tzamos, C.: Strong duality for a multiple-good monopolist (2014). <http://arxiv.org/pdf/1409.4150v1.pdf>
3. Daskalakis, C., Deckelbaum, A., Tzamos, C.: Strong duality for a multiple-good monopolist. In: The 16th ACM Conference on Economics and Computation, EC 2015, pp. 449–450. ACM, New York (2015)
4. Giannakopoulos, Y., Koutsoupias, E.: Duality and optimality of auctions for uniform distributions. In: The 15th ACM Conference on Economics and Computation, EC 2014, pp. 259–276. ACM, New York (2014)
5. Giannakopoulos, Y., Koutsoupias, E.: Selling two goods optimally. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015. LNCS, vol. 9135, pp. 650–662. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-47666-6\\_52](https://doi.org/10.1007/978-3-662-47666-6_52)
6. Manelli, A.M., Vincent, D.R.: Bundling as an optimal selling mechanism for a multiple-good monopolist. *J. Econ. Theory* **127**(1), 1–35 (2006)
7. Manelli, A.M., Vincent, D.R.: Multidimensional mechanism design: revenue maximization and the multiple-good monopoly. *J. Econ. Theory* **137**(1), 153–185 (2007)
8. Myerson, R.B.: Optimal auction design. *Math. Oper. Res.* **6**(1), 58–73 (1981)
9. Pavlov, G.: Optimal mechanism for selling two goods. *B.E. J. Theor. Econ.* **11**(1), 1–35 (2011)
10. Rochet, J.C., Choné, P.: Ironing, sweeping, and multidimensional screening. *Econometrica* **66**(4), 783–826 (1998). <http://www.jstor.org/stable/2999574>
11. Rochet, J.C.: A necessary and sufficient condition for rationalizability in a quasi-linear context. *J. Math. Econ.* **16**(2), 191–200 (1987)
12. Wang, Z., Tang, P.: Optimal mechanisms with simple menus. In: The 15th ACM Conference on Economics and Computation, pp. 227–240. ACM (2014)

# Anonymous Auctions Maximizing Revenue

Christos Tzamos<sup>1</sup>(✉) and Christopher A. Wilkens<sup>2</sup>

<sup>1</sup> Massachusetts Institute of Technology, Cambridge, USA

[tzamos@mit.edu](mailto:tzamos@mit.edu)

<sup>2</sup> Yahoo Research, San Francisco, USA

[cwilkens@yahoo-inc.com](mailto:cwilkens@yahoo-inc.com)

**Abstract.** Auctions like sponsored search often implicitly or explicitly require that bidders are treated fairly. This may be because large bidders have market power to negotiate equal treatment, because small bidders are difficult to identify, or for many other reasons. We study so-called *anonymous* auctions to understand the revenue tradeoffs and to develop simple anonymous auctions that are approximately optimal.

We begin with the canonical digital goods setting and show that the optimal anonymous, ex-post incentive compatible auction has an intuitive structure — imagine that bidders are randomly permuted before the auction, then infer a posterior belief about bidder  $i$ 's valuation from the values of other bidders and set a posted price that maximizes revenue given this posterior.

We prove that no anonymous mechanism can guarantee an approximation better than  $\Theta(n)$  to the optimal revenue in the worst case (or  $\Theta(\log n)$  for regular distributions) and that even posted price mechanisms match those guarantees. Understanding that the real power of anonymous mechanisms comes when the auctioneer can infer the bidder identities accurately, we show a tight  $\Theta(k)$  approximation guarantee when each bidder can be confused with at most  $k$  “higher types”. Moreover, we introduce a simple mechanism based on  $n$  target prices that is asymptotically optimal. Finally, we return to our original motivation and build on this mechanism to extend our results to  $m$ -unit auctions and sponsored search.

**Keywords:** Revenue maximization · Auction design · Anonymous mechanisms

## 1 Introduction

In 1981, Myerson elegantly derived the revenue-optimal way to sell a single item [13] — each buyer's bid is transformed through a personalized virtual valuation function and then submitted to a standard second-price auction. Myerson's auction leverages precise prior beliefs in order to identify the bidder who generates the highest marginal expected revenue, allowing the seller to discriminate among bidders and extract more money from those with a higher willingness to pay.

For all its mathematical beauty, Myerson’s optimal auction violates an inherently desirable property: *fairness*. One definition of fairness says that the auctioneer should not *a priori* discriminate among the auction’s participants. It is a property that may be both desirable and necessary — it is undeniably philosophically important in many applications; moreover, many settings lack a strong notion of identity, precluding explicit discrimination.

Sponsored search illustrates the practical importance and limitations of treating bidders equally *ex-ante*. A typical sponsored search auction run by Google, Bing, or Yahoo matches bidders to ad slots on a page of search results — higher slots get more clicks, so higher bidders get higher slots. Suppose that the search engine identifies a group of queries where the market is thin, so the top bid is much higher than the second one. The search engine would like to enforce a premium price for the top slot; however, this effectively requires discriminating against the highest bidder.<sup>1</sup> Unfortunately, *ex-ante* discrimination may not be possible. Advertisers who are large will desire and demand “fair” treatment; due to their size, they may have the negotiating power to get it. Advertisers who are small lack the clout to demand equality; however, they are plentiful and could copy their accounts, blending into the masses to avoid explicit discrimination. As a result, search platforms like Google, Bing, and Yahoo may be prohibited from such discrimination out of necessity.

In this paper, we study the value of discriminating among your opponents in advance. Myerson’s optimal auction critically requires that the seller know the identities of bidders *ex-ante*, so that he can price discriminate among them — our goal is to quantify tradeoff inherent in requiring *ex-ante fairness* in *ex-post* incentive compatible auctions.

*Anonymous Mechanism Design.* An *anonymous auction* treats all bidders equally *ex-ante*. While the auctioneer may know information about the kinds of bidders who will participate — even knowing precise prior beliefs about bidders’ values — this information cannot *ex-ante* be used to discriminate among them. Alternatively, one may say that the auctioneer knows precise priors but does not know which prior belongs to which bidder. Technically, an auction is anonymous if and only if it is symmetric in the sense that permuting bids will analogously permute allocations and prices.

Anonymity is a weaker property compared to uniform pricing and envy-freeness that require that all bidders pay the same price if they receive the same goods. To see the potential power of anonymous mechanisms, consider the following example: in a sponsored search auction for two identical slots, two bidders have values  $v_1 = \$2$  and  $v_2 = \$1$  for a click, and the auctioneer knows these values precisely. The optimal mechanism gives a slot to the first bidder at

---

<sup>1</sup> In some sense, the search engine would like to set a reserve price for the top slot. However, this must be carefully defined when no bidder meets the reserve price or when more than one bidder meets it; the decreasing price mechanism we discuss later in this paper may be considered a natural interpretation of setting different reserve prices for different slots in a sponsored search auction.

a price of  $p_1 = \$2$  per click and a slot to the second bidder at a price  $p_2 = \$1$  per click. This mechanism however is not anonymous since it discriminates among bidders based on their identity. What can anonymous mechanisms do in this setting? If a common price per click was used, it would result into revenue loss because either the price would be too high for bidder 2 to afford or too low compared to the value of bidder 1. On the other hand, the following anonymous mechanism can extract the optimal revenue:

- if one bidder bids \$2, the other needs to bid at least \$1 to win a slot;
- if one bidder bids \$1, the other needs to bid at least \$2 to win a slot.

We begin by characterizing the optimal anonymous auction that is ex-post incentive compatible and ex-post individually rational. We show that it has a simple intuition in the digital goods setting: (1) imagine that bidders are relabeled uniformly at random before participating in the auction, then (2) use  $v_{-i}$  to infer a posterior belief about  $v_i$  and (3) choose a posted price for bidder  $i$  that maximizes revenue given this posterior. This intuition generalizes beyond the digital goods setting when the inferred posterior is regular. Some simple cases bear mention here: if the auctioneer’s prior is the same for all bidders (an IID setting) or if it is impossible to confuse bidders, the optimal anonymous auction will correctly deduce everyone’s identity and coincide with the unconstrained optimal auction.

With a basic understanding of anonymous auctions in hand, we study the performance of anonymous digital-goods auctions; our results are not immediately encouraging. We begin with a single-price mechanism — a simple and naturally anonymous auction — and show that it offers only a  $\Theta(n)$  approximation in general and a  $\Theta(\log n)$  approximation when priors are regular. Moreover, we show that the above results are tight even for the class of all anonymous mechanisms: prior beliefs exist so that no anonymous auction can guarantee revenue approximation better than  $\Theta(n)$  to the revenue of Myerson’s optimal auction while if bidders’ values are known to be drawn from uniform distributions, we can prove a lower-bound of  $\Omega(\log n)$ . Together, these suggest that general anonymous mechanisms cannot achieve better asymptotic guarantees than pricing in general settings and can be very far from optimal.

Having shown that anonymity can hurt revenue substantially in the worst case, we ask whether there are particular conditions under which anonymous auctions perform well. Our characterization of the optimal mechanism gives us hope: if all bidders are almost identical or almost perfectly distinguishable, then the optimal anonymous mechanism should be close to the unconstrained optimal one. In order to formalize this observation, we consider  $k$ -ambiguous distributions where each bidder can be confused with at most  $k$  bidders with “higher ranked distributions” and show that anonymous mechanisms can guarantee a  $\Theta(k)$  approximation to the optimal revenue.

Moreover, we introduce the *decreasing price mechanism*, a simple mechanism that naturally generalizes single price mechanisms and matches the asymptotic guarantees of the best anonymous auction. Intuitively, the mechanism is succinctly defined by a set of  $n$  prices  $p_1 \geq \dots \geq p_n$ , where  $p_i$  is the price that

**Table 1.** Worst-case approximation bounds for the case of digital goods. The table compares the revenue of optimal non-anonymous mechanisms to the maximum revenue achievable by pricing mechanisms, the decreasing-price mechanism and optimal anonymous auctions.

| Distributions  | Single-price mechanism           | Decreasing-price mechanism/<br>Optimal anonymous auction |
|----------------|----------------------------------|--|
| IDENTICAL      | 1 [Myerson’s auction]            | 1 [Myerson’s auction]                                    |
| DETERMINISTIC  | $\Theta(\log n)$ [Theorem 4.2]   | 1 [Corollary 3.1]  |
| $k$ -AMBIGUOUS | $\Theta(k + \log n)$ [Lemma 5.1] | $\Theta(k)$ [Theorem 5.2]                                |
| MHR / REGULAR  | $\Theta(\log n)$ [Theorem 4.2]   | $\Theta(\log n)$ [Theorem 4.4]                           |
| GENERAL        | $\Theta(n)$ [Theorem 4.1]        | $\Theta(n)$ [Theorem 4.3]                                |

the  $i$ -th-highest bidder should pay. The decreasing price mechanism implements this idea with the minimal modifications required to maintain incentive compatibility. Notably, this auction has linear description complexity, whereas the description complexity of the true optimal anonymous mechanism may be exponential or even unbounded for continuous distributions since it might offer a wide range of different prices to a bidder depending on what others bid.

Finally, we return to sponsored search. As motivated above, a sponsored search platform may wish to charge a premium for certain slots based on the demand profile of a market. Without the ability to discriminate among bidders, the platform may be constrained to run an anonymous auction.<sup>2</sup> A slight modification to our decreasing price mechanism offers a way to do this and achieve the same  $\Theta(k)$  guarantee for  $k$ -ambiguous distributions in both sponsored search and  $m$ -unit auctions.

We provide a detailed comparison of pricing mechanisms and anonymous mechanisms with the optimal non-anonymous mechanisms in Table 1.

*Related Work.* Deb and Pai [7] also study the problem of designing a revenue maximizing mechanism under the anonymity constraint. They devise a set of allocation and payment functions such that in equilibrium bidders pay the Myerson virtual values of their corresponding distributions and the seller achieves revenue that matches the optimal revenue in the unrestricted case. Their results are only for a single item and their mechanisms are Bayesian IC and Bayesian IR. In contrast, attempting to get more robust and practical results, we require our mechanisms to be ex-post IC and ex-post IR.

Ashlagi [3] characterizes anonymous truth-revealing position auctions. He shows that under two different notions of anonymity, namely anonymity of the

<sup>2</sup> Many factors, such as click-through-rates (CTRs) and relevance scores, will break symmetry in a sponsored search auction. As discussed in Ashlagi [3], these can be handled in a variety of ways, e.g. by requiring symmetry among bidders with the same CTR or score. We follow Ashlagi and consider a simple model without such parameters to avoid these complexities.

allocation rule and utility symmetry, every truth-revealing position auction is a VCG position auction. His work applies to deterministic auctions and doesn't consider optimizing revenue.

A variety of problems in the optimal auction literature employ similar ideas to reach different ends. Hartline and Roughgarden [12] study simple mechanisms that maximize seller revenue for selling a single item. They show that when bidder distributions are regular, a second price auction with a single reserve — a simple anonymous mechanism — offers a constant fraction of the revenue that is achievable by Myerson's optimal auction [13]. This constant was recently improved in [1]. For multiple items, work on envy free auctions [5, 9–11] provides tight approximation guarantees for uniform pricing mechanisms compared to the Myerson optimal auction.

Also related to our work is literature on prior-independent mechanisms (e.g. [8, 16]) that use the other bids to infer the valuation of an agent. However, they assume that values are drawn I.I.D. from an unknown prior distribution. In contrast, anonymity will only be a significant constraint when values are non-I.I.D. and the optimal auction must discriminate among them. Optimal auctions for correlated bidders also use the bids of other agents  $\mathbf{v}_{-i}$  to infer a posterior over the bid of the  $i$ -th agent  $v_i$  (see e.g. [6, 15]). We will see that the optimal anonymous auction is closely related to the optimal general auction for a particular correlated prior.

## 2 Model and Preliminaries

A seller has  $m$  identical items to sell to  $n$  bidders. We will refer to the case where  $m = n$  as a *digital goods* setting. Each bidder  $i$  has a private valuation  $v_i$  for getting one item. The profile of agent valuations is denoted by  $\mathbf{v} = (v_1, \dots, v_n)$ . The valuations of the agents are drawn from a product distribution  $\mathbf{F} = F_1 \times \dots \times F_n$ .

A mechanism  $\mathcal{M} = (\mathcal{A}, \mathcal{P})$  consists of an allocation function  $\mathcal{A}$  and a pricing function  $\mathcal{P}$ . If agents choose strategies  $\mathbf{s} = (s_1, \dots, s_n)$ ,  $\mathcal{A}_i(\mathbf{s})$  indicates the probability that bidder  $i$  gets the item and  $\mathcal{P}_i(\mathbf{s})$  indicates the price that he pays.

A tuple of strategy functions  $s_1(\cdot), \dots, s_n(\cdot)$  is an ex-post Nash-Equilibrium if

$$\mathcal{A}_i(s_i(v_i), \mathbf{s}_{-i}(\mathbf{v}_{-i}))v_i - \mathcal{P}_i(s_i(v_i), \mathbf{s}_{-i}(\mathbf{v}_{-i})) \geq \mathcal{A}_i(s'_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))v_i - \mathcal{P}_i(s'_i, \mathbf{s}_{-i}(\mathbf{v}_{-i}))$$

for all agents  $i$ , agent valuations  $\mathbf{v}$  and alternative strategies  $s'_i$  for agent  $i$ .

We are interested in *anonymous* mechanisms  $\mathcal{M} = (\mathcal{A}, \mathcal{P})$ . A mechanism is anonymous if  $\mathcal{A}$  and  $\mathcal{P}$  are symmetric functions in the sense that permuting their arguments also permutes the resulting allocations and prices, i.e. for all permutations  $\pi = (\pi_1, \dots, \pi_n)$  and strategies  $\mathbf{s} = (s_1, \dots, s_n)$  it holds:

$$\mathcal{A}_i(s_1, \dots, s_n) = \mathcal{A}_{\pi_i}(s_{\pi_1}, \dots, s_{\pi_n}) \quad \text{and} \quad \mathcal{P}_i(s_1, \dots, s_n) = \mathcal{P}_{\pi_i}(s_{\pi_1}, \dots, s_{\pi_n}) \quad \text{for all agents } i$$

Anonymous mechanisms might have multiple ex-post Nash-Equilibria. However, we restrict our attention to *symmetric equilibria*, i.e. for every value  $v$ ,  $s_1(v) = s_2(v) = \dots = s_n(v) = s(v)$ . This assumption is natural as it only requires that agents know their own value and not necessarily the identity assigned to them by the mechanism designer. It also removes trivial solutions where a mechanism asks agents to report their identities together with their bids and charges them according to their reported identities<sup>3</sup>.

By the Revelation Principle for any symmetric ex-post Nash-Equilibrium  $s(\cdot)$  in an anonymous mechanism  $\hat{\mathcal{M}} = (\hat{\mathcal{A}}, \hat{\mathcal{P}})$  there exists an equivalent anonymous mechanism  $\mathcal{M} = (\mathcal{A}, \mathcal{P})$  that is direct, i.e. the strategy of every agent is to truthfully report his value  $s_i(v_i) = v_i$ . Mechanism  $\mathcal{M}$  has allocation function  $\mathcal{A}(v_1, \dots, v_n) = \hat{\mathcal{A}}(s(v_1), \dots, s(v_n))$  and price function  $\mathcal{P}(v_1, \dots, v_n) = \hat{\mathcal{P}}(s(v_1), \dots, s(v_n))$ . We can thus consider direct mechanisms. The ex-post Nash Equilibrium constraint translates to the following ex-post incentive compatibility constraint:

$$\mathcal{A}_i(\mathbf{v}) \cdot v_i - \mathcal{P}_i(\mathbf{v}) \geq \mathcal{A}_i(\mathbf{v}_{-i}, v'_i) \cdot v_i - \mathcal{P}_i(\mathbf{v}_{-i}, v'_i) \quad \text{for all } \mathbf{v}, v'_i, i \quad (\text{IC})$$

Moreover, the mechanism must satisfy ex-post individual rationality so that players participate voluntarily:

$$\mathcal{A}_i(\mathbf{v}) \cdot v_i - \mathcal{P}_i(\mathbf{v}) \geq 0 \quad \text{for all } \mathbf{v}, i \quad (\text{IR})$$

An additional property that is desirable from mechanisms is the following monotonicity guarantee.

**Definition 2.1 (Monotone Mechanisms).** *A mechanism is monotone if for all profiles  $\mathbf{v}$ , we have that  $\mathcal{A}_i(\mathbf{v}) \geq \mathcal{A}_j(\mathbf{v}) \Leftrightarrow v_i \geq v_j$  for all  $i, j$ .*

### 3 Optimal Anonymous Auctions

First, we study optimal anonymous auctions and show that they have a natural structure — informally, the mechanism uses the values of others  $\mathbf{v}_{-i}$  to infer a posterior belief  $h$  about bidder  $i$ 's value, then maximizes revenue in the standard way subject to the posterior beliefs  $h$  (maximizing virtual value and charging the associated single-parameter payments [2, 13]). In the special case of a digital goods auction, each bidder is offered the item at the optimal posted price for her inferred distribution  $h$ .

First, since anonymous mechanisms generate the same outcome when bidders are permuted, we observe the following:

**Observation 1.** *The optimal anonymous mechanism remains optimal if we randomly rename bidders before running the auction.*

<sup>3</sup> Truthful reporting by the agents can be made an ex-post Nash Equilibrium by canceling the auction completely if an agent misreported his identity and the mechanism receives an identity twice.



Thus, if the prior knowledge about the bidders was the correlated distribution

$$g(\mathbf{x}) = \frac{1}{n!} \sum_{\pi \in \Pi(n)} \prod_{i \in N} f_i(x_{\pi_i})$$

the optimal anonymous auction would still achieve exactly the same revenue. Moreover, in the setting where the prior knowledge is the distribution  $g$ , the anonymity constraint doesn't hurt the achievable revenue since the prior beliefs about the bidders are symmetric and we can assume without loss of generality that the optimal mechanism treats all bidders equally ex-ante.

**Observation 2.** *Suppose that prior beliefs  $\mathbf{F}$  are symmetric (possibly correlated). Then there exists a symmetric mechanism that maximizes revenue.*

These observations immediately lead to the following claim that allows us to reduce the problem of finding the optimal symmetric auction to optimizing with respect to a symmetric distribution of bidders:

*Claim.* Any mechanism that is optimal among ex-post IC and ex-post IR mechanisms for the symmetric distribution

$$g(\mathbf{x}) = \frac{1}{n!} \sum_{\pi \in \Pi(n)} \prod_{i \in N} f_i(x_{\pi_i})$$

can be transformed into a mechanism that is optimal among symmetric, ex-post IC, and ex-post IR auctions for the beliefs  $\mathbf{F}$  by relabeling bidders according to a uniformly random permutation.

Building on this claim, our characterization theorem for digital goods follows by characterizing the optimal auction for  $g$ :

**Theorem 3.1.** *The optimal anonymous digital goods auction offers bidder  $i$  a copy of the item at the revenue-maximizing price given  $h(v_i | \mathbf{v}_{-i})$ , the posterior belief about  $v_i$  given  $\mathbf{v}_{-i}$ .*

For mechanisms beyond digital goods, the  $k$ -lookahead mechanism of Ronnen [14] achieves at least half of the revenue of the optimal auction for correlated distributions. We can also apply a theorem of Roughgarden and Talgam-Cohen [15] to characterize the optimal auction for  $g$  for special cases where the inferred posterior  $h$  is regular — the resulting optimal mechanism will infer  $h$  and maximize virtual value with respect to  $h$ . We defer additional details to the full version of the paper.

*Proof.* Following Claim 3, it is equivalent to study the optimal auction for the correlated distribution  $g$ . We know from Myerson and others [2, 13], that a normalized mechanism  $\mathcal{M}$  will be ex-post IC if and only if  $\mathcal{A}_i$  is monotone in  $v_i$  and payments are given by  $\mathcal{P}(v) = v\mathcal{A}(v) - \int_0^v \mathcal{A}(z)dz$ .

We can thus write the expected revenue  $R_i$  from bidder  $i$  as

$$\begin{aligned} R_i &= \int_{\mathbb{R}_+^n} \mathcal{P}_i(\mathbf{v})g(\mathbf{v})d\mathbf{v} \\ &= \int_{\mathbb{R}_+^n} \left( v_i \mathcal{A}_i(\mathbf{v}) - \int_0^{v_i} \mathcal{A}_i(\mathbf{v}_{-i}, z) dz \right) g(\mathbf{v})d\mathbf{v} \\ &= \int_{\mathbb{R}_+^{n-1}} \int_{\mathbb{R}_+} g(\mathbf{v}_{-i}, v_i) \mathcal{A}_i(\mathbf{v}_{-i}, v_i) \left( v_i - \frac{\int_{v_i}^\infty g(\mathbf{v}_{-i}, z) dv_i}{g(\mathbf{v}_{-i}, v_i)} \right) dv_i d\mathbf{v}_{-i} \end{aligned}$$

If we let  $h(v_i|\mathbf{v}_{-i})$  denote the density of  $v_i$  that we can infer given  $\mathbf{v}_{-i}$ ,  $H$  the associated CDF, and  $\phi^{h|\mathbf{v}_{-i}}(v_i)$  its Myersonian virtual value, we have

$$h(v_i|\mathbf{v}_{-i}) = \frac{g(\mathbf{v}_{-i}, v_i)}{\int_0^\infty g(\mathbf{v}_{-i}, z) dz} \quad \text{and} \quad \phi^{h|\mathbf{v}_{-i}}(v_i) = v_i - \frac{1 - H(v_i|\mathbf{v}_{-i})}{h(v_i|\mathbf{v}_{-i})}$$

and can rearrange to get

$$R_i = \int_{\mathbb{R}_+^{n-1}} \left( \int_0^\infty g(\mathbf{v}_{-i}, z) dz \right) \int_{\mathbb{R}_+} h(v_i|\mathbf{v}_{-i}) \mathcal{A}_i(\mathbf{v}_{-i}, v_i) \phi^{h|\mathbf{v}_{-i}}(v_i) dv_i d\mathbf{v}_{-i} .$$

It remains to choose  $\mathcal{A}$ , which can be done in an arbitrary (monotone) way for digital goods. The inner integral  $\int h \mathcal{A}_i \phi dv_i$  is precisely the revenue when bidder  $i$  has value distributed according to  $h(v_i|\mathbf{v}_{-i})$ , so Myerson [13] tells us that the optimal allocation  $\mathcal{A}_i(\mathbf{v}_{-i}, v_i)$  is a posted price to bidder  $i$  that maximizes revenue given the distribution  $h$ .  $\square$

A few noteworthy extreme cases arise when the auctioneer can identify bidder  $i$  given only the bids  $\mathbf{v}_{-i}$ :

**Corollary 3.1.** *If the distributions  $f_i$  are point distributions (bidders’ values are known precisely to the auctioneer), have non-overlapping support, or are the same for all bidders, then the optimal anonymous mechanism coincides with Myerson’s optimal mechanism.*

In all three cases, the posterior distribution inferred from  $\mathbf{v}_{-i}$  is precisely  $f_i$ , therefore the auction precisely identifies each bidder and runs the optimal auction. Note that if bidders’ distributions have overlapping supports but are not identical, then the distribution inferred from the posterior will likely be different from  $f_i$  and the resulting algorithm will not correspond with Myerson’s optimal one.

These results suggest that anonymous mechanisms perform best when we can differentiate among the bidders; indeed, we will see that this is necessary. In Sect. 4, we show that the anonymity constraint substantially limits revenue even when distributions are discrete over  $n$  points and that assumptions like regularity of  $f_i$  are insufficient. In Sect. 5, we show that the performance degrades continuously with the auctioneer’s ability to differentiate among the bidders.

## 4 Worst-Case Approximations

We compare the revenue guarantees of single price and anonymous mechanisms and find that anonymous mechanisms can do no better in the worst case.

### 4.1 Single Price Mechanisms

We first look at how well single price mechanisms for  $m$ -unit auctions perform compared to the optimal. A single price mechanism allocates items to the  $m$  highest bidders with values exceeding  $p$  and charges them the maximum of  $p$  and the  $m+1$  highest bid.<sup>4</sup> It is easy to see that single price mechanisms can get at least a  $\frac{1}{n}$  fraction of the revenue by choosing as price the Myerson reserve price of a bidder's distribution chosen uniformly at random. However, such a linear approximation guarantee is unavoidable as we can also show a lower bound of  $\Omega(n)$  for the approximation. The proof is a simple extension of Proposition 5.1 in [1] and we give it here for completeness.

**Theorem 4.1 (Single Price for General Distributions [1]).** *For general distributions, a single price mechanism achieves a  $\Theta(n)$  approximation to the optimal revenue.*

*Proof.* We already saw that single price mechanisms achieve  $\frac{1}{n}$  fraction of the optimal revenue. Now, consider a case where each bidder  $i$  has a value of  $\frac{1}{\epsilon^i}$  with probability  $\epsilon^i$  and 0 otherwise for some constant  $\epsilon < \frac{1}{2}$ . We show that even for a single item the optimal auction achieves revenue  $\Omega(n)$  while any single price auction gets at most  $O(1)$  even if there is an unlimited supply of goods (digital goods).

First, consider a non-anonymous sequential posted price mechanism for a single item that asks every bidder  $i$  (from largest to smallest index) to pay  $\frac{1}{\epsilon^i}$  until the item is sold. The expected revenue it gets from bidder  $i$  is equal to  $\frac{1}{\epsilon^i} \epsilon^i \cdot \text{Prob}[\text{bidder } i \text{ reached}]$ . But the probability that bidder  $i$  is reached is at least  $1 - \sum_{j=i+1}^n \epsilon^j \geq 1 - \frac{\epsilon}{1-\epsilon} = \Omega(1)$ . Thus the total revenue achievable by this mechanism and hence by the optimal mechanism is  $\Omega(n)$ .

Finally, it is easy to see that even in a digital goods setting, a single price mechanism gives revenue at most  $\max_i \frac{1}{\epsilon^i} \sum_{j \geq i} \epsilon^j \leq \frac{1}{1-\epsilon} = O(1)$ .  $\square$

However, when all distributions are regular, we show that single price mechanisms perform much better.

**Theorem 4.2 (Single Price for Regular Distributions).** *For regular distributions, a single price mechanism achieves a  $\Theta(\log m)$  approximation to the optimal revenue.*

*Proof.* To prove the theorem we will apply Theorem 4.1 from [4] which states that running VCG with the median of each agent's distribution as a reserve price

<sup>4</sup> This is a regular VCG mechanism with a reserve price  $p$ .

(VCG-m) gives a 4-approximation to the optimal revenue. Therefore, it suffices to prove that the revenue of single price mechanisms is a  $\Theta(\log m)$  approximation to the revenue of VCG-m.

Let  $p_i$  be the median prices for each bidder and assume that  $p_1 \geq p_2 \geq \dots \geq p_n$ .

The revenue of VCG-m comes from 2 different sources: reserve prices, where a bidder is charged his reserve price, and competition between bidders, where a bidder is charged the bid of someone else.

More formally, let  $P_i$  be the random variable that gives the price agent  $i$  pays to the mechanism.  $P_i = 0$  whenever bidder  $i$  doesn't get an item. We can write the expected revenue of the mechanism as  $\sum_i E[P_i] = \sum_i E[P_i|P_i = p_i]Pr[P_i = p_i] + E[P_i|P_i > p_i]Pr[P_i > p_i]$ .

If more than half of the revenue comes from competition between bidders, i.e.  $\sum_i E[P_i|P_i > p_i]Pr[P_i > p_i] \geq \frac{1}{2} \sum_i E[P_i]$ , then setting a price 0 for all bidders and running a simple VCG gives a 2-approximation. This is because in VCG-m whenever someone doesn't get charged his reserve price  $p_i$ , he is charged the  $m+1$  largest bid that exceeds the reserve price. Therefore,  $\sum_i E[P_i|P_i > p_i]Pr[P_i > p_i]$  is at most  $m$  times the expectation of the  $m+1$  largest bid overall which is equal to the revenue of VCG with no reserve prices.

Otherwise, more than half of the revenue comes from charging the reserve prices to bidders, i.e.

$$\sum_i E[P_i|P_i = p_i]Pr[P_i = p_i] \geq \frac{1}{2} \sum_i E[P_i]$$

In this case, the revenue is at most equal to  $2 \sum_{i=1}^m p_i$ . Consider a mechanism that charges each price  $p_i$  with probability  $q_i = (iH_m)^{-1}$ . The revenue of this mechanism is  $\sum_{i=1}^m q_i p_i E[\# \text{ bids} \geq p_i]$ . However, we have that  $E[\# \text{ bids} \geq p_i] \geq i/2$  since each of the first  $i$  bidders has at least  $1/2$  of exceeding  $p_i$ . This gives a revenue of  $\sum_{i=1}^m (iH_m)^{-1} p_i (i/2) = \frac{\sum_{i=1}^m p_i}{2H_m}$  which is a  $4H_m$  approximation to  $2 \sum_{i=1}^m p_i$ .  $\square$

This bound is tight even for bidders coming from point distributions. Suppose that each bidder  $i$  has a value of  $1/i$ . The best single price gets revenue of 1 while the optimal mechanism gets revenue  $H_m = \Theta(\log m)$ .

## 4.2 General Anonymous Mechanisms

For general anonymous mechanisms, we show that even the best mechanism cannot get any better asymptotic guarantees than single price mechanisms for general distributions.

**Theorem 4.3 (Optimal Anonymous Mechanism for General Distributions).** *The optimal anonymous mechanism  $\mathcal{M}$  gives a  $\Theta(n)$  approximation to the optimal revenue for general distributions.*

*Proof.* We revisit the construction from the Theorem 4.1 but lower the probability that a bidder gets a high value even further. Each bidder  $i$  now has a value of  $\frac{1}{\epsilon^i}$  with probability  $\delta \epsilon^i$  and 0 otherwise. The optimal non-anonymous mechanism gets a revenue of  $\Omega(n\delta)$ . The optimal anonymous mechanism must charge the same price whenever there is only one bidder with a high bid. Let  $E$  be the event that at least two bidders value the item high. Given  $\neg E$ , the mechanism is identical to a single price mechanism. So the approximation of the optimal anonymous mechanism is upper bounded by  $O\left(\frac{\delta \frac{1}{1-\epsilon} + Pr[E]Rev[E]}{n\delta}\right) = O\left(\frac{1}{n} + \frac{Pr[E]Rev[E]}{n\delta}\right)$ . The theorem follows since  $\frac{Pr[E]Rev[E]}{n\delta}$  goes to 0 as  $\delta \rightarrow 0$  because  $Pr[E] = O(\delta^2)$ .  $\square$

Moreover, we can show that general anonymous mechanisms cannot beat the asymptotic guarantees that single price mechanisms achieve for regular distributions. In fact, we can show that this is true even for uniform distributions.

**Theorem 4.4 (Uniform Distributions Lower Bound).** *For uniform distributions, the best anonymous mechanism gets at least a  $\Theta(\log m)$  approximation to the optimal revenue.*

**Proof of Theorem 4.4.** To construct the lower bound instance, we consider the case where there are  $N = (2^n - 1)L$  agents and  $m = N$  items and exactly  $2^i L$  agents have distributions in  $U[0, 2^{-i}]$  for  $i \in \{0, \dots, n - 1\}$ . We can see that the optimal non-anonymous mechanism gets a revenue of  $\frac{Ln}{4}$  by charging each agent his monopoly price which is at the midpoint of his distribution.

We will now upper bound the revenue that the optimal anonymous mechanism achieves. To do this we consider an instance where a vector of values  $\mathbf{v}$  is reported.

Let  $b_i = \#\{j | v_j > 2^{-i}\}$ , i.e. the number of agents with values greater than  $2^{-i}$ . We will show that if all  $b_i$ 's are large, the optimal anonymous mechanism charges a very low price to each agent.

**Lemma 4.1.** *If  $b_i > (\frac{2}{3}2^i - 1)L + 1$  for all  $i \in \{1, \dots, n - 1\}$ , the optimal anonymous mechanism charges a price lower than  $2^{-(n-1)}$  to every agent.*

*Proof.* Since we are in a digital goods setting we can apply Theorem 3.1 and consider the distribution that the mechanism infers for an agent  $k$ 's value by looking at all bids of the other agents. The probability density of agent's  $k$  value at a point  $x$  given the bids  $\mathbf{v}_{-k}$  of the other agents is  $h(x | \mathbf{v}_{-k}) = \frac{1}{n!} \sum_{\pi \in \Pi(n)} \int \pi_k(x) \prod_{i \neq k} f_{\pi_i}(v_i)$ , which is proportional to the number of ways to match agents to probability distributions for the bid vector  $\mathbf{v}' = (\mathbf{v}_{-k}, x)$ .

We can compute the number of ways exactly in terms of  $b'_i = \#\{j | v'_j > 2^{-i}\}$  as  $\prod_{i=0}^{n-1} ((2^{i+1} - 1)L - b'_i)_{b'_{i+1} - b'_i}$  where the notation  $(a)_b \equiv a(a-1) \dots (a-b+1)$  denotes the falling factorial and  $b'_n$  is defined to be equal to  $N$ . This is because the  $b'_1$  agents that have values in  $[1/2, 1]$  can only be in the distributions  $U[0, 1]$  so there are  $L$  choices for distributions which means there are  $L(L-1) \dots (L-b'_1+1)$

ways to match them. For the  $b'_2 - b'_1$  agents that have values in  $[1/4, 1/2]$ , there are  $3L$  possible distributions ( $L$  that are  $U[0, 1]$  and  $2L$  that are  $U[0, 1/2]$ ) but  $b'_1$  of them are already taken so there are exactly  $(3L - b'_1)_{b'_2 - b'_1}$  choices over all and so on.

We now show that  $4h(x|v_{-k}) < h(y|v_{-k})$  for  $x \in (2^{-t}, 2^{-(t-1)})$ ,  $y \in (2^{-(t+1)}, 2^{-t})$  and  $1 \leq t \leq n - 1$ . That is the probability density at the interval  $(2^{-t}, 2^{-(t-1)})$  is at most a fourth of the probability density at the interval  $(2^{-(t+1)}, 2^{-t})$ . Let  $b'(x)$  and  $b'(y)$  be the corresponding  $b'$  parameters for  $x$  and  $y$  respectively. It is easy to see that  $b'_i(x) = b'_i(y)$  for  $i \neq t$  and that  $b'_t(x) = b'_t(y) + 1$ . We have that:

$$\begin{aligned}
 \frac{h(x|v_{-k})}{h(y|v_{-k})} &= \frac{\prod_{i=0}^{n-1} ((2^{i+1} - 1)L - b'_i(x))_{b'_{i+1}(x) - b'_i(x)}}{\prod_{i=0}^{n-1} ((2^{i+1} - 1)L - b'_i(y))_{b'_{i+1}(y) - b'_i(y)}} \\
 &= \frac{\prod_{i=t-1}^t ((2^{i+1} - 1)L - b'_i(x))_{b'_{i+1}(x) - b'_i(x)}}{\prod_{i=t-1}^t ((2^{i+1} - 1)L - b'_i(y))_{b'_{i+1}(y) - b'_i(y)}} \\
 &= \frac{(2^t - 1)L - b'_t(x)}{(2^{t+1} - 1)L - b'_t(y)} && \text{cancelling all identical terms} \\
 &< \frac{(2^t - 1)L - (\frac{2}{3}2^t - 1)L}{(2^{t+1} - 1)L - (\frac{2}{3}2^t - 1)L} && \text{since } b'_t(y) \geq b_t - 1 > (\frac{2}{3}2^t - 1)L \\
 &= \frac{\frac{1}{3}2^t}{\frac{4}{3}2^t} = \frac{1}{4}
 \end{aligned}$$

We now show that the optimal price for the inferred distribution is less than  $2^{-(n-1)}$ . Assume that this is not the case and the optimal price is  $p > 2^{-(n-1)}$ . We will show that by charging  $p/2$  we get strictly more revenue. We will prove by induction that  $Pr[x > p] < Pr[x > p/2]/2$  for  $p \in [2^{-(n-1)}, 2)$ . This is trivial to see if  $p \in [1, 2)$  since  $Pr[x > p] = 0$  while  $Pr[x > p/2] > 0$ . Assume that  $Pr[x > p] < Pr[x > p/2]/2$  for  $p \in [2^{-i}, 2^{-i+1})$ . Then for  $p \in [2^{-i-1}, 2^{-i})$  we have that:

$$\begin{aligned}
 Pr[x > p] &= Pr[x > 2^{-i}] + Pr[x \in (p, 2^{-i})] \\
 &< \frac{Pr[x > 2^{-i-1}]}{2} + Pr[x \in (p, 2^{-i})] && \text{by the induction hypothesis} \\
 &< \frac{Pr[x > 2^{-i-1}]}{2} + \frac{Pr[x \in (\frac{p}{2}, 2^{-i-1})]}{2} && \text{since } \frac{h(x|v_{-k})}{h(x/2|v_{-k})} < \frac{1}{4} \text{ for } x \in (p, 2^{-i}) \\
 &= Pr[x > p/2]/2
 \end{aligned}$$

We conclude that  $Pr[x > p] < Pr[x > p/2]/2$  which implies that  $pPr[x > p] < pPr[x > p/2]/2$ , i.e. the revenue we get by charging  $p$  is less than charging  $p/2$  if  $p > 2^{-(n-1)}$ .  $\square$

We show that for large  $L$  the conditions of Lemma 4.1 are satisfied with extremely high probability.

**Lemma 4.2.** *Let  $L = 2^{5n}$  and let  $E$  the event that  $b_i > (\frac{2}{3}2^i - 1)L + 1$  for all  $i$ . Then  $Pr[E] > 1 - ne^{-2^{n-2}}$ .*

*Proof.* Consider the expectation of  $b_i$ .

$$\begin{aligned} E[b_i] &= \sum_j Pr[v_j > 2^{-i}] = L2^0(1 - 2^{-i}) + L2^1(1 - 2^{-i+1}) + \dots + L2^{i-1}(1 - 2^{-1}) \\ &= L \left( 2^i - 1 - 2^i \sum_{j=1}^i 2^{-2j} \right) = L \left( 2^i - 1 - 2^i \frac{1 - 2^{-2i}}{3} \right) = L \left( \frac{2}{3}2^i - 1 + \frac{2^{-i}}{3} \right) \end{aligned}$$

We have that  $E[b_i](1 - 2^{-2n}) > 2^{5n} \left( \frac{2}{3}2^i - 1 + \frac{2^{-i}}{3} \right) - 2^{3n} \frac{2}{3}2^i > 2^{5n} \left( \frac{2}{3}2^i - 1 \right) + 1$ . Therefore,

$$\begin{aligned} Pr[b_i < (\frac{2}{3}2^i - 1)L + 1] &< Pr[b_i < E[b_i](1 - 2^{-2n})] \\ &\leq e^{-2^{-4n} E[b_i]/2} && \text{applying a Chernoff bound} \\ &\leq e^{-2^{n-2}} && \text{since } E[b_i] \geq L/2 = 2^{5n-1} \end{aligned}$$

By a union bound for all  $n$  possible values of  $i$  we get that  $Pr[E] > 1 - ne^{-2^{n-2}}$ . □

Therefore, the revenue of the optimal anonymous mechanism is at most  $N2^{-(n-1)} = L(2^n - 1)2^{-(n-1)} \leq 2L$  when event  $E$  happens and at most  $Ln$  otherwise since the total value of all agents is always at most  $Ln$ . Thus, the expected revenue is at most  $L(2 + n^2e^{-2^{n-2}})$ . Since the optimal non-anonymous mechanism achieves revenue  $Ln/4$ , the approximation ratio is  $n/8 + o(1)$ . Since the number of agents is at most  $N \leq 2^{6n}$ , we have that  $n \geq \log N/6$ . Thus the approximation ratio in terms of  $N$  is  $\frac{\log N}{48} + o(1) = \Theta(\log N) = \Theta(\log m)$  since  $m = N$  in the digital goods setting.

## 5 Anonymous Auctions with Limited Ambiguity

In the previous section, we showed that the best anonymous auction cannot offer better worst-case revenue guarantees than single price mechanisms, even when distributions are regular or have a monotone hazard rate. In this section, we explore a key property called limited ambiguity that separates anonymous mechanisms from single price mechanisms and demonstrates their power.

One natural model of limited ambiguity is that the auctioneer can roughly order the bidders in terms of their values. Two bidders who are close together in this ordering may be ordered incorrectly (their distributions overlap), but two bidders who are far apart in the ordering will have distributions that do not overlap. The following definition is a particular formalization of that idea:

**Definition 5.1.** Let  $[a_i, b_i]$  be the support of the distribution of agent  $i$  and assume without loss of generality that  $a_1 \geq a_2 \geq \dots \geq a_n$ . We say that the set of distributions is  $k$ -ambiguous if  $b_i < a_{i-1-k}$  for all  $i$ , i.e. a sample from the  $i$ -th distribution can be confused with at most  $k$  distributions ahead of it.

The extreme case where  $k = 0$  — i.e. bidders' values are drawn from distributions with disjoint supports — gives our first separation between general anonymous auctions and single price mechanisms. Consider the single point distribution  $1/i$  for each agent  $i$  — it is easy to see that the approximation ratio of any single price is  $\log m$ . In contrast, we showed that the optimal anonymous auction achieves the same revenue as the optimal non-anonymous auction in Sect. 3. For general  $k$ , Lemma 5.1 gives the following bound for pricing mechanisms. The proof of the lemma is given in the full version of the paper.

**Lemma 5.1.** *Single price mechanisms can get  $O(k + \log m)$  approximation to the optimal revenue.*

In this section, we will show that anonymous mechanisms can guarantee an approximation ratio of  $O(k)$  for  $k$ -ambiguous distributions, and that this is tight. We focus first on the case of digital goods, where  $m = n$ , and then extend to  $m < n$  as well as to sponsored search auctions.

To show that anonymous mechanisms can achieve an  $O(k)$  approximation to the optimal revenue, we construct a simple mechanism called the Decreasing Price Mechanism (DPM) that is efficiently defined by  $n$  prices. We will begin with a slight variation that is not ex-post incentive compatible to motivate the choice of mechanism.

**Definition 5.2 (Non-IC Decreasing Price Mechanism).** *The Non-IC Decreasing Price Mechanism is defined by a set of prices  $p_1 \geq p_2 \geq \dots \geq p_n$  and works as follows: incoming bids are sorted in decreasing order, then bidder  $i$  is offered an item at price  $p_i$ .*

This mechanism is both simple and anonymous, but unfortunately it is not ex-post IC, since a bidder can lower the price she pays simply by ranking lower in the ordering of bids (indeed, she can always get an item at price  $p_n$  simply by placing the lowest bid). We add two key ingredients to define our ex-post IC decreasing price mechanism.

The first ingredient we add limits a bidder's ability to win the item at a lower price: the auction only sells an item at price  $p_i$  if it has successfully sold items at all higher prices. Consequently, for example, bidder  $i + 1$  must be willing to pay  $p_i$  in order for bidder  $i$  to have a chance to win an item at a lower price. When the auction fails to sell an item at price  $p_i$  and therefore stops selling more items, we call this a “**drop**” event.

The second ingredient we add restores incentive compatibility: if a bidder could have won an item at a lower price by ranking lower in the bid order, then we automatically charge her the lower price instead. Observe that given our first modification, bidder  $i$  can win an item at a lower price  $p_i$  if and only if  $b_j \geq p_{j-1}$



for all  $j \in \{i + 1, \dots, l\}$ . We call this a “**chain**” effect since there is a chain of bidders with  $b_j \geq p_{j-1}$ .

These two additional ingredients are the intuition for our decreasing price mechanism:

**Definition 5.3 (Decreasing Price Mechanism).** *The Decreasing Price Mechanism (DPM) is defined by a set of prices  $p_1 \geq p_2 \geq \dots \geq p_n$  and works as follows:*

- Sort bids in decreasing order.
- Starting with  $i = 1$ , allocate items as long as  $b_i \geq p_i$ , then stop allocating items.
- Each winner  $i$  is charged  $p_{j(i)}$ , where  $j(i)$  is the smallest  $j \geq i$  such that exactly  $j$  bidders are bidding above  $p_j$ .

We note that single price mechanisms are a special case of DPM where all the prices  $p_1 = \dots = p_n = p$ . The following lemma shows several interesting properties of DPM.

**Lemma 5.2.** *The Decreasing Price Mechanism is anonymous, ex-post IR, ex-post IC, and monotone in the sense that if  $b_i > b_j$ , then  $\mathcal{A}_i(\mathbf{b}) \geq \mathcal{A}_j(\mathbf{b})$ .*

*Proof.* It is clear that the mechanism is anonymous because it ignores any initial labeling and relabels bidders in decreasing order of their bids. The auction is individually rational because a bidder only wins if  $b_i \geq p_i$  and pays a price  $p_{j(i)} \leq p_i$ . The claimed monotonicity property is also easy to see as the mechanism considers bids in decreasing order and allocates items only until it reaches the first bidder with  $b_i < p_i$ .

To see that the mechanism is ex-post IC, we look at an agent  $i$  and show that  $i$  cannot win an item at a lower price. Note that if  $i$  changes her bid to  $b'_i < p_{j(i)}$ , then there will be  $j(i) - 1$  bids  $\geq p_{j(i)}$  (there were exactly  $j(i)$  such bids before  $i$  changed her bid) and the auction must stop by the time it reaches bidder  $j(i)$ . Thus, the auction will not sell an item for less than  $p_{j(i)}$ , so  $i$  will not get an item. On the other hand, keeping other bids fixed, if  $i$  bids  $b_i \geq p_{j(i)}$ , there will be exactly  $j(i)$  bidders bidding  $\geq p_{j(i)}$ , so  $i$  cannot win at a price less than  $p_{j(i)}$ . □

We will now show that the decreasing price mechanism achieves an approximation ratio of  $O(k)$  for  $k$ -ambiguous distributions. To illustrate the significant ideas in the proof we will first show the statement for  $k = 1$  before proving the general case.

### 5.1 The Case of $k = 1$

For 1-ambiguous distributions, we prove the following theorem:

**Theorem 5.1.** *The optimal Decreasing Price Mechanism approximates the revenue of the optimal auction within a factor of 5 for 1-ambiguous distributions.*

*Proof.* The proof has two parts. First, we use a distribution over DPM pricing schemes to approximate the revenue contribution of agents 3 to  $n$ . This distribution will have expected revenue that is a 3-approximation to the welfare of those agents and therefore also to the revenue they contribute in the optimal auction. Second, we use our single price results to cover the revenue from the first two agents.

First, to cover the revenue contributions of agents 3 to  $n$ , DPM prices are chosen as follows (the parameters  $r_i$  will be chosen later):

$$p_i = \begin{cases} a_{i-1} & \text{with probability } r_i \\ a_i & \text{otherwise.} \end{cases}$$

Intuitively, choosing  $p_i = a_i$  is safe because  $v_i \geq a_i$ , whereas  $p_i = a_{i-1}$  extracts more revenue at the risk of triggering a drop event that prevents selling items to bidders  $> i$ . We take  $r_1 = 0$  so  $p_1 = a_1$ .

Let  $q_i$  be the probability that  $v_i \geq a_{i-1}$  and define  $q_1 = 0$ . We define  $c_i$ , the conditional likelihood of a chain effect, and  $d_i$ , the conditional likelihood of a drop event, as follows:

$$c_i \equiv \Pr[v_i \geq a_{i-1} \text{ and } p_i = a_i] = (1 - r_i)q_i$$

$$d_i \equiv \Pr[v_i < p_i] = r_i(1 - q_i)$$

By definition of the auction, agent  $i$  pays at least  $a_t$  for some  $t \geq i$  if and only if (a) all bidders  $j \leq i$  have  $v_j \geq p_j$  so that bidder  $i$  wins an item, and (b) there exists a  $j \in \{i + 1, \dots, t + 1\}$  such that exactly  $j$  bidders have bids  $b_j \geq p_j$ . Condition (a) is equivalent to saying that a drop event does not occur among the first  $i$  bidders and happens with probability  $\prod_{j=1}^i (1 - d_j)$ . Condition (b), assuming truthfulness and using 1-ambiguity, happens if and only if there is some  $j \in \{i + 1, \dots, t + 1\}$  such that either  $v_j < a_{j-1}$  or  $p_j = a_{j-1}$ , which happens precisely when  $j$  does not trigger a chain effect, so the likelihood that such a  $j$  exists is  $1 - \prod_{j=i+1}^{t+1} c_j$ .

Let  $x_t$  denote the expected number of bidders who pay  $a_t$  and  $y_t = \sum_{i=1}^t x_i$  the expected number who pay at least  $a_t$ . We can now write  $y_t$  as

$$y_t \geq \sum_{i=1}^t \Pr[\text{Agent } i \text{ pays at least } a_t] \geq \sum_{i=1}^t \left[ \left( 1 - \prod_{j=i+1}^{t+1} c_j \right) \prod_{j=1}^i (1 - d_j) \right].$$

To bound this sum, we relate the  $c_i$ 's and  $d_i$ 's with the following lemma:

**Lemma 5.3.** *We can choose  $r_i$  such that  $d_i \leq \rho$  and  $c_i \leq (1 - \sqrt{\rho})^2$  for any  $\rho \in [0, 1]$ .*

*Proof.* For any such  $\rho$  choose  $r_i = \min(\frac{\rho}{(1-q_i)}, 1)$ . We have that  $d_i = (1 - q_i)r_i \leq \rho$ . We also have that  $c_i = (1 - r_i)q_i$ . If  $r_i = 1$  then  $c_i = 0 \leq (1 - \sqrt{\rho})^2$ . Otherwise  $r_i = \frac{\rho}{(1-q_i)}$  and  $c_i = (1 - \frac{\rho}{(1-q_i)})q_i$  which achieves a maximum value at  $(1 - \sqrt{\rho})^2$  for  $q_i = 1 - \sqrt{\rho}$ .  $\square$

Applying this lemma with  $\rho = 1/i^2$  gives  $r_i$ 's such that  $d_i \leq 1/i^2$  and  $c_i \leq (1 - 1/i)^2$  for  $i \geq 2$ . This makes  $\prod_{j=1}^i (1 - d_j) \geq \prod_{j=2}^i (1 - 1/j^2) = (1 + 1/i)/2 \geq 1/2$ . Moreover,  $\prod_{j=i+1}^{t+1} c_j \leq \prod_{j=i+1}^{t+1} (1 - 1/j)^2 = (i/(t + 1))^2$ . Therefore,

$$y_t \geq \frac{1}{2} \sum_{i=1}^t \left[ 1 - \left( \frac{i}{t+1} \right)^2 \right] = \frac{1}{2} \left( t - \frac{t(t+1)(2t+1)}{6(t+1)^2} \right) \geq \frac{1}{2}(t - t/3) \geq t/3$$

The total expected revenue of the mechanism is  $\sum_{i=1}^n x_i a_i$ . Since  $y_t = \sum_{i=1}^t x_i \geq t/3$  for all  $t$ , it must be that  $\sum_{i=1}^n x_i a_i \geq \sum_{i=1}^n a_i/3$ . Moreover, since  $a_t > b_{t+2} \geq Rev[Agent_{t+2}]$ , it follows that  $\sum_{t=1}^n a_t/3 \geq \sum_{t=3}^n Rev[Agent_t]/3$ , i.e. the revenue is at least 1/3 of the optimal revenue generated by agents 3 to  $n$ .

It remains to handle the revenue contributed by the first two agents. To do so, we use the single price lemma that says that a single price  $p$  is a 2-factor approximation for 2 distributions. If we choose prices  $p_1 = \dots = p_n = p$  with probability 2/5 or the pricing scheme that is defined above with probability 3/5, we get an expected revenue of at least:

$$\frac{2}{5} \left( \frac{Rev[Agent_1] + Rev[Agent_2]}{2} \right) + \frac{3}{5} \left( \frac{\sum_{t=3}^n Rev[Agent_t]}{3} \right) = \frac{\sum_{t=1}^n Rev[Agent_t]}{5}$$

Since we are randomizing over DPM pricing schemes, there exists a single pricing scheme that achieves the necessary approximation. This completes the proof and shows a 5 approximation. □

### 5.2 The General Case

For general  $k$ -ambiguous distributions, the following theorem shows an  $O(k)$  approximation.

**Theorem 5.2.** *The Decreasing Price Mechanism achieves an approximation ratio of  $(3e^2 + 2)k$  for  $k$ -ambiguous distributions.*

The proof of this theorem is given in the full version of the paper and mimics the 1-ambiguous case. We split agents into blocks of size  $k$  such that an agent in block  $t$  cannot be confused with any agents in blocks  $< t - 1$ . Then, a technical lemma analogous to Lemma 5.3 bounds the drop and chain rates between blocks to achieve an  $O(k)$  approximation to the revenue from blocks 3 to  $n/k$ . Finally, a single price mechanism covers the revenue from the top two blocks.

### 5.3 Extension to $m$ -goods and Position Auctions

We extend the results of the previous section from digital goods, where we have an unlimited supply of identical goods, to the  $m$ -unit setting where we have  $m$  copies of a good and to position auctions.

**Definition 5.4 (Position Auction).** *In a position auction, there are  $m$  items for sale, each with a scale factor  $s_j \in [0, 1]$ . We assume that  $s_1 \geq s_2 \geq \dots \geq s_m$ . The utility of an agent  $i$  with value  $v_i$  that receives an item  $j$  and pays  $p$  is equal to  $s_j v_i - p$ .*

In sponsored search auctions, the items are slots on a page of search results and the scale factors correspond to the click through rate of each slot. We prove the following theorem in the full version of the paper.

**Theorem 5.3.** *In any  $m$ -good or position auction setting, there exists an anonymous mechanism that achieves an approximation of  $O(k)$  for  $k$ -ambiguous distributions.*

## 6 Conclusion

Anonymity imposes real constraints on an auction and, as we have seen, on the revenue it can achieve. In the worst case, we have shown that anonymous mechanisms are quite limited, and that the best anonymous mechanism cannot substantially beat a simple single price. The real advantage of an anonymous mechanism is directly related to the auctioneer’s ability to infer information about  $f_i$  and  $v_i$  from the bids of other advertisers,  $\mathbf{v}_{-i}$ , in essence circumventing the ex-ante anonymity requirement.

Our work leaves a few immediate open questions about anonymous auctions with limited ambiguity. We showed that anonymous auctions can achieve a  $\Theta(k)$  approximation for general  $k$ -ambiguous distributions. For single price mechanisms, we saw that the worst-case approximation improves from  $\Theta(n)$  to  $\Theta(\log n)$  when distributions are regular — *can we show an analogous  $\Theta(\log k)$  bound in the  $k$ -ambiguous setting when distributions are regular?* Another interesting research direction is to identify alternative metrics for measuring ambiguity. For example, *what can we say about the revenue from an anonymous auction when the differential entropy between  $f_i$  and the inferred posterior  $h$  is small?*

More broadly, our work suggests many general questions about anonymous mechanisms. *Can anonymous auctions achieve good approximations beyond the settings we have studied?* Interesting dependencies arise outside the digital goods setting because one bidder’s bid can affect the auctioneer’s inference about another bidder, affecting the outcome of the auction in a complicated way. Another question is one of computational complexity — *how difficult is it to compute the optimal anonymous auction?*

## References

1. Alaei, S., Hartline, J.D., Niazadeh, R., Pountourakis, E., Yuan, Y.: Optimal auctions vs. anonymous pricing. In: Guruswami, V. (ed.) IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17–20 October 2015, pp. 1446–1463. IEEE Computer Society (2015)

2. Archer, A., Tardos, É.: Truthful mechanisms for one-parameter agents. In: FOCS 2001: Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, p. 482. IEEE Computer Society, Washington, DC (2001)
3. Ashlagi, I.: A characterization of anonymous truth-revealing position auctions. Harvard University, Discussion Paper, Harvard Business School (2008)
4. Azar, P.D., Daskalakis, C., Micali, S., Weinberg, S.M.: Optimal and efficient parametric auctions. In: SODA, pp. 596–604 (2013)
5. Balcan, M.-F., Blum, A., Mansour, Y.: Item pricing for revenue maximization. In: Proceedings of the 9th ACM conference on Electronic commerce, pp. 50–59. ACM (2008)
6. Cremer, J., McLean, R.P.: Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* **53**(2), 345–361 (1985)
7. Deb, R., Pai, M.: Discrimination via symmetric auctions. *Am. Econ. J. Microecon.* (2015)
8. Dhangwatnotai, P., Roughgarden, T., Yan, Q.: Revenue maximization with a single sample. In: Proceedings of the 11th ACM Conference on Electronic Commerce, EC 2010, pp. 129–138. ACM, New York (2010)
9. Feldman, M., Fiat, A., Leonardi, S., Sankowski, P.: Revenue maximizing envy-free multi-unit auctions with budgets. In: Proceedings of the 13th ACM Conference on Electronic Commerce, pp. 532–549. ACM (2012)
10. Goldberg, A.V., Hartline, J.D.: Envy-free auctions for digital goods. In: Proceedings of the 4th ACM Conference on Electronic Commerce, pp. 29–35. ACM (2003)
11. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1164–1173. Society for Industrial and Applied Mathematics (2005)
12. Hartline, J.D., Roughgarden, T.: Simple versus optimal mechanisms. *SIGecom Exch.* **8**(1), 5:1–5:3 (2009)
13. Myerson, R.B.: Optimal auction design. *Math. Oper. Res.* **6**(1), 58–73 (1981)
14. Ronen, A.: On approximating optimal auctions. In: Proceedings of the 3rd ACM Conference on Electronic Commerce, pp. 11–17. ACM (2001)
15. Roughgarden, T., Talgam-Cohen, I.: Optimal and near-optimal mechanism design with interdependent values. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC 2013, pp. 767–784. ACM, New York (2013)
16. Segal, I.: Optimal pricing mechanisms with unknown demand. *Am. Econ. Rev.* **93**(3), 509–529 (2003)

# Revenue Maximizing Envy-Free Pricing in Matching Markets with Budgets

Riccardo Colini-Baldeschi<sup>1</sup>(✉), Stefano Leonardi<sup>2</sup>, and Qiang Zhang<sup>2</sup>

<sup>1</sup> Dipartimento di Economia e Finanza, LUISS, Rome, Italy  
rcolini@luiss.it

<sup>2</sup> Sapienza University of Rome, Rome, Italy  
leonardi@dis.uniroma1.it, qzhang@gmail.com

**Abstract.** We study envy-free pricing mechanisms in matching markets with  $m$  items and  $n$  budget constrained buyers. Each buyer is interested in a subset of the items on sale, and she appraises at some single-value every item in her preference-set. Moreover, each buyer has a budget that constraints the maximum affordable payment, while she aims to obtain as many items as possible of her preference-set. Our goal is to compute an envy-free pricing allocation that maximizes the revenue. This pricing problem is hard to approximate better than  $\Omega(\min\{n, m\}^{1/2-\epsilon})$  for any  $\epsilon > 0$ , unless  $P = NP$  [7]. The goal of this paper is to circumvent the hardness result by restricting ourselves to specific settings of valuations and budgets. Two particularly significant scenarios are: each buyer has a budget that is greater than her single-value valuation, and each buyer has a budget that is lower than her single-value valuation. Surprisingly, in both scenarios we are able to achieve a  $1/4$ -approximation to the optimal envy-free revenue.

## 1 Introduction

In this paper we study revenue maximization with envy-free pricing in matching markets. Imagine a seller that would like to sell  $m$  different items to  $n$  buyers. Every buyer  $i$  is interested in a subset of items  $S_i$  (the *preference-set*) and has a budget  $b_i$  that represents her maximum affordable payment. Moreover, every buyer  $i$  appraises each item in her preference-set at value  $v_i$ , and any other item has zero-value for her. The buyers are willing to get the largest number of items in their preference-sets. But every buyer  $i$  doesn't want to pay more than her value  $v_i$  for an item and more than her budget  $b_i$  for the whole set of obtained items. The seller has full knowledge on the buyer types and aims to compute an outcome that maximizes her revenue. The outcome is composed by a payment vector (a payment  $p_i$  for each buyer  $i$ ) and an allocation vector (a, potentially empty, set of items  $X_i$  for each buyer  $i$ ).

---

This work is partly supported by the EU FET project MULTIPLEX no. 317532 and the Google Focused Award on “Algorithms for Large-scale Data Analysis”.

Every buyer  $i$ , given a payment  $p_i$  and an allocation  $X_i$ , has a utility equal to  $v_i \cdot |X_i \cap S_i| - p_i$  if the payment is less than or equal to the budget, i.e.  $p_i \leq b_i$ . However, the utility becomes  $-\infty$  if the payment exceeds the budget, i.e.  $p_i > b_i$ .

Our goal is to design a pricing algorithm that is able to provide a good revenue to the seller and observes some fairness criterion for the buyers. In order to model fairness, we consider two notions: (i) the *individual rationality*, and (ii) the *envy-freeness*. The individual rationality is obtained if at the end of the auction no buyer will experience a negative utility. The envy-freeness is obtained if, given a pricing scheme (an assignment of prices to items or sets of items), every buyer obtains the most desired set of items.

Different notions of envy-freeness have been studied in literature, and a detailed discussion about them is deferred to Sect. 1.1. In classical economics, the standard envy-free definition embodies *bundle pricing* [12, 17]. In the bundle pricing scheme each (different) set of items has a (potentially) different price. Thus in a *bundle-price envy-free* allocation no buyer has an incentive to barter her bundle with the bundle of someone else. More recently, envy-freeness has been often studied in the more restrictive setting of *item pricing* [3, 5, 14]. In an item pricing scheme each (different) item has a (potentially) different price. Thus, the price for a bundle is obtained by summing the single items' prices contained in it. Consequently, in *item-price envy-free* allocations every buyer receives the bundle that maximizes her own utility. The difference is that the former definition of envy-freeness allows a buyer to envy only bundles that are assigned to someone else. Instead, the latter definition allows a buyer to be envious if the obtained bundle is not the utility-maximizer bundle (even if the utility-maximizer bundle is not allocated at all). In this paper we will use both definitions, and we will explicitly state which envy-free condition is satisfied for each provided algorithm. All our algorithms use as a benchmark the optimal bundle-price envy-free revenue.<sup>1</sup>

The envy-free revenue-maximization problem with budget constraints has been initially studied by Feldman et al. [10] in the multi-unit setting that considers one single kind of items (no matching constraints). The authors discussed the limitations of different pricing schemes, specifically they showed through separation examples that an item-pricing scheme cannot achieve more than  $O(1/m)$  fraction of the optimal envy-free bundle-price revenue on some specific examples. Thus, relying on a bundle pricing scheme, they provided a 2-approximation algorithm to the optimal envy-free bundle-price revenue for multi-unit setting with budgeted buyers. More recently, Branzei et al. [4] studied the same multi-unit setting in the context of item-pricing schemes, and they gave an FPTAS for the optimal envy-free item-price revenue (and an exact algorithm for special cases) using an item-pricing scheme.

The envy-free revenue-maximization problem in matching markets with budget constraints has been considered for the first time in Colini-Baldeschi et al. [7].

---

<sup>1</sup> A bundle pricing scheme is able to extract more revenue than an item pricing scheme. Thus, competing against the optimal bundle-price envy-free revenue is the hardest task in this context. See Sect. 1.1.

They proved that for  $n$  buyers and  $m$  items, the optimal revenue cannot be approximated by a polynomial time algorithm within  $\Omega(\min\{n, m\}^{1/2-\epsilon})$  for any  $\epsilon > 0$ , unless  $P = NP$ . In this paper we present a novel approach that allows us to circumvent the  $\Omega(\min\{n, m\}^{1/2-\epsilon})$  impossibility result for relevant special cases.

**Our Results.** We first consider the case where every buyer has a budget that is greater than or equal to her valuation. For example, in online advertising, it is typical that the advertisers' budgets are greater than their CPM valuations for a block of impressions, and the advertisers are interested in buying multiple blocks of impressions from different publishers. Second, complementary to the first case, we study the case where every buyer has a budget that is less than her valuation. In the two cases we are able to provide: (i) an individually-rational and item-price envy-free algorithm that achieves a 4-approximation to the optimal bundle-price envy-free revenue, when  $b_i \geq v_i$  for every buyer  $i$ , and (ii) an individually-rational and bundle-price envy-free algorithm that achieves a 4-approximation to the optimal bundle-price envy-free revenue, when  $b_i < v_i$  for every buyer  $i$ . The key ingredient needed to achieve these results is the design of an ascending price auction (inspired by the Ausubel auction [1]) for the matching market setting. The main idea is to use in the selling procedure proper arguments from matching theory ( $B$ -matchings and their properties) to compute an envy-free allocation of items to buyers.

Besides envy-freeness we are interested to secure the fairest pricing scheme to the buyers. The fix-price scheme (one single price for all the items) is the standard way to achieve fairness in setting with identical items. But, as expected, in setting with different items (like matching markets) a fixed-price scheme cannot obtain more than a logarithmic fraction of the optimal revenue regardless of computational complexity considerations (see [8]). Thus different prices for different items is a strict requirement to achieve a constant approximation. Nonetheless our aim is still to diversify the prices of the items as little as possible to avoid discrimination among buyers.

Since the ascending price auction approach can be adapted in either special cases. We conclude that the general inapproximability result of this problem provided in Colini-Baldeschi et al. [7] is the result of the interaction between the above mentioned buyer classes. Due to space limitations some proofs are omitted, but they are available in [8].

## 1.1 Related Work

*Envy-Free Pricing.* The notion of envy-freeness was initially introduced by Foley [12] and Varian [17]. The key property of an envy-free allocation is that no buyer has incentive to exchange her bundle-payment pair with the bundle-payment pair of another buyer. That is to say, given a set of bundles and the corresponding prices, every buyer obtains the bundle that maximizes her utility. This definition became the standard definition in the economics literature. More recently, the notion of envy-freeness has been deeply studied with a different



perspective that involves item-pricing instead of bundle-pricing. Indeed the key property of an envy-free allocation has been reshaped as follows: given a set of items and the corresponding (per-item) prices, every buyer obtains the bundle that maximizes her utility. Notice that the price of a bundle is automatically obtained from the sum of the items' prices contained in it. It is easy to see that this definition is more difficult to satisfy than the classical economics definition. The item-pricing version of envy-freeness has been considered in [3, 5, 14]. The envy-freeness was studied in the context of different pricing schemes by Feldman et al. [10]. Specifically, they focused their attention on: (i) the bundle-pricing scheme, where it is possible to specify different prices for different bundles, (ii) the item-pricing scheme, where different prices can be assigned to different items, and (iii) proportional-pricing scheme, that embody an item-pricing scheme plus the possibility to specify a maximum or a minimum size on the bundles that can be demanded by the buyers. Moreover, the authors were able to rank these pricing schemes with respect to two distinct criteria: the customer experience (how much the customers judge fair a pricing scheme) or the obtainable revenue (how much a pricing scheme is able to extract from the customers). In terms of customer experience the most desired pricing scheme is the item-pricing scheme, then the proportional-pricing scheme, and finally the bundle-pricing scheme. This is because the customers prefer the pricing schemes that allow less discrimination and are more uniform (uniformity is perceived as fairness). As expected, on the revenue perspective the ranking is reversed: the bundle-pricing scheme is able to produce the highest revenue, then proportional-pricing scheme is able to extract something less, and then the item-pricing scheme is the scheme that is able to extract less. Feldman et al. [10] have shown several separation examples that clearly state the limits of the different pricing schemes.

*Revenue-Maximization.* Moreover, Feldman et al. [10] proved that the problem of computing an optimal bundle-pricing envy-free revenue is NP-hard. Thus, they provided a 2-approximation algorithm for multi-unit setting with budget constraints, it is also proved that this result is tight. Colini-Baldeschi et al. [7] investigated the problem in the context of multi-unit fixed-price auctions with budget constraints and matching markets with budget constraints. Particularly relevant for our paper is the hardness result presented there. Indeed, they proved that the revenue maximizing envy-free problem in matching markets is  $\Omega(\min\{n, m\}^{1/2-\epsilon})$  inapproximable for any  $\epsilon > 0$ , unless  $P = NP$ . This is why we are forced to study (particularly relevant) special cases. The envy-free revenue maximization problem in multi-unit setting with budgets is also studied in [4]. They provided algorithms that approximate optimal social welfare and optimal revenue with an item-pricing scheme when buyers are price takers, and an impossibility result for price making buyers. The main difference between [7] and [4] is that the former considered bundle-pricing envy-freeness (pairwise envy-freeness), while the latter focused on the definition of item-pricing envy-freeness. Recently, the item-pricing envy-free problem for general valuations in multi-unit markets without budgets is considered in [16], and a dynamic programming algorithm is provided in such setting.

*Ascending Price Auctions.* Ascending price auctions were used in FCC spectrum auctions and were initially studied in [1, 2, 15]. Later, ascending price auctions were widely applied in the context of sponsored search auctions. Dobzinski et al. [9] designed a Pareto-optimal, incentive compatible ascending price auction for a multi-unit setting when buyers have budget constraints. Numerous subsequent papers extended this setting, see [6, 11, 13].

## 2 Preliminaries

An instance of the revenue-maximizing envy-free pricing problem in matching markets can be formally depicted by the tuple  $\mathcal{A} = \langle I, J, \mathbf{S}, \mathbf{v}, \mathbf{b} \rangle$ . There is a set of  $|I| = n$  buyers and a set of  $|J| = m$  different items in the market. Every buyer  $i \in I$  is interested in a set of items  $S_i \subseteq J$ , which we refer as the preference-set of buyer  $i$ . Buyers equally value the items in their preference-sets. Specifically, every buyer  $i$  has a valuation  $v_i \in \mathbb{R}_{>0}$  for each item  $j \in S_i$  and has a valuation of zero for any item  $j \notin S_i$ . Every buyer  $i$  has a budget  $b_i \in \mathbb{R}_{>0}$  that is the maximum payment she can afford.

An algorithm computes an outcome  $\langle \mathbf{X}, \mathbf{p} \rangle$  for every possible instance  $\mathcal{A}$ , where  $\mathbf{X} = \langle X_1, \dots, X_n \rangle$  is the allocation vector, and  $\mathbf{p} = \langle p_1, \dots, p_n \rangle$  is the payment vector. That is, for each buyer  $i$ ,  $X_i \subseteq J$  is the set of items allocated to buyer  $i$ , and  $p_i \in \mathbb{R}_{\geq 0}$  is the payment charged to buyer  $i$ . Moreover, we use  $\bar{p}_i$  to denote the per-item-price paid by a buyer  $i$ , i.e.,  $\bar{p}_i = \frac{p_i}{|X_i|}$ . Given an allocation  $X_i$  and a payment  $p_i$ , the utility  $u_i(X_i, p_i)$  of a buyer  $i$  is defined as  $v_i \cdot |X_i \cap S_i| - p_i$  if  $b_i \geq p_i$  and  $-\infty$  otherwise.

A feasible outcome  $\langle \mathbf{X}, \mathbf{p} \rangle$  must satisfy the following constraints: (i) *feasibility* (or *supply constraint*): for any pair of buyers  $i, i' \in I$ ,  $X_i \cap X_{i'} = \emptyset$ , and (ii) *individual rationality*: for any buyer  $i \in I$ ,  $u_i(X_i, p_i) \geq 0$ .

Furthermore, we incline toward our algorithm to produce envy-free outcomes. An outcome  $\langle \mathbf{X}, \mathbf{p} \rangle$  is:

- *envy-free*: if given an item-pricing vector  $\rho = \{\rho_1, \dots, \rho_m\}$  such that the price for the item  $j$  is  $\rho_j$ . Then  $p_i = \sum_{j \in X_i} \rho_j$ , and there is no bundle  $X' \subseteq J$  such that  $u_i(X', \sum_{j \in X'} \rho_j) > u_i(X_i, p_i)$ . And,
- *pairwise envy-free*: if given a set of proposed bundles  $\mathbf{X}$  such that every bundle  $X_i \in \mathbf{X}$  has a corresponding price  $p_i \in \mathbf{p}$ , and every bundle  $X' \notin \mathbf{X}$  has price equal to  $\infty$ . Then for every buyer  $i$  and every bundle  $X_j \in \mathbf{X}$ , buyer  $i$  prefers her own bundle, i.e.,  $u_i(X_i, p_i) \geq u_i(X_j, p_j)$ .

Given outcome  $\langle \mathbf{X}, \mathbf{p} \rangle$ , the revenue (i.e., the revenue of the algorithm on instance  $\mathcal{A}$ ) is the sum of the payments of all buyers, i.e.,  $\mathcal{R}(\mathbf{X}, \mathbf{p}) = \sum_{i \in I} p_i$ . Our goal is to design an envy-free algorithm that approximates the optimal envy-free revenue for every possible instance  $\mathcal{A}$ .

**Ascending Price Auction.** Our technique to design revenue-maximizing envy-free algorithms relies on the implementation of an ascending price auction. The standard implementation of an ascending price auction is as follows: (i) the

price (initialized at zero) is raised until some condition is met (usually referred as *selling condition*), (ii) when the selling condition is met an appropriate *selling procedure* is executed, and the auction goes back to (i) (if there is at least one buyer with positive demand). The novelty of the ascending price auction described in this paper is about the selling procedure. We will give a detailed description of the selling procedures in Sects. 3 and 4, but in either cases the selling procedure relies on the graph representation of the problems and their properties.

To facilitate our future presentation and analysis, here we introduce some necessary notations. For a price  $p$ , the *demand* of buyer  $i$  is defined as follows:

$$D_i(p) = \begin{cases} \min\{\lfloor \frac{b_i}{p} \rfloor, |S_i|\}, & \text{if } p \leq v_i \\ 0, & \text{if } p > v_i \end{cases} \tag{2.1}$$

Intuitively,  $D_i(p)$  is the number of the items that maximizes the utility of buyer  $i$  if all items in  $S_i$  are priced at  $p$ .

Given price  $p$ , we define three sets of buyers  $A^p$ ,  $Q^p$ , and  $I^p$ .  $A^p$  contains the buyers whose valuations are strictly greater than  $p$  and having positive demands, i.e.,  $A^p = \{i \in I | v_i > p \wedge D_i(p) > 0\}$ .  $Q^p$  contains the buyers whose valuations are equal to  $p$  and having positive demands, i.e.,  $Q^p = \{i \in I | v_i = p \wedge D_i(p) > 0\}$ . Finally, let  $I^p$  be the union of  $A^p$  and  $Q^p$ , i.e.,  $I^p = A^p \cup Q^p$ .

**Graph Representation.** Matching markets have a very intuitive bipartite graph representation. Consider a set of buyers  $I' \subseteq I$  and a set of items  $J' \subseteq J$  as two disjoint sets of nodes in a bipartite graph. Given a price  $p$ , there exists an edge between  $i \in I'$  and  $j \in J'$  if buyer  $i$  demands item  $j$  at price  $p$ , i.e.,  $(j \in S_i) \wedge \min\{v_i, b_i\} \geq p$ . More specifically, taking the notations established above, given a particular price  $p$  and a subset of items  $J' \subseteq J$ , we define a bipartite graph  $G^p = (I^p \cup J', E^p)$ , where  $E^p = \{(i, j) | i \in I^p, j \in J' \cap S_i\}$ . Similarly, we define  $\bar{G}^p = (A^p \cup J', \bar{E}^p)$  as the bipartite graph that only includes buyers in  $A^p$ . In this paper we refer to  $G^p$  and  $\bar{G}^p$  as *preference-graphs*.

Additionally, allocations in matching markets can be seen as matchings in preference-graphs, because an allocation essentially “maps” buyers to a subset of items. To formalize this idea, we introduce the concept of  $B$ -matching.

**Definition 2.1.** *Given a bipartite graph  $G^p = (I^p \cup J', E^p)$ , a  $B$ -matching  $\mathcal{M}(G^p)$  is a sub-graph of  $G^p$  such that every buyer is not matched to a number of items greater than her demand, i.e.,  $\forall i \in I^p, |\{j \in J' : (i, j) \in \mathcal{M}(G^p)\}| \leq D_i(p)$ , and every item is not matched to more than one buyer, i.e.,  $\forall j \in J^p, |\{i \in I^p : (i, j) \in \mathcal{M}(G^p)\}| \leq 1$ .*

Similarly, we use  $\mathcal{M}(\bar{G}^p)$  to denote the  $B$ -matching on graph  $\bar{G}^p$ . From now on, we will simply write matchings instead of  $B$ -matchings, and the notation  $\mathcal{M}(G^p)$  will refer to a *maximum matching* on  $G^p$ . By the matching’s definition, the allocations constructed from the matchings would satisfy the supply constraint and the budget constraint.

Finally, we introduce the concept of augmenting path. Given a preference-graph, an augmenting path starts with a buyer and ends with an unallocated item in the preference-graph. We mainly use augmenting paths to produce envy-free outcomes.

**Definition 2.2.** *Given a B-matching  $\mathcal{M}(G^p)$  (resp.  $\mathcal{M}(\bar{G}^p)$ ), an augmenting path  $\pi$  between buyer  $i$  and item  $j$  is a path  $\pi = \{i = y_1, z_1, y_2, z_2, \dots, y_h, z_h = j\}$  such that the following conditions hold:*

1.  $\forall k \in [1, \dots, h]$ ,  $y_k$  is a buyer. That is,  $y_k \in I^p$  (resp.  $y_k \in A^p$ ).
2.  $\forall k \in [1, \dots, h]$ ,  $z_k$  is an item. That is,  $z_k \in J'$ .
3.  $\forall k \in [1, \dots, h]$ ,  $z_k \in S_{y_k}$  and  $(y_k, z_k) \notin \mathcal{M}(G^p)$  (resp.  $(y_k, z_k) \notin \mathcal{M}(\bar{G}^p)$ ).
4.  $\forall k \in [1, \dots, h-1]$ ,  $(z_k, y_{k+1}) \in \mathcal{M}(G^p)$  (resp.  $(z_k, y_{k+1}) \in \mathcal{M}(\bar{G}^p)$ ).

In an augmenting path defined above, if item  $z_h$  is allocated to a buyer  $i \notin \{y_1, \dots, y_h\}$  at price  $p$ , we will also allocate some item to every buyer in  $\{y_1, \dots, y_h\}$ . We use this technique to achieve the envy-freeness among buyers.

### 3 An Ascending Price Algorithm for $b_i \geq v_i$

In this section we present an algorithm that obtains a 4-approximation to the optimal bundle-price envy-free revenue when all buyers have budgets that are equal to or greater than their valuations, i.e.,  $\forall i \in I, b_i \geq v_i$ . We remark that this is an item pricing algorithm, and it achieves a constant approximation with respect to the optimal bundle-price envy-free revenue.

The general idea is to implement an ascending price auction with selling conditions and selling procedures accurately designed to take advantage of this scenario. The implementation of the ascending price auction is described in Algorithm 1. At the beginning the price is set to zero. Then the algorithm increases the price until a proper selling condition is satisfied (line 4). When the selling condition is matched, the algorithm executes a selling procedure and assigns items and payments to the involved buyers. Notice that in Algorithm 1 there are two different selling procedures: COMPUTE-ALLOCATION-I described by Algorithm 2, and COMPUTE-ALLOCATION-II described by Algorithm 3. The selling condition and the two selling procedures will be described in the next subsections. The omitted proofs are available in [8].

**Selling Condition and Critical Prices.** Algorithm 1 uses a selling condition to catch the correct price at which a selling procedure can be executed. Intuitively, the notion of correct price relies on the following abstraction: nothing can be sold if the buyers' cumulative demand is too high (the sum of the buyers' demands is greater than the available items). In order to determine if the buyers' cumulative demand can be satisfied, we have to tackle two main difficulties: (i) in a matching markets setting the demand can be too high on a particular set of items but too low on a different set, and (ii) the demand functions are not continuous, thus we can have a cumulative demand that is too high at  $p$  and too low at  $p + \epsilon$ .

So, we want to detect each price  $p$  that is borderline between a too high and a too low cumulative demand. These prices are called *critical prices*. To detect critical prices, we have to compute two maximum matchings for each price  $p$ : a maximum matching at price  $p$  and a maximum matching at price  $p + \epsilon$ . If the size of the maximum matching at  $p$  is greater than the size of the maximum matching at  $p + \epsilon$ , then we know that on some set of items the cumulative demand will be too low at any price higher than  $p$ . So, we have to check if something should be sold at  $p$ . To be more formal, we define the set of critical prices as follows:

**Definition 3.1.** *Given a price  $p$ , let  $G^p = (I^p \cup J', E^p)$  where  $J' \subseteq J$  is the set of unsold items. The price  $p$  is critical if, for  $\epsilon$  small enough,  $|\mathcal{M}(G^p)| > |\mathcal{M}(G^{p+\epsilon})|$ , where  $\mathcal{M}(G^p)$  and  $\mathcal{M}(G^{p+\epsilon})$  are the maximum matchings in  $G^p$  and  $G^{p+\epsilon}$ , respectively.*

**Detailed Description and Selling Procedures.** For each critical price  $p$ , the Algorithm 1 checks why  $p$  is a critical price. It can be because the demand of some buyer in  $Q^p$  goes to zero when the price is increased above  $p$  (line 5). In this case, Procedure COMPUTE-ALLOCATION-I computes an envy-free partial assignment, where every item is sold at  $p$ . The other reason is that some buyer in  $A^p$  cannot afford the same amount of items at a slightly higher price (for budget limitations). In this case, Procedure COMPUTE-ALLOCATION-II computes an envy-free partial assignment, where every item is sold at  $p + \epsilon$ .

In either procedures the key element is the computation of an envy-free partial assignment. In order to describe how an envy-free partial assignment is computed, we introduce some further notation. Given a graph  $G^p = (I^p \cup J', E^p)$ ,

---

**Algorithm 1.** An ascending price algorithm for matching markets when  $b_i \geq v_i$

---

**Input:**  $\langle I, J, \mathbf{S}, \mathbf{v}, \mathbf{b} \rangle$

**Output:**  $\langle \mathbf{X}, \mathbf{p} \rangle$

```

1:  $p \leftarrow 0$ ; /* $p$  is the uniform price for all items and it is dynamically increasing*/
2:  $G^p = (I^p \cup J, E^p)$ ; /* $G^p$  is the preference-graph for unsold items at price  $p$  */
3: while  $|\mathcal{M}(G^p)| \neq \emptyset$  do
4:   Increase  $p$  until  $|\mathcal{M}(G^p)| > |\mathcal{M}(G^{p+\epsilon})|$ ;
5:   if  $|\mathcal{M}(G^p)| > |\mathcal{M}(\bar{G}^p)|$  then
6:      $M \leftarrow$ COMPUTE-ALLOCATION-I( $I, J, \mathbf{S}, \mathbf{v}, \mathbf{b}, p$ );
7:     for each edge  $(i, j) \in M$  do
8:        $J = J \setminus \{j\}, X_i = X_i \cup \{j\}, p_i = p_i + p, b_i = b_i - p$ ;
9:     end for
10:  else
11:     $M \leftarrow$ COMPUTE-ALLOCATION-II( $I, J, \mathbf{S}, \mathbf{v}, \mathbf{b}, p$ );
12:    for each edge  $(i, j) \in M$  do
13:       $J = J \setminus \{j\}, X_i = X_i \cup \{j\}, p_i = p_i + (p + \epsilon), b_i = b_i - (p + \epsilon)$ ;
14:    end for
15:    Remove all the items that are not demanded anymore.
16:  end if
17: end while

```

---

a maximum matching  $\mathcal{M}(G^p)$  on  $G^p$ , and a subset of items  $J' \subseteq J$ , we denote with  $\bar{J}^p \subseteq J'$  the set of items that are not matched in  $\mathcal{M}(G^p)$ , and with  $N(\bar{J}^p) \subseteq I^p$  the set of buyers that are connected with an augmenting path in  $\mathcal{M}(G^p)$  to some item in  $\bar{J}^p$ . Now, an envy-free partial assignment is computed as follows: (i) a maximum matching  $\mathcal{M}(G^p)$  on the graph  $G^p = (I^p \cup J, E^p)$  is computed, (ii) let  $\bar{J}^p \subseteq J'$  be the subset of items that are not matched in  $\mathcal{M}(G^p)$ , (iii) an envy-free partial assignment is the subgraph of  $\mathcal{M}(G^p)$  that involves only buyers in  $N(\bar{J}^p)$ . Restrict the allocation to the buyers in  $N(\bar{J}^p)$  is the key to obtain envy-freeness and a good revenue from a partial assignment. Indeed, we claim that if a buyer  $i$  is connected with an augmenting path to an item  $j$ , but that item  $j$  remains unmatched in a maximum matching, then the buyer  $i$  obtains as many items as she demands.

The next paragraphs will discuss the details of the two procedures. Procedure COMPUTE-ALLOCATION-I and Procedure COMPUTE-ALLOCATION-II implement the computation of envy-free partial assignment with different details, but the two procedures are similar in spirit.

**Compute-Allocation-I.** Algorithm 2 is executed when the price becomes equal to the valuation of some buyer, and it is not possible to sell the same amount of items at a slightly higher price. Notice that these buyers belong to the set  $Q^p$ . Moreover, recall that, by definition of envy-freeness, every buyer in  $Q^p$  cannot envy any other buyer that obtains items at price  $p$ . Let us denote with  $J_Q^p$  the set of available items in the preference sets of buyers in  $Q^p$ . Thus, the general idea is to compute an envy-free partial assignment such that most of the items in  $J_Q^p$  are assigned to the buyers in  $A^p$ . In succession, the items that remain unmatched will be assigned to the buyers in  $Q^p$ .

To achieve this, the algorithm computes a maximum matching  $\mathcal{M}(\bar{G}^p)$  with the minimal number of items in  $J_Q^p$  matched (line 3 in Algorithm 2). This can be done computing a maximal weight matching. In order to allocate the remaining items to buyers in  $Q^p$  and preserve the envy-freeness of the outcome, the

---

**Algorithm 2.** COMPUTE-ALLOCATION-I

---

- 1: **procedure** COMPUTE-ALLOCATION-I( $I, J, \mathbf{S}, \mathbf{v}, \mathbf{b}, p$ )
  - 2:   Let  $J_Q^p$  be the set of items in the preference-sets of buyers in  $Q^p$  at price  $p$ , i.e.  
 $J_Q^p = \bigcup_{i \in Q^p} S_i$ ;
  - 3:   Compute a maximum  $B$ -matching  $\mathcal{M}(\bar{G}^p)$  with minimum number of items in  $J_Q^p$  matched;
  - 4:   Let  $\bar{J}^p$  be the set of items that are not matched in  $\mathcal{M}(\bar{G}^p)$ ;
  - 5:   Let  $N(\bar{J}^p)$  be the set of buyers that are connected to an item in  $\bar{J}^p$  with an augmenting path in  $\mathcal{M}(\bar{G}^p)$ ;
  - 6:   Let  $M = \{(i, j) \in \mathcal{M}(\bar{G}^p) | i \in N(\bar{J}^p)\}$ ;
  - 7:   Assign items in  $\bar{J}^p$  to buyers in  $Q^p$  such that the allocation satisfies the supply constraint and budget constraint;
  - 8:   Include the assignment of  $\bar{J}^p$  to  $M$ ;
  - 9:   **return**  $M$
  - 10: **end procedure**
-

algorithm first allocates items to buyers in  $A^p$  who have augmenting paths to any remaining item in  $\mathcal{M}(\bar{G}^p)$ . Essentially, if an item is allocated to a buyer in  $Q^p$  at price  $p$ , every buyer  $i$  in  $A^p$  who is connected to the item by an augmenting path obtains  $\lfloor \frac{b_i}{p} \rfloor$  items and pays  $p \cdot \lfloor \frac{b_i}{p} \rfloor$ . Thus, buyer  $i$  will not envy the buyers in  $Q^p$ . Every envy-free partial assignment  $M \subseteq \mathcal{M}(\bar{G}^p)$  computed by procedure COMPUTE-ALLOCATION-I satisfies two important properties: (i) every buyer  $i \in M$  receives an amount of items exactly equal to her demand, i.e.,  $|\{(i, j) \in M\}| = D_i(p) \forall i \in M$ , and (ii) every item that is not assigned at the end of procedure COMPUTE-ALLOCATION-I is requested by at least one buyer after procedure COMPUTE-ALLOCATION-I. These properties are formally proved by the following lemmata.

**Lemma 3.1.** *Let  $\mathcal{M}(\bar{G}^p)$  be a maximum matching on graph  $\bar{G}^p$  and  $J_Q^p \neq \emptyset$ , then each buyer  $i \in N(\bar{J}^p)$  is matched to  $D_i(p)$  items in  $\mathcal{M}(\bar{G}^p)$ , where  $\bar{J}^p$  is the set of items not matched in  $\mathcal{M}(\bar{G}^p)$ .*

**Lemma 3.2.** *In Procedure COMPUTE-ALLOCATION-I, all items in  $\bar{J}^p$  can be allocated to the buyers in  $Q^p$  at price  $p$  per each.*

**Compute-Allocation-II.** Algorithm 3 is executed when it is not possible to sell the same amount of items at a slightly higher price, but no buyer in  $Q^p$  is responsible for that. Notice that in this case, the buyers that decrease their demands are in  $A^p$ . Namely, the buyers drop their demands because budgets and not valuations. Thus, there are no buyers in  $Q^p$  that are relevant for us. Consequently, it is enough to compute an envy-free partial assignment at price  $p+\epsilon$ . Similarly to the previous case, to preserve the envy-freeness of the outcome, the algorithm allocates items to buyers in  $A^p$  who have an augmenting path to an unmatched items in  $\mathcal{M}(G^{p+\epsilon})$ . Given the fact that all items are matched in  $\mathcal{M}(G^p)$ , if an items is not matched to any buyer in  $\mathcal{M}(G^{p+\epsilon})$ , then it implies that all buyers who have augmenting paths to this item must be fully matched in  $\mathcal{M}(G^{p+\epsilon})$ . In other words, those buyers do not have enough budgets to buy one more item. Otherwise, the size of  $\mathcal{M}(G^{p+\epsilon})$  can be increased. We prove this property in the following lemma.

---

**Algorithm 3.** COMPUTE-ALLOCATION-II

---

- 1: **procedure** COMPUTE-ALLOCATION-II( $I, J, \mathbf{S}, \mathbf{v}, \mathbf{b}, p$ )
  - 2:    Compute a maximum  $B$ -matching  $\mathcal{M}(G^{p+\epsilon})$ ;
  - 3:    Let  $\bar{J}^{p+\epsilon}$  be the set of items that are not matched in  $\mathcal{M}(G^{p+\epsilon})$ ;
  - 4:    Let  $N(\bar{J}^{p+\epsilon})$  be the set of buyers that are connected to an item  $\bar{J}^{p+\epsilon}$  with an augmenting path;
  - 5:    Let  $M = \{(i, j) \in \mathcal{M}(G^{p+\epsilon}) | i \in N(\bar{J}^{p+\epsilon})\}$ ;
  - 6:    Remove items in  $\bar{J}^{p+\epsilon}$  from  $J$ ;
  - 7:    **return**  $M$
  - 8: **end procedure**
-

**Lemma 3.3.** *Let  $\mathcal{M}(G^{p+\epsilon})$  be a maximum matching on graph  $G^{p+\epsilon}$ , and let  $\bar{J}^{p+\epsilon} \neq \emptyset$  be the set of items not matched in the maximum matching  $\mathcal{M}(G^{p+\epsilon})$ . Then for each buyers  $i \in N(\bar{J}^{p+\epsilon})$  the followings hold: (i)  $v_i > p$ , and (ii) buyer  $i$  is matched to  $D_i(p + \epsilon)$  items in  $\mathcal{M}(G^{p+\epsilon})$ .*

Notice that at the end of the procedure it is possible that some items remain unassigned but no buyer will demand them anymore. This is a potential problem for the revenue, because unassigned items can be translated in unused budgets. But, the next lemma shows that the number of unassigned items is bounded.

**Lemma 3.4.** *Let  $\mathcal{M}(G^{p+\epsilon})$  be a maximum matching on graph  $G^{p+\epsilon}$ , and let  $\bar{J}^{p+\epsilon}$  be the set of items not matched in the maximum matching  $\mathcal{M}(G^{p+\epsilon})$ , then  $|\bar{J}^{p+\epsilon}| \leq |N(\bar{J}^{p+\epsilon})|$ .*

**Main result.** Finally, we are ready to prove that the outcome computed by Algorithm 1 is envy-free and achieves a 4-approximation to the optimal bundle-price revenue.

We start with some auxiliary lemmata:

**Lemma 3.5.** *Let  $\langle \mathbf{X}, \mathbf{p} \rangle$  be the outcome obtained by Algorithm 1. If  $X_i \neq \emptyset$ , the buyer  $i$  obtains all the items in  $X_i$  at a unique price-per-item  $\bar{p}_i = \frac{p_i}{|X_i|}$ .*

**Lemma 3.6.** *If buyer  $i$  obtains  $X_i$  at price-per-item  $\bar{p}_i$ , then all the items assigned at a price  $p' < \bar{p}_i$  are not in her preference-set.*

**Lemma 3.7.** *If buyer  $i$  does not obtain any item, i.e.,  $X_i = \emptyset$ , then all items in  $S_i$  are sold at a price greater than or equal to  $v_i$ .*

Now we show that Algorithm 1 is envy-free.

**Theorem 3.1.** *The outcome  $\langle \mathbf{X}, \mathbf{p} \rangle$  produced by Algorithm 1 is envy-free.*

*Proof.* First, by Lemma 3.7, we know that the buyers that do not obtain any item do not envy anyone. Furthermore, by Lemma 3.5, we know that, for the rest of buyers, they obtain all items in  $X_i$  at a unique per-item-price  $\bar{p}_i = \frac{p_i}{|X_i|}$ . The rest of the proof is divided into two cases. (i)  $\bar{p}_i = v_i$ : The buyer  $i$  would not envy any buyer  $j$  that gets her bundle  $X_j$  at a price-per-item  $\bar{p}_j \geq \bar{p}_i$ . Moreover, by Lemma 3.6, buyer  $i$  is not interested in any item allocated at a price  $p < \bar{p}_i$ . This implies that buyer  $i$  cannot envy any buyer  $j$  who obtains the bundle at price-per-item  $p \leq \bar{p}_i$ . (ii)  $\bar{p}_i < v_i$ : Buyer  $i$  obtains  $X_i$  in either COMPUTE-ALLOCATION-I or COMPUTE-ALLOCATION-II. In both cases, by Lemmas 3.1 and 3.3, buyer  $i$  obtains  $D_i(\bar{p}_i) = |X_i|$  items. Hence, buyer  $i$  does not envy any buyer who obtains her bundle at price-per-item  $p \geq \bar{p}_i$ . Moreover, by Lemma 3.6, buyer  $i$  is not interested in any item allocated at a price  $p < \bar{p}_i$ . This implies that buyer  $i$  cannot envy any buyer  $j$  who obtains the bundle at price-per-item  $\bar{p}_j \leq \bar{p}_i$ . Thus we conclude that Algorithm 1 is envy-free.  $\square$

Now, we show that the outcome computed by Algorithm 1 is a 4-approximation to the optimal bundle-price envy-free revenue.

**Theorem 3.2.** *Algorithm 1 achieves a 4-approximation to the optimal envy-free revenue when all buyers have budgets that are at least their valuations.*



## 4 An Ascending Price Algorithm for $b_i < v_i$

In this section we present an algorithm that obtains a 4-approximation to the optimal bundle-price envy-free revenue in matching markets when all buyers have budgets less than their valuations, i.e.,  $\forall i \in I, b_i < v_i$ . Due to a separation example showed in Feldman et al. [10], we know that the optimal bundle-price envy-free revenue cannot be approximated within  $O(\frac{1}{m})$  using an item pricing scheme. Thus the algorithm presented in this section will embody a bundle pricing scheme, and it will produce a pairwise envy-free outcome. Before the description of the algorithm, we overload some notations. Given price  $p$ , let  $A^p$  contain the buyers whose budget are strictly greater than  $p$  and having positive demands, i.e.,  $A^p = \{i \in I | b_i > p \wedge D_i(p) > 0\}$ . Let  $Q^p$  contain the buyers whose budgets are equal to  $p$  and having positive demands, i.e.,  $Q^p = \{i \in I | b_i = p \wedge D_i(p) > 0\}$ . Finally, let  $I^p$  be the union of  $A^p$  and  $Q^p$ , i.e.,  $I^p = A^p \cup Q^p$ .

**Detailed Description.** The algorithm, which is referred to as Algorithm 4, shares the similar spirit as Algorithm 1 but possesses some tweaks on the actions performed at critical prices. The algorithm starts with an initial price  $p = 0$  for all items and keeps increasing the price for all items until the price becomes a

---

**Algorithm 4.** An ascending price algorithm for matching markets when  $v_i > b_i$

---

**Input:**  $\langle I, J, \mathbf{S}, \mathbf{v}, \mathbf{b} \rangle$

**Output:**  $\langle \mathbf{X}, \mathbf{p} \rangle$

```

1:  $p \leftarrow 0$ ;  $/*p$  is the price per-item and it is dynamically increasing*/
2:  $G^p = (I^p \cup J, E^p)$ ;  $/*G^p$  is the preference-graph for unsold items at price  $p$  */
3: while  $|\mathcal{M}(G^p)| \neq \emptyset$  do
4:   Increase  $p$  until  $|\mathcal{M}(G^p)| > |\mathcal{M}(G^{p+\epsilon})|$ ;
5:   Partition  $(N(\bar{J}^p) \cup Q^p)$  into sets of buyers  $Y_1, \dots, Y_k$ , where for any  $Y_t$  and  $Y_{t'}$ 
   there do not exist buyers  $y \in Y_t$  and  $y' \in Y_{t'}$  such that  $S_y \cap S_{y'} \neq \emptyset$ ;
6:   for each  $Y_t$  do
7:     if  $\sum_{i \in (Y_t \cap Q^p) \setminus \hat{I}} b_i < \sum_{i \in Y_t \cap N(\bar{J}^p)} b_i$  then
8:        $M \leftarrow \text{COMPUTE-ALLOCATION-II}(Y_t, J, \mathbf{S}, \mathbf{v}, \mathbf{b}, p + \epsilon)$ ;
9:       for each edge  $(i, j) \in M$  do
10:         $\hat{I} = \hat{I} \cup \{i\}$ ,  $J = J \setminus \{j\}$ ,  $X_i = X_i \cup \{j\}$ ,  $p_i = p_i + p$ ,  $b_i = b_i - p$ ;
11:       end for
12:     else
13:       if  $\exists$  a matching  $M$  s.t. every buyer in  $Y_t \setminus \hat{I}$  is matched to one item then
14:         Compute such a matching  $M$ ;
15:       else
16:          $M \leftarrow \text{COMPUTE-ALLOCATION-II}(Y_t, J, \mathbf{S}, \mathbf{v}, \mathbf{b}, p + \epsilon)$ ;
17:       end if
18:       for each edge  $(i, j) \in M$  do
19:         $\hat{I} = \hat{I} \cup \{i\}$ ,  $J = J \setminus \{j\}$ ,  $X_i = X_i \cup \{j\}$ ,  $p_i = p_i + p$ ,  $b_i = b_i - p$ ;
20:       end for
21:     end if
22:   end for
23: end while

```

---

critical price. The reason of a price being a critical price is a little different from the previous case. Since  $v_i > b_i$  for all buyers, a price becomes a critical price when it is equal to the budget of some buyer, or buyers cannot afford the same amount of items at higher prices. At each critical price, the algorithm divides buyers into different partitions (a partition is denoted by  $Y_t$ ). One property of these partitions is that no buyers from different partitions have a common item in their preference sets. It allows us to focus on each partition separately. Then the following actions are performed on each partition.

The algorithm compares the remaining budgets between buyers in  $(Y_t \cap Q^p) \setminus \hat{I}$  and buyers in  $Y_t \cap N(\bar{J}^p)$  where  $\hat{I}$  is the set of buyers who have obtained items at previous critical prices and still have budgets to demand more items. Recall that  $\bar{J}^p$  is the set of items that are not matched in  $\mathcal{M}(\bar{G}^p)$  and  $N(\bar{J}^p)$  is the set of buyers that are connected to an item in  $\bar{J}^p$  in  $\mathcal{M}(\bar{G}^p)$ . If the sum of the budgets of buyers in  $(Y_t \cap Q^p) \setminus \hat{I}$  is relatively small, then the algorithm “ignores” them (i.e. does not allocate them any item) but allocates items to buyers in  $Y_t \cap N(\bar{J}^p)$  at  $p + \epsilon$  each.

It can be achieved by the same as Procedure COMPUTE-ALLOCATION-II. It would extract at least half of the budgets of buyers in  $Y_t \cap N(\bar{J}^p)$ , which in turn is a good approximation to the optimal revenue from all buyers in  $Y_t \setminus \hat{I}$ . On the other case, when the sum of the budgets of buyers in  $(Y_t \cap Q^p) \setminus \hat{I}$  is relatively large, the algorithm checks if it is possible to give one item to every buyer in  $Y_t \setminus \hat{I}$ . If yes, the algorithm allocates one item to each of them. By doing so, the algorithm extracts all the budgets of buyers in  $(Y_t \cap Q^p) \setminus \hat{I}$  since their budgets are equal to the price. It gives us a good approximation to the optimal revenue extracted from buyers in  $Y_t$ . Otherwise, the algorithm “ignores” buyers in  $(Y_t \cap Q^p) \setminus \hat{I}$  and allocate items to buyers in  $Y_t \cap N(\bar{J}^p)$  at price  $p + \epsilon$  each. We show that it does not hurt the revenue since the optimal envy-free algorithm cannot extract any revenue from those buyers either.

**Main Result.** Our main result is the following. The proof is available in [8].

**Theorem 4.1.** *Algorithm 4 is pairwise envy-free and achieves a 4-approximation to the optimal revenue in envy-free outcomes when for each  $i \in I$   $b_i < v_i$ .*

## References

1. Ausubel, L.M.: An efficient ascending-bid auction for multiple objects. *Am. Econ. Rev.* **94**(5), 1452–1475 (2004)
2. Ausubel, L.M., Milgrom, P.R.: Ascending auctions with package bidding. *Front. Theor. Econ.* **1**, 1019–1019 (2002)
3. Balcan, M.-F., Blum, A., Mansour, Y.: Item pricing for revenue maximization. In: *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 50–59 (2008)
4. Brânzei, S., Filos-Ratsikas, A., Miltersen, P.B., Zeng, Y.: Envy-free pricing in multi-unit markets. *CoRR*, abs/1602.08719 (2016)
5. Briest, P., Krysta, P.: Single-minded unlimited supply pricing on sparse instances. In: *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, pp. 1093–1102 (2006)

6. Colini-Baldeschi, R., Henzinger, M., Leonardi, S., Starnberger, M.: On multiple keyword sponsored search auctions with budgets. In: Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.) ICALP 2012. LNCS, vol. 7392, pp. 1–12. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-31585-5\\_1](https://doi.org/10.1007/978-3-642-31585-5_1)
7. Colini-Baldeschi, R., Leonardi, S., Sankowski, P., Zhang, Q.: Revenue maximizing envy-free fixed-price auctions with budgets. In: Liu, T.-Y., Qi, Q., Ye, Y. (eds.) WINE 2014. LNCS, vol. 8877, pp. 233–246. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-13129-0\\_18](https://doi.org/10.1007/978-3-319-13129-0_18)
8. Colini-Baldeschi, R., Leonardi, S., Zhang, Q.: Revenue maximizing envy-free pricing in matching markets with budgets. CoRR, abs/1610.09221
9. Dobzinski, S., Lavi, R., Nisan, N.: Multi-unit auctions with budget limits. In: The Proceedings of the 49th Annual Symposium on Foundations of Computer Science, pp. 260–269 (2008)
10. Feldman, M., Fiat, A., Leonardi, S., Sankowski, P.: Revenue maximizing envy-free multi-unit auctions with budgets. In: Proceedings of the 13th ACM Conference on Electronic Commerce (EC), pp. 532–549 (2012)
11. Fiat, A., Leonardi, S., Saia, J., Sankowski, P.: Single valued combinatorial auctions with budgets. In: Proceedings of the 12th ACM Conference on Electronic Commerce (EC), pp. 223–232 (2011)
12. Foley, D.: Resource allocation and the public sector. *Yale Econ. Essays* **7**, 45–98 (1967)
13. Goel, G., Mirrokni, V.S., Leme, R.P.: Polyhedral clinching auctions and the adwords polytope. In: Proceedings of the 44th Symposium on Theory of Computing Conference, pp. 107–122 (2012)
14. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1164–1173 (2005)
15. Milgrom, P.: Putting auction theory to work: the simultaneous ascending auction. *J. Political Econ.* **108**, 245–272 (1998)
16. Monaco, G., Sankowski, P., Zhang, Q.: Revenue maximization envy-free pricing for homogeneous resources. In: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, 25–31 July 2015, pp. 90–96 (2015)
17. Varian, H.R.: Equity, envy, and efficiency. *J. Econ. Theory* **9**, 63–91 (1974)

# Conference Program Design with Single-Peaked and Single-Crossing Preferences

Dimitris Fotakis<sup>1</sup>(✉), Laurent Gourvès<sup>2</sup>, and Jérôme Monnot<sup>2</sup>

<sup>1</sup> School of Electrical and Computer Engineering,  
National Technical University of Athens, 15780 Athens, Greece  
fotakis@cs.ntua.gr

<sup>2</sup> Université Paris-Dauphine, PSL Research University,  
CNRS UMR [7243] LAMSADE, 75016 Paris, France  
{laurent.gourves, jerome.monnot}@dauphine.fr

**Abstract.** We consider the *Conference Program Design* (CPD) problem, a multi-round generalization of (the maximization versions of)  $q$ -Facility Location and the Chamberlin-Courant multi-winner election, introduced by (Caragiannis, Gourvès and Monnot, IJCAI 2016). CPD asks for the selection of  $kq$  items and their assignment to  $k$  disjoint sets of size  $q$  each. The agents receive utility only from their best item in each set and we want to maximize the total utility derived by all agents from all sets. Given that CPD is **NP**-hard for general utilities, we focus on utility functions that are either single-peaked or single-crossing. For general single-peaked utilities, we show that CPD is solvable in polynomial time and that Percentile Mechanisms are truthful. If the agent utilities are given by distances in the unit interval, we show that a Percentile Mechanism achieves an approximation ratio  $1/3$ , if  $q = 1$ , and at least  $(2q-3)/(2q-1)$ , for any  $q \geq 2$ . On the negative side, we show that a generalization of CPD, where some items must be assigned to specific sets in the solution, is **NP**-hard for dichotomous single-peaked preferences. For single-crossing preferences, we present a dynamic programming exact algorithm that runs in polynomial time if  $k$  is constant.

## 1 Introduction

Many problems in Social Choice deal with selecting  $q$  items (or candidates), from a given set of  $m$  items, based on the preferences of  $n$  agents. In more than a few, each agent derives utility from his best item in the solution and the objective is to maximize the total utility of the agents.

An instance of this general setting is the classical  $q$ -Facility Location problem, where we place  $q$  facilities in a metric space, based on the locations suggested by  $n$  agents. Each agent uses his nearest facility in the solution and the objective is to minimize the total distance of the agents to their facilities. Facility Location has studied extensively as an optimization problem. In Social Choice, the relevant literature mostly focuses on strategic agents with single-peaked preferences

---

Supported by the projects ANR CoCoRiCo-CoDec and PSL MULTIFAC.

over the possible facility locations. The goal is to characterize the class of truthful mechanisms and to determine the best approximation ratio achievable by truthful mechanisms when the agent preferences are quantified by distances on the real line (see e.g., [10, 14, 16, 18] and the references therein).

A different specimen appears in the context of multi-winner elections. In the model introduced by Chamberlin and Courant [3], we form a committee of  $q$  representatives so as to minimize the committee’s “misrepresentation” with respect to a set of  $n$  agents. Similarly to Facility Location, each agent is associated with the committee member that represents him best, and we want to minimize the total “misrepresentation cost” of the agents. The winner determination problem for the multi-winner election of Chamberlin-Courant has received significant attention recently, with **NP**-hardness results and approximation algorithms for general agent preferences and polynomial-time algorithms for restricted preferences, such as single-peaked or single-crossing (see e.g., [1, 2, 15, 17]).

In this work, we study the *Conference Program Design* problem, or CPD in short, which was recently introduced by Caragiannis *et al.* [2] and can be regarded as a generalization of (the maximization versions of)  $q$ -Facility Location and the Chamberlin-Courant election. An instance of CPD consists of a set of  $m$  items  $X = \{x_1, \dots, x_m\}$ , a set of  $n$  agents  $L = \{1, \dots, n\}$ , each with a utility function  $u_\ell : X \rightarrow \mathbb{R}_{\geq 0}$ , and two positive integers  $k$  and  $q$ . A feasible solution  $\mathcal{S} = \{S_1, \dots, S_k\}$  is a collection of  $k$  pairwise disjoint subsets of  $X$  (or *slots*) such that each slot  $S_i$  contains at most  $q$  items. The agents derive utility only from their most preferred item in each slot and have additive utilities for different slots. Hence, the utility of an agent  $\ell$  for a solution  $\mathcal{S}$  is  $u_\ell(\mathcal{S}) = \sum_{i=1}^k \max_{x \in S_i} u_\ell(x)$ . We want to maximize the total utility of all agents, which is  $U(\mathcal{S}) = \sum_{\ell \in L} u_\ell(\mathcal{S})$  for any given solution  $\mathcal{S} = \{S_1, \dots, S_k\}$ . We underline that although a greater total utility could be achieved by assigning some items to multiple slots, we require that the slots  $S_1, \dots, S_k$  should be pairwise disjoint.

*Example.* We consider 5 items  $\{x_1, x_2, x_3, x_4, x_5\}$ , 3 agents and  $k = q = 2$ . The utility functions of the agents are  $u_1 = (4, 3, 5, 1, 2)$ ,  $u_2 = (1, 2, 3, 9, 2)$  and  $u_3 = (6, 1, 4, 0, 7)$  (the  $i$ -th coordinate in  $u_j$  denotes the utility of agent  $j$  for item  $x_i$ ). The total utility of the solution  $\mathcal{S} = (\{x_1, x_2\}, \{x_3, x_4\})$  is  $U(\mathcal{S}) = (u_1(x_1) + u_1(x_3)) + (u_2(x_2) + u_2(x_4)) + (u_3(x_1) + u_3(x_3)) = 30$ . ♦

The name of Conference Program Design is motivated by the possibility of regarding each item as a conference talk. The conference has  $q$  parallel sessions and  $k$  time slots. In each slot  $S_i$ , at most  $q$  talks are given and each agent can attend only one of them. We assume that every agent attends the talk that maximizes his utility in each slot. More generally, CPD should be regarded as an abstraction of multi-round multi-winner elections, where the set of winners in different rounds must be disjoint, each agent is represented by his most preferred winner in each round, and the utility functions of the agents are additive with respect to their representatives in different rounds (see also [2]).

**Previous Work.** CPD incorporates both  $q$ -Facility Location and the election of Chamberlin-Courant (for  $k = 1$ ): each item is a facility/candidate and the utilities are the opposite of the distance/misrepresentation costs. Since the multi-

winner election of Chamberlin-Courant and  $q$ -Facility Location [11, 15] are known to be **NP**-hard for general cost functions, CPD is **NP**-hard for general utilities. Interestingly, Caragiannis *et al.* [2] proved that CPD remains **NP**-hard (and hard to approximate) in the special case where agent utilities are either 0 or 1 (a.k.a. *uniformly dichotomous* preferences), all items fit in the solution, i.e.,  $m = kq$ , and either  $k = 2$  or  $q = 3$ . The only case where CPD is known to be polynomially solvable is for  $q = 2$ , by a reduction to maximum matching. Based on a natural Integer Linear Programming formulation, [2] obtained polynomial-time approximation algorithms for CPD with general utilities, with ratios  $1 - 1/e$ , if  $q$  is a constant, and  $1/e - 1/e^2$ , if  $q$  is part of the input.

However, many positive results are known for natural special cases of  $q$ -Facility Location and of the Chamberlin-Courant election, especially for the line metric and for single-peaked or single-crossing preferences. Specifically,  $q$ -Facility Location and its *fault tolerant* version, where each agent must connect to  $k$  different facilities, are polynomially solvable on the line [12]. As for the approximability of  $q$ -Facility Location on the line by truthful deterministic mechanisms, the Median Mechanism is optimal for  $q = 1$  [14, 16], the 2-Extremes Mechanism achieves a best possible approximation ratio of  $n - 2$  for  $q = 2$  [10, 16], and the Percentile Mechanisms comprise the only known general class of truthful deterministic mechanisms for all  $q \geq 2$  [18], but their worst-case approximation ratio cannot be bounded in terms of  $n$  and  $q$  [10] (all these mechanisms are actually known to be group strategyproof). However, all these results on the approximability of Facility Location by truthful mechanisms are about cost minimization and assume that a facility can be placed at any point on the real line. So they are not directly relevant for CPD, where we want to maximize the total utility and the item locations are restricted by the input. In a recent work, Feldman *et al.* [9] characterized the approximability of 1-Facility Location on the line metric when the potential facility locations are restricted by the input.

For the Chamberlin-Courant election, it is reasonable to assume that the agent preferences on the candidates are consistent with a placement of the candidates on a societal axis. The line metric is a special case of two popular classes of structured preferences in Social Choice, namely *single-peaked* and *single-crossing* preferences. Recent work presents polynomial-time exact algorithms for the winner determination problem of the Chamberlin-Courant election when the agent preferences are either single-peaked [1] or single-crossing [17].

**Contribution and Techniques.** Motivated by the many interesting positive results for  $q$ -Facility Location and for the Chamberlin-Courant election when the agent preferences either are determined by the line metric or are single-peaked or single-crossing, we investigate the algorithmic properties of the Conference Program Design problem for such preferences. We give an almost complete picture for CPD with single-peaked preferences and show that CPD with single-crossing preferences is polynomially solvable if the number of slots  $k$  is constant.

An interesting observation is that for single-peaked utility functions, the best  $k$  items of any agent occupy consecutive positions on the societal axis (Proposition 1). Therefore, for any set  $M$  of items,  $|M| \leq kq$ , a simple greedy

assignment of the items to slots ensures that any agent can derive utility from his best  $k$  items in  $M$ . So, we can focus on the item selection aspect of CPD for single-peaked preferences. Combining this observation with a generalization of the Linear Programming approach of [12], in Sect. 3, we show that CPD can be solved in polynomial-time for general single-peaked utility functions (Theorem 1).

In Sects. 4 and 5, we study the approximability of CPD with single-peaked preferences by truthful mechanisms. To achieve truthfulness, we exploit the idea of Percentile Mechanisms, which are known to be group strategyproof for single-peaked preferences [18, Theorem 1]. To optimize the approximation guarantees, we apply Percentile Mechanisms to the set of all tuples consisting of  $k$  consecutive items on the societal axis. We show that the extension of any single-peaked utility function on items to a utility function on tuples of  $k$  consecutive items is also single-peaked (Lemma 1). Consequently, this variant of Percentile Mechanisms is truthful (Theorem 2, we can also show that it is group strategyproof).

We analyze the approximation for the special case of linear preferences where the items and the agents lie in the unit interval  $[0, 1]$  and the utility of an agent  $\ell$  located at  $v_\ell$  for an item  $j$  located at  $x_j$  is  $u_\ell(x_j) = 1 - |v_\ell - x_j|$ . The restriction to the unit interval is wlog., since all our results hold for any interval length  $B$ , provided that the utility functions are  $u_\ell(x_j) = B - |x_j - v_\ell|$ . We first observe that if  $k = q = 1$ , the optimal solution is not truthful and any deterministic truthful mechanism must have an approximation ratio at most  $5/7$ . For  $q = 1$  and any  $k \geq 1$ , we show that the approximation ratio of the  $1/2$ -Percentile Mechanism is  $1/3$  (Lemmas 3 and 4). For any  $q \geq 2$  and  $k \geq 1$ , we show that if the number of agents is a multiple of  $q$ , the approximation ratio of the  $(\frac{1}{2q}, \frac{3}{2q}, \dots, \frac{2q-1}{2q})$ -Percentile Mechanism is at least  $(2q-3)/(2q-1)$  (Theorem 3) and at most  $(2q-1)/(2q-1/q)$ . Interestingly, the approximation ratio tends to 1, as  $q$  increases. If the number of agents  $n$  is not a multiple of  $q$ , we obtain an approximation ratio of  $(2q-3)/(2q-1) - O(q/n)$  (Theorem 4).

To the best of our knowledge, this is the first analysis of the approximation ratio of Percentile Mechanisms for linear preferences. As for the proof technique, for the general case where  $q \geq 2$ , we introduce the notion of the *width* of a subset of agents, which allows to bound the approximation ratio by analyzing independently the approximation ratio of non-overlapping groups with  $n/q$  agents each.

Nevertheless, single-peaked preferences are not enough to make CPD polynomially solvable if some items need to be assigned to specific slots. Using a reduction from PRECOLORING EXTENSION, which is known to be **NP**-complete in unit interval graphs [13], we show that this generalization of CPD is **NP**-hard if the agent utilities are single-peaked and either 0 or 1 for each item (a.k.a. *dichotomous single-peaked* preferences, see Theorem 5, in Sect. 6).

Finally, in Sect. 7, we extend the dynamic programming approach applied in [17] to the Chamberlin-Courant election with single-crossing preferences and show that CPD with single-crossing preferences can be solved in  $O(m(nq)^{k+1})$  time (Theorem 6). An interesting open question is whether CPD with single-crossing preferences is polynomially solvable if  $k$  is part of the input.

## 2 Notation and Preliminaries

CPD is introduced in Sect. 1. We introduce here some additional notation and terminology. For any integer  $p \geq 1$ , we let  $[p] = \{1, \dots, p\}$ . We write  $x \succ_\ell x'$  to denote that an agent  $\ell$  prefers item  $x$  to item  $x'$ , which happens iff  $u_\ell(x) > u_\ell(x')$ . In such cases, we write that  $\succ_\ell$  is the preference order induced by the utility function  $u_\ell$ . We always break ties in an arbitrary fixed deterministic way.

The best item of an agent  $\ell$  in a set  $Y \subseteq X$  is  $Y$ 's most valuable item to  $\ell$ , i.e.,  $\arg \max_{y \in Y} u_\ell(y)$ . We define the second,  $\dots$ , the  $k$ -th best item of an agent  $\ell$  in  $Y$  similarly. Given a set of items  $Y \subseteq X$ , and assuming that  $k = 1$ , i.e., that each agent uses a single item, we let  $u_\ell(Y) = \max_{y \in Y} \{u_\ell(y)\}$  denote the utility of an agent  $\ell$  for his best item in  $Y$ , and let  $U(Y) = \sum_{\ell=1}^n u_\ell(Y)$  denote the total utility derived by the agents from  $Y$ . Similarly, we let  $U(x) = \sum_{\ell=1}^n u_\ell(x)$  denote the total utility derived by the agents from an item  $x \in X$ .

**Conference Program Design with Item Preselection.** In Sect. 6, we consider a natural generalization of CPD, where a specified subset of items  $X' \subseteq X$  must appear in the final solution and the assignment of the items in  $X'$  to slots is fully specified by the input. We call this variant *Conference Program Design with Item Preselection*, or PRE-CPD, in short.

More formally, in addition to the input of CPD, the input of PRE-CPD includes a subset  $X' \subseteq X$  of items and a mapping  $g : X' \rightarrow [k]$ . A solution  $\mathcal{S}$  is a collection of  $k$  disjoint subsets  $S_1, \dots, S_k$  of  $X$ , such that each  $S_i$  contains at most  $q$  items and  $g^{-1}(i) = \{x \in X' : g(x) = i\} \subseteq S_i$ . In particular, we assume  $|g^{-1}(i)| \leq q$ . Thus, CPD corresponds to PRE-CPD with  $X' = \emptyset$ .

**Approximation Ratio.** An algorithm achieves an *approximation ratio* of  $\rho \in (0, 1]$ , if for any instance  $I$  of CPD, the solution  $\mathcal{S}$  computed by the algorithm satisfies  $U(\mathcal{S}) \geq \rho U(\mathcal{S}^*)$ , where  $\mathcal{S}^*$  denotes the optimal solution to instance  $I$ .

**Truthfulness.** A mechanism  $A$  for CPD is *truthful* (or *strategyproof*) if for any pair of instances  $I$  and  $I'$  that differ in the utility function of any single agent  $\ell$ , with  $u_\ell$  denoting  $\ell$ 's utility in  $I$ , we have that  $u_\ell(\mathcal{S}) \geq u_\ell(\mathcal{S}')$ , where  $\mathcal{S}$  (resp.  $\mathcal{S}'$ ) is the solution of  $A$  on instance  $I$  (resp.  $I'$ ).

A mechanism  $A$  for CPD is *group strategyproof* if for any pair of instances  $I$  and  $I'$  that differ in the utilities of any nonempty subset  $L' \subseteq L$  of agents, with  $u_1, \dots, u_n$  denoting the utility functions in  $I$ , there exists an agent  $\ell \in L'$  such that  $u_\ell(\mathcal{S}) \geq u_\ell(\mathcal{S}')$ , where  $\mathcal{S}$  (resp.  $\mathcal{S}'$ ) is the solution of  $A$  on  $I$  (resp.  $I'$ ).

**Single-Peaked Preferences.** A societal axis is a linear order  $\sqsupset$  over  $X$ . An agent's preference order  $\succ$  is *consistent* with  $\sqsupset$ , if for each three items  $x_a, x_b, x_c \in X$ ,  $((x_a \sqsupset x_b \sqsupset x_c) \vee (x_c \sqsupset x_b \sqsupset x_a)) \Rightarrow (x_a \succ x_b \Rightarrow x_b \succ x_c)$ . We say that a utility function  $u_\ell$  of an agent  $\ell \in L$  is *single-peaked* wrt axis  $\sqsupset$ , if the preference order  $\succ_\ell$  induced by  $u_\ell$  is consistent with  $\sqsupset$ . An instance of CPD is *single-peaked* (or has single-peaked utilities or preferences) wrt axis  $\sqsupset$ , if the utility functions  $u_\ell$  of all agents  $\ell \in L$  are single-peaked wrt axis  $\sqsupset$ . An instance of CPD is *single-peaked* if it is single-peaked wrt some societal axis.



One can determine in polynomial time whether a set of utility functions  $u_1, \dots, u_n$  is single-peaked (see e.g., [8]). For instances of CPD with single-peaked utilities wrt axis  $\sqsupset$ , we always index the items according to  $\sqsupset$ , i.e., we have that  $x_1 \sqsupset x_2 \sqsupset \dots \sqsupset x_m$ . We sometimes abuse the notation and use  $x_i \sqsupseteq x_j$  to denote that either  $x_i$  precedes  $x_j$  in  $\sqsupset$  or  $x_i = x_j$ .

For instances of CPD with single-peaked preferences wrt some axis  $\sqsupset$ , we say that two items  $x_i$  and  $x_j$  are consecutive, if there is no other item  $x'$  such that  $x_i \sqsupset x' \sqsupset x_j$  or  $x_j \sqsupset x' \sqsupset x_i$ . This definition extends to any number of items.

E.g., let us consider 4 items  $x_1, x_2, x_3, x_4$  and 5 agents with preferences:

|   |   |   |
|---|---|---|
| 1 : $x_1 \succ_1 x_2 \succ_1 x_3 \succ_1 x_4$ | 2 : $x_2 \succ_2 x_1 \succ_2 x_3 \succ_2 x_4$ | 3 : $x_2 \succ_3 x_3 \succ_3 x_1 \succ_3 x_4$ |
| 4 : $x_3 \succ_4 x_2 \succ_4 x_4 \succ_4 x_1$ | 5 : $x_3 \succ_5 x_4 \succ_5 x_2 \succ_5 x_1$ |   |

This set of preferences is single-peaked wrt the societal axis  $x_1 \sqsupset x_2 \sqsupset x_3 \sqsupset x_4$ . In this example, the items e.g.,  $x_1, x_2$  and  $x_3$  are consecutive.

**Optimal Item Allocation.** For instances with single-peaked preferences wrt axis  $\sqsupset$ , we can allocate any set of items  $M$ ,  $|M| = kq$ , to slots  $S_1, \dots, S_k$  in a greedy manner, so that each slot gets  $q$  items and the utility of each agent  $\ell$  is  $\max_{S \subseteq M, |S|=k} \sum_{x \in S} u_\ell(x)$ , i.e., equal to the maximum utility that  $\ell$  can derive from  $M$ . Specifically, we arrange the items in  $M$  according to  $\sqsupset$ , so that  $x_1 \sqsupset x_2 \sqsupset \dots \sqsupset x_{kq}$ , and let each slot  $S_i = \{x_i, x_{i+k}, \dots, x_{i+(q-1)k}\}$ . This ensures that any  $k$  items consecutive in  $\sqsupset$  are assigned to  $k$  different slots.

**Proposition 1.** *Let  $X$  be a set of  $m$  items arranged as  $x_1 \sqsupset x_2 \sqsupset \dots \sqsupset x_m$ , according to the societal axis  $\sqsupset$ , and let  $u : X \rightarrow \mathbb{R}_{\geq 0}$  be any utility function that is single-peaked wrt  $\sqsupset$ . Then, for any  $k \in [m]$ , the maximum utility obtained from  $k$  items in  $X$  is achieved by considering  $k$  consecutive items in  $\sqsupset$ , i.e.,  $\max_{S \subseteq X, |S|=k} \sum_{x_p \in S} u(x_p) = \max_{x_j \in X} \sum_{p=j}^{k+j-1} u(x_p)$ .*

Proposition 1 implies that for instances with single-peaked utilities, we can assume a greedy allocation of the items to slots and focus on the item selection aspect of CPD. Hence, given any set of items  $M \subseteq X$ , with  $|M| \leq kq$ , we avoid referring to any particular allocation of  $M$ . Moreover, we assume that  $|X| > kq$ , since otherwise, CPD is easily solvable.

**Linear Preferences.** An interesting special case of single-peaked preferences are *linear preferences* (or *linear utilities*), where both the items and the agents lie in  $[0, 1]$  and the utility of an agent  $\ell$  for an item  $j$  is a linear decreasing function of their distance. For such instances, we assume that the items are located at  $0 \leq x_1 < x_2 < \dots < x_m \leq 1$  and the agents are located  $0 \leq v_1 \leq v_2 \leq \dots \leq v_n \leq 1$ . The utility of an agent  $\ell$  for an item  $x_j$  is  $u_\ell(x_j) = 1 - |x_j - v_\ell|$ .

**Single-Crossing Preferences.** A profile of preferences is single-crossing if there exists an ordering of the agents, say  $\pi : [n] \rightarrow L$ , such that for every pair of items  $x_i, x_j \in X$ , either all the agents rank  $x_i$  and  $x_j$  in the same way, or there is an index  $t_{ij} \in \{1, \dots, n\}$  such that agents  $\pi(1)$  to  $\pi(t_{ij})$  all agree to rank  $x_i$  and  $x_j$

in the same way, and agents  $\pi(t_{ij} + 1)$  to  $\pi(n)$  all agree to rank  $x_i$  and  $x_j$  in the opposite way. So, either all the agents agree on the relative positions of two given items, or there is a dichotomy  $L_1, L \setminus L_1$  such that both  $L_1$  and  $L \setminus L_1$  contain consecutive agents with respect to ordering  $\pi$ . One can determine whether a preference profile is single-crossing in polynomial time [5].

E.g., let us consider 4 items  $x_1, x_2, x_3, x_4$  and 5 agents with preferences:

|   |   |   |
|---|---|---|
| 1 : $x_1 \succ_1 x_2 \succ_1 x_3 \succ_1 x_4$ | 2 : $x_1 \succ_2 x_2 \succ_2 x_4 \succ_2 x_3$ | 3 : $x_1 \succ_3 x_4 \succ_3 x_2 \succ_3 x_3$ |
| 4 : $x_4 \succ_4 x_1 \succ_4 x_2 \succ_4 x_3$ | 5 : $x_4 \succ_5 x_1 \succ_5 x_3 \succ_5 x_2$ |   |

These preferences are single-crossing, where  $\pi$  is the identity permutation of  $L$ .

### 3 CPD with Single-Peaked Preferences

**Theorem 1.** *CPD with single-peaked preferences is solvable in polynomial time.*

*Proof (sketch).* By Proposition 1, we can assume a greedy allocation of the selected items to the slots and focus on the item selection aspect of CPD. Hence, we can consider a simplified Integer Linear Programming formulation of CPD.

$$\begin{aligned}
 \text{(SLP)} \quad & \text{maximize} \quad \sum_{\ell \in N} \sum_{x_i \in X} z_{\ell i} \cdot u_{\ell}(x_i) \\
 & \text{subject to :} \quad y_i - z_{\ell i} \geq 0, \forall \ell \in N, x_i \in X \tag{1} \\
 & \quad \sum_{x_i \in X} z_{\ell i} \leq k, \forall \ell \in N \tag{2} \\
 & \quad \sum_{x_i \in X} y_i \leq k \cdot q \tag{3} \\
 & \quad y_i, z_{\ell i} \in \{0, 1\}, \forall \ell \in N, x_i \in X
 \end{aligned}$$

In (SLP), each variable  $y_i$  indicates whether an item  $x_i$  is included in the solution and each variable  $z_{\ell i}$  indicates whether an agent  $\ell$  derives utility from an item  $x_i$ . (1) ensures that an agent  $\ell$  derives utility from an item  $x_i$  only if  $x_i$  is included in the solution. (2) ensures that every agent derives utility from at most  $k$  items. (4) ensures that at most  $kq$  items are selected in the solution.

The optimum of (SLP) is equal to the optimal total utility. Let us denote by (R-SLP) the relaxation of (SLP) where the constraints  $y_i, z_{\ell i} \in \{0, 1\}$  are replaced by  $y_i, z_{\ell i} \in [0, 1]$ . Thus, the optimal value of (R-SLP) is no less than the value of (SLP). We solve (R-SLP) and let  $X' = \{x_i \in X : \exists \ell \in L \text{ with } z_{\ell i} > 0\}$ . We say that the usage of an item  $x_i \in X'$  by an agent  $\ell$  is *full* when  $z_{\ell i} = y_i$ , *null* when  $z_{\ell i} = 0$  and *intermediate* when  $0 < z_{\ell i} < y_i$ .

We can show that for any agent  $\ell$  and any two consecutive items  $x_a, x_b \in X'$  with  $x_a \succ_{\ell} x_b$ , the intermediate or null usage of  $x_a$  by  $\ell$  implies a null usage of  $x_b$  by  $\ell$ . Therefore, the items of  $X'$  for which a given agent has a non-null

usage are consecutive. Moreover, within this set of consecutive items, only the two extreme items can be used in an intermediate way. Using this observation, we modify  $X'$  as done in [12, Sect. 3], in the context of fault tolerant  $q$ -Facility Location on the line. Working as in [12], we write a new Linear Program (FLP) for the modified instance such that (i) the optimum of (FLP) is as good as the optimum of (R-SLP); and (ii) (FLP) satisfies the *consecutive ones* property, and thus, has an integral optimal solution. Obtaining an optimal selection of  $kq$  items from the solution of (FLP), we allocate the selected items using the optimal greedy allocation described in Sect. 2.  $\square$

## 4 A Truthful Mechanism for Single-Peaked Preferences

Next, we present a truthful mechanism for CPD with single-peaked preferences. Given a set  $X$  of  $m$  items arranged as  $x_1 \sqsupset x_2 \sqsupset \dots \sqsupset x_m$  on axis  $\sqsupset$ , we consider the set  $\mathcal{X} = \{C_1, \dots, C_{m-k+1}\}$  of  $k$ -tuples of consecutive items, where  $C_i = (x_i, \dots, x_{k+i-1})$  for each  $i \in [m - k + 1]$ . These  $k$ -tuples can be naturally arranged on  $\sqsupset$ , as  $C_1 \sqsupset C_2 \sqsupset \dots \sqsupset C_{m-k+1}$ , according to their first coordinate.

For each agent  $\ell$  and each  $k$ -tuple  $C_i$ , we let  $\bar{u}_\ell(C_i) = \sum_{j=i}^{k+i-1} u_\ell(x_j)$  be the utility of agent  $\ell$  for the items in  $C_i$ . Then, utilities  $\bar{u}_\ell(C_1), \dots, \bar{u}_\ell(C_{m-k+1})$  define the preference relation of agent  $\ell$  on the set  $\mathcal{X}$  of  $k$ -tuples of consecutive items. We can show that if  $u_\ell$  is single-peaked on  $X$ ,  $\bar{u}_\ell$  is single-peaked on  $\mathcal{X}$ .

**Lemma 1.** *Let  $u : X \rightarrow \mathbb{R}_{\geq 0}$  be a single-peaked utility function wrt  $x_1 \sqsupset x_2 \sqsupset \dots \sqsupset x_m$  and let  $\bar{u} : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$  be its extension on the set  $\mathcal{X} = \{C_1, \dots, C_{m-k+1}\}$  of  $k$ -tuples of consecutive items, where  $\bar{u}(C_i) = \sum_{j=i}^{k+i-1} u(x_j)$  for each  $C_i \in \mathcal{X}$ . Then,  $\bar{u}$  is single-peaked wrt  $C_1 \sqsupset C_2 \sqsupset \dots \sqsupset C_{m-k+1}$ .*

**Percentile Mechanism.** In an  $(\alpha_1, \dots, \alpha_q)$ -Percentile Mechanism for CPD, with  $0 \leq \alpha_1 < \dots < \alpha_q \leq 1$ , each agent  $\ell$  casts a vote for his best  $k$ -tuple  $C^\ell = \arg \max_{C \in \mathcal{X}} \{\bar{u}_\ell(C)\}$ . For each  $k$ -tuple  $C_i$ , we let  $\text{cnt}(C_i)$  denote the number of agents voting for  $C_i$ , i.e.,  $\text{cnt}(C_i) = \{\ell \in L : C_i = C^\ell\}$ . The mechanism selects  $q$  tuples from  $\mathcal{X}$ . For each  $j \in [q]$ , the  $k$ -tuple  $C_i \in \mathcal{X}$  is selected as the  $j$ -th tuple of the  $(\alpha_1, \dots, \alpha_q)$ -Percentile Mechanism if  $\sum_{p=1}^{i-1} \text{cnt}(C_p) < \alpha_j n \leq \sum_{p=1}^i \text{cnt}(C_p)$ .

Let  $C(1), \dots, C(q)$  be the  $k$ -tuples selected by the Percentile Mechanism, in the order of selection, and let  $M = \bigcup_{j=1}^q C(j)$  be the set of items selected by the mechanism. It may be  $|M| < kq$ , since  $C(1), \dots, C(q)$  do not need to be disjoint. The items in  $M$  are assigned greedily to slots, as explained in Sect. 2. By Proposition 1, the greedy allocation is optimal and ensures that the utility  $u_\ell(M)$  of each agent  $\ell$  from the outcome of the mechanism is best possible.  $\blacklozenge$

Using Lemma 1, we can now show that Percentile Mechanisms are truthful.

**Theorem 2.** *For any tuple  $(\alpha_1, \dots, \alpha_q)$ , with  $0 \leq \alpha_1 < \dots < \alpha_q \leq 1$ , the  $(\alpha_1, \dots, \alpha_q)$ -Percentile Mechanism is truthful for the Conference Program Design problem with single-peaked preferences.*

*Proof (sketch).* The greedy allocation of the items in  $M$  to slots ensures that all agents get a maximum utility from  $M$ . Thus, they do not have any incentive to manipulate the greedy assignment. We can also show that the agents cannot change the item selected by the mechanism in their favor. The intuition is the same as the intuition in the proofs that Generalized Median and Percentile Mechanisms [14,18] are truthful for agents with single-peaked preferences. If an agent  $\ell$  lies and votes for a  $k$ -tuple  $C'$  on the left (resp. on the right) of  $C^\ell$ , this could only push a  $k$ -tuple  $C(j)$  selected by the mechanism further on the left (resp. on the right) of  $C^\ell$ . Since agent  $\ell$  has single-peaked preferences over  $\mathcal{X}$ , such a change is not profitable for him. In fact, working as in the proof of [18, Theorem 1], we can show that Percentile Mechanisms for CPD with single-peaked preferences are group strategyproof.  $\square$

## 5 The Approximation Ratio for Linear Preferences

In this section, we analyze the approximation ratio of Percentile Mechanisms for the special case of linear preferences. The items are located at  $0 \leq x_1 < x_2 < \dots < x_m \leq 1$  and the agents are located at  $0 \leq v_1 \leq v_2 \leq \dots \leq v_n \leq 1$ . The utility of an agent  $\ell$  for an item  $x_j$  is  $u_\ell(x_j) = 1 - |x_j - v_\ell|$ .

**The Approximation Ratio for Selecting a Single Item.** We start with the case where  $k = q = 1$ . In contrast to 1-Facility Location, where the Median Mechanism is optimal (see e.g., [16]), the approximation ratio for this special case of CPD is  $1/3$  and we can show that any deterministic truthful mechanism has approximation ratio at most  $5/7$ .

In the  $1/2$ -Percentile Mechanism, each agent  $\ell$  votes for his best item, i.e. for the item  $x_j$  that minimizes  $|v_\ell - x_j|$ . We recall that  $\text{cnt}(x_j)$  is the number of agents that vote for  $x_j$ . Then, the  $1/2$ -Percentile Mechanism selects the item  $x_i$  that satisfies  $\sum_{j=1}^{i-1} \text{cnt}(x_j) < n/2 \leq \sum_{j=1}^i \text{cnt}(x_j)$ . So, because it cannot select the location of the median agent, the  $1/2$ -Percentile Mechanism selects the item closest to the location  $v_{\text{med}}$  of the median agent. The analysis of the approximation ratio is based on the following.

**Lemma 2.** *Let  $0 \leq v_1 \leq \dots \leq v_n \leq 1$  be  $n$  agent locations in  $[0, 1]$  and let  $v_{\text{med}}$  be the location of the median agent. For any items  $z, y \in [0, 1]$  with  $|v_{\text{med}} - z| \leq |v_{\text{med}} - y|$ ,  $U(y) \leq 3U(z)$ .*

We can now determine the approximation ratio for the case where  $k = q = 1$ .

**Lemma 3.** *If  $k = q = 1$ , the  $1/2$ -Percentile Mechanism achieves an approximation ratio of  $1/3$ .*

*Proof.* For the lower bound on the approximation ratio, we apply Lemma 2 with the item selected by the mechanism as  $z$  and the item selected by the optimal solution as  $y$ . Since  $|v_{\text{med}} - z| \leq |v_{\text{med}} - y|$ , Lemma 2 immediately implies that the approximation ratio of the  $1/2$ -Percentile Mechanism is at least  $1/3$ .

To conclude the proof, we present a class of instances where the mechanism has an approximation ratio of  $1/3 + \epsilon$ , for any  $\epsilon > 0$ . Such instances consist of  $n/2$  agents located at  $1/2 - \epsilon$ , where  $\epsilon > 0$  is arbitrarily small, and  $n/2$  agents located at 1, and of 2 items, one at 0 and the other at 1. The optimal solution selects the item at 1 and has a total utility of  $3n/4 - n\epsilon/2$ . The 1/2-Percentile Mechanism selects the item at 0 and has a total utility of  $n/4 + n\epsilon/2$ .  $\square$

**The Approximation Ratio for Singleton Slots.** We now use Lemma 2 and show that for  $q = 1$  and any  $k \geq 1$ , the approximation ratio of the 1/2-Percentile Mechanism is  $1/3$ . In this case, each agent  $\ell$  votes for his best  $k$ -tuple of consecutive items. The mechanism selects the  $k$ -tuple  $C_i$  that satisfies  $\sum_{j=1}^{i-1} \text{cnt}(C_j) < n/2 \leq \sum_{j=1}^i \text{cnt}(C_j)$ . Therefore, the  $k$ -tuple  $C_i$  selected by the mechanism is the best  $k$ -tuple of the median agent, i.e.,  $C_i = \arg \max_{C_j \in \mathcal{X}} \{\bar{u}_{\lfloor (n+1)/2 \rfloor}(C_j)\}$ .

**Lemma 4.** *If  $q = 1$ , for any  $k \geq 2$ , the 1/2-Percentile Mechanism on  $k$ -tuples of consecutive items achieves an approximation ratio of  $1/3$ .*

*Proof.* Let  $Z = \{z_1, \dots, z_k\}$  be the solution of the mechanism and let  $Y$  be the optimal solution. Since  $z_1, \dots, z_k$  are consecutive in  $[0, 1]$  and correspond to the  $k$  items closest to the location  $v_{\text{med}}$  of the median agent, we can arrange the items in  $Y$  as  $y_1, \dots, y_k$  so that for each  $j \in [k]$ ,  $|z_j - v_{\text{med}}| < |y_j - v_{\text{med}}|$ . Hence, Lemma 2 implies that for each pair of items  $z_j$  and  $y_j$ ,  $U(y_j) \leq 3U(z_j)$ . Since the optimal utility is  $U(Y) = \sum_{j=1}^k U(y_j)$  and the mechanism's utility is  $U(Z) = \sum_{j=1}^k U(z_j)$ , the approximation ratio is at least  $1/3$ .

Moreover, for any  $k \geq 2$ , we can generalize the tight example in the proof of Lemma 3. To this end, we consider the same agent locations and  $2k$  items,  $k$  of them are essentially collocated at 0 and  $k$  of them are essentially collocated at 1. One can verify that the approximation ratio of the 1/2-Percentile Mechanism for this class of instances can be arbitrarily close to  $1/3$ .  $\square$

**The Approximation Ratio for the General Case.** We proceed to bound the approximation of the  $(\frac{1}{2q}, \frac{3}{2q}, \dots, \frac{2q-1}{2q})$ -Percentile Mechanism for agents with linear preferences. The tight example in the proof of Lemma 3 shows that the distances  $|y - z|$  and  $v_n - v_1$  essentially determine the approximation ratio. This motivates us to introduce the notion of the *width* for a subset of agents.

We let  $L$  be a set of  $n$  agents with locations  $0 \leq v_1 \leq \dots \leq v_n \leq 1$  in  $[0, 1]$ , let  $z \in [0, 1]$  be an item, and let  $Y \subseteq [0, 1]$  be a nonempty set of items. Assuming that  $L$ ,  $z$  and  $Y$  are fixed, we denote  $y_l = \arg \max_{y \in Y \cup \{z\}} u_1(y)$  and  $y_r = \arg \max_{y \in Y \cup \{z\}} u_n(y)$  the leftmost and the rightmost items in  $Y \cup \{z\}$  used by some agent in  $L$ . The *width*  $\beta(L, z, Y)$  of the agent set  $L$  with respect to the item  $z$  and to the set  $Y$  is defined as:

$$\beta(L, z, Y) = \begin{cases} 0 & \text{if } Y \cap [y_l, y_r] \subseteq \{z\} \\ \max\{v_n - \min\{z, v_1\}, \max\{z, v_n\} - v_1\} & \text{otherwise} \end{cases}$$

Namely, if the only useful item in  $Y \cup \{z\}$  is  $z$ , the width is 0. Otherwise, the width of  $L$  is either  $v_n - v_1$ , if  $z \in [v_1, v_n]$ , or  $v_n - z$ , if  $z < v_1$ , or  $z - v_1$ , if  $z > v_n$ .

We can show that when a set of agents is partitioned into groups that occupy non-overlapping intervals in  $[0, 1]$ , the total width of all groups is at most 2.

**Lemma 5.** *Let  $L$  be a set of  $n$  agents partitioned into groups  $L^1, \dots, L^q$ , where each group consists of agents at consecutive locations. For any  $j \in [q]$ , let  $v_{\text{med}}^j$  be the location of the median agent in group  $L^j$ , and for any set  $Z$  of items, let  $z^j = \arg \min_{z \in Z} |v_{\text{med}}^j - z|$ . For any set  $Z$  with at most  $q$  items and any set  $Y$  of items with  $|v_{\text{med}}^j - z^j| \leq \min_{y \in Y} |v_{\text{med}}^j - y|$ , let  $\beta^j$  denote the width of group  $L^j$  with respect to  $z^j$  and  $Y$ . Then,  $\sum_{j=1}^q \beta^j \leq 2$ .*

The following lemma determines the approximation ratio for a group of agents  $L$  that use the same item  $z$ , as a function of the width  $\beta$ . Note that using  $\beta = 1$  and  $Y = \{y\}$ , we can obtain Lemma 2 as a special case of Lemma 6.

**Lemma 6.** *Let  $L$  be a set of  $n$  agents located at  $0 \leq v_1 \leq \dots \leq v_n \leq 1$  and let  $v_{\text{med}}$  be the location of the median agent in  $L$ . For any item  $z$  and any set of items  $Y$  such that  $|v_{\text{med}} - z| \leq \min_{y \in Y} |v_{\text{med}} - y|$ ,  $U(Y) \leq \frac{4-\beta}{4-3\beta} U(z)$ , where  $\beta \in [0, 1]$  is the width of  $L$  with respect to  $z$  and  $Y$ .*

*Proof (sketch).* We use integer division by 2 and deal with both even and odd  $n = |L|$ . Since we are interested in the ratio of  $U(Y)/U(z)$ , we focus on the set of useful items (for the agents in  $L$ ) in  $Y$ . Specifically, we assume that  $Y = (Y \cap [y_l, y_r]) \cup \{z\}$ . In case where  $Y = \{z\}$ , the lemma holds trivially, because  $\beta = 0$  and  $U(Y) = U(z)$ . So, from now on, we assume that  $\{z\} \subset Y$ .

We let  $y = \arg \min_{y' \in Y \setminus \{z\}} |v_{\text{med}} - y'|$  and consider the case where  $z < y$  (the case where  $z > y$  is symmetric). So,  $\beta = \max\{v_n - v_1, v_n - z\}$  (if  $z > y$ ,  $\beta = \max\{v_n - v_1, z - v_1\}$ ). We denote  $\delta = (y - z)/2$ . We distinguish two cases depending on whether  $z \geq v_1$  or  $z < v_1$ .

We first consider the case where  $z \geq v_1$ . For convenience, we let  $\gamma = z - v_1$ . In this case,  $\beta = v_n - v_1$ . Wlog., we assume that  $y \leq v_n$  and that  $\gamma + 2\delta \leq \beta$ . (These inequalities can be enforced if we add to  $Y$  an artificial item at  $v_n$ , which does not change the value of  $\beta$ , can only increase  $U(Y)$  and does not change  $U(z)$ .) We let  $n_1$  be the number of agents located in  $[v_1, z]$ ,  $n_2$  (resp.  $L_2$ ) denote the number (resp. the set) of agents located in  $[z, z + \delta]$ , and  $n_3$  denote the number of agents in  $(z + \delta, v_n]$ . Since  $z < y$  and  $|v_{\text{med}} - z| \leq |v_{\text{med}} - y|$ , the median agent is located in  $[v_1, z + \delta]$ . Therefore,  $n_3 \leq n/2$ . Moreover, we assume that  $n_1 \leq n/2$  (i.e., we assume that  $v_{\text{med}} \geq z$ ). Otherwise, the median agent is located on the left of  $z$  and this case is similar to the case where  $y < z$ . We have that

$$U(Y) \leq n - \sum_{j \in L_2} (z - v_j), \tag{4}$$

because each agent  $j \in L_2$  has utility at most  $1 - (z - v_j)$  for his best item in  $Y$ , while all the remaining agents have utility at most 1 for  $y$ . Similarly,

$$U(z) \geq n - n_1\gamma - \sum_{j \in L_2} (z - v_j) - n_3(\beta - \gamma), \tag{5}$$

because  $n_1$  agents have utility at least  $1 - (z - v_1) = 1 - \gamma$  for  $z$ , each agent  $j \in L_2$  has utility  $1 - (z - v_j)$  for  $z$ , and  $n_3$  agents have utility at least  $1 - (v_n - z) = 1 - (\beta - \gamma)$  for  $z$ . Using (5), we can show that for any  $\alpha \geq 1$ ,

$$\alpha U(z) + \sum_{j \in L_2} (z - v_j) \geq \alpha n - \frac{3\alpha - 1}{4} \beta n \quad (6)$$

Then, using  $\alpha = (4 - \beta)/(4 - 3\beta)$ , we obtain that  $(3\alpha - 1)\beta/4 = \alpha - 1$ . Combining this equation with (6), we conclude that  $\frac{4 - \beta}{4 - 3\beta} U(z) \geq n - \sum_{j \in L_2} (z - v_j) \geq U(Y)$ .

The analysis for the case where  $z < v_1$  follows exactly the same steps, but it is simpler, since we have  $n_1 = \gamma = 0$  in this case.  $\square$

The following provides a lower bound on the approximation ratio in case where the number of agents is a multiple of  $q$ . The proof of Theorem 3 is based on the analysis in the proof of Lemma 6 and on Lemma 5.

**Theorem 3.** *For any integers  $k \geq 1$  and  $q \geq 2$ ,  $m > qk$  items and  $qn$  agents, the approximation ratio of the  $(\frac{1}{2q}, \frac{3}{2q}, \dots, \frac{2q-1}{2q})$ -Percentile Mechanism for CPD instances with linear preferences is at least  $(2q - 3)/(2q - 1)$ .*

*Proof (idea).* The mechanism partitions the agents into groups  $L^j$ ,  $j \in [q]$ , with  $n$  consecutive agents each. Given the optimal set of items  $Y_i$  assigned to each slot  $i$ , we can determine, for each group  $L^j$ , an item  $z_i^j \in M$  (different for each slot  $i$ ) so that  $|z_i^j - v_{\text{med}}^j| \leq \min_{y \in Y_i} |y - v_{\text{med}}^j|$ , where  $v_{\text{med}}^j$  is the median of  $L^j$ . Using the width  $\beta_i^j$  of group  $L^j$  wrt  $z_i^j$  and  $Y_i$ , we can lower bound the total utility of the agents in  $L^j$  for item  $z_i^j$  by an inequality similar to (6). Then, we sum up all these inequalities and use  $\sum_{j=1}^q \beta_i^j \leq 2$ , by Lemma 5. The approximation ratio follows by an upper bound similar to (4) on the total optimal utility.  $\square$

There are instances with  $nq$  agents and  $k = 1$  where the approximation ratio of the Percentile Mechanism tends to  $(2q - 1)/(2q - 1/q)$ . E.g., for some odd integer  $n \geq 3$ , we consider  $(n - 1)/2$  agents at 0,  $n - 1$  agents at each point  $i/q$ ,  $i \in [q - 1]$ ,  $(n - 1)/2$  agents at 1, and a single agent at each point  $(2i + 1)/(2q)$ ,  $i = 0, \dots, q - 1$ . We have  $2q$  items located at points  $i/q$ ,  $i \in [q]$ , and at points  $(2i + 1)/(2q)$ ,  $i = 0, \dots, q - 1$ . The optimal solution is to select the items at  $i/q$ , for a total utility of roughly  $(n - 1)q - (n - 1)/(2q)$ . The mechanism selects the items at  $(2i + 1)/(2q)$ ,  $i = 0, \dots, q - 1$ , for a total utility of roughly  $nq - (n - 1)/2$ .

If the number of agents is not a multiple of  $q$ , we obtain a slightly weaker approximation ratio. The proof is similar to the proof of Theorem 3.

**Theorem 4.** *For any  $k \geq 1$  and  $q \geq 2$ , any  $m > qk$  and any number of agents  $|L| \geq q + 1$ , the approximation ratio of the  $(\frac{1}{2q}, \frac{3}{2q}, \dots, \frac{2q-1}{2q})$ -Percentile Mechanism for CPD instances with linear preferences is at least  $(2q - 3 - 3q/|L|)/(2q - 1 - q/|L|) = (2q - 3)/(2q - 1) - O(q/|L|)$ .*

## 6 Conference Program Design with Item Preselection

In this section, we show that despite CPD being polynomially solvable for single-peaked preferences (Theorem 1), PRE-CPD is NP-hard for single-peaked preferences, even with the additional restriction of dichotomous preferences.

The agent preferences are *dichotomous* if each agent “likes” a subset of items and “dislikes” the remaining ones. This induces a preorder with two indifference classes for every agent. In general, this implies  $u_i(x) \in \{a, b\}$  where  $a, b$  are two nonnegative reals satisfying  $a < b$ . It is called *approval-based utility* when  $a = 0$  and  $b = 1$ . Dichotomous preferences have received attention by the community of Computational Social Choice, especially in the case of committee selection rules for voters [6, 7] or in judgment aggregation [4]. In our setting, dealing with approval utilities is not restrictive with respect to algorithmic complexity issues, after a rescaling. So, we can assume that  $u_\ell(x) \in \{0, 1\}$  for  $\ell \in L$  and  $x \in X$ .

The preferences are *dichotomous single-peaked* if they are both single-peaked and dichotomous. Equivalently, the items are located on a line and each agent  $\ell$  corresponds to a closed interval  $I_\ell$ , where  $u_\ell(x) = 1$  if  $x \in I_\ell$  and  $u_\ell(x) = 0$  otherwise. This is also known as *Voter Interval* in Voting Theory [6] or *Single-Plateauedness* in majority judgments [4]. We can show that:

**Theorem 5.** *PRE-CPD is NP-hard for dichotomous single-peaked preferences.*

## 7 CPD with Single-Crossing Preferences

In this section, we consider CPD with single-crossing preferences. Wlog., we assume that the preference profile is single-crossing for the identity permutation of the agents and that agent 1 prefers  $x_i$  to  $x_j$  if and only if  $i < j$ .

We extend the dynamic programming approach applied to the Chamberlin-Courant election in [17]. We exploit the contiguous blocks property of the optimal solution of Chamberlin-Courant with single-crossing preferences [17, Lemma 5], which directly extends to CPD. For a slot  $S_j$  of a solution to CPD such that  $x_i \in S_j$ , we let  $L(j, i)$  be the set of agents who consider  $x_i$  as their best item in  $S_j$ . The *contiguous blocks property* for CPD states that for every  $j \in [k]$  and  $x_i \in S_j$ , either  $L(j, i) = \emptyset$  or there are two indices,  $t_{ji}$  and  $t'_{ji}$ , such that  $t_{ji} \leq t'_{ji}$  and  $L(j, i) = \{t_{ji}, t_{ji} + 1, \dots, t'_{ji}\}$ . Moreover, for each  $i < i'$  such that  $L(j, i) \neq \emptyset$  and  $L(j, i') \neq \emptyset$ , it holds that  $t'_{ji} < t_{ji'}$ . Namely, an item is considered as the most preferred in a slot by a set of consecutive agents and such sets of agents who prefer different items of the same slot do not overlap with each other.

**Theorem 6.** *A dynamic programming algorithm solves every single-crossing instance of CPD in  $O(m(nq)^{k+1})$  time.*

*Proof.* Let  $U(j, (i_1, t_1), \dots, (i_k, t_k))$  be the maximum total utility if we use items from set  $X_j = \{x_1, \dots, x_j\}$  only, and in each slot  $S_p$ , only the agents  $1, \dots, i_p$  are considered and only  $t_p$  items are used. The function  $U$  is defined for all  $j = 0, \dots, m$  and for all tuples  $(i_p, t_p)$  with  $t_1 + \dots + t_k \leq j$ . If  $j = 0$ ,  $X_0 = \emptyset$ .



We start with  $U(0, (i_1, t_1), \dots, (i_k, t_k)) = 0$ , for all pairs  $(i_1, t_1), \dots, (i_k, t_k)$ . For each  $j \geq 0$ , the next item  $x_{j+1}$  either is not selected (provided that  $t_1 + \dots + t_k \leq j$ ), in which case  $U(j + 1, (i_1, t_1), \dots, (i_k, t_k)) = U(j, (i_1, t_1), \dots, (i_k, t_k))$ , or it is assigned to some slot  $S_p$ , in which case  $U(j + 1, (i_1, t_1), \dots, (i_k, t_k)) =$

$$\max_{0 \leq \ell \leq i_p} \left\{ U(j, (i_1, t_1), \dots, (\ell, t_p - 1), \dots, (i_k, t_k)) + \sum_{\nu=\ell+1}^{i_p} u_\nu(x_{j+1}) \right\}$$

Therefore, for each  $j \geq 0$  and each fixed  $(i_1, t_1), \dots, (i_k, t_k)$ , with  $t_1 + \dots + t_k \leq j + 1$ ,  $U(j + 1, (i_1, t_1), \dots, (i_k, t_k))$  can be defined recursively as follows:

$$\max \left\{ U(j, (i_1, t_1), \dots, (i_k, t_k)) \right. \\ \left. \max_{1 \leq p \leq k} \max_{0 \leq \ell \leq i_p} \left\{ U(j, \dots, (\ell, t_p - 1), \dots) + \sum_{\nu=\ell+1}^{i_p} u_\nu(x_{j+1}) \right\} \right\}$$

in case where  $t_1 + \dots + t_k \leq j$ , or

$$\max_{1 \leq p \leq k} \max_{0 \leq \ell \leq i_p} \left\{ U(j, \dots, (\ell, t_p - 1), \dots) + \sum_{\nu=\ell+1}^{i_p} u_\nu(x_{j+1}) \right\}$$

in case where  $t_1 + \dots + t_k = j + 1$ . The optimal solution is given by  $U(m, (n, k), \dots, (n, k))$ . The number of values that we need to compute is  $O(m(nq)^k)$  and the total running time is  $O(m(nq)^{k+1})$ .  $\square$

## References

1. Betzler, N., Slinko, A., Uhlmann, J.: On the computation of fully proportional representation. *J. Artif. Intell. Res. (JAIR)* **47**, 475–519 (2013)
2. Caragiannis, I., Gourvès, L., Monnot, J.: Achieving proportional representation in conference programs. In: *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*, pp. 144–150 (2016)
3. Chamberlin, J.R., Courant, P.N.: Representative deliberations and representative decisions: proportional representation and the Borda rule. *Am. Polit. Sci. Rev.* **77**(3), 718–733 (1983)
4. Dietrich, F., List, C.: Majority voting on restricted domains. *J. Econ. Theory* **145**(2), 512–543 (2010)
5. Elkind, E., Faliszewski, P., Slinko, A.M.: Clone structures in voters’ preferences. In: *Proceedings of the 13th ACM Conference on Electronic Commerce (EC 2012)*, pp. 496–513 (2012)
6. Elkind, E., Lackner, M.: Structure in dichotomous preferences. In: *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pp. 2019–2025 (2015)
7. Erdélyi, G., Lackner, M., Pfandler, A.: Manipulation of  $k$ -approval in nearly single-peaked electorates. In: Walsh, T. (ed.) *ADT 2015. LNCS (LNAI)*, vol. 9346, pp. 71–85. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-23114-3\\_5](https://doi.org/10.1007/978-3-319-23114-3_5)

8. Escoffier, B., Lang, J., Öztürk, M.: Single-peaked consistency and its complexity. In: Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008), pp. 366–370. IOS Press (2008)
9. Feldman, M., Fiat, A., Golomb, I.: On voting and facility location. In: Proceedings of the 17th ACM Conference on Economics and Computation (EC 2016), pp. 269–286 (2016)
10. Fotakis, D., Tzamos, C.: On the power of deterministic mechanisms for facility location games. *ACM Trans. Econ. Comput.* **2**(4), 15:1–15:37 (2014)
11. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York (1979)
12. Hajiaghayi, M., Hu, W., Li, J., Li, S., Saha, B.: A constant factor approximation algorithm for fault-tolerant  $k$ -median. In: Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA 2014), pp. 1–12 (2014)
13. Marx, D.: Precoloring extension on unit interval graphs. *Discrete Appl. Math.* **154**(6), 995–1002 (2006)
14. Moulin, H.: On strategy-proofness and single-peakedness. *Public Choice* **35**, 437–455 (1980)
15. Procaccia, A., Rosenschein, J., Zohar, A.: On the complexity of achieving proportional representation. *Soc. Choice Welfare* **30**(3), 353–362 (2008)
16. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. *ACM Trans. Econ. Comput.* **1**(4), 18 (2013)
17. Skowron, P., Yu, L., Faliszewski, P., Elkind, E.: The complexity of fully proportional representation for single-crossing electorates. *Theoret. Comput. Sci.* **569**, 43–57 (2015)
18. Sui, X., Boutilier, C., Sandholm, T.: Analysis and optimization of multi-dimensional percentile mechanisms. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013), pp. 367–374 (2013)

# Truthful Facility Assignment with Resource Augmentation: An Exact Analysis of Serial Dictatorship

Ioannis Caragiannis<sup>1</sup>, Aris Filos-Ratsikas<sup>2(✉)</sup>,  
Søren Kristoffer Stiil Frederiksen<sup>3</sup>, Kristoffer Arnsfelt Hansen<sup>3</sup>, and Zihan Tan<sup>4</sup>

<sup>1</sup> Department of Computer Engineering and Informatics, University of Patras,  
Patras, Greece

`caragian@ceid.upatras.gr`

<sup>2</sup> Department of Computer Science, University of Oxford, Oxford, UK

`Aris.Filos-Ratsikas@cs.ox.ac.uk`

<sup>3</sup> Department of Computer Science, Aarhus University, Aarhus, Denmark

`sorensf@gmail.com, arnsfelt@cs.au.dk`

<sup>4</sup> Department of Computer Science, University of Chicago, Chicago, USA

`zihantan@uchicago.edu`

**Abstract.** We study the *truthful facility assignment* problem, where a set of agents with private most-preferred points on a metric space are assigned to facilities that lie on the metric space, under capacity constraints on the facilities. The goal is to produce such an assignment that minimizes the social cost, i.e., the total distance between the most-preferred points of the agents and their corresponding facilities in the assignment, under the constraint of truthfulness, which ensures that agents do not misreport their most-preferred points.

We propose a *resource augmentation framework*, where a truthful mechanism is evaluated by its worst-case performance on an instance with enhanced facility capacities against the optimal mechanism on the same instance with the original capacities. We study a well-known mechanism, Serial Dictatorship, and provide an exact analysis of its performance. Among other results, we prove that Serial Dictatorship has approximation ratio  $g/(g-2)$  when the capacities are multiplied by any integer  $g \geq 3$ . Our results suggest that even a limited augmentation of the resources can have wondrous effects on the performance of the mechanism and in particular, the approximation ratio goes to 1 as the augmentation factor becomes large. We complement our results

---

Ioannis Caragiannis was partially supported by a Caratheodory research grant E.114 from the University of Patras. Aris Filos-Ratsikas was partially supported by the COST Action IC1205 on “Computational Social Choice” and by the ERC Advanced Grant 321171 (ALGAME). Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Kristoffer Arnsfelt Hansen acknowledge support from the Danish National Research Foundation and The National Science Foundation of China (under the grant 61361136003) for the Sino-Danish Center for the Theory of Interactive Computation and from the Center for Research in Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council.

with bounds on the approximation ratio of Random Serial Dictatorship, the randomized version of Serial Dictatorship, when there is no resource augmentation.

## 1 Introduction

We study the *facility assignment problem*, in which there is a set of agents and a set of *facilities* with finite capacities; facilities are located on a metric space at points  $F_i$  and each agent has a most-preferred point  $A_i$ , which is her private information. The goal is to produce an *assignment* of agents to facilities, such that no capacity is exceeded and the sum of distances between agents and their assigned facilities, *the social cost*, is minimized. A *mechanism* is a function that elicits the points  $A_i$  from the agents and outputs an assignment. We will be interested in *truthful* mechanisms, i.e., mechanisms that do not incentivize agents to misreport their most-preferred locations and we will be aiming to find mechanisms that achieve a social cost as close as possible to that of the optimal assignment when applied to the true points  $A_i$  of the agents. Our setting has various applications such as assigning patients to personal GPs, vehicles to parking spots, children to schools and pretty much any matching environment where there is some notion of distance involved.

Our work falls under the umbrella of *approximate mechanism design without money*, a term coined by Procaccia and Tennenholtz [16] to describe problems where some objective function is optimized under the hard constraints imposed by the requirement of truthfulness. The standard measure of performance for truthful mechanisms is the *approximation ratio*, which for our objective, is the worst-case ratio between the social cost of the truthful mechanism in question over the minimum social cost, calculated over all input instances of the problem.

However, it is arguably unfair to compare the performance of a mechanism that is severely limited by the requirement of truthfulness to that of an omnipotent mechanism that operates under no restrictions and has access to the real inputs of the agents, without giving the truthful mechanism any additional capabilities. This is even more evident in general settings, where strong impossibility results restrict the performance of all truthful mechanisms to be rather poor. The need for a departure from the worst-case approach has been often advocated in the literature, but the suggestions mainly involve some average case analysis or experimental evaluations.

Instead, we will adopt a different approach, that has been made popular in the field of online algorithms and competitive analysis [13, 17]; the approach suggests enhancing the capabilities of the mechanism operating under some very limiting requirement (such as truthfulness or lack of information) before comparing to the optimal solution. Our main conceptual contribution is the adoption of a *resource augmentation approach to approximate mechanism design*. In the resource augmentation framework, we evaluate the performance of a truthful mechanism on an input with additional resources, when compared to the optimal solutions for the set of original resources. For our problem, we consider the

social cost achievable by a truthful mechanism on some input with augmented facility capacities against the optimal assignment under the original capacities given as input.

More precisely, let  $I$  be an input instance to the facility assignment problem and let  $I_g$  be the same instance where each capacity has been multiplied by some integer constant  $g$ , that we call the *augmentation factor*. Then, the *approximation ratio with augmentation  $g$*  of a truthful mechanism  $M$  is the worst-case ratio of the social cost achievable by  $M$  on  $I_g$  over the social cost of the optimal assignment on  $I$ , over all possible inputs of the problem. The idea is that if the ratio achievable by a mechanism with small augmentation is much better when compared to the standard approximation ratio, it might make sense to invest in additional resources. At the same time, such a result would imply that the set of “bad” instances in the worst-case analysis is rather pathological and not very likely to appear in practice. To the best of our knowledge, this is the first time that such a resource augmentation framework has been explicitly proposed in algorithmic mechanism design.

## 1.1 Our Results

As our main contribution, we study the well-known truthful mechanisms for assignment problems, Serial Dictatorship (SD) and Random Serial Dictatorship (RSD). For SD, we provide an *exact analysis*, obtaining tight bounds on the approximation ratio of the mechanism for all possible augmentation factors  $g$ . Specifically, we prove that when  $n$  is the number of agents, while without any augmentation, the approximation ratio of SD is  $2^n - 1$ , the approximation ratio with augmentation factor  $g = 2$  is exactly  $\log(n + 1)$  whereas for  $g \geq 3$ , the approximation ratio is  $g/(g - 2)$ , i.e., a small constant. In particular, our results imply that as the augmentation factor becomes large, the approximation ratio of SD with augmentation goes to 1 and the convergence is rather fast. Our results for SD improve and extend some results in the field of online algorithms [12].

To prove the approximation ratios for all augmentation factors, we use an interesting technique based on linear programming. Specifically, we first provide a directed graph interpretation of the assignment produced by SD and the optimal assignment, and then prove that the worst-case instances appear on  $g$ -trees, i.e., trees where (practically) every vertex has exactly  $g$  successors. Then, we formulate the problem of calculating the worst ratio on such trees as a linear program and bound the ratio by obtaining feasible solutions to its dual. Such a solution can be seen as a “path covering” of the assignment graph and we obtain the bounds by constructing appropriate path coverings of low cost.

We also consider randomized mechanisms and the very well-known Random Serial Dictatorship mechanism. We prove that for augmentation factor 1 (i.e., no resource augmentation), the approximation ratio of the mechanism is between  $n^{0.26}$  and  $n$ ; the result suggests that even a small augmentation ( $g = 2$ ) is a more powerful tool than randomization.

## 1.2 Related Work

Assignment problems are central in the literature of economics and computer science. The literature on one-sided matchings dates back to the seminal paper by Hylland and Zeckhauser [10] and includes many very influential papers [5, 18] in economics as well as a rich recent literature in computer science [2, 8, 9, 15]. Serial Dictatorships (or their randomized counterparts) have been in the focus of much of this literature, mainly due to their simplicity and the fragile nature of truthfulness, which makes it quite hard to construct more involved truthful mechanisms. In a celebrated result, Svensson [18] characterized a large class of truthful mechanisms by serial dictatorships. Random Serial Dictatorship has also been extensively studied [1, 15] and recently it was proven [8] that is asymptotically the best truthful mechanism for one-sided matchings under the general cardinal preference domain.

The facility assignment problem can be interpreted as a matching problem; somewhat surprisingly, matching problems in metric spaces have only recently been considered in the mechanism design literature. Emek et al. [7] study a setting very closely related to ours, where the goal is to find matchings on metric spaces, but they are interested in how well a mechanism that produces a stable matching can approximate the cost of the optimal matching. In a conceptually similar work, Anshelevich and Shreyas [3] study the performance of *ordinal* matching mechanisms on metric spaces, when the limitation is the lack of information. The fundamental difference between those works and ours is that we consider truthful mechanisms and bound their performance due to the truthfulness requirement; to the best of our knowledge, this is the first time where truthful mechanisms have been considered in a matching setting with metric preferences. Another difference between our work and the aforementioned papers is that they do not consider resource augmentation and only bound the performance of mechanisms on the same set of resources.<sup>1</sup> However, given the generality of the augmentation framework, the same idea could be applied to their settings. In that sense, our paper proposes a *resource augmentation approach to algorithmic mechanism design* that could be adopted in most resource allocation and assignment settings.

As we mentioned earlier, the idea of resource augmentation was popularized by the field of online algorithms and competitive analysis and is tightly related to the literature on *weak adversaries* where an online competitive algorithm is compared to the adversary that uses a smaller number of resources. The idea for this approach originated in the seminal paper by Sleator and Tarjan [17] and has been adopted by others ever since [14, 19]; the term “resource augmentation” was explicitly introduced by Kalyanasundaram and Pruhs [13].

Most closely related to our problem is the *online transportation problem* [12] (also known as the minimum online metric bipartite matching). In particular, results about the greedy algorithm in the online transportation problem imply bounds for the facility assignment problem. However, contrary to [12], our analysis is *exact*, i.e. our results involve no asymptotics. Furthermore, compared to the

<sup>1</sup> With the exception of the bi-criteria result in [3].

related result in [12], we remark that our analysis is substantially different due to the use of linear programming; our primal-dual technique could be applicable for greedy assignment mechanisms on other resource augmentation settings, beyond the problem studied here. For a detailed discussion of the connection between the two settings, the reader might refer to the full version of this paper.

Finally, there is some resemblance between our problem and the facility location problem [16] that has been studied extensively in the literature of approximate mechanism design, in the sense that in both settings, agents specify their most preferred positions on a metric space. Note that the settings are fundamentally different however, since in the facility location problem, the task is to identify the appropriate point to locate a facility whereas in our setting, facilities are already in place and we are looking for an assignment of agents to them.

## 2 Preliminaries

In the *facility assignment* problem, there is a set  $N = \{1, \dots, n\}$  of agents and a set  $M = \{1, \dots, m\}$  of *facilities*, where agents and facilities are located on a metric space, equipped with a distance function  $d$ . Each facility has a *capacity*  $c_i \in \mathbb{N}_+$ , which is the number of agents that the facility can accommodate. We assume that  $\sum_{i=1}^m c_i \geq n$ , i.e., all agents can be accommodated by some facility. Each agent has a most preferred position  $A_i$  on the space and his cost  $d_i(j)$  from facility  $j$  is the distance  $d(A_i, F_j)$  between  $A_i$  and the position  $F_j$  of the facility. Let  $A = (A_1, \dots, A_n)$  be a vector of preferred positions and call it a *location profile*. Let  $F = (F_1, \dots, F_m)$  be the corresponding set of points of the facilities. A pair of agents' most preferred points and facility points  $(A, F)$  is called an *instance* of the facility assignment problem and is denoted by  $I$ .

The locations of the facilities are known but the location profiles are not known; agents are asked to report them to a central planner, who then decides on an *assignment*  $S$ , i.e., a pairing of agents and facilities such that no agent is assigned to more than one facility and no facility capacity is exceeded. Let  $S_i$  be the restriction of the assignment to the  $i$ 'th coordinate, i.e., the facility to which agent  $i$  is assigned in  $S$  and let  $\mathcal{S}$  be the set of all assignments. The *social cost* of an assignment  $S$  on input  $I$  is the sum of the agents' costs from their facilities assigned by  $S$  i.e.,  $\sum_{i=1}^n d_i(S_i)$ . A deterministic mechanism maps instances to assignments whereas a randomized mechanism maps instances to probability distributions over assignments.

A mechanism is *truthful* if no agent has an incentive to misreport his most preferred location. Formally, this is guaranteed when for every location profile  $A$ , any report  $A'_i$ , and any reports  $A_{-i}$  of all agents besides agent  $i$ , it holds that  $d_i(S_i) \geq d_i(S'_i)$ , where  $S = M(I)$  and  $S' = M(I')$ , with  $I = (A, F)$  and  $I' = ((A'_i, A_{-i}), F)$ . For randomized mechanisms, the corresponding notion is *truthfulness-in-expectation*, where an agent can not decrease her expected distance from the assigned facilities by deviating, i.e., it holds that  $\mathbb{E}_{S \sim D}[d_i(S_i)] \geq \mathbb{E}_{S \sim D'}[d_i(S_i)]$ , where  $D$  and  $D'$  are the probability distributions output by the mechanism on inputs  $I$  and  $I'$  respectively. A stronger notion of

truthfulness for randomized mechanisms is that of *universal truthfulness*, which guarantees that for every realization of randomness, there will not be any agent with an incentive to deviate. Alternatively, one can view a universally truthful mechanism as a mechanism that runs a deterministic truthful mechanism at random, according to some distribution.

As our main conceptual contribution, we will consider a *resource augmentation* framework where the minimum social cost of any assignment will be compared with the social cost achievable by a mechanism on a location profile with augmented facility capacities. Given an instance  $I$ , we will use the term *g-augmented instance* to refer to an instance of the problem where the input is  $I$  and the facility of each capacity has been multiplied by  $g$ . We will denote that instance by  $I_g$  and we will call  $g$  the *augmentation factor* of  $I$ . For example, when  $g = 2$ , we will compare the minimum social cost with the social cost of a mechanism on the same inputs but with double capacities.

For the facility assignment problem, the optimal mechanism computes a minimum cost matching (which can be computed using an algorithm for maximum weight bipartite matching) and it can be easily shown that it is not truthful; in order to achieve truthfulness, we have to output suboptimal solutions. As performance measure, we define the *approximation ratio with augmentation* of a mechanism  $M$  as

$$ratio_g(M) = \sup_I \frac{SC_M(I_g)}{SC_{OPT}(I)}$$

where  $SC_M(I_g) = \sum_{i=1}^n d_i(M(I_g)_i)$  is the social cost of the assignment produced by mechanism  $M$  on input instance  $I$  with augmentation factor  $g$  and  $SC_{OPT}(I)$  is the minimum social cost of any assignment on  $I$  i.e.,  $SC_{OPT}(I) = \min_{S \in \mathcal{S}} \sum_{i=1}^n d_i(S_i)$ . For randomized mechanisms, the definitions involve the expected social cost and are very similar. Obviously, if we set  $g = 1$ , we obtain the standard notion of the approximation ratio for truthful mechanisms [16]. For consistency with the literature, we will denote  $ratio_1(M)$  by  $ratio(M)$ .

We will be interested in two natural truthful mechanisms that assign agents to facilities in a greedy nature. A *serial dictatorship* (SD) is a mechanism that first fixes an ordering of the agents and then assigns each agent to his most preferred facility, from the set of facilities with non-zero residual capacities. Its randomized counterpart, *Random Serial Dictatorship* (RSD), is the mechanism that first fixes the ordering of agents uniformly at random and then assigns them to their favorite facilities that still have capacities left. In other words, RSD runs one of the  $n!$  possible serial dictatorships uniformly at random and hence it is universally truthful.

### 3 Approximation Guarantees for Serial Dictatorships

In this section we provide our main results, the upper bounds on the approximation ratio with augmentation of Serial Dictatorship, for all possible augmentation factors. In Sect. 4, we state the theorem that ensures that the bounds proven here



are tight. At the end of the section, we also consider Random Serial Dictatorship, when there is no resource augmentation.

**Theorem 1.** *The approximation ratio of SD with augmentation factor  $g$  in facility assignment instances with  $n$  agents is*

1.  $\text{ratio}(SD) \leq 2^n - 1$ ,
2.  $\text{ratio}_2(SD) \leq \log(n + 1)$ ,
3.  $\text{ratio}_g(SD) \leq \frac{g}{g-2}$  when  $g \geq 3$ .

In order to prove the theorem,<sup>2</sup> we first need to introduce a different interpretation of the assignment produced by SD and the optimal assignment, in terms of a directed graph. We begin with a roadmap of the proof of Theorem 1.

1. We show how to represent an instance of facility assignment together with an optimal solution and a solution computed by the SD mechanism as a directed graph and argue that the instances in which the SD mechanism has the worst approximation ratio are specifically structured as directed trees.
2. We observe that the cost of the SD mechanism in these instances is upper-bounded by the objective value of a maximization linear program defined over the corresponding directed trees.
3. We use duality to upper-bound the objective value of this LP by the value of a feasible solution for the dual LP. This reveals a direct relation of the approximation ratio of the SD mechanism to a graph-theoretic quantity defined on a directed tree, which we call the cost of a path covering.
4. Our last step is to prove bounds on this quantity; these might be of independent interest and could find applications in other contexts.

Consider an instance  $I$  of facility assignment. Recall the interpretation of the problem as a metric bipartite matching and note that without loss of generality, each facility can be assumed to have capacity 1 and  $m \geq n$ . Unless otherwise specified, agents and facilities are identified by the integers in  $[n]$  and  $[m]$ , respectively.

Now, let  $O$  be any assignment on input  $I$ , and let  $S$  be an assignment returned by the SD mechanism when applied on the instance  $I_g$  (where each facility has capacity  $g$ ). We use a directed graph to represent the triplet  $I$ ,  $O$ , and  $S$  as follows. The graph has a node for each facility. Each directed edge corresponds to an agent. A directed edge from a node corresponding to facility  $j_1$  to a node corresponding to facility  $j_2$  indicates that the agent corresponding to the edge is assigned to facility  $j_1$  in  $O$  and facility  $j_2$  in  $S$ . Observe that there is at most one edge outgoing from each node; this edge corresponds to the agent that is assigned to the facility corresponding to the node in solution  $O$ . Furthermore, a

---

<sup>2</sup> We point out here that statement 1 and a weaker version of statement 2 in Theorem 1 can be obtained as corollaries of results in the literature for the online transportation problem (see [11, 12]). However, we will prove the three statements of Theorem 1 as part of our more general framework.

node may have up to  $g$  incoming edges, corresponding to agents assigned to the facility by the SD mechanism.

Representations as directed  $g$ -trees are of particular importance. A *directed  $g$ -tree*  $T$  is an acyclic directed graph that has a root node  $r$  of in-degree 1 and out-degree 0, leaves with in-degree 0 and out-degree 1, and intermediate nodes with in-degree  $g$  and out-degree 1. We now show that it suffices to restrict our attention to directed  $g$ -trees as graph representations of instances in which the SD mechanism achieves its worst performance.

**Lemma 1.** *Given a instance  $I$  with  $n$  agents, an optimal solution  $O$  for  $I$  and a solution  $S$  consistent with the SD mechanism when applied to instance  $I_g$ , there is another instance  $I'$  with at most  $n$  agents, with an optimal solution  $O'$  and a solution  $S'$  consistent with the application of the SD mechanism on the instance  $I'_g$  such that the representation graph of the triplet  $(I', O', S')$  is a directed  $g$ -tree and such that*

$$\frac{\text{cost}(S, I_g)}{\text{cost}(O, I)} \leq \frac{\text{cost}(S', I'_g)}{\text{cost}(O', I')}.$$

*Proof.* Let  $o_i$  and  $s_i$  denote the facility to which agent  $i$  is connected in assignments  $O$  and  $S$ , respectively. We say that agent  $i$  is *optimal* if  $o_i = s_i$ . We say that agent  $i$  is *greedy* if  $s_i \neq o_i$  and less than  $g$  agents are assigned to facility  $o_i$  when SD decides the assignment of agent  $i$ . This means that  $d(A_i, F_{s_i}) \leq d(A_i, F_{o_i})$ . We say that agent  $i$  is *blocked* if  $g$  agents are already assigned to facility  $o_i$  when SD decides the assignment of agent  $i$ .

Starting from  $(I, O, S)$ , we construct a new triplet  $(I', O', S')$  as follows:

- First, we remove all optimal agents. This corresponds to removing loops from the representation graph.
- Then, we repeat the following process as long as there exists a blocked agent  $i$  that is connected under  $S$  to a facility  $j$  that is the optimal facility of a greedy agent. In this case, we introduce a new facility  $j'$  at point  $F_{j'}$  such that  $d(A_i, F_{j'}) = d(A_i, F_j)$  and  $d(F_{j'}, X) = d(A_i, F_{j'}) + d(A_i, X)$  for every other point  $X$  of the space. The second equality guarantees that the set of all points corresponding to locations of agents and facilities that have survived and the newly introduced point  $F_{j'}$  is a metric. This can easily be achieved by placing the new facility  $j'$  such that it coincides with  $j$  on the metric space. We assign agent  $i$  to facility  $j'$  instead of  $j$ ; by the first equality above, this is consistent to the definition of the SD mechanism. In the representation graph, this step adds a new node corresponding to the new facility  $j'$  and modifies the directed edge corresponding to blocked agent  $i$  so that it is directed to the new node.
- Then, we remove all greedy agents that are not connected under  $S$  to optimal facilities of blocked agents together with their optimal facilities.
- Then, for each facility  $j$  that is used by  $t \geq 2$  agents  $i_1, i_2, \dots, i_t$  in  $S$  but is not used by any agent in  $O$ , we remove facility  $j$  and introduce  $t$  new facilities  $j_1, j_2, \dots, j_t$  such that  $d(A_{i_k}, j_k) = d(A_{i_k}, j)$  for  $k = 1, \dots, t$  and  $d(X, j_k) = d(X, A_{i_k}) + d(A_{i_k}, j_k)$  for every other point  $X$  of the space.

Again, the second equality guarantees that the set of all points corresponding to locations of agents and facilities that have survived and the newly introduced points  $F_{j_1}, \dots, F_{j_t}$  is a metric. For  $k = 1, \dots, t$ , we assign agent  $i_k$  to facility  $j_k$ ; by the first equality above, this is consistent to the definition of the SD mechanism. In the representation graph, this step adds  $t$  nodes corresponding to the new facilities  $j_1, \dots, j_t$  and, for  $k = 1, \dots, t$ , it modifies the directed edge corresponding to blocked agent  $i_k$  so that it is directed to the new node  $j_k$ , and removes node corresponding to facility  $j$ .

- Finally, we remove any facility that is not used by any of the non-removed agents in any of the two solutions.

We denote by  $I'$  the resulting instance and by  $O'$  the restriction of  $O$  to the survived agents. Also,  $S'$  is the assignment obtained by the modification of  $S$  and considering the survived agents only. We remark that the representation graph of  $(I', O', S')$  is a forest of directed  $g$ -trees. Indeed, the optimal facility of a greedy agent is not used by any agent in  $S'$ ; the corresponding node is a leaf in the representation graph. Now, assume that the representation graph contains a directed cycle; this should consist of directed edges corresponding to blocked agents. By the definition above, this would mean that, for every agent  $j$  in this cycle, the assignment of all agents that were assigned by the SD mechanism to the optimal facility  $o_j$  took place before the assignment of agent  $j$  to a facility; this yields a contradiction and no such cycle exists. The optimal facility of a blocked agent has out-degree 1 and in-degree  $g$ . Nodes with zero out-degree have degree exactly 1; these are nodes corresponding to the newly added facilities and serve as roots of the directed  $g$ -trees.

Let  $R$  be the set of (greedy and optimal) agents removed and observe that  $d(A_i, F_{s_i}) \leq d(A_i, F_{o_i})$  for each such agent  $i \in R$ . Hence, it is

$$\begin{aligned} \frac{\text{cost}(S, I_g)}{\text{cost}(O, I)} &= \frac{\sum_{i \in [n]} d(A_i, F_{s_i})}{\sum_{i \in [n]} d(A_i, F_{o_i})} \leq \frac{\sum_{i \in [n]} d(A_i, F_{s_i}) - \sum_{i \in R} d(A_i, F_{s_i})}{\sum_{i \in [n] \setminus R} d(A_i, F_{o_i}) - \sum_{i \in R} d(A_i, F_{o_i})} \\ &= \frac{\sum_{i \in [n] \setminus R} d(A_i, F_{s'_i})}{\sum_{i \in [n] \setminus R} d(A_i, F_{o'_i})}. \end{aligned}$$

Clearly, if the representation of triplet  $I', O', S'$  consists of more than one  $g$ -trees, there is an instance  $I''$  and assignments  $O''$  and  $S''$  corresponding to the restriction of  $(I', O', S')$  in one of the  $g$ -trees which satisfies  $\frac{\text{cost}(S, I_g)}{\text{cost}(O, I)} \leq \frac{\text{cost}(S'', I''_g)}{\text{cost}(O'', I'')}$ . If  $O''$  is indeed an optimal solution for instance  $I''$ , the proof is complete. Otherwise, we repeat the whole process using instance  $I''$  as  $I$ , solution  $O$  to be the optimal solution for instance  $I''$ , and the SD solution  $S''$  until the solution  $O''$  obtained is optimal for the  $g$ -tree instance obtained at the final step (this condition will eventually be satisfied as the optimal cost decreases in each application of the process). By setting  $\tilde{I} = I''$ ,  $\tilde{O} = O''$ , and  $\tilde{S} = S''$  will then yield the triplet with the desired characteristics.  $\square$

So, in the following, we will focus on triplets  $(I, O, S)$  of a facility assignment instance  $I$  with at most  $n$  agents, with an optimal solution  $O$ , and with an SD

solution  $S$  for instance  $I_g$  that have a graph representations as a directed  $g$ -tree  $T$ . Below, we use  $\mathcal{P}$  to denote the set of all paths that originate from leaves. Given an edge  $e$  of a  $g$ -tree, we use  $\mathcal{P}_e$  (respectively,  $\tilde{\mathcal{P}}_e$ ) to denote the set of all paths that originate from a leaf and cross (respectively, terminate with) edge  $e$ . We always use  $e_r$  to denote the edge incident to the root of a  $g$ -tree.

Our next observation is that  $\text{cost}(S, I_g)$  is upper-bounded by the objective value of the following linear program.

$$\begin{aligned} & \text{maximize } \sum_{e \in T} z_e \\ & \text{subject to: } z_e - \sum_{a \in p \setminus \{e\}} z_a \leq \sum_{a \in p} d(A_a, F_{o_a}), e \in T, p \in \tilde{\mathcal{P}}_e \\ & z_e \geq 0, e \in T \end{aligned}$$

To see why, interpret variable  $z_e$  as the distance of agent corresponding to edge  $e$  of  $T$  to the facility it is connected to under assignment  $S$ . Then, clearly, the objective  $\sum_{e \in T} z_e$  represents  $\text{cost}(S, I_g)$ . Now, how high can  $\text{cost}(S, I)$  be? The LP essentially answers this question (partially, because it does not use all constraints of the SD mechanism but sufficiently for our purposes). In particular, the LP takes into account the fact that the distance of agent  $e$  to the facility to which it is connected in  $S$  is not higher than the distance from the agent to any leaf facility in its subtree; this follows by the definition of the SD mechanism since leaf facilities are by definition available throughout the execution of the SD mechanism. Indeed, consider agent  $e$  and a path  $p \in \tilde{\mathcal{P}}_e$ . Since agent  $e$  is connected to facility  $s_e$  under SD and not to the facility corresponding to the leaf from which path  $p$  originates from, this means that the distance  $d(A_e, F_{s_e})$  is not higher than the distance of  $A_e$  from the location of the facility corresponding to that leaf. Since  $d$  is a metric, this distance is at most  $d(A_e, F_{o_e}) + \sum_{a \in p \setminus \{e\}} d(F_{s_a}, F_{o_a}) \leq d(A_e, F_{o_e}) + \sum_{a \in p \setminus \{e\}} (d(A_a, F_{s_a}) + d(A_a, F_{o_a}))$ . So, the constraint associated with path  $p \in \tilde{\mathcal{P}}_e$  in the LP captures the inequality  $d(A_e, F_{s_e}) \leq d(A_e, F_{o_e}) + \sum_{a \in p \setminus \{e\}} d(F_{s_a}, F_{o_a}) \leq d(A_e, F_{o_e}) + \sum_{a \in p \setminus \{e\}} (d(A_a, F_{s_a}) + d(A_a, F_{o_a}))$ , by replacing  $d(A_e, F_{s_e})$  with  $z_e$  and  $d(A_a, F_{s_a})$  with  $z_a$  and rearranging the terms.

By duality, the cost  $\text{cost}(I_g, S)$  of solution  $S$  is upper-bounded by the objective value of the dual linear program, defined as follows:

$$\begin{aligned} & \text{minimize } \sum_{p \in \mathcal{P}} x_p \sum_{e \in p} d(A_e, F_{o_e}) \\ & \text{subject to: } \sum_{p \in \mathcal{P}_{e_r}} x_p \geq 1 \\ & \sum_{p \in \tilde{\mathcal{P}}_e} x_p - \sum_{p \in \mathcal{P}_e \setminus \tilde{\mathcal{P}}_e} x_p \geq 1, e \in T, e \neq e_r \\ & x_p \geq 0, p \in \mathcal{P} \end{aligned}$$

Actually, for any feasible solution  $x$  of the dual LP,  $\text{cost}(S, I_g)$  is upper bounded by the quantity  $\sum_{p \in \mathcal{P}} x_p \sum_{e \in p} d(A_e, F_{o_e})$ . We will refer to any assignment  $x$  over

the paths of  $\mathcal{P}$  that satisfies the constraints of the dual LP as a *path covering* of the directed  $g$ -tree  $T$  and will denote its cost by  $c(x) = \max_{e \in T} \sum_{p \in \mathcal{P}_e} x_p$ . We repeat these definitions for clarity:

**Definition 1.** Let  $T$  be a directed tree. A function  $x : \mathcal{P} \rightarrow \mathbb{R}^+$  is called a path covering of  $T$  if the following conditions hold:

- $\sum_{p \in \mathcal{P}_{e_r}} x_p \geq 1$  for the edge  $e_r$  incident to the root of  $T$ ;
- $\sum_{p \in \tilde{\mathcal{P}}_e} x_p - \sum_{p \in \mathcal{P}_e \cap \mathcal{P}_f} x_p \geq 1$  if  $e \neq e_r$  and  $f$  denotes the parent edge of  $e$ .

The cost  $c(x)$  of  $x$  is equal to  $\max_{e \in T} \sum_{p \in \mathcal{P}_e} x_p$ .

**Lemma 2.** Let  $g \geq 2$  be an integer,  $I$  be a facility assignment instance with an optimal solution  $O$ ,  $S$  be a solution of the SD mechanism when applied on instance  $I_g$ , so that the triplet  $(I, O, S)$  is represented as a directed  $g$ -tree  $T$  which has a path covering  $x$ . Then,  $\text{cost}(S, I_g) \leq c(x) \cdot \text{cost}(O, I)$ .

*Proof.* Using the interpretation of the variables of the primal LP, duality, and the definition of the cost of path covering  $x$ , we have that

$$\begin{aligned} \text{cost}(S, I_g) &= \sum_{e \in T} z_e \leq \sum_{p \in \mathcal{P}} x_p \sum_{e \in \mathcal{P}_p} d(A_e, F_{o_e}) = \sum_{e \in T} d(A_e, F_{o_e}) \cdot \sum_{p \in \mathcal{P}_e} x_p \\ &\leq c(x) \cdot \sum_{e \in T} d(A_e, F_{o_e}) = c(x) \cdot \text{cost}(O, I) \end{aligned}$$

as desired. □

In order to establish the upper bounds in Theorem 1, it remains to show that path coverings with low cost do exist; this is what we do in the next three lemmas. We start with the Lemma for no augmentation. The proof of the lemma is omitted due to lack of space.

**Lemma 3.** Let  $T$  be a 1-tree. Then, there is a path covering of  $T$  of cost  $2^n - 1$ .

In the following, we identify path coverings of low cost for the case of  $g \geq 3$  and  $g = 2$ . The next two lemmas complete the part of Theorem 1 that regards the upper bounds.

**Lemma 4.** Let  $g \geq 3$  be an integer and  $T$  be a  $g$ -tree. Then, there is a path covering of  $T$  of cost  $\frac{g}{g-2}$ .

*Proof.* We prove the lemma using the following assignment  $x$ : for every path  $p$  of length  $\ell$ , we set  $x_p = \frac{1}{g-2} g^{2-\ell}$  if it contains an edge that is adjacent to the root and  $x_p = \frac{g-1}{g-2} g^{1-\ell}$  otherwise.

We will first show that  $\sum_{p \in \mathcal{P}_e} x_p = \frac{g}{g-2}$  for every edge  $e$  using induction. We will do so by visiting the edges in a bottom-up manner (i.e., an edge will be visited only after its child-edges have been visited) and prove that the equality for edge  $e$  using the information that the equality holds for its child-edges. As

the basis of our induction, consider an edge  $e$  that is adjacent to a leaf at depth  $\ell \geq 1$  from the root. If  $\ell = 1$ , this means that the tree consists of a single edge and there is a single path  $p$  with  $x_p = \frac{g}{g-2}$ . If  $\ell \geq 2$ , then the paths that contain edge  $e$  are those who end at each ancestor of the leaf adjacent to  $e$ . Hence,

$$\sum_{p \in \mathcal{P}_e} x_p = \sum_{i=1}^{\ell-1} \frac{g-1}{g-2} g^{1-i} + \frac{1}{g-2} g^{2-\ell} = \frac{g}{g-2}.$$

Now, let us focus on a non-leaf edge  $e$  and assume that  $\sum_{p \in \mathcal{P}_{e_i}} x_p = \frac{g}{g-2}$  for each child-edge  $e_i$  (for  $i \in [g]$ ) of  $e$  (this is the induction hypothesis). Let  $u$  be the node to which edges  $e$  and  $e_i$  with  $i \in [g]$  are incident. The set of paths in  $\mathcal{P}_e$  consists of the following disjoint sets of paths: for each edge  $e_i$  and for each path  $p \in \tilde{\mathcal{P}}_{e_i}$ , set  $\mathcal{P}_e$  contains all super-paths of  $p$ , i.e., paths originating from the leaf-node reached by  $p$  and ending at each ancestor of node  $u$ ; we use the notation  $\text{sup}(p)$  to denote the set of super-paths of  $p$ . Observe that, the definition of  $x$  implies that a super-path  $q$  of  $p$  that is longer than  $p$  by  $j$  has  $x_q = \frac{1}{g-1} g^{1-j} x_p$  if  $q$  is adjacent to the root and  $x_q = g^{-j} x_p$  otherwise. Hence, assuming that node  $u$  is at depth  $\ell \geq 1$  from the root, we have that

$$\begin{aligned} \sum_{p \in \mathcal{P}_e} x_p &= \sum_{i=1}^g \sum_{p \in \tilde{\mathcal{P}}_{e_i}} \sum_{q \in \text{sup}(p)} x_q = \left( \sum_{j=1}^{\ell-1} g^{-j} + \frac{1}{g-1} g^{1-\ell} \right) \sum_{i=1}^g \sum_{p \in \tilde{\mathcal{P}}_{e_i}} x_p \\ &= \frac{1}{g-1} \left( \sum_{i=1}^g \sum_{p \in \mathcal{P}_{e_i}} x_p - \sum_{p \in \mathcal{P}_e} x_p \right), \end{aligned}$$

which yields  $\sum_{p \in \mathcal{P}_e} x_p = \frac{g}{g-2}$  as desired, since  $\sum_{p \in \mathcal{P}_{e_i}} x_p = \frac{g}{g-2}$  by the induction hypothesis.

It remains to show feasibility. Clearly,  $\sum_{p \in \mathcal{P}_e} x_p = \frac{g}{g-2} \geq 1$  if  $e$  is adjacent to the root. Otherwise, consider an edge  $e$ , its parent edge  $f$ , and their common endpoint  $u$ . Assuming that  $u$  is at depth  $\ell$  from the root (and using definitions and observations we used above), we have

$$\sum_{p \in \mathcal{P}_e \cap \mathcal{P}_f} x_p = \sum_{p \in \tilde{\mathcal{P}}_e} \sum_{q \in \text{sup}(p)} x_q = \left( \sum_{j=1}^{\ell-1} g^{-j} + \frac{1}{g-1} g^{1-\ell} \right) \sum_{p \in \tilde{\mathcal{P}}_e} x_p = \frac{1}{g-1} \sum_{p \in \tilde{\mathcal{P}}_e} x_p,$$

which, together with the fact that  $\frac{g}{g-2} = \sum_{p \in \mathcal{P}_e} x_p = \sum_{p \in \mathcal{P}_e \cap \mathcal{P}_f} x_p + \sum_{p \in \tilde{\mathcal{P}}_e} x_p$  yields  $\sum_{p \in \mathcal{P}_e \cap \mathcal{P}_f} x_p = \frac{1}{g-2}$  and  $\sum_{p \in \tilde{\mathcal{P}}_e} x_p = \frac{g-1}{g-2}$  and, consequently,  $\sum_{p \in \tilde{\mathcal{P}}_e} x_p - \sum_{p \in \mathcal{P}_e \cap \mathcal{P}_f} x_p = 1$  as desired.  $\square$

Finally, we state the lemma for augmentation factor  $g = 2$ . The proof is omitted due to lack of space.

**Lemma 5.** *Let  $T$  be an  $N$ -node 2-tree. Then, there is a path covering of  $T$  of cost at most  $\log N$ .*

We have shown that the performance of SD significantly improves even with a small augmentation factor. A natural next question is to study its randomized counterpart, RSD. Could randomization help in achieving much better ratios? In the following, we state an approximation guarantee for RSD, when there is no resource augmentation. The proof is omitted due to lack of space.

**Theorem 2.** *The approximation ratio of RSD without resource augmentation is  $\text{ratio}(RSD) \leq n$ .*

## 4 Lower Bounds

In this section, we provide lower bounds on the approximation ratio with augmentation of the two mechanisms that we study. Interestingly, the constructed instances are all on a simple metric space, the *real line metric*.

**Theorem 3.** *The approximation ratio of Serial Dictatorship with augmentation factor  $g$  in facility assignment instances with  $n$  agents is*

1.  $\text{ratio}(SD) \geq 2^n - 1$
2.  $\text{ratio}_2(SD) \geq \log(n + 1)$
3.  $\text{ratio}_g(SD) \geq \frac{g}{g-2} - \delta$  for any  $\delta > 0$  when  $g \geq 3$ .

*The approximation ratio of Random Serial Dictatorship is at least  $\text{ratio}(RSD) \geq n^{0.26}$  (without resource augmentation).*

We omit the proof of the theorem due to lack of space. The instances that provide the lower bounds as well as the proofs are included in the full version of the paper.

## 5 Discussion

We proposed a resource augmentation framework for algorithmic mechanism design, where a mechanism, severely limited by the need for truthfulness is given some additional allocative power before being compared to the optimal mechanism, which operates under no restrictions. The framework is applicable to other related problems as well; for example, the bi-criteria algorithms of [3] can be seen as instances of resource augmentation. The framework can also be applied to broader settings where the loss in performance is due to restrictions other than truthfulness, such as fairness [6], stability [7] or ordinality [4,8]; all the problems in those papers can be studied through the resource augmentation lens. It is not hard to imagine that similar notions like the *price of fairness* [6], could be redefined in terms of resource augmentation.

For the facility assignment problem, we took a positive step in the study of Random Serial Dictatorship, proving approximation ratio bounds when there is no augmentation. It seems like an interesting technical question to obtain (tight) bounds for RSD and for different augmentation factors. It would also be

meaningful to consider augmentation factors smaller than 2; note that a similar construction to the one in our main lower bound can be used to show that additive factors can not achieve significantly improved approximations. Finally, it makes sense to consider other truthful mechanisms, beyond the greedy ones. In the full version, we actually prove that for two facilities and no resource augmentation, the approximation ratio of SD is 3, which is optimal among all truthful mechanisms, even randomized ones.

## References

1. Abdulkadiroğlu, A., Sönmez, T.: Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* **66**, 689–701 (1998)
2. Anshelevich, E., Das, S.: Matching, cardinal utility, and social welfare. *ACM SIGE-Com Exch.* **9**(1), 4 (2010)
3. Anshelevich, E., Sekar, S.: Blind, greedy, random: algorithms for matching and clustering using only ordinal information. In: *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 390–396 (2016)
4. Aziz, H., Chen, J., Filos-Ratsikas, A., Mackenzie, S., Mattei, N.: Egalitarianism of random assignment mechanisms. In: *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2016)
5. Bogomolnaia, A., Moulin, H.: A new solution to the random assignment problem. *J. Econ. Theory* **100**, 295–328 (2001)
6. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: The efficiency of fair division. *Theory Comput. Syst.* **50**(4), 589–610 (2012)
7. Emek, Y., Langner, T., Wattenhofer, R.: The price of matching with metric preferences. In: Bansal, N., Finocchi, I. (eds.) *ESA 2015. LNCS*, vol. 9294, pp. 459–470. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-48350-3\\_39](https://doi.org/10.1007/978-3-662-48350-3_39)
8. Filos-Ratsikas, A., Frederiksen, S.K.S., Zhang, J.: Social welfare in one-sided matchings: random priority and beyond. In: Lavi, R. (ed.) *SAGT 2014. LNCS*, vol. 8768, pp. 1–12. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-44803-8\\_1](https://doi.org/10.1007/978-3-662-44803-8_1)
9. Guo, M., Conitzer, V.: Strategy-proof allocation of multiple items between two agents without payments or priors. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 881–888 (2010)
10. Hylland, A., Zeckhauser, R.: The efficient allocation of individuals to positions. *J. Polit. Econ.* **87**(2), 293–314 (1979)
11. Kalyanasundaram, B., Pruhs, K.: Online weighted matching. *J. Algorithms* **14**(3), 478–488 (1993)
12. Kalyanasundaram, B., Pruhs, K.: The online transportation problem. *SIAM J. Discrete Math.* **13**(3), 370–383 (2000)
13. Kalyanasundaram, B., Pruhs, K.: Speed is as powerful as clairvoyance. *J. ACM* **47**(4), 617–643 (2000)
14. Koutsoupias, E.: Weak adversaries for the k-server problem. In: *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 444–449 (1999)
15. Krysta, P., Manlove, D., Rastegari, B., Zhang, J.: Size versus truthfulness in the house allocation problem. In: *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pp. 453–470 (2014)



16. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. *ACM Trans. Econ. Comput.* **1**(4), Article No. 18 (2013)
17. Sleator, D.D., Tarjan, R.E.: Amortized efficiency of list update and paging rules. *Commun. ACM* **28**(2), 202–208 (1985)
18. Svensson, L.-G.: Strategy-proof allocation of indivisible goods. *Soc. Choice Welfare* **16**(4), 557–567 (1999)
19. Young, N.: The k-server dual and loose competitiveness for paging. *Algorithmica* **11**(6), 525–541 (1994)

# Putting Peer Prediction Under the Micro(economic)scope and Making Truth-Telling Focal

Yuqing Kong<sup>1</sup>(✉), Katrina Ligett<sup>2,3</sup>, and Grant Schoenebeck<sup>1</sup>

<sup>1</sup> University of Michigan, Ann Arbor, USA  
{yuqkong, schoeneb}@umich.edu

<sup>2</sup> California Institute of Technology, Pasadena, USA  
katrina@caltech.edu

<sup>3</sup> Hebrew University, Jerusalem, Israel

**Abstract.** Peer-prediction [19] is a (meta-)mechanism which, given any proper scoring rule, produces a mechanism to elicit private information from self-interested agents. Formally, truth-telling is a strict Nash equilibrium of the mechanism. Unfortunately, there may be other equilibria as well (including uninformative equilibria where all players simply report the same fixed signal, regardless of their true signal) and, typically, the truth-telling equilibrium does not have the highest expected payoff. The main result of this paper is to show that, in the symmetric binary setting, by tweaking peer-prediction, in part by carefully selecting the proper scoring rule it is based on, we can make the truth-telling equilibrium focal—that is, truth-telling has higher expected payoff than any other equilibrium.

Along the way, we prove the following: in the setting where agents receive binary signals we (1) classify all equilibria of the peer-prediction mechanism; (2) introduce a new technical tool for understanding scoring rules, which allows us to make truth-telling pay better than any other informative equilibrium; (3) leverage this tool to provide an optimal version of the previous result; that is, we optimize the gap between the expected payoff of truth-telling and other informative equilibria; and (4) show that with a slight modification to the peer-prediction framework, we can, in general, make the truth-telling equilibrium focal—that is, truth-telling pays more than any other equilibrium (including the uninformative equilibria).

**Keywords:** Information elicitation · Peer prediction · Crowdsourcing

---

Y. Kong—Supported by National Science Foundation Career Award 1452915 and CCF Award 1618187.

K. Ligett—Supported in part NSF grants 1254169 and 1518941, US-Israel Binational Science Foundation Grant 2012348, the Charles Lee Powell Foundation, a subcontract through the Brandeis project, a grant from the HUIJ Cyber Security Research Center, and a startup grant from Hebrew University's School of Computer Science.

G. Schoenebeck—Supported by National Science Foundation Career Award 1452915 and Algorithms in the Field Award 1535912.

## 1 Introduction

From Facebook.com’s “What’s on your mind?” to Netflix’s 5-point ratings, from innumerable survey requests in one’s email inbox to Ebay’s reputation system, user feedback plays an increasingly central role in our online lives. This feedback can serve a variety of important purposes, including supporting product recommendations, scholarly research, product development, pricing, and purchasing decisions. With increasing requests for information, agents must decide where to turn their attention. When privately held information is elicited, sometimes agents may be intrinsically motivated to both participate and report the truth. Other times, self-interested agents may need incentives to compensate for costs associated with truth-telling and reporting: the effort required to complete the rating (which could lead to a lack of reviews), the effort required to produce an accurate rating (which might lead to inaccurate reviews), foregoing the opportunity to submit an inaccurate review that could benefit the agent in future interactions [11] (which could, e.g., encourage negative reviews), or a potential loss of privacy [8] (which could encourage either non-participation or incorrect reviews).

To overcome a lack of (representative) reviews, a system could reward users for reviews. However, this can create perverse incentives that lead to inaccurate reviews. If agents are merely rewarded for participation, they may not take time to answer the questions carefully, or even meaningfully.

To this end, explicit reward systems for honest ratings have been developed. If the ratings correspond to objective information that will be revealed at a future date, this information can be leveraged (e.g., via prediction markets) to incentive honesty. In this paper, we study situations where this is not the case: the ratings cannot be independently verified either because no objective truth exists (the ratings are inherently subjective) or an objective truth exists, but is not observable.

In such cases, it is known that correlation between user types can be leveraged to elicit truthful reports by using side payments [1, 2, 4, 5]. Miller et al. [19] propose a particular such (meta-)mechanism for truthful feedback elicitation, known as *peer prediction*. Given any proper scoring rule (a simple class of payment functions we describe further below), and a prior where each agent’s signal is “stochastically relevant” (informative about other agents’ signals), the corresponding peer prediction mechanism has truth-telling as a strict Bayesian-Nash equilibrium.

There is a major problem, however: alternative, non-truthful equilibria may have higher payoff for the agents than truth-telling. This is the challenge that our work addresses.

*Our Results.* The main result of this paper is to show that by tweaking peer prediction, in part by specially selecting the proper scoring rule it is based on, we can make the truth-telling equilibrium focal—that is, truth-telling has higher expected payoff than any other equilibrium.

Along the way we prove the following: in the setting where agents receive binary signals we (1) classify all equilibria of the peer prediction mechanism; (2) introduce a new technical tool, the best response plot, and use it to show that we can find proper scoring rules so the truth-telling pays more, in expectation, than any other informative equilibrium; (3) we provide an optimal version of the previous result, that is we optimize the gap between the expected payoff of truth-telling and other informative equilibrium; and (4) we show that with a slight modification to the peer prediction framework, we can, in general, make the truth-telling equilibrium focal—that is, truth-telling pays more than any other equilibrium (including the uninformative equilibria).

The main technical tool we use is a best response plot, which allows us to easily reason about the payoffs of different equilibria. We first prove that no asymmetric equilibria exist. The naive approach then would be to simply plot the payoffs of different symmetric strategies. However, for even the simplest proper-scoring rules, these payoff curves are paraboloid, and hence difficult to analyze directly. The best response plot differs from this naive approach in two ways: first, instead of plotting the strategies of agents explicitly, the best response plot aggregates the results of these actions; second, instead of plotting the payoffs of all agents, the best response plot analyzes the payoff of one distinguished agent which, given the strategies of the remaining agents, plays her best response. This makes the plot piece-wise linear for all proper scoring rules, which makes analysis tractable. We hope that the best response plot will be useful in future work using proper scoring rules.

## 1.1 Related Work

Since the seminal work of Miller et al. introducing peer prediction [19], a host of results in closely related models have followed (see, e.g., [9, 11, 12, 14]), primarily motivated by opinion elicitation in online settings where there is no objective ground truth.

Recent research [7] indicates that individuals in lab experiments do not always truth-tell when faced with peer prediction mechanisms; this may in part be related to the issue of equilibrium multiplicity. Gao et al. [7] ran studies over Mechanical Turk using two treatments: in the first they compensated the participants according to peer prediction payments, and in the second they gave them a flat reward for participation. In their work, the mechanism had complete knowledge of the prior. The participants responded truthfully more often when the payoffs were fixed than in response to the peer prediction payments. However, it should be noted that the task the agents were asked to perform took little effort (report the received signal), and the participants were not primed with any information about the truthful equilibrium of the peer prediction mechanism (they were only told the payoffs)—an actual surveyor would have incentive to prime the participants to report truthfully.

The most closely related work is a series of papers by Jurca and Faltings [12, 14], which studies collusion between the reporting agents. In a weak model

of collusion, the agents may be able to coordinate ahead of time (before receiving their signals) to select the equilibrium with the highest payoff. Jurca and Faltings use techniques from algorithmic mechanism design to design a mechanism where, in most situations, the only symmetric *pure* Nash equilibria are truth-telling. They explicitly state the challenge of analysing mixed-Nash equilibrium as an open question, and show challenges to doing this in their algorithmic mechanism design framework [12, 14]. Our techniques, in contrast, allow us to analyse all Nash equilibria of the peer prediction mechanism including both mixed-strategy and asymmetric equilibria. Instead of eliminating equilibria, we enforce that they have a lower expected payoff than truth-telling. Additionally, the algorithmic mechanism design framework used by Jurca and Faltings sacrifices “the simplicity of specifying the payments through closed-form scoring rules” [12] that was present in the peer prediction paper. Our work recovers a good deal of that simplicity.

Jurca and Faltings further analyze other settings where colluding agents can make transfer payments, or may collude after receiving their signals. In particular, they again use automated mechanism design to show that in the case where agents coordinate after receiving their signals that even without transfer payments, there will always be multiple equilibria; in this setting, they pose the question of whether the truth-telling equilibrium can be endowed with the highest expected payoff. We do not deal with this setting explicitly, but in the settings we consider, we show that even in the face of multiple equilibria, we can ensure that the truth-telling equilibrium has the highest expected payoff and no other equilibrium is paid the same with truth-telling.

In a different paper [11], Jurca and Faltings show how to minimize payments in the peer prediction framework. Their goal is to discover how much “cost” is associated with a certain marginal improvement of truth-telling over lying. In this paper, they also consider generalizations of peer prediction, where more than one other agent’s report is used as a reference. Our work takes this to the extreme (as did [8] before us) using *all* of the *other* agents’ reports as references.

A key motivation of one branch of the related work is removing the assumption that the mechanism knows the common prior [3, 6, 10, 13, 15, 18, 20–22, 24, 25]. Dasgupta and Ghosh [3], Kamble et al. [15], Kong and Schoenebeck [16], Shnayder et al. [23] study a different setting where agents are asked to answer several a priori similar questions. Our results can be applied even when there is just a single questions (thus we do not need to assume any relation between questions). Kamble et al. [15]’s mechanism applies to both homogeneous and heterogeneous population but requires a large number of a priori similar tasks. However, Kamble et al. [15]’s mechanism contains non-truthful equilibria that are paid higher than truth-telling. Dasgupta and Ghosh [3]’s mechanism has truth-telling as the equilibrium with the highest payoff, but contains a non-truthful equilibrium that is paid as much as truth-telling. Prelec [20] shows that in his Bayesian Truth Serum (BTS), truth-telling maximizes each individual’s expected “Information-score” across all equilibria. However, this guarantee is not strict, and requires the number of agents to be infinite, even to just have truth-telling be an equilibrium.

Moreover, it is hard to classify the equilibria or optimize mechanism in Prelec’s setting. Another drawback of BTS is that it requires agents to report “prediction” while our mechanism only requires agents to report a single signal. Radanovic and Faltings [21]’s mechanism solves this drawback but that mechanism is in a sensing scenario and needs to compare the information of an sensor’s local neighbours with the information of global sensors while our mechanism does not require this local/global structure. Moreover, like BTS, Radanovic and Faltings [21]’s mechanism does not have the strictness guarantee and requires the number of agents to be infinite even to have truth-telling as an equilibrium. In addition, Lambert and Shoham [18] provide a mechanism such that no equilibrium pays more than truth-telling, but here all equilibria pay the same amount; and while truth-telling is a Bayesian Nash equilibrium, unlike in peer prediction it generally is not a strict Bayesian Nash equilibrium. *Minimal Truth Serum (MTS)* [22] is a mechanism where agents have the option to report or not report their predictions, and also lacks analysis of non-truthful equilibria. MTS uses a typical zero-sum technique such that all equilibria are paid equally.

Equilibrium multiplicity is clearly a pervasive problem in this literature. While our present work only applies to the classical peer prediction mechanism, it provides an important step in addressing equilibrium multiplicity, and a new toolkit for reasoning about proper scoring rules.

*Subsequent Work.* Kong and Schoenebeck [17] show analogous results in the setting where mechanism does not know the prior; however, they also prove that results as strong as those in this paper are impossible in that setting.

## 2 Preliminaries, Background, and Notation

### 2.1 Game Setting

Consider a setting with  $n$  agents  $A$ . If  $A' \subseteq A$ , we let  $-A'$  denote  $A \setminus A'$ . Each agent  $i$  has a private signal  $\sigma_i \in \Sigma$ . We consider a game in which each agent  $i$  reports some signal  $\hat{\sigma}_i \in \Sigma$ . Let  $\sigma$  denote the vector of signals and  $\hat{\sigma}$  denote the vector of reports. Let  $\sigma_{-i}$  and  $\hat{\sigma}_{-i}$  denote the signals and reports excluding that of agent  $i$ ; we regularly use the  $-i$  notation to exclude an agent  $i$ .

We would like to encourage truth-telling, namely that agent  $i$  reports  $\hat{\sigma}_i = \sigma_i$ . To this end, agent  $i$  will receive some payment  $\nu_i(\hat{\sigma}_i, \hat{\sigma}_{-i})$  from our mechanism. In this paper, the game will be *anonymous*, in that each player’s payoffs will depend only on the player’s own report and the *fraction* of other players giving each possible report  $\in \Sigma$ , and not on the identities of those players.

**Assumption 1 (Binary Signals).** *We will refer to the case when  $\Sigma = \{0, 1\}$  as the **binary signal** setting, and we focus on this setting in this paper.*

**Assumption 2 (Symmetric Prior).** *We assume throughout that the agents’ signals  $\sigma$  are drawn from some joint **symmetric prior**  $Q$ : a priori, each agent’s signal is drawn from the same distribution. We in fact only leverage a weaker assumption, that  $\forall \sigma, \sigma'$ , and  $\forall i \neq j$  and  $k \neq l$ , we have  $\Pr[\sigma_j = \sigma' | \sigma_i = \sigma] = \Pr[\sigma_l = \sigma' | \sigma_k = \sigma]$ .*

That is, the inference your signal lets you draw about others' signals does not depend on your identity or on the identity of the other agent.

Given the prior  $Q$ , for  $\sigma \in \Sigma$ , let  $q(\sigma)$  be the fraction of agents that an agent expects will have  $\sigma_j = \sigma$  *a priori*. Let

$$q(\sigma'|\sigma) := \Pr[\sigma_j = \sigma' | \sigma_i = \sigma]$$

(where  $j \neq i$ ) be the fraction of other agents that a user  $i$  expects have received signal  $\sigma'$  given that he has signal  $\sigma$ .

**Assumption 3 (Signals Positively Correlated).** *We assume throughout that the prior  $Q$  is **positively correlated**, namely that  $q(\sigma|\sigma) > q(\sigma)$ , for all  $\sigma \in \Sigma$ .*

That is, once a player sees that his signal is  $\sigma$ , this strictly increases his belief that others will have signal  $\sigma$ , when compared with his prior. Notice that even after an agent receives his signal, he may still believe that he is in the minority. Thus, simply encouraging agent agreement is not sufficient to incentivize truthful reporting.

**Assumption 4 (Signal Asymmetric Prior).** *An additional assumption we will often use is that the prior is **signal asymmetric**. For binary signals, as we consider in this paper, this simply means that  $q(0) \neq q(1)$ .*

For a richer signal space, intuitively, a signal asymmetric prior is one that changes under a relabeling of the signals, so that lying can potentially be distinguishable from truth-telling.

We say that an agent plays *response*  $\sigma \rightarrow \hat{\sigma}$ , if the agent reports signal  $\hat{\sigma}$  when he receives signal  $\sigma$ . Let  $X$  be the set of all responses (e.g.  $X = \{0 \rightarrow 0, 0 \rightarrow 1, 1 \rightarrow 0, 1 \rightarrow 1\}$  when  $\Sigma = \{0, 1\}$ ). In a *pure-strategy* an agent chooses a response for each  $\sigma \in \Sigma$ , and thus there are  $|\Sigma|^{|\Sigma|}$  possible pure strategies. Let  $S$  be the set of pure strategies and let  $s_i \in S$  denote a pure-strategy for agent  $i$ . We will also consider mixed strategies  $\theta_i$ , where agent  $i$  randomizes over pure strategies. Here we write

$$\theta_i(\sigma'|\sigma) := \Pr[\hat{\sigma}_i = \sigma' | \sigma_i = \sigma].$$

A strategy profile  $\theta = (\theta_1, \dots, \theta_n)$  consists of a strategy for each agent.

We can think of each  $\theta$  as a linear transformation from a distribution over received signals to a distribution of reported signals. Given a set of agents  $A' \subset A$ , we define

$$\theta'_{A'}(\sigma'|\sigma) := E_{i \leftarrow A'}[\theta_i(\sigma'|\sigma)]$$

where  $i \leftarrow A'$  means  $i$  is chosen uniformly at random from  $A'$ . When discussing *symmetric strategy profiles* where all players use the same strategy, we will often abuse notation and use notation for one agent's strategy to denote the entire strategy profile.

A *Bayesian Nash equilibrium* consists of a strategy profile  $\theta = (\theta_1, \dots, \theta_n)$  such that no player wishes to change his strategy, given the strategies of the other

players and the information contained the prior and his signal: for all  $i$  and for all alternative strategies  $\theta'_i$  for  $i$ ,  $\mathbb{E}[\nu_i(\boldsymbol{\theta})] \geq \mathbb{E}[\nu_i(\theta'_i, \boldsymbol{\theta}_{-i})]$ , where the expectations are over the realizations of the randomized strategies and the prior  $Q$ . We call such an equilibrium **focal** if it provides a strictly larger payoff, in expectation, to each agent, than any other equilibrium and **weakly focal** if it provides a larger payoff (maybe not strictly).

Given a symmetric prior  $Q$  and strategy profile  $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$ , we define

$$\hat{q}_j(\sigma'|\sigma) := \Pr[\hat{\sigma}_j = \sigma' | \sigma_i = \sigma] = \sum_{\sigma'' \in \Sigma} q(\sigma''|\sigma)\theta_j(\sigma'|\sigma'')$$

for  $i \neq j$ . Intuitively,  $\hat{q}_j(\sigma'|\sigma)$  is the probability of player  $j$  reporting  $\sigma'$ , given that another player  $i$  sees signal  $\sigma$ ; note that this probability does not depend on the identity of  $i$ , by symmetry of the prior. Given a set of agents  $A' \subset A$ , we define

$$\hat{q}'_{A'}(\sigma'|\sigma) := E_{j \leftarrow A'} \hat{q}_j(\sigma'|\sigma)$$

where  $j \leftarrow A'$  means  $j$  is chosen uniformly at random from  $A'$  (again assuming that the implicit reference agent  $i \notin A'$ ). If  $\boldsymbol{\theta} = (\theta, \dots, \theta)$  is symmetric, we simplify our notation to  $\hat{q}(\sigma'|\sigma)$  because the referenced set of agents does not matter.

In the binary signal setting when  $\boldsymbol{\theta}$  is symmetric, we have:

$$\hat{q}(1|0) = \theta(1|0)q(0|0) + \theta(1|1)q(1|0) \tag{1}$$

$$\hat{q}(1|1) = \theta(1|0)q(0|1) + \theta(1|1)q(1|1) \tag{2}$$

Additionally, we observe that  $q(1|b) = 1 - q(0|b)$ ,  $\theta_i(1|b) = 1 - \theta_i(0|b) \forall i$ , and  $\hat{q}(1|b) = 1 - \hat{q}(0|b)$ . Note that we will typically use  $b$  instead of  $\sigma$  to refer to binary signals (bits).

There are four pure strategies for playing the game in the binary signal setting: always 0, always 1, truth-telling, lying:

$$S = \left\{ \begin{pmatrix} 0 \rightarrow 0 \\ 1 \rightarrow 0 \end{pmatrix}, \begin{pmatrix} 0 \rightarrow 1 \\ 1 \rightarrow 1 \end{pmatrix}, \begin{pmatrix} 0 \rightarrow 0 \\ 1 \rightarrow 1 \end{pmatrix}, \begin{pmatrix} 0 \rightarrow 1 \\ 1 \rightarrow 0 \end{pmatrix} \right\} = \{\mathbf{0}, \mathbf{1}, \mathbf{T}, \mathbf{F}\}.$$

We will denote mixed strategies as  $\begin{pmatrix} 0 \rightarrow \theta(1|0) \\ 1 \rightarrow \theta(1|1) \end{pmatrix}$ .

### 2.2 Proper Scoring Rules

A scoring rule  $PS : \Sigma \times \Delta_\Sigma \rightarrow \mathbb{R}$  takes in signal  $\sigma \in \Sigma$  and a distribution over signals  $\delta_\Sigma \in \Delta_\Sigma$  and outputs a real number. A scoring rule is *proper* if, whenever the first input is drawn from a distribution  $\delta_\Sigma$ , then the expectation of  $PS$  is maximized by  $\delta_\Sigma$ . A scoring rule is called *strictly proper* if this maximum is unique. We will assume throughout that the scoring rules we use are strictly proper. By slightly abusing notation, we can extend a scoring rule to be  $PS : \Delta_\Sigma \times \Delta_\Sigma \rightarrow \mathbb{R}$  by simply taking  $PS(\delta_\Sigma, \delta'_\Sigma) = \mathbb{E}_{\sigma \leftarrow \delta_\Sigma}(\sigma, \delta'_\Sigma)$ .

In the case of scoring rules over binary signals, a distribution can be represented by a number in the unit interval, denoting the probability placed on the signal 1. In the binary signal setting, then, we extend proper scoring rules to be defined on  $[0, 1] \times [0, 1]$ .



*Example 1 (Example of Proper Scoring Rule).* The Brier Scoring Rule for predicting a binary event is defined as follows. Let  $I$  be the indicator random variable for the binary event to be predicted. Let  $q$  be the predicted probability of the event occurring. Then:

$$B(I, q) = 2I \cdot q + 2(1 - I) \cdot (1 - q) - q^2 - (1 - q)^2.$$

Note that if the event occurs with probability  $p$ , then the expected payoff of reporting a guess  $q$  is (abusing notation slightly):

$$B(p, q) = 2p \cdot q + 2(1 - p) \cdot (1 - q) - q^2 - (1 - q)^2 = 1 - 2(p - 2p \cdot q + q^2)$$

This is (uniquely) maximized when  $p = q$ , and so the Brier scoring rule is a strictly proper scoring rule. Note also that  $B(p, q)$  is a linear function in  $p$ . Hence, if  $p$  is drawn from a distribution, we have:  $\mathbb{E}_p[B(p, q)] = B(\mathbb{E}[p], q)$ , and so this is also maximized by reporting  $q = \mathbb{E}[p]$ .

### 2.3 Peer Prediction

Peer Prediction [19] with  $n$  agents receiving positively correlated binary signals  $\mathbf{b}$ , with symmetric prior  $Q$ , consists of the following mechanism  $\mathcal{M}(\hat{\mathbf{b}})$ :

1. Each agent  $i$  reports a signal  $\hat{b}_i$ .
2. Each agent  $i$  is uniformly randomly matched with an individual  $j \neq i$ , and is then paid  $PS(\hat{b}_j, q(1|\hat{b}_i))$ , where  $PS$  is a proper scoring rule.

That is, agent  $i$  is paid according to a proper scoring rule, based on  $i$ 's prediction that  $\hat{b}_j = 1$ , where  $i$ 's prediction is computed as either  $q(1|0)$  or  $q(1|1)$ , depending on  $i$ 's report to the mechanism. This can be thought of as having agent  $i$  bet on what agent  $j$ 's reported signal will be.

Notice that if agent  $j$  is truth-telling, then the Bayesian agent  $i$  would also be incentivized to truth-tell (strictly incentivized, if the proper scoring rule is strict). Agent  $i$ 's expected payoff (according to his own posterior distribution) for reporting his true type  $b_i$  has a premium compared to reporting  $-b_i$  of:

$$PS(\hat{b}_j, q(1|b_i)) - PS(\hat{b}_j, q(1|-b_i)) \geq 0$$

(strictly, for strict proper scoring rules) because we know that the expectation of  $PS(\hat{b}_j, \cdot)$  is (uniquely) maximized at  $q(1|b_i)$ . Now we introduce a convenient way to represent peer prediction mechanism.

**Definition 1 (Payoff Function Matrix).** *Each agent  $i$  who reports  $\hat{b}_i$  and is paired with agent  $j$  who reports  $\hat{b}_j$ , will be paid  $h_{\hat{b}_j, \hat{b}_i}$ . Then the peer prediction mechanism can be naturally represented as a  $2 \times 2$  matrix:*

$$\begin{pmatrix} h_{1,1} & h_{1,0} \\ h_{0,1} & h_{0,0} \end{pmatrix} = \begin{pmatrix} PS(1, q(1|1)) & PS(1, q(1|0)) \\ PS(0, q(1|1)) & PS(0, q(1|0)) \end{pmatrix}$$

which we call the payoff function matrix.

An example of a peer-prediction setting is included in the full version.

While truth-telling is always an equilibrium of the peer prediction mechanism, as we will see, it is not the only equilibrium. Two more equilibria are to always play 0 or always play 1. In Sect. 3.1, we further investigate equilibria of the peer prediction game. Based on the analysis of these multiple equilibria, we will develop a **modified peer prediction mechanism**, wherein players are paid according to the peer prediction based on a carefully-designed proper scoring rule, modulo some punishment imposed on the all playing 0 or all playing 1 strategy profiles. This modified mechanism will make the truth-telling equilibrium focal.

### 3 Summary of Main Results

In this section, we introduce our modified peer prediction mechanism and sketch the main theorem of this paper, that for almost any symmetric prior, there exists a modified peer prediction mechanism such that truth-telling is the focal equilibrium. Recall, we use the term *focal* to refer to an equilibrium with expected payoff strictly higher than that of any other Bayesian Nash equilibrium.

#### 3.1 Our Modified Peer Prediction Mechanism MPPM

Recall that modified peer prediction mechanism is the mechanism wherein players are paid according to peer prediction based on a carefully-designed proper scoring rule, modulo some punishment imposed on the all playing 0 or all playing 1 strategy profiles. So our approach differentiates between two types of equilibria:

**Definition 2 (Informative Strategy).** *We call always reporting 1 and always reporting 0 uninformative strategies; we call all other strategies (equilibria) informative.*

*Designing the Optimal Peer Prediction Mechanism.* We start to describe our modified peer prediction mechanism MPPM. We use two steps to design our MPPM. First we define the PPM:

**Definition 3.** *Given any binary, symmetric, positively correlated, and signal asymmetric prior  $Q$ , with  $q(1|1) > q(0|0)$  (the  $q(0|0) < q(1|1)$  case is analogous), we first design our peer prediction mechanism  $PPM(Q)$  and represent it as a payoff function matrix (See Definition 1).  $PPM(Q)$  depends on the region that  $Q$  belongs to, we defer the definitions of regions  $R_1, R_2, R_3$  to full version but provide Fig. 1 here to illustrate them.*

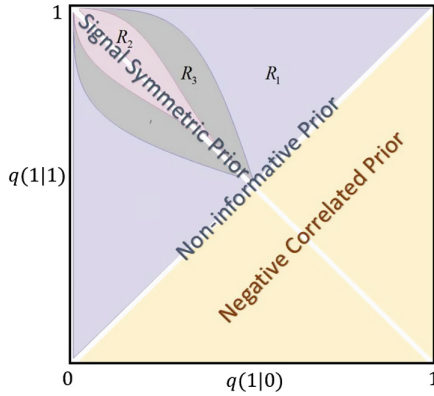
1. If  $Q \in R_1$ , then  $PPM(Q) = \mathcal{M}_1(Q)$
2. If  $Q \in R_2$ , then  $PPM(Q) = \mathcal{M}_2(Q)$
3. If  $Q \in R_3$ , then we pick a small number  $\epsilon > 0$  and  $PPM(Q, \epsilon) = \mathcal{M}_3(Q, \epsilon)$

where

$$\mathcal{M}_1(Q) = \begin{pmatrix} \zeta(Q) & 0 \\ 0 & 1 \end{pmatrix}, \mathcal{M}_2(Q) = \begin{pmatrix} 1 & 0 \\ 0 & \eta(Q) \end{pmatrix}, \mathcal{M}_3(Q, \epsilon) = \begin{pmatrix} \zeta(Q, \epsilon) & \delta(Q, \epsilon) \\ 0 & 1 \end{pmatrix}$$

and

$0 \leq \zeta(Q), \eta(Q) \leq 1$  are constants that only depend on common prior  $Q$ .  $0 \leq \zeta(Q, \epsilon), \delta(Q, \epsilon) \leq 1$  are constants that only depend on common prior  $Q$  and  $\epsilon > 0$ .<sup>1</sup>



**Fig. 1.** The regions  $R_1, R_2, R_3$  are good “priors” where we can make truth-telling focal when the number of agents is sufficient large. The white diagonals are “bad” priors we cannot make truth-telling focal. In the top-right to bottom-left diagonal,  $q(1|0) = q(1|1)$ , so the private signal does not have any information. We call this diagonal the set of non-informative priors. In the top-left to bottom-right diagonal,  $q(0|0) = q(1|1)$  (actually we can see  $q(0|0) = q(1|1)$  iff  $q(0) = q(1)$  via some calculations). This diagonal is the set of signal symmetric priors. The yellow region is the set of the negative correlated priors. (Color figure online)

Note that actually  $PPM(Q)$  is a quite simple mechanism. We use region  $R_1$  as example: if the prior belongs to region  $R_1$ , for every  $i$ , agent  $i$  will receive 0 payment if the agent paired with agent  $i$ , call him agent  $j$ , reports a different signal than him. If both agent  $i$  and agent  $j$  report 1, agent  $i$  will receive a payment of  $0 \leq \zeta(Q) \leq 1$ , if both agent  $i$  and agent  $j$  report 0, agent  $i$  will receive payment of 1.

Actually for regions  $R_1, R_2$ , the  $PPM(Q)$  we define here is the optimal peer prediction mechanism in that it maximizes the advantage of truth-telling over the informative equilibria which have the second largest expected payoff over all Peer-prediction mechanisms with payoffs in  $[0, 1]$ . For region  $R_3$ , the optimal peer prediction mechanism does not exist, but the advantage of the  $PPM(Q, \epsilon)$  we define approaches the optimal advantage as  $\epsilon$  goes to 0.

<sup>1</sup> Explicit statement in full version.

**Definition 4.** We define  $\Delta^*(Q)$  to be the supremum of the advantage of truth-telling over the informative equilibria which have the second largest expected payoff over all Peer-prediction mechanisms with payoffs in  $[0, 1]$ .

*Add Punishment.* In our  $PPM(Q)$ , an uninformative strategy can still obtain the highest payoff. For example, in mechanism  $\mathcal{M}_1$ , agents will receive maximal payment 1 by simply always reporting 0.

Our final  $MPPM(Q)$  Mechanism is the same as the  $PPM(Q)$  except that we add a punishment designed to hurt the all 0 or all 1 equilibria.

**Definition 5.** Our Modified Peer-Prediction Mechanism  $MPPM(Q)$  (or  $MPPM(Q, \epsilon)$  has payoffs identical to  $PPM(Q)$  (or  $MPPM(Q, \epsilon)$ ) except that, in the event all the other agents play all  $\mathbf{0}$  or all  $\mathbf{1}$ , it will issue an agent a punishment of  $p = \frac{1-t}{2(1-\epsilon_Q)} + \frac{\Delta^*(Q)}{2\epsilon_Q}$  where  $\epsilon_Q$  is the maximum probability that a fixed set of  $n - 1$  agents receive the same signal (either all  $\mathbf{0}$  or all  $\mathbf{1}$ );  $t$  is the expected of payoff of truth-telling  $\mathbf{T}$  in the  $PPM(Q)$ , and  $\Delta^*(Q)$  is as defined in Definition 4.

To make truth-telling focal, we would like to impose a punishment to the agents if everyone reports the same signal. However, such a punishment may distort the equilibria of the mechanism. To avoid this, we punish an agent by  $p$  when all the other agents report the same signal. Because an agent’s strategy does not influence his punishment, his marginal benefit for deviation remains the same and so the equilibrium remain the same. However, while all  $\mathbf{0}$  and all  $\mathbf{1}$  remain equilibrium, in them,  $MPPM(Q)$  will punish each agent by  $p$ .

A difficulty arises: if the number of agents is too small like 2 or 3, it is possible (and even probable) that all agents report their true signals, yet are still punished by the  $MPPM(Q)$  mechanism. Punishments like this might distort the payoffs among the informative equilibrium. However, if  $\epsilon_Q$  (the probability that  $n - 1$  agents receives the same signal) is sufficient small, this is no longer a problem. For most reasonable priors, as the number of agents increases,  $\epsilon_Q$  will go to zero. Formally we will need that the number of agents is large enough such that  $\epsilon_Q < \frac{\Delta^*(Q)}{1-t+\Delta^*(Q)}$ .

If the number of agents is too small such that  $\epsilon_Q \leq \frac{\Delta^*(Q)}{1-t+\Delta^*(Q)}$ , we cannot show that  $MPPM(Q)$  has truth-telling as a focal equilibrium.

In particular, we can see if  $\epsilon_Q \rightarrow 0$  (say as the number of agents increases), then at some point, truth-telling will be focal. We know that such a limit is necessary because, for example, with two agents making truth-telling focal is impossible.

Note that if the prior tells us the probability of a 1 event is concentrated far away from 0 and 1, the number of agents we need to make truth-telling focal will be very small since uninformative equilibria (all 1 and all 0) are far away from truth-telling.


**Theorem 5.** (Main Theorem (Informal)) Let  $Q$  be a binary, symmetric, positively correlated and signal asymmetric prior, and let  $\epsilon_Q$  be the maximum probability that a fixed set of  $n - 1$  agents receive the same signal (either all  $\mathbf{0}$  or all  $\mathbf{1}$ ). Then

1. In our PPM, truth-telling has the largest expected payoff among all informative equilibria. Moreover, over the space of Peer-Prediction mechanisms, our PPM(Q) maximizes the advantage truth-telling has over the informative equilibrium which have the second largest expected payoff, over all Peer-prediction mechanisms with payoffs in [0, 1] for regions R<sub>1</sub>, R<sub>2</sub> and PPM(Q, ε) approaches the maximal advantage for region R<sub>3</sub> as ε goes to 0.
2. There exists a constant ξ<sub>q(1|1),q(1|0)</sub> which only depends on q(1|1) and q(1|0) such that, if ε<sub>Q</sub> < ξ, our MPPM(Q) makes truth-telling focal.

Now we list all equilibria of the peer prediction mechanism in the below theorem (Fig. 2).

**Definition 6.** For a prior Q, proper scoring rule PS, and a binary signal space, we define q\* to be the fraction of other agents reporting 1 that would make an agent indifferent between reporting 0 or 1, i.e.,

$$q^* := \{p \mid PS(p, q(b|1)) = PS(p, q(b|0)), 0 \leq p \leq 1\}.$$



**Fig. 2.** Illustration of the 7 equilibria of a peer prediction mechanism under a specific scoring rule (see the full version). Note that to the right of the dashed red line where  $\hat{q}(1|0) = q^*$ , the best response is to increase  $\theta(1|0)$ ; to the left of the dashed red line, the best response is to decrease  $\theta(1|0)$ ; and on the line an agent is indifferent. Similarly, above the dotted blue line where  $\hat{q}(1|1) = q^*$ , the best response is to increase  $\theta(1|1)$ ; below the dotted blue line, the best response is to decrease  $\theta(1|1)$ ; and on the line an agent is indifferent. (Color figure online)

**Theorem 6.** Let Q be a symmetric and positively correlated prior on  $\{0, 1\}^n$ , and let M be a peer-prediction mechanism run with a strictly proper scoring rule with break-even q\* (Definition 6). Then there are no asymmetric equilibria. All equilibria are symmetric and depend only on q\*; they are

$$0, 1, T, Q^* \triangleq \begin{pmatrix} 0 \rightarrow q^* \\ 1 \rightarrow q^* \end{pmatrix},$$

$$\mathbf{T0} \triangleq \begin{pmatrix} 0 \rightarrow 0 \\ 1 \rightarrow \frac{q^*}{q(1|1)} \end{pmatrix}, \mathbf{T1} \triangleq \begin{pmatrix} 0 \rightarrow \frac{q^* - q(1|0)}{q(0|0)} \\ 1 \rightarrow 1 \end{pmatrix}$$

and also conditionally include

$$\mathbf{F} \text{ if } q(0|1) \leq q^* \leq q(0|0) \tag{3}$$

$$\mathbf{F1} \triangleq \begin{pmatrix} 0 \rightarrow 1 \\ 1 \rightarrow \frac{q^* - q(0|1)}{q(1|1)} \end{pmatrix} \text{ if } q(0|1) \leq q^* \tag{4}$$

$$\mathbf{F0} \triangleq \begin{pmatrix} 0 \rightarrow \frac{q^*}{q(0|0)} \\ 1 \rightarrow 0 \end{pmatrix} \text{ if } q^* \leq q(0|0) \tag{5}$$

Due to the space limitation, we defer all proofs to our full version (see <https://arxiv.org/abs/1603.07319>).

**Acknowledgments.** We thank David Parkes and Paul Resnick for helpful suggestions.

## References

1. Crémer, J., McLean, R.P.: Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* **53**(2), 345–361 (1985)
2. Crémer, J., McLean, R.P.: Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica* **56**(6), 1247–1257 (1988)
3. Dasgupta, A., Ghosh, A.: Crowdsourced judgement elicitation with endogenous proficiency. In: Proceedings of the 22nd International Conference on World Wide Web, pp. 319–330. International World Wide Web Conferences Steering Committee (2013)
4. d’Aspremont, C., Gérard-Varet, L.A.: Bayesian incentive compatible beliefs. *J. Math. Econ.* **10**(1), 83–103 (1982)
5. d’Aspremont, C., Grard-Varet, L.A.: Incentives and incomplete information. *J. Publ. Econ.* **11**(1), 25–45 (1979)
6. Faltings, B., Jurca, R., Pu, P., Tran, B.D.: Incentives to counter bias in human computation. In: Second AAAI Conference on Human Computation and Crowdsourcing (2014)
7. Gao, X.A., Mao, A., Chen, Y., Adams, R.P.: Trick or treat: putting peer prediction to the test. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation, pp. 507–524. ACM (2014)
8. Ghosh, A., Ligett, K., Roth, A., Schoenebeck, G.: Buying private data without verification. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation, pp. 931–948. ACM (2014)
9. Goel, S., Reeves, D.M., Pennock, D.M.: Collective revelation: a mechanism for self-verified, weighted, and truthful predictions. In: Proceedings of the 10th ACM Conference on Electronic Commerce (EC 2009) (2009)
10. Jurca, R., Faltings, B.: Incentives for answering hypothetical questions. In: Proceedings of the 1st Workshop on Social Computing and User Generated Content (SC 2011). ACM (2011)

11. Jurca, R., Faltings, B.: Minimum payments that reward honest reputation feedback. In: Proceedings of the 7th ACM Conference on Electronic Commerce (EC 2006) (2006)
12. Jurca, R., Faltings, B.: Collusion-resistant, incentive-compatible feedback payments. In: Proceedings of the 8th ACM Conference on Electronic Commerce, pp. 200–209. ACM (2007)
13. Jurca, R., Faltings, B.: Incentives for expressing opinions in online polls. In: Proceedings of the 9th ACM Conference on Electronic Commerce (EC 2008) (2008)
14. Jurca, R., Faltings, B.: Mechanisms for making crowds truthful. *J. Artif. Int. Res.* **34**(1), 209 (2009)
15. Kamble, V., Shah, N., Marn, D., Parekh, A., Ramachandran, K.: Truth serums for massively crowdsourced evaluation tasks. arXiv preprint [arXiv:1507.07045](https://arxiv.org/abs/1507.07045)
16. Kong, Y., Schoenebeck, G.: A framework for designing information elicitation mechanisms that reward truth-telling. ArXiv e-prints, May 2016
17. Kong, Y., Schoenebeck, G.: Equilibrium selection in information elicitation without verification via information monotonicity. ArXiv e-prints, March 2016
18. Lambert, N., Shoham, Y.: Truthful surveys. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 154–165. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-92185-1\\_23](https://doi.org/10.1007/978-3-540-92185-1_23)
19. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting informative feedback: the peer-prediction method. *Manag. Sci.* **51**(9), 1359–1373 (2005)
20. Prelec, D.: A Bayesian truth serum for subjective data. *Science* **306**(5695), 462–466 (2004)
21. Radanovic, G., Faltings, B.: Incentive schemes for participatory sensing. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, pp. 1081–1089. International Foundation for Autonomous Agents and Multiagent Systems (2015)
22. Riley, B.: Minimum truth serums with optional predictions. In: Proceedings of the 4th Workshop on Social Computing and User Generated Content (SC14) (2014)
23. Shnayder, V., Agarwal, A., Frongillo, R., Parkes, D.C.: Informed truthfulness in multi-task peer prediction. ArXiv e-prints, March 2016
24. Witkowski, J., Parkes, D.C.: Peer prediction without a common prior. In: Proceedings of the 13th ACM Conference on Electronic Commerce, pp. 964–981. ACM (2012)
25. Witkowski, J., Parkes, D.C.: Learning the prior in minimal peer prediction. In: Proceedings of the 3rd Workshop on Social Computing and User Generated Content at the ACM Conference on Electronic Commerce, p. 14. Citeseer (2013)

# Truthful Mechanisms for Matching and Clustering in an Ordinal World

Elliot Anshelevich and Shreyas Sekar<sup>(✉)</sup>

Rensselaer Polytechnic Institute, Troy, NY, USA  
eanshel@cs.rpi.edu, sekars@rpi.edu

**Abstract.** We study truthful mechanisms for matching and related problems in a partial information setting, where the agents' true utilities are hidden, and the algorithm only has access to ordinal preference information. Our model is motivated by the fact that in many settings, agents cannot express the numerical values of their utility for different outcomes, but are still able to rank the outcomes in their order of preference. Specifically, we study problems where the ground truth exists in the form of a weighted graph of agent utilities, but the algorithm can only elicit the agents' private information in the form of a preference ordering for each agent induced by the underlying weights. Against this backdrop, we design truthful algorithms to approximate the true optimum solution with respect to the hidden weights. Our techniques yield universally truthful algorithms for a number of graph problems: a 1.76-approximation algorithm for Max-Weight Matching, 2-approximation algorithm for Max  $k$ -matching, a 6-approximation algorithm for Densest  $k$ -subgraph, and a 2-approximation algorithm for Max Traveling Salesman as long as the hidden weights constitute a metric. Our results are the first non-trivial truthful approximation algorithms for these problems, and indicate that in many situations, we can design robust algorithms even when the agents may lie and only provide ordinal information instead of precise utilities.

## 1 Introduction

In recent years, the field of algorithm design has been marked by a steady shift towards newer paradigms that take into the account the behavioral aspects and communication bottlenecks pertaining to self-interested agents. In contrast to traditional algorithms that are assumed to have complete information regarding the inputs, mechanisms that interact with autonomous individuals commonly assume that the input to the algorithm is controlled by the agents themselves. In this context, a natural constraint that governs the process by which the algorithm elicits inputs from these agents is *truthfulness*: agents cannot improve upon the resulting outcome by misreporting the inputs. Another constraint that has recently gained traction in optimization problems on weighted graphs (where the agents correspond to the nodes) is that of *ordinality*: here, each agent can only submit a preference list of their neighbors ranked in the order of the edge



weights. The need for algorithms that are both truthful and ordinal arises in a number of important settings; however, it is well known that it is impossible to obtain optimum solutions even when the algorithm is required to satisfy only one of these two constraints.

In this work, we study the design of approximation algorithms for popular graph optimization problems including matching, clustering, and team formation with the goal of understanding the *combined price of truthfulness and ordinality*. To be more specific, we consider the above optimization problems on a weighted graph whose vertices represent the agents, and where the edge weights (that correspond to agent utilities) are private to the agents constituting that edge, and pose the following natural question: “*How does a computationally efficient, truthful algorithm that only has access to each agent’s edge weights in the form of preference rankings perform in comparison to an optimal algorithm that has full knowledge of the weighted graph?*”.

*Truthfulness in an Ordinal World.* Mechanisms that are either truthful or ordinal have received extensive attention across the spectrum of optimization problems. However, non-trivial algorithms that satisfy both of these considerations exist only for very specific settings [1, 11]. For instance, the price of ordinality (also referred to as *distortion*) is well understood for a number of applications such as voting [2, 5], matching [4, 13], facility location [11], and subset selection [4, 7]. The common thread in all of these settings where the (input) information is often held by the users is that it may be impossible or prohibitively expensive for the agents to express their full utilities to the mechanism; the same agents may incur a smaller overhead if they communicate preference lists over the other users or candidates in the system. Our main contention in this paper is that in exactly the same types of settings, it is reasonable to expect strategic agents to lie about their preferences if it improves their resulting utilities. Motivated by this, we study ordinal algorithms that are also truthful. Even though such mechanisms are clearly less powerful than their ‘ordinal but not necessarily truthful’ counterparts, our high level-level contribution is that for several well-studied graph maximization problems, one can obtain solutions that are only a constant factor away from the (social welfare of the) optimum, omniscient solution.

*Model and Problem Statements.* The high-level model in this paper is the same as the one in [4], with the addition of truthfulness as a constraint. The common setting for all the problems studied in this work is an undirected, complete weighted graph  $G$  whose nodes are the set of self-interested agents  $\mathcal{N}$  with  $|\mathcal{N}| = N$ . We use  $w(x, y)$  to denote the weight of the edge  $(x, y)$  in the graph for  $x, y \in \mathcal{N}$ . All of the optimization problems studied in this work involve selecting a subset of edges from  $G$  that obey some condition, with the objective of maximizing the weight of the edges chosen.

**Max  $k$ -Matching.** Compute the maximum weight matching consisting of exactly  $k$  edges. We refer to the  $k = \frac{N}{2}$  case as the *Weighted Perfect Matching* problem.

**$k$ -Sum Clustering.** Given an integer  $k$ , partition the nodes into  $k$  disjoint sets  $(S_1, \dots, S_k)$  of equal size in order to maximize  $\sum_{i=1}^k \sum_{x,y \in S_i} w(x,y)$ . (It is assumed that  $N$  is divisible by  $k$ ). When  $k = N/2$ ,  $k$ -sum clustering reduces to the weighted perfect matching problem.

**Densest  $k$ -subgraph.** Given an integer  $k$ , compute a set  $S \subseteq \mathcal{N}$  of size  $k$  to maximize the weight of the edges inside  $S$ .

**Max TSP.** In the maximum traveling salesman problem, the objective is to compute a tour  $T$  (cycle that visits each node in  $\mathcal{N}$  exactly once) to maximize  $\sum_{(x,y) \in T} w(x,y)$ .

A crucial but reasonably natural assumption that we make in this work is that the edge weights satisfy the *triangle inequality*, i.e., for  $x, y, z \in \mathcal{N}$ ,  $w(x,y) \leq w(x,z) + w(z,y)$ . For the specific kind of the problems that we study, the metric structure occurs in a number of well-motivated environments such as: (i) social networks, where the property captures a specific notion of friendship, (ii) Euclidean metrics: each agent is a point in a metric space which denotes her skills or beliefs, and (iii) edit distances: each agent could be represented by a string over a finite alphabet (for e.g., a gene sequence) and the graph weights represent the edit or Levenshtein distances [19]. The reader is asked to refer to [4] for additional details on these specific applications and a mathematical treatment of friendship in social networks.

Our framework and problem set models a multitude of interesting applications, and not surprisingly, all of the problems described above (with the metric assumption) have been the subject of a dense body of algorithmic work [4, 12, 14, 16]. In many of these applications, it becomes imperative that the algorithm provide good approximation guarantees even in the absence of precise numerical information regarding the graph weights. For instance, one can imagine partitioning a set of wedding guests to form a table assignment ( $k$ -sum clustering) or selecting a diverse team of agents in order to tackle a complex task (dense subgraph).

**Algorithmic Framework.** In this work, we are interested in the design of algorithms that are both ordinal and truthful. Suppose that for any one of the above problems, we are given an instance described by a weighted graph; then an algorithm  $\mathcal{A}$  for this problem is said to be ordinal if it has access only to a vector of preference orderings induced by the graph weights. That is, the input to this algorithm consists of a set of  $N$  preference orderings reported by each of the agents, where the preference list corresponding to agent  $i \in \mathcal{I}$  is a ranking over the agents in  $\mathcal{N} - \{i\}$  such that  $\forall j, k \in \mathcal{N}$ , if  $i$  prefers  $j$  to  $k$ , then  $w(i,j) \geq w(i,k)$ .

The algorithm is truthful if no single agent can improve their utility by submitting a preference ordering different from the ‘true ranking’ induced by the graph weights. Here, the utility of each agent  $i$  is simply the total weight of the edges incident to  $i$  which are chosen. These utilities have a natural interpretation with respect to the problems considered in this work. For instance, for matching problems, an agent’s utility corresponds to her affinity or weight to the agent to whom she is matched, and for densest subgraph as well as  $k$ -sum clustering,

the utility is her aggregate weight to the agents in the same team or cluster. Our objective in this paper is to design mechanisms that maximize the overall social welfare, i.e., the sum of the utilities of all the agents. Thus, the goal is to select a maximum-weight set of edges while knowing only ordinal preferences (instead of the true weights  $w$ ), with even the ordinal preferences possibly being misrepresented by the self-interested agents.

Finally,  $\mathcal{A}$  is said to be an ordinal  $\alpha$ -approximation algorithm for  $\alpha \geq 1$  if for any given instance along with the graph weights, the total objective value of the maximum weight solution with respect to the instance weights is at most a factor  $\alpha$  times the value of the solution returned by  $\mathcal{A}$ , when the input corresponds to the preference rankings induced by the weights. In other words, such algorithms produce solutions which are always a factor  $\alpha$  away from optimum, *without actually knowing* what the weights  $w$  are. We conclude by pointing out that despite the extensive body of work on all of the problems described previously, hardly any of the proposed mechanisms satisfy either truthfulness or ordinality (see Related Work for exceptions), motivating the need for a new line of algorithmic thinking.

**Our Contributions.** Our main results are summarized in Table 1. All of the non-matching problems that we study are NP-Hard even in the full information setting [12, 15, 18]. Our truthful ordinal algorithms provide constant approximation factors for a variety of problems in this setting, showing that even if only ordinal information is presented to the algorithm, and even if the agents can lie about their preferences, we can still form solutions efficiently with close to optimal utility. Note that as seen in Table 1, in [4] the authors already gave ordinal approximation algorithms for matching problems: those algorithms were not truthful, however, and achieving non-trivial approximation bounds while always giving players incentive to tell the truth requires significant additional work. For example, even the natural, greedy 2-approximation algorithm for Max  $k$ -matching from [4] is *not* truthful.

**Table 1.** Approximation factors provided in this paper by both truthful and non-truthful ordinal algorithms. (\*) A bicriteria result for Densest  $k$ -subgraph where the set size is relaxed to  $\beta k$  but the approximation factor is improved from 4 to  $\frac{4}{\beta^2}$  for  $\beta \geq 1$ .

| Problem                   | Our results      |                                  |
|---------------------------|------------------|----------------------------------|
|                           | Truthful ordinal | Non-truthful ordinal             |
| Weighted perfect matching | 1.7638           | 1.6 [4]                          |
| Max $k$ -matching         | 2                | 2 [4]                            |
| $k$ -Sum clustering       | 2                | 2                                |
| Densest $k$ -subgraph     | 6                | $(\frac{4}{\beta^2}, \beta)$ (*) |
| Max TSP                   | 2                | 1.88                             |

In addition to considering truthful mechanisms, we also develop new approximation algorithms for the setting where the agents are not able to lie, and thus the algorithm knows their true preference ordering. By dropping the truthfulness constraint, we are able to obtain better approximation factors for clustering, densest subgraph, and max TSP. The improved results are enabled by more involved algorithmic techniques that invariably sacrifice truthfulness; they establish a clear separation between the performance of an unconstrained ordinal algorithm and one that is required to be truthful. Owing to space constraints, we do not present these theorems in this version of the paper. The algorithms for the non-truthful versions of the problems can be found in the full version of this paper.

**Techniques.** Our proof techniques involve carefully stitching together *greedy*, *random*, and *serial dictatorship* based solutions. Understandably, and perhaps unavoidably for ordinal settings, the algorithmic paradigms that form the bedrock for our mechanisms are rather simple. However, beating the guarantees obtained by a naive application of these techniques involves a more intricate understanding of the interplay between the various approaches. For instance, our algorithm for the weighted perfect matching problem involves mixing between two simple 2-approximation algorithms (greedy, random) to achieve a 1.764-guarantee: towards this end, we establish new tradeoffs between greedy and random matchings showing that when one is far away from the optimum solution, the other one must provably be close to optimum.

## 1.1 Related Work

Broadly speaking, the truthful mechanisms in our work fall under the umbrella of ‘mechanism design without money’ [1, 6, 10, 13, 17], a recent line of work on designing strategyproof mechanisms for settings like ours, where monetary transfers are irrelevant. A majority of the papers in this domain deal with mechanisms that elicit agent utilities, specifically for one-sided matchings, assignments and facility location problems that are somewhat different from the graph problems we are interested in. The notable exceptions are the recent papers on truthful, ordinal mechanisms for one-sided matchings [6, 13] and general allocation problems [1]. While [13] looks at normalized agent utilities and shows that no ordinal algorithm can provide an approximation factor better than  $\Theta(\sqrt{N})$ , [6] considers *minimum* cost metric matching under a resource augmentation framework. The main differences between our work and these two papers are (1) we consider two-sided matching instead of one-sided, as well as other clustering problems, as well as non-truthful algorithms with better approximation factors, and (2) we consider maximization objectives in which users attempt to maximize their utility instead of minimize their cost. The latter may seem like a small difference, but it completely changes the nature of these problems, allowing us to create many different truthful mechanisms and achieve *constant-factor approximations*. Finally, [1] looks at the problem of allocating goods to buyers in a ‘fair fashion’. In that paper, the focus is on maximizing a popular non-linear

objective known as the *maximin share*, which is incompatible with our objective of social welfare maximization. That said, an interesting direction is to see if our techniques extend to other objectives.

As discussed in the Introduction, this paper improves on several results from [4]. In [4], the authors focused on the problem of maximum-weight matching for the non-truthful setting, with the main result being an ordinal 1.6-approximation algorithm. In the current paper, we greatly extend the techniques from [4] so that they may be applied to other problems in addition to matching. Moreover, we introduce several new techniques for this setting in order to create *truthful* algorithms; such algorithms require a somewhat different approach and make much more sense for many of the settings that we are interested in. Other than [6], these are the first known truthful algorithms for matching and clustering with metric utilities.

Our work is similar in motivation to the growing body of research studying settings where the voter preferences are induced by a set of hidden utilities [2, 3, 5, 7, 8, 11]. The voting protocols in these papers are essentially ordinal approximation algorithms, albeit for a very specific problem of selecting the utility-maximizing candidate from a set of alternatives.

## 2 Preliminaries

### 2.1 Truthful Ordinal Mechanisms

As mentioned previously, we are interested in designing incentive-compatible mechanisms that elicit ordinal preference information from the users, i.e., mechanisms where agents are incentivized to truthfully report their preferences in order to maximize their utility. We now formally define the notions of truthfulness pertinent to our setting. Throughout the rest of this paper, we will use  $P_i$  to represent the private ordinal preference of agent  $i$  (i.e., one that is induced by the weights  $w(i, j)$ ), and  $s_i$  to represent the preference ordering that agent  $i$  submits to the mechanisms (which will be equal to  $P_i$  if  $i$  tells the truth).

**Definition 1** (*Truthful Mechanism*). A deterministic mechanism  $\mathcal{M}$  is said to be truthful if for every  $i \in \mathcal{N}$ , all  $\mathbf{s}_{-i}, s'_i$ , we have that  $u_i(P_i, \mathbf{s}_{-i}) \geq u_i(s'_i, \mathbf{s}_{-i})$ , where  $u_i$  is the utility guaranteed to agent  $i$  by the mechanism.

**Definition 2** (*Universally Truthful Mechanisms*). A randomized mechanism is said to be universally truthful if it is a probability distribution over truthful deterministic mechanisms.

Informally, in a universally truthful mechanism, a user is incentivized to be truthful even when she knows the exact realization of the random variables involved in determining the mechanism. All of the algorithms in this work are *universally* truthful, not just in expectation. The reader is asked to refer to [9] for a useful discussion on the types of randomized mechanisms.

## 2.2 Approaches for Designing Truthful Matching Mechanisms

As a concrete first step towards designing truthful ordinal mechanisms, we introduce three high-level algorithmic paradigms that will form the backbone of all the results in this work. These paradigms are based on the popular algorithmic notions of *Greedy*, *Serial Dictatorship*, and *Uniformly Random*. For each of these paradigms, we develop approaches towards designing truthful mechanisms for the maximum matching problem. In Sects. 3 and 4, we develop more sophisticated truthful mechanisms that build upon the simple paradigms presented here, leading to improved approximation factors. All of the missing proofs from this section can be found in the full version of our paper.

**Greedy via Undominated Edges.** Our first algorithm is the ordinal analogue of the classic greedy matching algorithm, that has been extensively applied across the matching literature. In order to better understand this algorithm, we first define the notion of an *undominated edge*.

**Definition 3** (*Undominated Edge*). *Given a set  $E$  of edges,  $(x, y) \in E$  is said to be an undominated edge if for all  $(x, a)$  and  $(y, b)$  in  $E$ ,  $w(x, y) \geq w(x, a)$  and  $w(x, y) \geq w(y, b)$ .*

We make two simple observations here regarding undominated edges based on which we define Algorithm 1.

1. Every edge set  $E$  has at least one undominated edge. In particular, any maximum weight edge in  $E$  is obviously an undominated edge.
2. Given an edge set  $E$ , one can efficiently find at least one undominated edge *using only the ordinal preference information*.

```

 $M := \emptyset$ ,  $T$  is the valid set of edges initialized to the complete graph on  $\mathcal{N}$ ;
while  $T$  is not empty do
  | pick an undominated edge  $e = (x, y)$  from  $T$  and add it to  $M$ ;
  | remove all edges containing  $x$  or  $y$  from  $T$ ; if  $|M| = k$ ,  $T = \emptyset$ .
end

```

**Algorithm 1.** Greedy Algorithm for Max  $k$ -Matching

It is not difficult to see that this algorithm gives a 2-approximation for Max-Weight Perfect Matching, and is truthful for that case. Unfortunately, for Max  $k$ -Matching with smaller  $k$ , it is no longer truthful, and thus none of the algorithms that use Greedy as a subroutine (such as the algorithms from [4]) are truthful.

**Proposition 4.** *Algorithm 1 is truthful for the Max  $k$ -Matching problem only when  $k = \frac{N}{2}$ .*

*Proof.* We need to prove that for any given strategy profile adopted by the other players  $\mathbf{s}_{-i}$ , player  $i$  maximizes her utility when she is truthful, i.e., when  $s_i = P_i$ . Our proof will proceed via contradiction and will make use of the

following fundamental property: *if Algorithm 1 (for some input) matches agent  $i$  to  $j$  during some iteration, then both  $i$  and  $j$  prefer each other to every other agent that is unmatched during the same round.*

We introduce some notation: suppose that  $M$  denotes the matching output by Algorithm 1 for input  $(P_i, \mathbf{s}_{-i})$ , and for every  $x \in \mathcal{N}$ ,  $m(x)$  is the agent to whom  $x$  is matched to under  $M$ . Let  $e_j$  be the edge added to the matching  $M$  in round  $j$  of Algorithm 1, denote the round in which  $i$  is matched to  $m(i)$  as round  $k$ . Assume to the contrary that for input  $(s'_i, \mathbf{s}_{-i})$ ,  $i$  is matched to an agent she prefers more than  $m(i)$ . Let the altered matching be referred to as  $M'$ , and let  $m'(x)$  be the agent who  $x$  is matched with in  $M'$ .

We begin by proving the following claim: *For each  $j < k$ , we have that  $e_j \in M'$ .* In other words, all the edges which are included into  $M$  before  $i$  is matched by Algorithm 1 must appear in both matchings no matter what  $i$  does. Once we prove this claim, we are done, since  $e_k$  is the highest-weight edge from  $i$  to any node not in  $e_1, \dots, e_{k-1}$ , so  $i$  maximizes its utility by telling the truth and receiving utility equal to the weight of  $e_k$ .

To prove the claim above, we proceed by induction. Note that if  $k = 1$ , then  $i$  is trivially truthful, since  $m(i)$  is its top choice in the entire graph. Now suppose that we have shown the claim for edges  $e_1, \dots, e_{j-1}$ . Let  $e_j = (x, y)$ , and without loss of generality suppose that  $x$  is matched in our algorithm constructing  $M'$  before  $y$ . At the time that  $x$  is matched with  $m'(x)$ , it must be that  $m'(x)$  is the top choice of  $x$  from all available nodes. But, by the definition of our algorithm,  $y$  is the top choice of  $x$  that is not contained in  $e_1, \dots, e_{j-1}$ . Since  $m'(x)$  is not contained in  $e_1, \dots, e_{j-1}$  due to our inductive hypothesis, this means that  $x$  prefers  $y$  over  $m'(x)$ , and since  $y$  is not matched yet, this means that  $x$  and  $y$  will become matched together in  $M'$ . Thus,  $e_j$  is in  $M'$  as well. This completes the proof of truthfulness for  $k = \frac{N}{2}$ .  $\square$

Can we use a similar approach to design algorithms for the other problems that we are interested in? For  $k$ -sum clustering and Densest  $k$ -subgraph, one can follow the approach taken in [4, 14], and use the above matching as an intermediate to compute 4-approximations for the above problems. For Max TSP, we can directly leverage the above algorithm by maintaining  $M$  as a (forest of) path(s) instead of a matching in order to obtain a 2-approximate Hamiltonian tour. Unfortunately, as we show in the full version, these approaches **do not** lead to truthful algorithms at all.

**Serial Dictatorship.** Another popular approach to compute incentive compatible matchings (albeit usually for one-sided matchings [6, 13]) is serial dictatorship, which we formally define below for our two-sided matching setting.

**Proposition 5.** *Algorithm 2 is universally truthful for the Max  $k$ -Matching problem for all  $k$ .*

Serial dictatorship is among the most prominent of algorithms to feature in this work: our primary approximation algorithms for Max  $k$ -matching and Max TSP involve randomized versions of serial dictatorship.

```

 $M := \emptyset$ ,  $T$  is the set of available agents initialized to  $\mathcal{N}$ ;
while  $T$  is not empty do
    | pick an available agent  $x$  arbitrarily from  $T$ ;
    | let  $y$  denote  $x$ 's most preferred agent in  $T - \{x\}$ ; add  $(x, y)$  to  $M$ ;
    | remove all edges containing  $x$  or  $y$  from  $T$ ; if  $|M| = k$ ,  $T = \emptyset$ .
end
    
```

**Algorithm 2.** Serial Dictatorship for Max  $k$ -Matching

**Randomness.** A much simpler approach that is completely oblivious to the input preferences involves selecting a solution uniformly at random. Such an algorithm (described in Algorithm 3) is obviously truthful. Many of the techniques in this paper rely on carefully combining these three types of algorithms in order to produce good approximation factors while retaining truthfulness.

```

 $M := \emptyset$ ,  $T$  is the valid set of edges initialized to the complete graph on  $\mathcal{N}$ ;
while  $T$  is not empty do
    | pick an edge  $e = (x, y)$  from  $T$  uniformly at random and add it to  $M$ ;
    | remove all edges containing  $x$  or  $y$  from  $T$ ; if  $|M| = k$ ,  $T = \emptyset$ .
end
    
```

**Algorithm 3.** Random Algorithm for Max  $k$ -matching

### 3 Truthful Mechanisms for Matching

#### 3.1 Weighted Perfect Matching

So far, we have looked at two simply approaches for designing truthful mechanisms (Greedy and Random) for the weighted perfect matching problem, both of which yield 2-approximations [4] to the optimum matching. Can we do any better? In [4], the authors use a complex interleaving of greedy and random approaches to extract a *non-truthful* 1.6-approximation algorithm. In this paper, we instead present a simpler algorithm and rather surprising result: a simple random combination of Algorithms 1 and 3 results in a 1.764-approximation to the optimum matching. The main insight driving this result is the fact that the random and greedy approaches are in some senses complementary to each other, i.e., on instances where the approximation guarantee for the greedy algorithm is close to 2, the random algorithm performs much better.

**Theorem 6.** *The following algorithm is a universally truthful mechanism for the weighted perfect matching problem that obtains a 1.7638-approximation to the optimum matching.*

**Greedy-Random Mix Algorithm for Weighted Perfect Matching.** *With probability  $\frac{3}{7}$ , return the output of Algorithm 1 for  $k = \frac{N}{2}$  and with probability  $\frac{4}{7}$ , return the output of Algorithm 3 for  $k = \frac{N}{2}$ .*



*Proof Sketch:* Although the algorithm is exceedingly simple, the proof of the approximation factor is quite involved (see full version for the proof). The high-level argument proceeds as follows.

Suppose that  $GR$  is the output of the greedy algorithm for the given instance,  $OPT$  is the weight of the maximum-weight matching, and  $w(RD)$  is the expected weight of the random matching for the same instance. Begin by dividing the graph into two sets as follows. Define  $T$  to be the set of nodes which are included in the top (i.e., highest-weight)  $N/4$  edges of  $GR$ , and let  $B$  be the rest of the nodes. It is not difficult to show that the edges of  $GR$  in  $T$  have weight at least  $\frac{OPT}{2}$ ; suppose that the weight of the edges of  $GR$  in  $B$  equals  $xOPT$  for some  $x \geq 0$ . This means that  $w(GR) \geq \frac{1}{2}OPT + xOPT$ . The main part of the proof consists of proving the following claim, which essentially shows that when the greedy algorithm performs poorly, the randomized algorithm must do well.

**Claim 7.** *The weight of the random matching is always at least*

$$E[w(RD)] \geq \frac{5}{8}OPT - x(1 - \frac{3}{2}x)OPT.$$

Moreover, when  $x \leq \frac{1}{8}$ , the following is a tighter lower bound for the random matching:  $E[w(RD)] \geq \frac{5}{8}OPT - x(1 - 2x)OPT$ .

Once this claim is proven, the desired approximation bound of our algorithm follows from simple algebra. Proving this claim, however, requires forming some non-trivial machinery to analyze the quality of random metric matchings compared with maximum-weight matching. For any set of nodes  $S$ , define  $w(S)$  to be the total weight of all edges in  $S$ , and  $w(S_1, S_2)$  to be the weight of all edges between  $S_1$  and  $S_2$ . By heavily using the triangle inequality, we know that

$$OPT \leq \frac{4}{N}w(B) + \frac{2}{N}w(T, B) \quad \text{and} \quad \frac{N \cdot OPT}{4} \leq 2w(T).$$

Since  $N \cdot w(RD) = w(T) + w(B) + w(T, B)$ , this tells us that  $w(RD) \geq \frac{5}{8}OPT - \frac{w(B)}{N}$ . Most of the work from this point on is to obtain an upper bound on  $w(B)$  in terms of  $xOPT$ . The main idea involves splitting  $B$  into two parts  $B_1$  and  $B_2$ , where  $B_1$  consists of the nodes that make up the top  $\frac{xN}{2}$  edges in  $B$  with respect to  $GR$ . Suppose that the weight of the greedy edges in  $B_1$  equals  $\alpha xOPT$ , where  $\alpha$  is a measure of how ‘concentrated’ the heaviest edges in the bottom half of  $GR$  are. Now, what if  $\alpha$  is not very large: in this case, the weight of the greedy edges in  $B$  are somewhat evenly distributed across  $B_1$  and  $B_2$ , and the random matching performs quite well on such instances. If  $\alpha$  is high, then the concentration is uneven and several edges in  $GR$  (namely, those inside  $B_2$ ) have a small weight. Here, we show that the random matching performs well owing to the high-weight edges from  $T \cup B_1$  to  $B_2$ . Applying these insights to the following generic lower bound for the weight of the random matching allows us to complete the proof of Claim 7 and hence, the theorem.

$$N \cdot w(RD) \geq w(T) - w(B_2) + \frac{1}{2}[w(T, B_1) + (|B| + |B_2|)OPT - w(B_1, B_2)]. \square$$

### 3.2 Max $k$ -Matching

We now move on to the more general Max  $k$ -matching problem, where the objective is to compute a maximum weight matching consisting only of  $k \leq \frac{N}{2}$  edges. Our previous results do not carry over to this problem. While we know from [4] that the greedy algorithm is half-optimal, one can easily construct examples where this is not truthful. On the other hand, the random matching algorithm is truthful but its approximation factor can be as large as  $\frac{N}{k}$ . Our main result in this section is based on the *Random Serial Dictatorship* algorithm that in some sense combines the best of greedy and random into a single algorithm. Such algorithms have received attention for other matching problems [6, 13]; ours is the first result showing that these algorithms can approximate the optimum matching up to a small constant factor for metric settings. Specifically, while serial dictatorship is usually easy to analyze, our algorithm greatly exploits the randomness to select good edges *in expectation*.

**Definition:** *Random Serial Dictatorship* is the same algorithm as Serial Dictatorship (Algorithm 2), except the agents  $x$  from  $T$  are picked uniformly at random.

**Theorem 8.** *Random serial dictatorship is a universally truthful mechanism that provides a 2-approximation for the Max  $k$ -matching problem.*

## 4 Truthful Mechanisms for Other Problems

We remark that the proofs of all of the results in this section are available in the full version of the paper.

### 4.1 Densest $k$ -Subgraph

In this section we present our truthful, ordinal algorithm for Densest  $k$ -subgraph, which requires techniques somewhat different from the ones outlined in Sect. 2. While “conventional” approaches such as Greedy and Serial Dictatorship *do* lead to good approximations for this problem, they are not truthful, whereas random approaches are truthful but result in poor worst-case approximation factors. We combat this problem with a somewhat novel approach that combines the best of both worlds by designing a semi-oblivious algorithm that has the following property: if agent  $i$  is included in the solution, then changing her preference ordering  $s_i$  does not affect the mechanism’s output.

**Theorem 9.** *Algorithm 4 is a universally truthful mechanism that yields a 6-approximation for the Densest  $k$ -Subgraph problem.*

To see why this is truthful, note that for any particular choice of the anchor agent  $a$ , the only case in which  $a$ ’s preference ordering makes a difference is when  $a$  is definitely not added to the final team. Therefore, by lying  $a$  cannot influence her utility in the event that she is actually chosen.

```

 $S := \emptyset$ ,  $T$  is the set of available agents initialized to  $\mathcal{N}$ ;
while  $|S| < k$  do
    pick an anchor agent  $a$  and another node  $x$ , both uniformly at random from
     $T$ ;
    let  $b$  denote  $a$ 's most preferred agent in  $T - \{a, x\}$ ;
    with probability  $\frac{1}{2}$ , add  $a, x$  to  $S$ , and set  $T = T - \{a, x\}$ ;
    with probability  $\frac{1}{2}$ , add  $b, x$  to  $S$  and set  $T = T - \{a, b, x\}$ ;
end
    
```

**Algorithm 4.** Hybrid Algorithm for Densest  $k$ -Subgraph

### 4.2 A 2-approximation Algorithm for $k$ -Sum Clustering

In the literature, the  $k$ -sum clustering problem has only been studied in a full information setting, sometimes amidst the class of dispersion problems [14]. The best known approximation algorithm for this is a 2-approximation that uses the optimum matching as an intermediate. Instead, we give a *much simpler* algorithm with the same factor that is completely oblivious to the input, and is therefore truthful. Specifically, we prove that simply choosing the clusters uniformly at random is enough to provide a 2-approximate solution in expectation. Although the analysis of the algorithm involves new upper bounds on the optimum solution, it is still not difficult, so we include this result mostly for completeness.

### 4.3 Max Traveling Salesman Problem

The max traveling salesman problem has received a lot of attention in the literature despite not being as popular as the minimization variant, and has seen a plethora of algorithms for both the metric and the non-metric versions [15, 16]. Such algorithms usually work by looking at the optimum matching and cycle cover and cleverly interspersing the two solutions to form a Hamiltonian cycle. In adapting this approach to our setting, we would be bottlenecked by the best possible ordinal algorithms for the above two problems. Instead, we take a direct approach towards computing a tour and show that a simple algorithm based on Serial Dictatorship results in a 2-approximation factor.

**Theorem 10.** *Algorithm 5 is a universally truthful mechanism that provides a 2-approximation to the optimum tour. Moreover, the algorithm provides a  $(2+\epsilon)$ -approximation, where  $\epsilon \rightarrow 0$  as  $N \rightarrow \infty$ , even when the edge weights do not obey the metric assumption.*

It is easy to see that this algorithm is truthful: when an agent  $i$  is asked for its preferences, the first edge of  $T$  incident to agent  $i$  has already been decided, so  $i$  cannot affect it. Thus, to form the second edge of  $T$  incident to  $i$ , it may as well specify its most-preferred edge. Note that the randomization in the first step is *essential*: if we had selected the first edge based on the input preferences, then the first node could improve its utility by lying, and the algorithm would no longer be strategy-proof.

```

Initialize  $T$  to be a random edge from the complete graph on  $\mathcal{N}$ ;
Let  $S$  be the set of available agents initialized to  $\mathcal{N}$ ;
while  $S \neq \emptyset$  do
    pick one of the end-points of  $T$ , say  $x$  ;
    let  $y$  denote  $x$ 's most preferred agent in  $S$ ; add  $(x, y)$  to  $T$  and remove  $y$ 
    from  $S$ ;
end
Complete  $T$  to form a Hamiltonian cycle;

```

**Algorithm 5.** Serial Dictatorship for Max TSP

## 5 Conclusion

In this paper we study ordinal algorithms, i.e., algorithms which are aware only of preference orderings instead of the hidden weights or utilities which generate such orderings. Perhaps surprisingly, our results indicate that for many problems including Matching,  $k$ -sum clustering, Densest Subgraph, and Traveling Salesman, ordinal algorithms perform almost as well as algorithms which know the underlying metric weights, *even when the agents involved can lie about their preferences*. This indicates that for settings involving strategic agents where it is expensive, or impossible to obtain the true numerical weights or utilities, one can use ordinal mechanisms without much loss in welfare.

How do these algorithms stand in comparison to unconstrained ordinal algorithms that do not obey truthfulness? In the full version of this paper, we present non-truthful, ordinal algorithms for the same set of problems including a 4-approximation algorithm for Densest subgraph and a 1.88-approximation algorithm for Max TSP. In conjunction with the ordinal 1.6-approximation algorithm for perfect matching from [4], the improved approximation factors indicate a clear separation between the two classes of algorithms. On the surface, the improvement is not surprising since in many settings, truthfulness often places strong constraints on the set of allowed algorithms and techniques; indeed, all of our truthful mechanisms are derived using the three simple techniques outlined in Sect. 2. That said, given the absence of matching lower bounds in this work, the resolution of the gap between these two classes of algorithms is perhaps the most important question that is yet to be addressed.

The full version of this paper is publicly available on arXiv and can be accessed at <https://arxiv.org/abs/1610.04069>.

**Acknowledgements.** This work was supported in part by NSF awards CCF-1527497 and CNS-1218374.

## References

1. Georgios, A., Birmpas, G., Markakis, E.: On truthful mechanisms for maximin share allocations. In: Proceedings of IJCAI (2016)
2. Anshelevich, E., Bhardwaj, O., Postl, J.: Approximating optimal social choice under metric preferences. In: Proceedings of AAAI (2015)

3. Anshelevich, E., Postl, J.: Randomized social choice functions under metric preferences. In: Proceedings of IJCAI (2016)
4. Anshelevich, E., Sekar, S.: Blind, greedy, random: algorithms for matching and clustering using only ordinal information. In: Proceedings of AAAI (2016)
5. Boutilier, C., Caragiannis, I., Haber, S., Tyler, L., Procaccia, A.D., Sheffet, O.: Optimal social choice functions: a utilitarian view. *Artif. Intell.* **227**, 190–213 (2015)
6. Caragiannis, I., Filos-Ratsikas, A., Frederiksen, S.K.S., Hansen, K.A., Tan, Z.: An exact analysis of serial dictatorship. In: Proceedings of WINE, Truthful Facility Assignment with Resource Augmentation (2016)
7. Caragiannis, I., Nath, S., Procaccia, A.D., Shah, N.: Subset selection via implicit utilitarian voting. In: Proceedings of IJCAI (2016)
8. Caragiannis, I., Procaccia, A.D.: Voting almost maximizes social welfare despite limited communication. *Artif. Intell.* **175**(9–10), 1655–1671 (2011)
9. Dobzinski, S., Nisan, N., Schapira, M.: Truthful randomized mechanisms for combinatorial auctions. *J. Comput. Syst. Sci.* **78**(1), 15–25 (2012)
10. Dughmi, S., Ghosh, A.: Truthful assignment without money. In: Proceedings of ACM (2010)
11. Feldman, M., Fiat, A., Golomb, I.: On voting and facility location. In: Proceedings of ACM EC (2016)
12. Thomas, A.F., Khellaf, M.: A class of bounded approximation algorithms for graph partitioning. *Networks* **20**(2), 181–195 (1990)
13. Filos-Ratsikas, A., Frederiksen, S.K.S., Zhang, J.: Social welfare in one-sided matchings: random priority and beyond. In: Lavi, R. (ed.) SAGT 2014. LNCS, vol. 8768, pp. 1–12. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-44803-8\\_1](https://doi.org/10.1007/978-3-662-44803-8_1)
14. Hassin, R., Rubinstein, S., Tamir, A.: Approximation algorithms for maximum dispersion. *Oper. Res. Lett.* **21**(3), 133–137 (1997)
15. Rao Kosaraju, S., Park, J.K., Stein, C.: Long tours and short superstrings. In: Proceedings of FOCS (1994)
16. Kowalik, L., Mucha, M.: Deterministic  $7/8$ -approximation for the metric maximum TSP. *Theoret. Comput. Sci.* **410**(47), 5000–5009 (2009)
17. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. *ACM Trans. Econ. Comput.* **1**(4), 18 (2013)
18. Sekharipuram, S.R., Rosenkrantz, D.J., Tayi, G.K.: Heuristic and special case algorithms for dispersion problems. *Oper. Res.* **42**(2), 299–310 (1994)
19. Troncoso-Pastoriza, J.R., Katzenbeisser, S., Utku Celik, M.: Privacy preserving error resilient DNA searching through oblivious automata. In: Proceedings of ACM CCS (2007)

# Computer-Aided Verification for Mechanism Design

Gilles Barthe<sup>1</sup>, Marco Gaboardi<sup>2</sup>, Emilio Jesús Gallego Arias<sup>3</sup>,  
Justin Hsu<sup>4</sup>(✉), Aaron Roth<sup>4</sup>, and Pierre-Yves Strub<sup>1</sup>

<sup>1</sup> IMDEA Software Institute, Madrid, Spain

<sup>2</sup> University at Buffalo, SUNY, Buffalo, USA

<sup>3</sup> MINES ParisTech, Paris, France

<sup>4</sup> University of Pennsylvania, Philadelphia, USA

email@justinh.su

**Abstract.** We explore techniques from computer-aided verification to construct formal proofs of incentive properties. Because formal proofs can be automatically checked, agents do not need to manually check the properties, or even understand the proof. To demonstrate, we present the verification of a sophisticated mechanism: the generic reduction from Bayesian incentive compatible mechanism design to algorithm design given by Hartline, Kleinberg, and Malekian. This mechanism presents new challenges for formal verification, including essential use of randomness from both the execution of the mechanism and from the prior type distributions.

## 1 Introduction

Recent years have seen a surge of interest in mechanism design, as researchers explore connections between computer science and economics. This fruitful collaboration has produced many sophisticated mechanisms, including mechanism deployed in high-stakes auctions. Many mechanisms satisfy properties that *incentivize* agents to behave in a straightforward and easily modeled manner; the gold standard properties are *dominant strategy truthful* (in settings of complete information) and *Bayesian incentive compatible* (in settings of incomplete information). While existing mechanisms are impressive achievements, their increasing complexity raises two concerns.

The first concern is correctness. As mechanisms become more sophisticated, proofs of their incentive properties have also grown in complexity, sometimes involving delicate reasoning about randomization or tedious case analysis. Complex mechanisms are also more prone to implementation errors. The second concern is more subtle. At its heart, mechanism design is algorithm design together with a predictive model of how agents will decide to behave. Unlike algorithm

---

The full version of this paper is available at <https://arxiv.org/abs/1502.04052>.

design, where correctness can be verified in a vacuum, the success of a mechanism requires a realistic behavioral model of the participants. How will agents behave when faced with a complex mechanism?

Different behavioral models assume different answers to this question. At one extreme, we may assume that agents will coordinate to play a Nash equilibrium of the game and we can study concepts like the *price of anarchy* [8, 22]. However, Nash equilibria are generally not unique, requiring coordination and communication to achieve [16]. Even if information is centralized, equilibria can be computationally hard to find [12]. Assuming that agents play at a Nash equilibrium may be unrealistic unless agents possess strong computational resources.

At the other extreme, we may ask for mechanisms which are dominant strategy truthful or Bayesian incentive compatible. In such mechanisms, agents can do no better than truthfully reporting their type, even in the worst case or in expectation over the other agents' types. These solution concepts assume little about the bidders: When interacting with truthful mechanisms, agents do not have to engage in complicated counter-speculation, communication, or computation—they merely have to tell the truth!

However, even with mechanisms that are dominant strategy truthful or Bayesian incentive compatible, participating agents must still *believe* that the mechanism is truthful. For complicated mechanisms this is no small matter, as the incentive properties may require significant domain expertise to verify. We are not the first to raise these concerns. When designing the FCC auction for reallocating radio spectrum, Milgrom and Segal [20] advocated an “*obviously* strategy-proof” mechanism; formalizing this notion is an ongoing area of investigation [19]. However, some useful mechanisms are just too complex to be obvious. Instead of restricting mechanisms, can we give users evidence for the incentive properties?

In this work, we consider using *formal proofs* as certificates. Formal proofs bear a resemblance to pen-and-paper proofs, but they are constructed in a rigorous fashion: They use a formal syntax, have a precise interpretation as a mathematical proof, and can be built with a rich palette of computer-assisted proof-construction tools. Compared to pen-and-paper proofs, the major benefit of formal proofs is that once constructed, they can be checked independently and fully automatically by a *proof checker* program.

Several previous works have explored formal methods for verifying mechanisms; Kerber et al. [18] provide an extensive survey. Broadly speaking, prior work falls into two groups. *Automated* approaches check properties via extensive search, guided by intelligent heuristics. These techniques are more suited to verifying simpler properties of mechanisms, perhaps instantiated on a specific input; properties like BIC lie beyond the reach of existing approaches.

More manual (sometimes called *interactive*) techniques divide the verification task into two separate stages. In the first stage, the formal proof is *constructed*. This step typically involves human assistance, perhaps encoding the mechanism in a specific form or constructing a formal proof. With the help of the human, these techniques can prove rich properties like BIC and support the level of

generality that is typical of existing proofs—say, for an arbitrary number of agents, or for any type space. In the second stage, the formal proof is *checked*. This step proceeds fully automatically: a proof checking program verifies that the formal proof is constructed correctly. This neat division of the verification task is a natural fit for mechanism design. We could imagine that the mechanism designer—a sophisticated party who is intimately familiar with the details of the proof—has the resources and knowledge to construct the formal proof. This proof could then be transmitted to the agents, who can automatically check the proof with no knowledge of the details.

The main difference between manual techniques is in the amount of human labor for proof construction, the most challenging phase. Existing verification approaches formalize the proof at a level that is far more detailed than existing proofs on paper, requiring extensive expertise in formal methods. Furthermore, existing works focus on general correctness properties—the output of a mechanism should be a partition, the prices should be non-negative, etc., rather than incentive properties.

In our work, we look to combine the best of both worlds: enabling a high level of automation during proof construction, while supporting formal proofs that can capture rich incentive properties. To demonstrate our approach, the primary technical contribution of our paper is a challenging case study: a formal proof of Bayesian incentive compatibility (BIC) for the generic reduction from algorithm design to mechanism design by Hartline, Kleinberg, and Malekian [17]. This example is an attractive proof-of-concept for several reasons.

1. Both the reduction and the proof of Bayesian incentive compatibility are complex. The mechanism is far from obviously strategy proof—indeed, the proof is a research contribution first published at SODA 2011.
2. It is a general reduction, so certifying its correctness once certifies the incentive properties for any instantiation of the reduction.
3. It relies on truthfulness of the Vickrey-Clarke-Groves (VCG) mechanism. As part of our efforts, we provide the first formal verification of truthfulness for this classical mechanism.
4. It employs randomization both within the algorithm and within the agent behavior—agent types are drawn from the known Bayesian prior.

The formal proofs bear a resemblance to the original proof, both easing formalization and making the proofs more accessible to the mechanism design community.

To formalize the proofs, we adapt techniques from program verification. We view incentive properties as a property of the mechanism and the agent’s payoff function, both expressed as programs. Formal verification has developed sophisticated tools and techniques for verifying program properties, but general-purpose tools require significant manual work. Verifying even moderately complex mechanisms seems well beyond the reach of current technology. To ease the task, we view incentive properties as *relational properties*: statements about the relationship between the outputs in two runs of the same program. Specifically, consider the program which calculates an agent’s payoff under the mechanism and assume



agents play their true value in the first run, while an agent may deviate arbitrarily in the second run. If the output in the first run is at least the output in the second run, then the mechanism is incentive compatible.

With this point of view, we can use tools specialized for relational properties. Such tools are significantly easier to use and have achieved notable successes for verifying proofs from differential privacy and cryptography. We use `HOARe2`, a recently-developed programming language that can express and check relational properties [4]. `HOARe2` has been used to verify differential privacy and basic truthfulness in simple mechanisms under complete information, like the fixed price auction and the random sampling mechanism of Goldberg et al. [13] for digital goods.

Our work goes significantly beyond prior efforts in several respects. First, the mechanism we verify is significantly more complex than previously considered mechanisms, and we analyze all uses of the reduction, rather than just a single instance. Second, the mechanism operates in the partial information setting, so the proof requires careful reasoning about randomization (from both the mechanism and from the prior distribution on types).

The main strength of our approach lies in the high degree of automation during *proof construction*. Once the mechanism and payoff functions have been encoded as programs, and once we have supplied some annotations, we can construct most of the formal proof automatically with the aid of automated solvers. However, there are a handful of particularly complex steps that `HOARe2` fails to automatically prove. To finish the proof, we manually build a formal proof for these missing pieces using `EasyCrypt`, a proof assistant for relational properties, and `Coq`, a general purpose proof assistant.<sup>1</sup>

*Related Work.* Closely related to our work, a recent paper by Caminati et al. [7] uses the theorem prover Isabelle to verify basic properties of the celebrated Vickrey-Clarke-Groves (VCG) mechanism. They consider general auction properties: the prices should be non-negative, VCG should produce a partition of goods, etc. Moreover, their framework can be used to automatically produce a correct, executable implementation of the mechanism. While their work demonstrates that formal verification can be applied to verify properties of mechanisms, their results are limited in two respects. First, they do not consider incentive properties, arguably the properties at the heart of mechanism design. Second, they apply general techniques from computer-aided verification that are not specifically tailored to mechanism design, requiring substantial effort to produce the machine-checked proof. Our work uses verification techniques that are particularly suited for incentive properties.

In the extended version we provide a primer on formal verification and discuss related work; a recent survey by Kerber et al. [18] provides a comprehensive review of formal methods for verifying mechanism design properties. The algorithmic game theory literature has for the most part ignored the problem of

---

<sup>1</sup> Our formal proofs, along with code for the `HOARe2` tool, are available online: <https://github.com/ejgallego/HOARe2/tree/master/examples/bic>.

*verifying* incentive properties, with a few notable exceptions. Recently, Brânzei and Procaccia [6] define *verifiably truthful mechanisms*. Informally, such a mechanism is selected from a fixed family of mechanisms such that for every truthful mechanism in that family, a certificate showing truthfulness can be found in polynomial time. Brânzei and Procaccia [6] consider mechanisms represented as decision trees and show that for the one-dimensional facility location problem, truthfulness for mechanisms in this class can be efficiently verified by linear programming. In contrast, we investigate significantly more complex mechanisms in exchange for forgoing worst-case polynomial time complexity.

Mu’alem [21] considers the problem of *property testing* for truthfulness in single parameter domains, which reduces to testing for a variant of monotonicity. Mu’alem [21] gives a tester that can test whether there exist payments that guarantee that truthful reporting is a dominant strategy with probability  $1 - \epsilon$ , given a  $\text{poly}(1/\epsilon)$  number of arbitrary evaluations of an allocation rule and assuming agents have uniformly random valuations. In contrast, we assume direct access to the code specifying the auction instead of merely black box access to the allocation rule, and we achieve verification of exact truthfulness, not just approximate truthfulness. We are also able to verify mechanisms in more complex settings, e.g., arbitrary type spaces, randomized mechanisms, and arbitrary priors.

Our work is also related to the literature on automated mechanism design, initiated by Conitzer and Sandholm [11] (see Sandholm [23] or Conitzer [10, Chapter 6] for an introduction). In broad strokes, automated mechanism design seeks to generate truthful mechanisms which optimize the designer’s objectives. This is often accomplished by solving explicitly for the distribution on outcomes defining a mechanism using a mixed integer linear program encoding the incentive constraints and objective, an NP hard problem that can often be solved efficiently on typical instances [10]. Automated mechanism design targets a more difficult problem than we do: it seeks not just to *verify* the truthfulness of a given mechanism, but to *optimize* over all truthful mechanisms. However, these techniques have some limitations: they produce explicit representations of mechanisms requiring size exponential in the number of bidders, and they use an explicit integer linear program, requiring a finite type space. In contrast, by only requiring full automation for proof verification and not proof construction, we are able to use a much more sophisticated toolkit—including symbolic manipulation, not just numeric optimization—and verify significantly more complex mechanisms that can have infinite outcome and type spaces.

## 2 Main Example: RSM

As our main proof of concept, we verify that the Replica-Surrogate-Matching (RSM) mechanism due to Hartline et al. [17] is Bayesian incentive compatible. The RSM mechanism reduces mechanism design to algorithm design: given an algorithm  $A$  that takes in agents’ reported types and selects an outcome, the RSM mechanism turns  $A$  into a Bayesian incentive compatible mechanism. Accordingly, our formal proof will carry over to any instantiation of RSM. We

first review the original proof by Hartline et al. [17]. Then, we describe our verification process, from pseudocode to a fully verified mechanism.

Let’s begin with the standard notion of Bayesian incentive compatibility. We assume there are  $n$  agents, each with a *type*  $t_i$  drawn from some set of types  $T$ . Furthermore, we have access to a distribution  $\mu$  on types, the *prior*. A *mechanism* is a (possibly randomized) function from the inputs—one per agent—to a single *outcome*  $o$  from set  $O$ , and a real-valued *payment*  $p_i$  for each agent. Without loss of generality, we will assume that the agents each report a type from  $T$  as their input. Agents have a valuation  $v(t, o)$  for type  $t$  and outcome  $o$ . Agents have *quasi-linear utility*: their utility for outcome  $o$  and payment  $p$  is  $v(t, o) - p$ . We will write  $(s, t_{-i})$  for the vector obtained by inserting  $s$  into the  $i$ th slot of  $t$ . Then, we want to check the following property.

**Definition 1.** A mechanism  $M$  is Bayesian incentive compatible (BIC) if for every agent  $i$  and types  $t_i, t'_i$ , we have

$$\mathbb{E}_{t_{-i} \sim \mu^{n-1}} [v(t_i, M(t_i, t_{-i})) - p_i(t_i, t_{-i})] \geq \mathbb{E}_{t_{-i} \sim \mu^{n-1}} [v(t_i, M(t'_i, t_{-i})) - p_i(t_i, t_{-i})].$$

The expectation is taken over the types  $t_{-i}$  of the other agents (drawn independently from  $\mu$ ) and any randomness used by the mechanism.

### 2.1 The RSM Mechanism

Now, let’s consider the mechanism: the RSM mechanism in the “idealized model” by Hartline et al. [17]. We will first recapitulate their proof, before explaining in detail how we verify it.

Hartline et al. RSM is a construction for turning an *algorithm*  $A : T^n \rightarrow O$  into a BIC mechanism. The process is easy to describe: each agent individually transforms their type  $t_i$  to a *surrogate type*  $s_i$  by applying the Replica-Surrogate-Matching procedure  $R$ . This procedure also produces a payment  $p_i$  for the agent. Then, the surrogates  $s$  are fed into the algorithm  $A$ , which produces the final outcome.

1. Pick  $i$  uniformly at random from  $[m]$ ;
2. Build a *replica type profile*  $\mathbf{r}$  by taking  $m - 1$  samples from  $\mu$  for  $\mathbf{r}_{-i}$ , setting  $r_i = t$ ;
3. Build a *surrogate type profile*  $\mathbf{s}$  by taking  $m$  samples from  $\mu$ ;
4. Build a bipartite graph with nodes  $\mathbf{r}, \mathbf{s}$ , and edges with weight

$$w(\mathbf{r}, \mathbf{s}) = \mathbb{E}_{t_{-i} \sim \mu^{n-1}} [v(\mathbf{r}, A(\mathbf{s}, t_{-i}))];$$

5. Run the VCG procedure on the generated graph, and return the surrogate  $s$  that is matched to the replica in slot  $i$ , and the appropriate payment  $p$ .

**Fig. 1.** Procedure  $R$  with parameter  $m$

The procedure  $R$  is described in Fig. 1. Let  $m$  be an integer parameter—the number of replicas. Given input type  $t$ , we take  $m - 1$  independent samples from

$\mu$ , the  $(r)$ eplicas. We then take  $m$  independent samples from  $\mu$ , the  $(s)$ urrogates. Finally, we select an index  $i$  uniformly at random from  $[m]$ , and place the original type  $t$  in the  $i$ th “slot” of the replicas  $\mathbf{r}$ . We will consider the replicas as “buyers”, and the surrogates as “goods”, and assign a numeric “value” for every pair of buyer and good. The value of replica  $r$  for surrogate  $s$  is set to be

$$w(r, s) = \mathbb{E}_{t_{-i} \sim \mu^{n-1}}[v(r, A(s, t_{-i}))], \quad (1)$$

that is, the expected utility of an agent with true type  $r$  reporting type  $s$ . Finally, RSM runs the well-known Vickrey-Clarke-Groves mechanism [9, 14, 24] to match each replica with a surrogate in this market. The output is the surrogate matched to replica in slot  $i$  (the original type  $t$ ), along with the payment charged.

*The Original Proof.* The proof of BIC from Hartline et al. [17] proceeds in two steps. First, they show that  $R$  is *distribution preserving*.

**Lemma 1 (Hartline et al. [17]).** *Sampling a type  $t \sim \mu$  as input to  $R$  gives the same distribution ( $\mu$ ) on the surrogates output.*

*Proof.* When  $R$  constructs the list of buyers before applying VCG, the distribution over buyers is simply  $m$  independent samples from  $\mu$ , no matter the value of  $i$ . So, we can delay sampling  $i$  and selecting the surrogate until after running VCG (via the principle of deferred decision). VCG produces a perfect matching of replicas to surrogates, and the surrogates are also  $m$  independent samples from  $\mu$ . So, sampling a random replica  $i$  and returning the matched surrogate is an unbiased sample from  $\mu$ .

With the lemma in hand, Hartline et al. [17] show that RSM is BIC.

**Theorem 1 (Hartline et al. [17]).** *The RSM mechanism is BIC.*

*Proof.* Consider bidder  $i$  with type  $t_i$ , and fix the randomness for bidder  $i$ . In the VCG procedure of  $R$ , the value of  $i$ 's replica for surrogate  $s$  is  $w(t_i, s)$ : the expected utility for submitting  $s$  to  $A$  while having true type  $t_i$ , assuming that all other inputs to  $A$  are drawn from  $\mu$ .

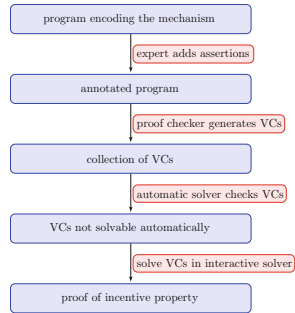
In the RSM mechanism, the other inputs to  $A$  are computed by sampling a type  $t_j \sim \mu$ , and taking the surrogate produced by  $R(t_j)$ . By Lemma 1, the distribution over surrogates is  $\mu$ . Therefore,  $w(t_i, s)$  is bidder  $i$ 's expected utility in the RSM mechanism for ending up matched to  $s$ . Since VCG is incentive compatible, bidder  $i$  has no incentive to deviate to any other bid  $t'_i$ . By taking expectation over the randomness of  $i$ , we get the result.

Crucially, Theorem 1 relies on the truthfulness property of the VCG mechanism. We have also verified this property but we postpone our discussion to the extended version; the verification of RSM is more interesting.

### 3 Verifying RSM

Now that we have seen the mechanism, we present our verification step by step.

1. We write the RSM mechanism as a program in the HOARE<sup>2</sup> programming language.
2. We annotate the program with assertions expressing the BIC property, and some additional intermediate facts (lemmas).
3. The tool automatically generates the *verification conditions* (VCs), which imply BIC.
4. The tool uses automatic solvers to check the verification conditions; they may fail to prove some assertions.
5. Finally, we prove the remaining verification conditions by using an interactive prover.



The outcome of these five steps is a formal proof that the RSM mechanism enjoys the BIC property. In the following, we will combine the description of different steps in the same subsection.

#### Step 1: Modeling the Mechanism

To express RSM as a program, we will code a single agent’s utility function when running the RSM mechanism, when all the other agents report truthfully and have types drawn from  $\mu$ . Remembering that we consider truthfulness as a *relational property*, we will then reason about what happens when the agent reports truthfully, compared to what happens when the agent deviates.

We model types and outcomes as drawn from (unspecified) sets  $T$  and  $O$ , and we assume an algorithm `alg` mapping  $T^n \rightarrow O$ . We will consider what happens when the first bidder deviates. This is without loss of generality: if  $j$  deviates, we can consider the RSM mechanism with `alg` replaced by a version `alg'` that first rotates the  $j$ th bidder to the first slot, when proving BIC for the first bidder under `alg'` implies BIC for the  $j$ th bidder under  $A$ . For the values, we will assume an arbitrary valuation function `value` mapping  $T \times O \rightarrow \mathbb{R}$ . In the code, we will write `mu` for the prior distribution  $\mu$ .

Let’s begin by coding the RSM transformation  $R$ , which transforms an agent’s type into a surrogate type and a payment. It will be convenient to separate the randomness from  $R$ . We encode  $R$  as a deterministic function `Rsmdet`, which takes as input the agent number `j`, the random coins `coins`, and the input type `report`. We will have `Rsmdet` take an additional parameter `truety` representing an agent’s true type. This variable does not show up in the code—RSM does not have access to this information—but will be useful later for expressing Bayesian incentive compatibility as a relational property. We will model the slot as a natural number.

In the extended version we will discuss our treatment of VCG in more detail, but it is enough to know that VCG takes a list of buyers and a list of goods. VCG will output a permutation of goods (representing the assignment), and a corresponding list of payments. In Fig. 2, bolded words are keywords and primitive operations of HOARE<sup>2</sup>. For a brief explanation, line (2) names the three components of `coins`: the replicas `rs-i`, the surrogates `ss`, and the slot `i`; line (3) puts the agent's input type `report` in the proper slot for the replicas; line (4) calls VCG on the list of buyers `vcgbuyers` produced at line (3) and the list of surrogates `ss` as goods; and line (5) returns the surrogate and payment.

```

1 def Rsmdet(j, coins, truety, report) =
2   (rs-i, ss, i) = coins;
3   vcgbuyers = (report, rs-i);
4   (surrs, pays) = Vcg(vcgbuyers, ss);
5   return (surrs[j], pays[j])

```

Fig. 2. Defining RSM

```

1 def Expwts(j, r, s) =
2   sample others-j = mun-1;
3   algInput = (s, others-j);
4   outcome = alg(algInput);
5   return expect_num { value(r, outcome) }

```

Fig. 3. Defining weights

The `Expwts` function in Fig. 3 implements the  $w$  function from Eq. (1), with the additional parameter  $j$  to indicate the agent. In Fig. 3, line (2) samples  $n - 1$  types `others-j` from  $\mu$  for the other agents. These are the types on which the expectation is taken in Eq. (1). Line (4) uses the algorithm `alg` to compute the outcome `outcome` when the agent  $j$  report type `s`. Finally, the `expect_num` on line (5) takes the expected value of the distribution over reals defined by evaluating the value function `value` on the true type `r` and on the randomized outcome of the `alg`.

To check the BIC property, we will code the expected utility for the first bidder and then check that it is maximized by truthful reporting. To break down the code, we will suppose that the function takes in a list of functions `othermoves` that transform each of the other bidder's type.

```

1 def Util(othermoves, myty, mybid) =
2   return (expect rsmcoins Helper)
3
4 def Helper(coins) =
5   (mysur, mypay) =
6   Rsmdet(1, coins, myty, mybid);
7   myval = expect_num {
8     for i = 1 .. n - 1:
9       sample othersurs[i] =
10        (sample otherty = mu;
11         othermoves[i](otherty));
12        algInput = (mysur, othersurs);
13        outcome = alg(algInput);
14        value(myty, outcome) };
15   return (myval - mypay)

```

Fig. 4. Defining utility

```

1 def Others(j, t) =
2   sample coins = rsmcoins;
3   (s, p) = Rsmdet(j, coins, t);
4   return s
5
6 def MyUtil(ty, bid) = Util(Others, ty, bid)

```

Fig. 5. Defining other reports

The distribution `rsmcoins` defines the distribution over the coins to  $R$ , i.e., sampling the replicas `r-i`, the surrogates `s`, and the coin `i`. We encoded this distribution in HOARE<sup>2</sup>, but we elide it for lack of space. In the code in Fig. 4,

on line (2) we take expectation of the function `Helper` over the distribution `rsmcoins`, with `expect`. In `Helper`, we then call `RsmDET` on line (6) to compute the surrogate and payment for the agent, passing 1 since we are calculating the utility for the first agent. We sample the other agents' types and transform them on lines (9–11), and we take expectation of the first agent's value for the outcome on lines (7–14). Finally, we subtract off the payment on line (15), giving the final utility for the first agent.

To complete our modeling of RSM, in Fig. 5 we plug in `Others` into the utility function: it simply takes an agent number and a type as input, samples the coins from `rsmcoins`, and returns the surrogate from calling `RsmDET`. So far, we have just written code describing how to implement the RSM mechanism and how to calculate the utility for a single bidder. Now, we express the BIC property as a property about this program and check it with HOARE<sup>2</sup>.

## Step 2: Adding Assertions

We specify properties in HOARE<sup>2</sup> by annotating variable and functions with assertions of the form  $\{x :: Q \mid \phi\}$ , read as “ $x$  is an element of set  $Q$  and satisfies the logical formula  $\phi$ ”. These assertions serve two purposes: (1) they express facts to be proved about the code and (2) they assert mathematical facts about primitive operations like `expect` and `expect_num`. The system will then formally verify that the first kind of annotations are correct, while assuming the assertions of the second kind as axioms.

A key feature of HOARE<sup>2</sup> is that the assertion  $\phi$  is *relational*: it can refer to two copies of each variable  $x$ , denoted  $x_1$  and  $x_2$ . Roughly, we may make assertions about two runs of the same program where in the first program we use variables  $x_1$ , and in the second run we use variables  $x_2$ . For instance, truthfulness corresponds to the following assertions:

$$\begin{array}{ll} \{ty :: T \mid ty_1 = ty_2\} & \text{(true type is equal on both runs)} \\ \{bid :: T \mid bid_1 = ty_1\} & \text{(bid is the true type in the first run)} \\ \{utility :: \mathbb{R} \mid utility_1 \geq utility_2\} & \text{(utility is higher in the first run)} \end{array}$$

Our goal is to check these assertions for the function `MyUtil`, which computes an agent's utility in expectation over the other types. Along the way we will use several intermediate facts, encoded as assertions in HOARE<sup>2</sup>. Assertions on primitive operations, like `expect` and `expect_num`, are the axioms. Assertions on larger chunks of code are proved correct from the assertions on the subcomponents.

*Monotonicity of Expectation.* Since the BIC property refers to *expected* utility, we use an expectation operation `expect` when computing an agent's utility (line (2) of the `Util` code). To show BIC, we need a standard fact about *monotonicity* of expected value: for functions  $f \leq g$ ,  $\mathbb{E}[f] \leq \mathbb{E}[g]$  taken over the same distribution. This can be encoded with an annotation for `expect`:

$$\text{distr} \{c :: C \mid c_1 = c_2\} \rightarrow \{f :: C \rightarrow \mathbb{R} \mid \forall x. f_1(x) \leq f_2(x)\} \rightarrow \{e :: \mathbb{R} \mid e_1 \leq e_2\}.$$

This annotation indicates that `expect` is a function that takes two arguments and returns a real number. In each of the three components, the annotation before the bar specifies the type of the value: The first argument is a distribution over  $C$ , the second argument is a real-valued function  $C \rightarrow \mathbb{R}$ , and the return value is a real number. The logical formulas after the pipe describe how two runs of the expectation function are related. The first component states that in the two runs, the distributions are the same. The second component states that the function  $f$  in the first run is pointwise less than  $f$  in the second run. The final component asserts that the expected value—a real number—is less on the first run than on the second run.

If think of the distribution as being over the coins `rsmcoins`, this fact allows us to prove deterministic truthfulness for each setting of the coins, then take expectation over the coins in order to show truthfulness in expectation. This is what we need to prove for the BIC property, and is precisely the first step in the original proof of Theorem 1.

*Distribution Preservation.* When we consider a single agent, truthful bidding may not be BIC for arbitrary transformations of the other agents' types (`othermoves` in the `Util` code). As indicated by Lemma 1, we also need the transformation to be distribution preserving: the output distribution on surrogates must be the same as the distribution on input types.

Much as we did above, we can capture this property with appropriate annotations. While we have so far used rather simple formulas  $\phi$  that only mention variables in  $\{x :: T \mid \phi\}$ , in general the formulas  $\phi$  can describe assertions about programs. We can annotate the `othermoves` argument to `Util` to require distribution independence:

$$\{\text{othermoves} : \text{list } (T \rightarrow \text{distr } T) \mid \forall j \in [n]. (\text{sample } \text{ot} = \text{mu}; \text{othermoves}[j](\text{ot})) = \text{mu}\}$$

To read this, `othermoves` is a list of functions  $f_j$  that take a type and return a distribution on types, such that if we sample a type from  $\text{mu}$  and feed it to  $f_j$ , the resulting distribution (including randomness over the initial choice of type) is equal to  $\text{mu}$ . In other words, this asserts the distribution preservation property of Lemma 1 for each of the other agent's transformations.

*Facts about VCG.* Recall that `Vcg` takes a list of bidders and a list of goods, and produces a permutation of the goods and a list of payments as output. In our case, the bidders and goods are both represented as types in  $T$ , so we can annotate the `Vcg` as:

$$\{\text{buys} :: \text{list } T\} \rightarrow \{\text{goods} :: \text{list } T\} \rightarrow \{(\text{alloc}, \text{pays}) :: \text{list } T \times \text{list } \mathbb{R} \mid \text{vcgTruth} \wedge \text{vcgPerm}\}.$$

The two assertions `vcgTruth` and `vcgPerm` reflect two facts about VCG. The first is that VCG is incentive compatible; this can be encoded like we have already seen, with a slight twist: We require that VCG is IC for a deviation by *any*



player rather than just the first player, since the possibly deviating player may be in any slot. More precisely, we define the formula

$$\begin{aligned} \text{vcgTruth} &:= \forall j \in [m]. (\text{bids}_{-j,1} = \text{bids}_{-j,2}) \implies \text{Expwts}(j, \text{bids}_1[j], \text{alloc}_1[j]) - \text{pays}_1[j] \\ &\geq \text{Expwts}(j, \text{bids}_1[j], \text{alloc}_2[j]) - \text{pays}_2[j]. \end{aligned}$$

We treat the bid in the first run ( $\text{bids}_1[j]$ ) as the true type, and the bid on the second run ( $\text{bids}_2[j]$ ) as a possible deviation—this is why we evaluate the  $j$ th bidder’s expected utility using the same true type in the two runs. The second fact we use is that VCG *matches* buyers to the goods. In fact, since the number of goods (surrogates) and the number of buyers (replicas) are equal, VCG produces a perfect matching. We express this by asserting that VCG outputs an assignment that is a permutation of the goods:

$$\text{vcgPerm} := \text{isPerm goods}_1 \text{ alloc}_1 \wedge \text{isPerm goods}_2 \text{ alloc}_2.$$

We verify these properties for a general version of VCG. The verification follows much like the current verification; we discuss the details in the extended version.

### Step 3: Handling Proof Obligations

After providing the annotations, HOARE<sup>2</sup> is able to automatically check most of the annotations with *SMT solvers*—fully automated solvers that check the validity of logical formulas. Such solvers are a staple of modern formal verification. While the underlying problem is often undecidable, modern solvers employ sophisticated heuristics that can efficiently handle large formulas in practice.

We are able to use SMT solvers to automatically check all but three proof obligations; for these three facts the SMT solvers time out without finding a proof. The first two are uninteresting, and we manually construct the formal proof using the Coq proof assistant. The last obligation is more interesting: it corresponds to Lemma 1. Concretely, when we define an agent’s expected utility

$$\text{def MyUtil}(\text{ty}, \text{bid}) = \text{Util}(\text{Others}, \text{ty}, \text{bid}),$$

recall that `Util` asserts that `Others` is distribution preserving. This is precisely Lemma 1, and the automated solvers fail to prove this automatically.

To handle this assertion we use a more manual tool called EasyCrypt [2, 3], a proof assistant that allows the user prove equivalence of two programs  $A$  and  $B$  by manually transforming the source code of  $A$  until the source code is identical to  $B$ . We prove that `Others` is equivalent to the program that simply samples from  $\mu$  by transforming the code for `Others` (including the code sampling the coins of the mechanism, `rsmcoins`) in several stages. We present the code in Fig. 6 with two replicas, for simplicity.

The proof boils down to showing that each step transforms a program to an equivalent program. Our starting point is `stage1`, the program that samples

```

def stage1 =
  sample ot = mu;
  Others(ot)

def stage2 =
  sample ot = mu;
  sample r' = mu;
  sample s1 = mu;
  sample s2 = mu;
  sample i = flip;

  if i then
    (r1,r2) = (ot,r');
  else
    (r1,r2) = (r',ot);

  bs = (r1,r2);
  gs = (s1,s2);

  (ss,ps) = Vcg(bs,gs);
  (o1,o2) = ss;

  if i then o1 else o2

def stage3 =
  sample ot = mu;
  sample r' = mu;
  sample s1 = mu;
  sample s2 = mu;

  (r1,r2) = (ot,r');

  bs = (r1,r2);
  gs = (s1,s2);

  (ss,ps) = Vcg(bs,gs);
  (o1,o2) = ss;

  sample i = flip;
  if i then o1 else o2

def stage4 =
  sample s1 = mu;
  sample s2 = mu;
  sample i = flip;
  if i then s1
    else s2

```

**Fig. 6.** Code transformations to prove Lemma 1.

an agent’s type from `mu` and runs `Others` on the sampled value. Unfolding the definition of `Others`, `Rsmdet`, `rsmcoins` and including the code that puts the agent’s input type in the proper slot for the replicas, we obtain program `stage2`. From there, the main step is to show that we don’t need to place the replicas in a random order before calling `Vcg`. Then, we can move the sampling for `i` down past the `Vcg` call, giving `stage3`. Finally, using the fact that the output assignment `ss` of `Vcg` is a permutation of the goods (`s1`, `s2`), we obtain the program `stage4` and conclude that this is equivalent to taking a single sample from `mu`. This chain of transformations has been verified with `EasyCrypt`.

## 4 Perspective

Now that we have presented our verification of the RSM mechanism, what have we learned and what does formal verification have to offer mechanism design going forward? In our experience, while formal verification of game theoretic mechanisms is by no means trivial, verification tools are maturing to a point where practical verification of complex mechanisms can be envisioned. Our verification of RSM, for instance, involved only coding the utility function and adding annotations, most of which can be checked automatically. The most time-consuming part was manually proving the last few assertions.

At the same time, the range of mechanisms that can be verified is less clear. There is an art to encoding a mechanism in the right way, and some mechanisms are easier to verify than others. Since we are trying to verify proofs, the crucial factor is the complexity of the *proof* rather than the complexity of the mechanism. Clean proofs where, each step reasons about localized parts of the program, are more amenable to verification; proof patterns—like universal truthfulness—also help.

In sum, formal verification can manage the increasing complexity of mechanisms by formally proving incentive properties for everyone—mechanism designers, mechanism users, and even mechanism programmers. We believe that the tools to verify one-shot mechanisms are already here. So, we propose a challenge:

Try using tools like HOARE<sup>2</sup> to verify your own mechanisms, putting formal verification techniques to the test. We hope that one day soon, verification for mechanisms will be both easy and commonplace.

**Acknowledgments.** We thank the anonymous reviewers for their careful reading; their suggestions have significantly improved this work. This work was partially supported by NSF grants TWC-1513694, CNS-1237235, CNS-1565365 and a grant from the Simons Foundation (#360368 to Justin Hsu).

## References

1. Barrett, C., Sebastini, R., Seshia, S.A., Tinelli, C.: Satisfiability modulo theories. In: Handbook of satisfiability, vol. 185. IOS press (2009)
2. Barthe, G., Grégoire, B., Héraud, S., Béguelin, S.Z.: Computer-aided security proofs for the working cryptographer. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 71–90. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-22792-9\\_5](https://doi.org/10.1007/978-3-642-22792-9_5)
3. Barthe, G., Dupressoir, F., Grégoire, B., Kunz, C., Schmidt, B., Strub, P.-Y.: EasyCrypt: a tutorial. In: Aldini, A., Lopez, J., Martinelli, F. (eds.) FOSAD 2012-2013. LNCS, vol. 8604, pp. 146–166. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10082-1\\_6](https://doi.org/10.1007/978-3-319-10082-1_6)
4. Barthe, G., Gaboardi, M., Gallego Arias, E.J., Hsu, J., Roth, A., Strub, P.-Y.: Higher-order approximate relational refinement types for mechanism design and differential privacy. In: ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), Mumbai, India, pp. 55–68 (2015). <http://arxiv.org/abs/1407.6845>
5. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006). doi:[10.1007/11761679\\_25](https://doi.org/10.1007/11761679_25)
6. Brânzei, S., Procaccia, A.D.: Verifiably truthful mechanisms. In: ACM SIGACT Innovations in Theoretical Computer Science (ITCS), Princeton, New Jersey (2014)
7. Caminati, M.B., Kerber, M., Lange, C., Rowat, C.: Sound auction specification and implementation. In: ACM SIGecom Conference on Economics and Computation (EC), Portland, Oregon, pp. 547–564 (2015). <http://doi.acm.org/10.1145/2764468.2764511>
8. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: ACM SIGACT Symposium on Theory of Computing (STOC), Baltimore, Maryland, pp. 67–73. ACM (2005). <http://dl.acm.org/citation.cfm?id=1060600>
9. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* **11**(1), 17–33 (1971)
10. Conitzer, V.: Computational aspects of preference aggregation. PhD thesis, IBM (2006)
11. Conitzer, V., Sandholm, T.: Complexity of mechanism design. In: Conference on Uncertainty in Artificial Intelligence (UAI), Edmonton, Alberta, pp. 103–110. Morgan Kaufmann Publishers Inc. (2002)
12. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. *SIAM J. Comput.* **39**(1), 195–259 (2009)
13. Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M., Wright, A.: Competitive auctions. *Games Econ. Behav.* **55**(2), 242–269 (2006)

14. Groves, T.: Incentives in teams. *Econ. J. Econ. Soc.* **41**(4), 617–631 (1973)
15. Halevi, S.: A plausible approach to computer-aided cryptographic proofs. *Cryptology ePrint Archive*, Report 2005/181 (2005). <https://eprint.iacr.org/2005/181.pdf>
16. Hart, S., Mansour, Y.: The communication complexity of uncoupled Nash equilibrium procedures. In: *ACM SIGACT Symposium on Theory of Computing (STOC)*, San Diego, California, pp. 345–353. ACM (2007)
17. Hartline, J.D., Kleinberg, R., Malekian, A.: Bayesian incentive compatibility via matchings. In: *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, San Francisco, California, pp. 734–747. SIAM (2011)
18. Kerber, M., Lange, C., Rowat, C.: An introduction to mechanized reasoning. *CoRR*, abs/1603.02478, 2016. <http://arxiv.org/abs/1603.02478>
19. Li, S.: Obviously strategy-proof mechanisms. *SSRN Electron. J.* <http://dx.doi.org/10.2139/ssrn.2560028>
20. Milgrom, P., Segal, I.: Deferred acceptance auctions and radio spectrum reallocation (2014). <http://www.as.huji.ac.il/sites/default/files/DA%20Heuristic%20Auctions%20June-2014.pdf>
21. Mu’alem, A.: A note on testing truthfulness. In: *Electronic Colloquium on Computational Complexity (ECCC)*, No. 130 (2005)
22. Roughgarden, T.: *Selfish Routing and the Price of Anarchy*, vol. 174. MIT Press, Cambridge (2005)
23. Sandholm, T.: Automated mechanism design: a new application area for search algorithms. In: Rossi, F. (ed.) *CP 2003*. LNCS, vol. 2833, pp. 19–36. Springer, Heidelberg (2003). doi:10.1007/978-3-540-45193-8\_2
24. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Finan.* **16**(1), 8–37 (1961)

# Smoothness for Simultaneous Composition of Mechanisms with Admission

Martin Hoefer, Thomas Kesselheim, and Bojana Kodric<sup>(✉)</sup>

MPI Informatik and Saarland University, Saarbrücken, Germany  
{martin.hoefer,thomas.kesselheim,bojana.kodric}@mpi-inf.mpg.de

**Abstract.** We study social welfare of learning outcomes in mechanisms with admission. In our repeated game there are  $n$  bidders and  $m$  mechanisms, and in each round each mechanism is available for each bidder only with a certain probability. Our scenario is an elementary case of simple mechanism design with incomplete information, where availabilities are bidder types. It captures natural applications in online markets with limited supply and can be used to model access of unreliable channels in wireless networks. If mechanisms satisfy a smoothness guarantee, existing results show that learning outcomes recover a significant fraction of the optimal social welfare. These approaches, however, have serious drawbacks in terms of plausibility and computational complexity. Also, the guarantees apply only when availabilities are stochastically independent among bidders. In contrast, we propose an alternative approach where each bidder uses a single no-regret learning algorithm and applies it in all rounds. This results in what we call *availability-oblivious* coarse correlated equilibria. It exponentially decreases the learning burden, simplifies implementation (e.g., as a method for channel access in wireless devices), and thereby addresses some of the concerns about Bayes-Nash equilibria and learning outcomes in Bayesian settings. Our main results are general composition theorems for smooth mechanisms when valuation functions of bidders are lattice-submodular. They rely on an interesting connection to the notion of correlation gap of submodular functions over product lattices.

## 1 Introduction

Truthful mechanism design is a central challenge at the intersection of economics and computer science, but many fundamental techniques are only very rarely used in practice. For example, sponsored search auctions are used on a daily basis and generate billions of dollars in revenue, but they are based on simple and non-truthful procedures to allocate ads on search result pages. In contrast, truthful mechanisms often involve heavy algorithmic machinery, complicated allocation techniques, or other hurdles to easy and transparent implementation.

---

This work is supported by DFG through Cluster of Excellence MMCI.

A recent trend is to study non-truthful and conceptually “simple” mechanisms for allocation in markets and their inherent loss in system performance. The idea is to analyze the induced game among the bidders and bound the quality of (possibly manipulated) outcomes in equilibrium. In a seminal paper, Syrgkanis and Tardos [22] propose a general technique for bounding social welfare of these equilibria, based on a so-called “smoothness” technique. These guarantees apply even to mixed Bayes-Nash equilibria in environments with composition of mechanisms. For example, in a combinatorial auction we might not sell all items via a complicated truthful mechanism, but instead sell each item simultaneously via simple individual single-item auctions. Such a mechanism is obviously not truthful, since bidders are not even able to express their valuations for all subsets of items. However, if bidders have complement-free XOS valuations, the (expected) social welfare of allocations in a mixed Bayes-Nash equilibrium turns out to be a constant-factor approximation of the optimal social welfare.

While this is a fundamental insight into non-truthful mechanisms, it is not well-understood how this result extends under more realistic conditions. In particular, there has been recent concern about the plausibility and computational complexity of exact and approximate Bayes-Nash equilibria [5]. For more general Bayesian concepts based on no-regret learning strategies in repeated games, there are two natural approaches – either bidder types are drawn newly with bids, or types are drawn only once initially. While the latter is not really in line with the idea of incomplete information (bidders could communicate their type in the course of learning, see [5]), the former is in general hard to obtain. Also, the composition theorem applies only if bidders’ types are drawn independently.

In this paper, we study a variant of simultaneous composition of mechanisms and show how to avoid the drawbacks of the Bayesian approach. Our scenario is motivated by limited availability or admission: Suppose bidders try to acquire items in a repeated online market, in which  $m$  items are sold simultaneously via, say, first-price auctions. However, in each round only some of the items are actually available for purchase. This scenario can be phrased in the Bayesian framework when bidder  $i$ ’s type is given by the set of items available to him. To obtain an equilibrium in the Bayesian sense, each bidder would have to consider a complicated bid vector and satisfy an equilibrium condition for each of the possible  $2^m$  subsets of items.

In contrast, here we assume that bidders do not even get to know (or are not able to account for) *their own availabilities* before making bids in each round. We assume they learn with no-regret strategies in a way that is *oblivious* to their own and all other bidders’ availabilities. Thereby, bidders arrive at what one might term an *availability-oblivious* coarse-correlated equilibrium – a bid distribution not tailored to the specific availabilities of bidders, which can be computed (approximately) in polynomial time. Our main result is that for a large class of valuation functions, we can apply smoothness ideas in this framework and prove bounds that mirror the guarantees above. The guarantees apply even if some bidders learn obliviously and others follow a Bayes-Nash bidding strategy.

In particular, we cover a broad domain with simultaneous composition of weakly smooth mechanisms in the sense of [22] when bidders have lattice-submodular valuations. Our study covers cases where availabilities are correlated among bidders and provides lower bounds for combinatorial auctions with item-bidding and XOS valuations. As a part of our analysis, we use the concept of correlation gap from [1] for submodular functions over product lattices.

## 1.1 Our Contribution

We assume that every mechanism satisfies a weak smoothness bound (for more details see Sect. 2 below) with parameters  $\lambda, \mu_1, \mu_2 \geq 0$ . It is known that for each individual mechanism, this implies an upper bound of  $(\max(1, \mu_1) + \mu_2)/\lambda$  on the price of anarchy for no-regret learning outcomes and Bayes-Nash equilibria. Furthermore, the same bound also applies for outcomes of multiple simultaneous mechanisms that are tailored to availabilities, i.e., not oblivious.

In Sect. 3 we consider smoothness for oblivious learning and composition with independent availabilities, where in each round  $t$ , each mechanism  $j$  is available to each bidder  $i$  independently with probability  $q_{i,j}$ . Our smoothness bound involves the above parameters and the correlation gap of the class of valuation functions. In particular, if valuations  $v_i$  come from a class  $\mathcal{V}$  with a correlation gap of  $\gamma(\mathcal{V})$ , the price of anarchy becomes  $\gamma(\mathcal{V}) \cdot (\max(1, \mu_1) + \mu_2)/\lambda$ .

Our construction uses smoothness of simultaneous composition from [22]. However, since learning is oblivious, the deviations establishing smoothness must be independent of availability. Here we use correlation gap to relate the value for independent deviations to that of type-dependent Bayesian deviations. Correlation gap is a notion originally defined for submodular set functions in [2]. It captures the worst-case ratio between the expected value of independent and correlated distributions over elements with the same marginals. We use an extension of this notion from [1] to Cartesian products of outcome spaces such as product lattices. For the class  $\mathcal{V}$  of monotone lattice-submodular valuations, we prove a correlation gap of  $\gamma(\mathcal{V}) = e/(e-1)$ , which simplifies and slightly extends previous results.

In Sect. 4, we analyze oblivious learning for composition with correlated availabilities in the form of “everybody-or-nobody” – each mechanism is either available to all bidders or to no bidder. The probability for availability of mechanism  $j$  is  $q_j$ , and availabilities are independent among mechanisms. In this case, we simulate independence by assuming that each bidder draws random types and outcomes for himself. We also consider distributions where outcomes are drawn independently according to the marginals from the optimal correlated distribution over outcomes. While these two distributions are directly related via correlation gap, the technical challenge is to show that there is a connection to the value obtained by the bidder. For lattice-submodular functions, we show a smoothness bound that implies a price of anarchy of  $4e/(e-1) \cdot (\max(1, \mu_1) + \mu_2)/\lambda^2$ .

For neither of the results is it necessary that all bidders follow our oblivious-learning approach. We only require that bidders have no regret compared to this

strategy. This is also fulfilled if some or all bidders determine their bids based on the actually available items rather than in the oblivious way.

Finally, in Sect. 5 we show a lower bound for simultaneous composition of single-item first-price auctions with general XOS valuation functions. The correlation gap for such functions is known to be large [2], but this does not directly imply a lower bound on the price of anarchy for oblivious learning. We provide a class of instances where the price of anarchy for oblivious learning becomes  $\Omega((\log m)/(\log m \log m))$ . This shows that for XOS functions it is impossible to generalize the constant price of anarchy for single-item first-price auctions.

Our results have additional implications beyond auctions for the analysis of regret learning in wireless networks. We discuss these in the full version [15, Appendix A].

## 1.2 Related Work

Closely related to our work are combinatorial auctions with item bidding, where multiple items are being sold in separate auctions. Bidders are generally interested in multiple items. However, depending on the bidder, some items may be substitutes for others. As the auctions work independently, bidders have to strategize in order to buy not too many items simultaneously. In a number of papers [3, 7, 12, 14] the efficiency of Nash and Bayes-Nash equilibria has been studied. It has been shown that, if the single items are sold in first or second price auctions and if the valuation functions are XOS or subadditive, the price of anarchy is constant. Limitations of this approach are shown in [8, 21].

Many of these proofs follow a similar pattern, namely showing smoothness. This concept has been introduced by [19, 20] to analyze correlated and Bayes-Nash equilibria of general games. In [22] it was adjusted to mechanisms, and it was shown that simultaneous or sequential composition of smooth mechanisms is again smooth. Combinatorial auctions with item bidding are an example of a simultaneous composition. To show smoothness of the combined mechanism, it is thus enough to show smoothness of each single auction. Other examples of smooth mechanisms are position auctions with generalized second price [6, 18] and greedy auctions [16]. The smoothness approach for mixed Bayes-Nash equilibria shown in [22] is, in fact, slightly more general and continues to hold for variants of Bayesian correlated equilibrium [13].

The complexity of finding such equilibria has been studied only very recently. It has been shown in [5, 11] that equilibria are hard to find in some settings. In contrast, in [10] a different auction format is studied that yields good bounds on social welfare for equilibria that can be found more easily. Although similar in spirit, our approach is different – it shows that in some scenarios agents can reduce the computational effort and still obtain reasonably good states with existing mechanisms.

As such, our approach is closer to recent work [9] that shows hardness results for learning full-information coarse-correlated equilibria in simultaneous single-item second-price auctions with unit-demand bidders. As a remedy, a form of so-called no-envy learning is proposed, in which bidders use a different form of



bidding that enables convergence in polynomial time. While achieving a general no-regret guarantee against all possible bid vectors is hard, we note here that our approach based on smoothness requires only a guarantee with respect to bids that are derived directly from the XOS representation of the bidder valuation. As such, bidders can obtain the guarantees required for our results in polynomial time. Conceptually, we here treat a different problem – the impact of availabilities, and more generally, different bidder types on learning outcomes in repeated mechanism design.

A model with dynamic populations in games has recently been considered in [17]. Each round a small portion of players are replaced by others with different utility functions. When players use algorithms that minimize a notion called adaptive regret, smoothness conditions and the resulting bounds on the price of anarchy continue to hold if there are solutions which remain near-optimal over time with a small number of structural changes. Using tools from differential privacy, these conditions are shown for some special classes of games, including first-price auctions with unit-demand or gross-substitutes valuations. In contrast, our scenario is orthogonal, since we consider much more general classes of mechanisms and allow changes in each round for possibly all players. However, our model of change captures the notion of availability and therefore is much more specific than the adversarial approach of [17].

The notion of correlation gap was defined and analyzed for stochastic optimization in [1, 2]. The notion was used in [23] for analyzing revenue maximization with sequential auctions, which is very different from our approach.

## 2 Model and Preliminaries

There are  $n$  bidders that participate in  $m$  simultaneous mechanisms. Each mechanism  $j \in [m]$  is a pair  $M_j = (f_j, p_j)$ , consisting of an outcome function and payment functions. More formally, function  $f_j: B_j \rightarrow \mathcal{X}_j$  maps every bid vector  $b_{\cdot,j}$  on mechanism  $j$  into an outcome space  $\mathcal{X}_j$ . The function  $p_j = (p_{1,j}, \dots, p_{n,j})$  defines a payment for each bidder. That is, depending on the bid vector,  $p_{i,j}: B_j \rightarrow \mathbb{R}_{\geq 0}$  defines the non-negative payment for bidder  $i$  in mechanism  $j$ .

We consider a repeated framework with oblivious learning in a simultaneous composition of mechanisms with availabilities. There are  $T$  rounds and in each round the bidders participate in  $m$  simultaneous mechanisms. In round  $t = 1, \dots, T$ , each bidder places a bid  $b_{i,j}^t$  for each mechanism, the mechanism determines the outcome and the payments, and bidder  $i$  has a utility function  $u_i(b^t) = v_i(f(b^t)) - p_i(b^t)$ , where  $v_i$  is a valuation function over vectors of outcomes and  $p_i = \sum_j p_{i,j}(b^t)$ . In addition, in each round we assume that each mechanism is available to each bidder with a certain probability. We let the Bernoulli random variable  $A_{i,j} = 1$  if mechanism  $j$  is available to bidder  $i$ . Due to availability, the mechanisms must also be applicable when only subsets of bidders are placing bids. For this reason, it will be convenient to assume that the outcome space for mechanism  $j \in [m]$  is  $\mathcal{X}_j = \mathcal{X}_{1,j} \times \dots \times \mathcal{X}_{n,j}$  and  $x_j \in \mathcal{X}_j$  is  $x_j = (x_{i,j})_{i \in [n]}$ . We assume that each bidder, for whom the mechanism is not available, must place a bid of “0”. If bidder  $i$  bids 0 for mechanism

$j$ , we assume  $f_j(0, b_{-i,j}) = \perp_{i,j}$ , where  $\perp_{i,j}$  is a “losing” outcome, and payment  $p_{i,j}(0, b_{-i,j}) = 0$ . For convenience, we will denote by  $f = (f_j)_{j \in [m]}$  the composed mechanism and by  $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_m$  its outcome space.

*Oblivious Learning.* We assume *oblivious learning* – each bidder runs a single no-regret learning algorithm and uses the utility of every round as feedback, no matter how the availability in each round turns out. In hindsight, the average history of play for oblivious learning becomes an availability-oblivious variant of coarse-correlated equilibrium [4]. Hence, the outcomes of oblivious learning are captured by the coarse-correlated equilibria in the following one-shot game: First, all bidders simultaneously place a bid for every mechanism. They know only the probability distribution of the availabilities. Only after they placed their bids, the availability of each mechanism for each bidder is determined at random.

**Definition 1.** *An availability-oblivious coarse-correlated equilibrium is a distribution over bid vectors  $b$  (independent of  $A$ ) such that, in expectation over all availabilities, it is not beneficial for any bidder  $i$  to switch to another bid  $b'_i$ . For each  $i$  and each  $b'_i$ , we have  $\mathbf{E}[u_i(b'_i, b_{-i})] \leq \mathbf{E}[u_i(b)]$ .*

Indeed, our results also hold for a larger class of equilibria, in which a subset of bidders might not be oblivious to availabilities. For our guarantees, it is enough to consider distributions over bidding strategies  $b$  which might depend on  $A$  such that, in expectation over all availabilities, it is not beneficial for any bidder  $i$  to switch to another bid  $b'_i$ . For each  $i$  and each  $b'_i$ , we have  $\mathbf{E}[u_i(b'_i, b_{-i})] \leq \mathbf{E}[u_i(b)]$ . Note that both ordinary coarse-correlated equilibria and availability-oblivious ones fulfill this property.

We bound the performance of these equilibria by deriving suitable smoothness bounds.

*Smoothness.* We assume that each mechanism  $j$  satisfies *weak smoothness* as defined in [22]. For any valuations  $v_{i,j} : \mathcal{X}_j \rightarrow \mathbb{R}^{\geq 0}$  there are (possibly randomized) deviations<sup>1</sup>  $b'_{i,j}$  for each  $i \in [n]$  such that for all bid vectors  $b_{\cdot,j}$

$$\begin{aligned} & \mathbf{E} \left[ \sum_{i \in [n]} v_{i,j}(f_j(b'_{i,j}, b_{-i,j})) - p_{i,j}(b'_{i,j}, b_{-i,j}) \right] \\ & \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - \mu_1 \cdot \sum_{i \in [n]} p_{i,j}(b_{\cdot,j}) - \mu_2 \sum_{i \in [n]} h_{i,j}(b_{i,j}, f_j(b_{\cdot,j})), \end{aligned} \quad (1)$$

where  $h_{i,j}(b_{i,j}, x_j) = \max_{b_{-i,j} : f_j(b_{\cdot,j}) = x_j} p_{i,j}(b_{\cdot,j})$ . For intuition, assume that (1) holds with  $\mu_2 = 0$ . Consider a learning outcome with a no-regret guarantee where

<sup>1</sup> In slight contrast to [22], we here assume that the smoothness deviations of a bidder do not depend on his own current bid. This serves to simplify our exposition and can be incorporated into our analysis.

every bidder  $i$  can gain at most  $\epsilon$  in any fixed deviation, i.e.,  $\mathbf{E}[v_{i,j}(f_j(b_{\cdot,j})) - p_{i,j}(b_{\cdot,j})] \geq \mathbf{E}[v_{i,j}(f_j(b'_{i,j}, b_{-i,j})) - p_{i,j}(b'_{i,j}, b_{-i,j})] - \epsilon$ . Applying (1) pointwise

$$\sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j})) - p_{i,j}(b_{\cdot,j})] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - \mu_1 \cdot \sum_{i \in [n]} \mathbf{E}[p_{i,j}(b_{\cdot,j})] - n\epsilon,$$

which implies for social welfare

$$\sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) + (1 - \mu_1) \cdot \sum_{i \in [n]} \mathbf{E}[p_{i,j}(b_{\cdot,j})] - n\epsilon.$$

Every bidder  $i$  can stay away from the market and payments are non-negative, so  $0 \leq \mathbf{E}[p_{i,j}(b_{\cdot,j})] \leq \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] + \epsilon$  and

$$\max(1, \mu_1) \sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - (n + \mu_1)\epsilon.$$

Thus, for  $\epsilon \rightarrow 0$ , the price of anarchy tends to  $\max(1, \mu_1)/\lambda$ . More generally, (1) implies a bound on the price of anarchy of  $(\mu_2 + \max(1, \mu_1))/\lambda$  for many equilibrium concepts. If  $\mu_2 > 0$ , then the bound relies on an additional no-overbidding assumption, which directly transfers to our results. For details see [22].

*Valuation Functions.* Our main results apply for the class of monotone lattice-submodular valuations. Suppose for every mechanism  $j$  the set  $\mathcal{X}_{i,j}$  of possible outcomes for bidder  $i$  forms a lattice  $(\mathcal{X}_{i,j}, \succeq_{i,j})$  with a partial order  $\succeq_{i,j}$ . Bidder  $i$  has a *lattice-submodular valuation*  $v_i$  if and only if it is submodular on the product lattice  $(\mathcal{X}_i, \succeq_i)$  of outcomes for bidder  $i$ :  $\forall x_i, \tilde{x}_i \in \mathcal{X}_i : v_i(x_i \vee \tilde{x}_i) + v_i(x_i \wedge \tilde{x}_i) \leq v_i(x_i) + v_i(\tilde{x}_i)$ . In the paper, we concentrate on distributive lattices, for which this definition is equivalent to the diminishing marginal returns property:

$$\forall z_i \succeq_i y_i \in \mathcal{X}_i \implies \forall t \in \mathcal{X}_i : v_i(t \vee y_i) - v(y_i) \geq v_i(t \vee z_i) - v(z_i).$$

Lattice-submodular functions generalize submodular set functions but are a strict subclass of XOS functions. Bidder  $i$  has an *XOS valuation*  $v_i$  if and only if there are additive functions  $v_i^1, v_i^2, \dots$  with  $v_i^{k_i}(x_i) = \sum_j v_j^{k_i}(x_{ij})$  for every  $x_{i,j} \in \mathcal{X}_{i,j}$  and  $v_i(x_i) = \max_{k_i} v_i^{k_i}(x_i)$ .

### 3 Composition with Independent Admission

We first consider simultaneous composition of smooth mechanisms with independent availabilities. Here, all random variables  $A_{i,j}$  are independent, and we let  $q_{i,j} = \Pr[A_{i,j} = 1]$ .

**Definition 2.** *Let  $v$  be a valuation function on a product lattice, coming from a class of valuation functions  $\mathcal{V}$ . Given vectors  $x^1, \dots, x^k$  and numbers  $\alpha_1, \dots, \alpha_k \in [0, 1]$  such that  $\sum_{j=1}^k \alpha_j = 1$ , determine another vector  $y$  at random by setting component  $y_i$  to  $x_i^j$  independently with probability  $\alpha_j$ . Then, the smallest  $\gamma$  s.t.  $\sum_{j=1}^k \alpha_j v(x^j) \leq \gamma \cdot \mathbf{E}[v(y)]$  is the correlation gap of class  $\mathcal{V}$ .*

**Theorem 1.** *Suppose bidder valuations are monotone and come from a class  $\mathcal{V}$  with a correlation gap of  $\gamma(\mathcal{V})$ . The price of anarchy for oblivious learning for simultaneous composition of weakly  $(\lambda, \mu_1, \mu_2)$ -smooth mechanisms with valuations from  $\mathcal{V}$  and fully independent availability is at most  $\gamma(\mathcal{V}) \cdot (\mu_2 + \max(1, \mu_1))/\lambda$ .*

Before the proof of the main theorem of this section, we note that in the full version [15, Appendix C.1] we also prove an upper bound of  $e/(e-1)$  on the correlation gap of lattice-submodular valuations with diminishing marginal returns. This result slightly generalizes the result of [1] from composition of totally ordered sets to arbitrary product lattices.

**Lemma 1 (Correlation Gap on a Product Lattice).** *Let  $v$  be a function with diminishing marginal returns on a product lattice. Given vectors  $x^1, \dots, x^k$  and numbers  $\alpha_1, \dots, \alpha_k \in [0, 1]$  such that  $\sum_{j=1}^k \alpha_j = 1$ , determine another vector  $y$  at random by setting component  $y_i$  to  $x_i^j$  independently with probability  $\alpha_j$ . Then  $\mathbf{E}[v(y)] \geq (1 - \frac{1}{e}) \sum_{j=1}^k \alpha_j v(x^j)$ .*

From here, we arrive at the following corollary of the main theorem.

**Corollary 1.** *The price of anarchy for oblivious learning for simultaneous composition of weakly  $(\lambda, \mu_1, \mu_2)$ -smooth mechanisms with monotone lattice-submodular valuations and fully independent availability is at most  $e/(e-1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda$ .*

*Proof of Theorem 1.* We will prove the theorem by defining an availability-oblivious (randomized) deviation  $b'_i$  for each player  $i$  such that the following inequality will hold for any (not necessarily availability-oblivious) bidding strategy  $b$ :

$$\begin{aligned} & \sum_i \mathbf{E}[u_i(b'_i, b_{-i})] \\ & \geq \frac{1}{\gamma(\mathcal{V})} \cdot \lambda \cdot \sum_i \mathbf{E}[v_i(x^*)] - \mu_1 \sum_i \mathbf{E}[p_i(b)] - \mu_2 \sum_i \mathbf{E}[h_i(b_i, f(b))], \quad (2) \end{aligned}$$

where  $x^*$  denotes the (random) optimal outcome. From this inequality, whose form is in fact exactly that of the smoothness condition (1), the claim of the theorem follows as described in Sect. 2.

In more detail, to attain the aforementioned inequality, we will relate each player's utility for deviating to  $b'_i$  to the utility he could achieve if he was allowed to see and react upon the availabilities. In that case, he could simply use the smoothness deviation tailored to the specific availability profile  $A_i = (A_{i,1}, \dots, A_{i,m})$  that he is encountering. We denote this non-oblivious smoothness deviation by  $b_i^{A_i}$ . Because the global mechanism is a simultaneous composition of  $(\lambda, \mu_1, \mu_2)$ -smooth mechanisms, it is again  $(\lambda, \mu_1, \mu_2)$ -smooth. Therefore we know that the non-oblivious deviations  $b_i^{A_i}$  do exist, and they satisfy the smoothness inequality (1) by definition.

We proceed to define, for each player  $i$ , the availability-oblivious deviation  $b'_i$ . First, bidder  $i$  assumes for himself a reduced valuation function  $\bar{v}_i = \alpha \cdot v_i$ , for some appropriate  $\alpha$  to be chosen later. The deviation  $b'_i$  is a composition of component-wise independent deviations  $b'_{i,j}$ , i.e.  $b'_i = (b'_{i,1}, \dots, b'_{i,m})$  where each  $b'_{i,j}$  is chosen independently. To arrive at  $b'_{i,j}$ , bidder  $i$  assumes that mechanism  $j$  is available to him and draws all other availabilities independently according to probabilities  $q_{i',j'}$ . This means that he draws availabilities for all other players on all mechanisms and also his own availabilities on all mechanisms other than  $j$ . Now he has a full availability profile, and therefore he can consider the non-oblivious smoothness deviation. He observes the  $j$ -th component of this smoothness deviation and sets  $b'_{i,j}$  to be equal to it. Note that  $b'_{i,j}$  will be applied only with the probability that mechanism  $j$  is in fact available to bidder  $i$ , i.e. with probability  $q_{i,j}$ .

Next, we want to compare  $u_i(b'_i, b_{-i})$  and  $u_i(b_i^{A_i}, b_{-i})$ . Let us focus on the valuation  $v_i(f(b'_i, b_{-i}))$  first. The non-oblivious smoothness deviation  $b_i^{A_i}$  is a vector whose components are correlated. More precisely, to form this bid we observe  $A_i$ , sample the availabilities  $A_{-i}$  and bids  $b_{-i}$  of other players, and take the optimal allocation  $x^*$  for the resulting availability profile  $A$ . Then, we determine the  $\ell$  for which  $\bar{v}_i(x_i^*) = \sum_j \bar{v}_{i,j}^\ell(x_{i,j}^*)$  and use  $\bar{v}_{i,j}^\ell$  for determining  $b_{i,j}^{A_i}$  (note that  $A_i$  can be regarded as bidder  $i$ 's type in a Bayesian sense, for more details see [22]). Therefore, the components of  $b_i^{A_i}$  are correlated through the common choice of  $\ell$ . Our deviation  $b'_i$  is assembled by setting  $b'_{i,j} = (b_{i,j}^{A_i})_{k_j}$  independently for each  $j$ .

Formally, let  $r_{i,j}^\ell$  denote the conditional probability that the optimum yields an outcome vector  $x^*$  that attains its maximum value for bidder  $i$  in  $\bar{v}_i^\ell$ , given that  $A_{i,j} = 1$ . Then, the marginal probability of observing  $b_{i,j}^{A_i} = (b_{i,j}^{A_i})_\ell$  is  $r_{i,j}^\ell q_{i,j}$ . In  $b'_i$  we pick  $\ell$  independently for each mechanism with probability  $r_{i,j}^\ell$ , which yields a combined probability of  $r_{i,j}^\ell q_{i,j}$  for availability and deviation. Thus,  $b'_i$  simulates the marginal probabilities of outcomes in  $b_i^{A_i}$ , i.e.,  $\Pr[f_j(b'_i, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}] = \Pr[f_j(b_i^{A_i}, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}]$  for all  $y_{i,j} \in \mathcal{X}_{i,j}$ , for each  $j \in [m]$ . Hence, for fixed  $A_{-i}, b_{-i}$ , the two expected valuations  $\mathbf{E}[v_i(f(b'_i, b_{-i})) \mid A_{-i}, b_{-i}]$  and  $\mathbf{E}[v_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i}]$  are related via correlation gap.

Thus, setting  $\alpha = 1/\gamma(\mathcal{V})$  and  $\bar{v}_i(x) = 1/\gamma(\mathcal{V}) \cdot v_i(x)$  we get

$$\begin{aligned} \mathbf{E}[v_i(f(b'_i, b_{-i})) \mid A_{-i}, b_{-i}] &= \sum_{y \in \mathcal{X}} v_i(y) \cdot \Pr[f(b'_i, b_{-i}) = y \mid A_{-i}, b_{-i}] \\ &= \sum_{y \in \mathcal{X}} v_i(y) \cdot \prod_j \Pr[f_j(b'_i, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}] \\ &\geq \frac{1}{\gamma(\mathcal{V})} \cdot \sum_{y \in \mathcal{X}} v_i(y) \cdot \Pr[f(b_i^{A_i}, b_{-i}) = y \mid A_{-i}, b_{-i}] \end{aligned}$$

$$\begin{aligned}
 &= \frac{1}{\gamma(\mathcal{V})} \cdot \mathbf{E} \left[ v_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i} \right] \\
 &= \mathbf{E} \left[ \bar{v}_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i} \right].
 \end{aligned}$$

In addition, because payments are simply additive across mechanisms, it is straightforward to see that for every bidder  $i$

$$\mathbf{E} [p_i(b'_i, b_{-i}) \mid A_{-i}, b_{-i}] = \mathbf{E} [p_i(b_i^{A_i}, b_{-i}) \mid A_{-i}, b_{-i}].$$

This allows to apply the smoothness bound for Bayesian mechanisms with independent types from [22] to derive

$$\begin{aligned}
 &\sum_i \mathbf{E} [u_i(b'_i, b_{-i})] \\
 &= \sum_i \mathbf{E} [v_i(f(b'_i, b_{-i}))] - \mathbf{E} [p_i(b'_i, b_{-i})] \\
 &\geq \sum_i \mathbf{E} [\bar{v}_i(f(b_i^{A_i}, b_{-i}))] - \mathbf{E} [p_i(b_i^{A_i}, b_{-i})] \\
 &\geq \lambda \cdot \sum_i \mathbf{E} [\bar{v}_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] \\
 &= \frac{\lambda}{\gamma(\mathcal{V})} \cdot \sum_i \mathbf{E} [v_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))]
 \end{aligned}$$

This proves the desired smoothness guarantee and implies the theorem. □

### 4 Composition with Everybody-or-Nobody Admission

We consider the case in which at each point in time each mechanism is either available to all bidders or to none. We let  $A_j = A_{i,j}$  for all  $i \in [n]$  and  $q_j = \Pr [A_j = 1]$ . Note that all  $A_j$  are assumed to be independent.

Let the social optimum be denoted by  $x^*$ . We assume that  $x_j^* = \perp_j$  if  $A_j = 0$ . Otherwise,  $x^*$  might have different values, depending on the availabilities of other mechanisms. Let us denote the possible outcomes by  $x_j^1, x_j^2, \dots$  and let  $r_j^\ell := \Pr [x_j^* = x_j^\ell \mid A_j = 1]$ . That is,  $r_j^\ell$  is the marginal probability of  $x_j^\ell$  conditioned on  $j$  being available. Theorem 2 formulates our main result in this section.

**Theorem 2.** *The price of anarchy for oblivious learning for simultaneous composition of weakly  $(\lambda, \mu_1, \mu_2)$ -smooth mechanisms with monotone lattice-submodular valuations and everybody-or-nobody admission is at most  $4e/(e - 1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda^2$ .*

*Proof.* We will prove that, for each bidder  $i$  and each mechanism  $j$  there are randomized deviation strategies  $b'_{i,j}$  that are independent of the availabilities

such that the following smoothness guarantee holds against any (potentially non-oblivious) bidding strategy  $b$ :

$$\begin{aligned} & \sum_i \mathbf{E} [u_i(b'_i, b_{-i})] \\ & \geq \left(1 - \frac{1}{e}\right) \frac{\lambda^2}{4} \sum_i \mathbf{E} [v_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))]. \end{aligned}$$

From this guarantee the claim of the theorem again follows as described in Sect. 2.

To define  $b'_{i,j}$ , every bidder  $i$  draws two vectors  $z^i$  and  $\tilde{t}^i$  at random as follows. He sets  $z^i_j$  to  $x^{\ell}_j$  with probability  $r^{\ell}_j/\alpha$ , where  $\alpha = 2/\lambda$ , and to  $\perp_j$  with the remaining probability. Furthermore, he sets  $\tilde{t}^i_j$  to  $x^{\ell}_j$  with probability  $q_j r^{\ell}_j$  and to  $\perp_j$  with the remaining probability. These draws are performed independent of any availabilities. Observe that for each  $i$ , we have  $\mathbf{E} [\sum_{i'} v_{i'}(\tilde{t}^i)] \geq (1 - \frac{1}{e}) \mathbf{E} [\sum_{i'} v_{i'}(x^*)]$  by Lemma 1.

Due to the random draws, each bidder  $i'$  defines functions  $w^{i'}_{i,j}: \Omega_j \rightarrow \mathbb{R}$  for each bidder  $i$  and each mechanism  $j$ . Function  $w^{i'}_{i,j}$  maps an outcome of mechanism  $j$ , denoted by  $y_j$ , to a real number as follows

$$w^{i'}_{i,j}(y_j) = v_i(\tilde{t}^i_1, \dots, \tilde{t}^i_{j-1}, y_j \wedge z^{i'}_j, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}^i_1, \dots, \tilde{t}^i_{j-1}, \perp_j, \dots, \perp_m).$$

Note that these functions do not necessarily reflect the actual value any outcome might have. They are only used to define the deviation strategy: bidder  $i'$  pretends all bidders  $i$ , including himself, would have valuations  $w^{i'}_{i,j}$  for the outcome of mechanism  $j$ . This gives him a deviation strategy  $b'_{i',j}$  by setting  $b'_{i',j} = b^*_{i',j}(w^{i'}_{1,j}, \dots, w^{i'}_{n,j})$  as defined by the smoothness of mechanism  $j$ .

The proofs for the following three lemmas are presented in the full version [15, Appendix C.2, C.3, C.4].

**Lemma 2.** *For every bidder  $i$  and deviating bids  $b'_{i,j} = b^*_{i,j}(w^i_{1,j}, \dots, w^i_{n,j})$ ,*

$$\mathbf{E} [v_i(f(b'_i, b_{-i}))] \geq \sum_j \mathbf{E} [w^i_{i,j}(f_j(b'_{i,j}, b_{-i}))] - \frac{1}{\alpha(\alpha + 1)} \mathbf{E} [v_i(\tilde{t}^i)].$$

**Lemma 3.** *For the adjusted functions  $w$  we can apply smoothness to obtain*

$$\begin{aligned} & \sum_i \sum_j \mathbf{E} [w^i_{i,j}(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] \\ & \geq \lambda \sum_i \sum_j q_j \mathbf{E} [w^1_{i,j}(z^1_j)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))]. \end{aligned}$$

**Lemma 4.** *For function  $w^1$ , random vectors  $z^1_j$  and  $\tilde{t}^1$ , and every mechanism  $j$*

$$\sum_j q_j \mathbf{E} [w^1_{i,j}(z^1_j)] = \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}^1)].$$

The bound from Lemma 3 has striking similarities to the smoothness bound (1). However, it is expressed in terms of the functions  $w_{i,j}^{i'}$ , rather than the actual valuation functions  $v_i$ . The other two Lemmas show that, in expectation, these functions are close enough to the functions  $v_i$  so that this bound actually suffices to prove the main result:

$$\begin{aligned}
 \sum_i \mathbf{E} [u_i(b'_i, b_{-i})] &= \sum_i \mathbf{E} \left[ v_i(f(b'_i, b_{-i})) - \sum_j p_{i,j}(b'_{i,j}, b_{-i}) \right] \\
 &\geq \sum_i \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] - \frac{1}{\alpha(\alpha + 1)} \sum_i \mathbf{E} [v_i(\tilde{t}^i)] \\
 &\hspace{20em} \text{(by Lemma 2)} \\
 &\geq \lambda \sum_i \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] \\
 &\quad - \frac{1}{\alpha(\alpha + 1)} \sum_i \mathbf{E} [v_i(\tilde{t}^1)] \\
 &\hspace{20em} \text{(by Lemma 3)} \\
 &= \sum_i \left( \frac{\lambda}{\alpha} - \frac{1}{\alpha(\alpha + 1)} \right) \mathbf{E} [v_i(\tilde{t}^i)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))]. \\
 &\hspace{20em} \text{(by Lemma 4)}
 \end{aligned}$$

By setting  $\alpha = \frac{2}{\lambda}$

$$\begin{aligned}
 \sum_i \mathbf{E} [u_i(b'_i, b_{-i})] &\geq \frac{\lambda^2}{4} \sum_i \mathbf{E} [v_i(\tilde{t}^i)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] \\
 &\geq \left( 1 - \frac{1}{e} \right) \frac{\lambda^2}{4} \sum_i \mathbf{E} [v_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))].
 \end{aligned}$$

The last step follows from Lemma 1.

Note that technically the mechanism could be randomized itself. Our results extend to this case in a straightforward way.

## 5 A Lower Bound for General XOS Functions

In this section we consider combinatorial auctions with item bidding and first-price auctions. We can apply the previous analysis, since for each bidder the outcomes form a trivial 2-element lattice – winning an item is the supremum outcome, not winning is the infimum outcome. In the analysis, observe that each bidder determines a random allocation of items according to the probabilities in the optimum. Based on these allocations, bidders determine the valuations  $w_{i,j}^{i'}$ , which in turn form the basis for the deviation. The first-price auction with



general bidding space is  $(1 - 1/e, 1, 0)$ -smooth [22]. If valuation functions are submodular, the composition theorems can be applied to yield the following corollary.

**Corollary 2.** *The price of anarchy for oblivious learning for simultaneous composition of single-item first-price auctions with monotone submodular valuations and fully independent availability is at most  $1/(1-1/e)^2$ ; for everybody-or-nobody admission it is at most  $4/(1 - 1/e)^3$ .*

For more general XOS valuations, we prove a lower bound that with oblivious bidding we will not be able to show a guarantee based on the smoothness parameters – even for a single bidder, so the bound applies without assumptions on correlation among bidders. The proof can be found in the full version [15, Appendix C.5].

**Theorem 3.** *In a simultaneous composition of discrete first-price single-item auctions with  $m$  items and XOS valuations, the price of anarchy for pure Nash equilibria with oblivious bidding can be as large as  $\Omega((\log m)/(\log \log m))$ , while each single mechanism is weakly  $(1/2, 1, 0)$ -smooth.*

## 6 Conclusion

In this paper, we have studied an oblivious variant for no-regret learning in repeated games with incomplete information and proved a composition theorem for smooth mechanisms. The bounds show that even if bidders apply learning algorithms independently of their types, they can still obtain outcomes that approximate the optimal social welfare within a small ratio.

Our primary motivation are changes over time on the supply side. That is, bidders value items the same at all times but are constrained when they can buy them. A different interpretation that leads to the same model is when bidders value items differently from time to time. Here the valuation for a bundle has the special structure that it is given by the value of a fixed submodular function evaluated on the intersection of this bundle with a random set.

There is potential to generalize this approach to other interesting settings. For example, one could consider general independent types, where the complete availability-vector of a single bidder is drawn from a bidder-specific distribution, and for each bidder this is done independently. In the full version [15, Appendix B], we give a partial answer and show how our techniques can be extended to the following case. Consider simultaneous single-item auctions with unit-demand valuations, i.e.,  $v_i(S) = \max_{j \in S} v_{i,j}$ . The distribution over valuations is such that for each item the value  $v_{i,j}$  is independently drawn from a distribution of small support. Independent availabilities can be captured in this setting by setting  $v_{i,j}$  to a fixed value or to 0 with the respective probabilities.

## References

1. Agrawal, S.: Optimization under uncertainty: bounding the correlation gap. Ph.D. thesis, Stanford University (2011)
2. Agrawal, S., Ding, Y., Saberi, A., Ye, Y.: Correlation robust stochastic optimization. In: Proceedings of 21st Symposium on Discrete Algorithms (SODA), pp. 1087–1096 (2010)
3. Bhawalkar, K., Roughgarden, T.: Welfare guarantees for combinatorial auctions with item bidding. In: Proceedings of 22nd Symposium on Discrete Algorithms (SODA), pp. 700–709 (2011)
4. Blum, A., Mansour, Y.: Learning, regret minimization, and equilibria. In: Nisan, N., Tardos, É., Roughgarden, T., Vazirani, V. (eds.), *Algorithmic Game Theory*, chapter 4. Cambridge University Press (2007)
5. Cai, Y., Papadimitriou, C.: Simultaneous Bayesian auctions and computational complexity. In Proceedings of the 15th ACM Conference on Economics and Computation (EC 2014), pp. 895–910 (2014)
6. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., Lucier, B., Leme, R.P., Tardos, É.: On the efficiency of equilibria in generalized second price auctions. *J. Econom. Theory* **156**, 343–388 (2015)
7. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian combinatorial auctions. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) *ICALP 2008. LNCS*, vol. 5125, pp. 820–832. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-70575-8\\_67](https://doi.org/10.1007/978-3-540-70575-8_67)
8. Christodoulou, G., Kovács, A., Sgouritsa, A., Tang, B.: Tight bounds for the price of anarchy of simultaneous first price auctions. *ACM Trans. Econom. Comput.* **4**(2), 9 (2016)
9. Daskalakis, C., Syrgkanis, V.: Learning in auctions: regret is hard, envy is easy. In: Proceedings of 57th Symposium on Foundations of Computer Science (FOCS) (2016, to appear)
10. Devanur, N., Morgenstern, J., Syrgkanis, V., Weinberg, M.: Simple auctions with simple strategies. In: Proceedings of the 16th ACM Conference on Economics and Computation (EC 2015), pp. 305–322 (2015)
11. Dobzinski, S., Fu, H., Kleinberg, R.D.: On the complexity of computing an equilibrium in combinatorial auctions. In: Proceedings of 26th Symposium on Discrete Algorithms (SODA), pp. 110–122 (2015)
12. Feldman, M., Fu, H., Gravin, N., Lucier, B.: Simultaneous auctions are (almost) efficient. In: Proceedings of 45th Symposium on Theory of Computing (STOC), pp. 201–210 (2013)
13. Hartline, J.D., Syrgkanis, V., Tardos, É.: No-regret learning in repeated Bayesian games. In: Proceedings of 28th Conference on Advanced Neural Information Processing Systems (NIPS), pp. 3043–3051 (2015)
14. Hassidim, A., Kaplan, H., Mansour, Y., Nisan, N.: Non-price equilibria in markets of discrete goods. In Proceedings of 12th Conference on Electronic Commerce (EC), pp. 295–296 (2011)
15. Hoefer, M., Kesselheim, T., Kodric, B.: Smoothness for simultaneous composition of mechanisms with admission. *CoRR*, abs/1509.00337 (2015)
16. Lucier, B., Borodin, A.: Price of anarchy for greedy auctions. In: Proceedings of 21st Symposium Discrete Algorithms (SODA), pp. 537–553 (2010)
17. Lykouris, T., Syrgkanis, V., Tardos, É.: Learning and efficiency in games with dynamic population. In Proceedings of 27th Symposium Discrete Algorithms (SODA), pp. 120–129 (2016)

18. Paes Leme, R., Tardos, É.: Pure and Bayes-Nash price of anarchy for generalized second price auction. In: Proceedings of 51st Symposium Foundations of Computer Science (FOCS), pp. 735–744 (2010)
19. Roughgarden, T.: Intrinsic robustness of the price of anarchy. *Comm. ACM* **55**(7), 116–123 (2012)
20. Roughgarden, T.: The price of anarchy in games of incomplete information. In: Proceedings of 13th Conference on Electronic Commerce (EC), pp. 862–879 (2012)
21. Roughgarden, T.: Barriers to near-optimal equilibria. In: Proceedings of 55th Symposium on Foundations of Computer Science (FOCS), pp. 71–80 (2014)
22. Syrgkanis, V., Tardos, É.: Composable and efficient mechanisms. In: Proceedings of 45th Symposium Theory of Computing (STOC), pp. 211–220 (2013)
23. Yan, Q.: Mechanism design via correlation gap. In: Proceedings of 22nd Symposium Discrete Algorithms (SODA), pp. 710–719 (2011)

# Motivating Time-Inconsistent Agents: A Computational Approach

Susanne Albers and Dennis Kraft<sup>(✉)</sup>

Department of Computer Science, Technical University of Munich, Munich, Germany  
{albers,kraftd}@in.tum.de

**Abstract.** We study the complexity of motivating time-inconsistent agents to complete long term projects in a graph-based planning model as proposed by Kleinberg and Oren [5]. Given a task graph  $G$  with  $n$  nodes, our objective is to guide an agent towards a target node  $t$  under certain budget constraints. The crux is that the agent may change its strategy over time due to its present-bias. We consider two strategies to guide the agent. First, a single reward is placed at  $t$  and arbitrary edges can be removed from  $G$ . Secondly, rewards can be placed at arbitrary nodes of  $G$  but no edges must be deleted. In both cases we show that it is NP-complete to decide if a given budget is sufficient to guide the agent. For the first setting, we give complementing upper and lower bounds on the approximability of the minimum required budget. In particular, we devise a  $(1 + \sqrt{n})$ -approximation algorithm and prove NP-hardness for ratios greater than  $\sqrt{n}/3$ . Finally, we argue that the second setting does not permit any efficient approximation unless  $P = NP$ .

**Keywords:** Approximation algorithms · Behavioral economics · Computational complexity · Planning and scheduling · Time-inconsistency

## 1 Introduction

In this paper we study the phenomenon of *time-inconsistent behavior* from a computational perspective. Time-inconsistency is a fundamental problem in behavioral economics and has many examples in every day life including academia. For instance, consider a referee who agrees to evaluate a scientific proposal. Despite good intentions, the referee gets distracted and never submits a report. Or consider a student who enrolls in a course. After completing the first homework assignments, the student drops out without earning any credit. In general, these situations have a reoccurring pattern: An agent makes a plan to complete a set of tasks in the future, but changes the plan at a later point in time. This behavior is sometimes the result of unforeseen circumstances. However, in many cases the plan is changed or abandoned even if the circumstances stay the same. This paradox behavior of *procrastination* and *abandonment* is well-known in the

---

S. Albers and D. Kraft—Work supported by the European Research Council, Grant Agreement No. 691672.

field of behavioral economics and might severely affect the performance of agents in an economic or social domain, see e.g. [1, 9, 11].

A sensible explanation for time-inconsistent behavior is that agents assign disproportionately greater value to current cost than to future expenses. For example, consider a simple *car wash problem* in which Alice commissions Bob to wash her car. Each day Bob can either do the chore or postpone it to the next day. However, the longer he waits, the dirtier the car gets. On day  $i$  cleaning the car incurs a cost of  $i/50$  while the cost of waiting another day is 0. After completing the task, Bob will receive a reward of 1 from Alice. Because Bob is present-biased, he perceives any current cost according to its true value, but discounts future costs and rewards by a factor of  $\beta \in [0, 1]$ . On day  $i$  he compares the cost of washing the car right away, which is  $i/50$ , to his perceived cost of washing it on the next day, which is  $\beta(i+1)/50$ . Suppose that  $\beta = 1/3$ . Because  $i/50 > \beta(i+1)/50$ , he procrastinates with good intentions of doing the job on the following day. On day  $i = 50$ , Bob's perceived cost for washing the car on the next day or any of the following days is at least  $\beta(50+1)/50$ . This exceeds his perceived reward of  $\beta$  and therefore he abandons the project.

**Previous Work:** There exists an extensive body of work on time-inconsistent behavior in the economic literature, cf. again [1, 9, 11]. In particular, the car wash problem as stated above is a special case of *quasi-hyperbolic discounting* [7]. We build on work by Kleinberg and Oren, who proposed a graph-based model that captures time-inconsistent behavior in general planning problems [5]. Since its introduction, their work has sparked an active line of research at the intersection of economics, mathematics and computer science, see e.g. [4, 6].

We will give a formal definition of our model in Sect. 2. Essentially, it consists of a directed acyclic task graph  $G$  with  $n$  nodes. Each node represents a certain state of the project, whereas the edges are tasks necessary to transition between states. The workload of individual tasks is modeled by edge costs. To complete the project, an agent with bias factor  $\beta \in [0, 1]$  must move from a designated source  $s$  to a target  $t$ . As a motivation, rewards are placed on the nodes of  $G$ . When located at some node of  $G$ , the agent considers all possible paths to  $t$ . However, because of its time inconsistency, the agent only evaluates the cost of incident edges accurately. All other costs and rewards are discounted by a factor of  $\beta$ . Let  $P$  be a path that minimizes the agent's *perceived net cost*. If this cost is at most 0, the agent traverses the first edge of  $P$  and then reassesses its plan. Otherwise the agent abandons the project. A graph in which the agent always reaches  $t$  is called *motivating*.

In this paper, we will take the perspective of a project designer, whose main objective is budget-efficiency. In other words, we try to minimize the reward we must spend to get the project completed. In general, various strategic arrangements can be made to increase budget-efficiency. Because the aim of such arrangements is to commit the agent to finish the project, they are also called *commitment devices* [2]. Consider, for instance, the car wash example. As the project designer Alice can introduce a deadline to keep Bob from procrastinating. As we will show in Sect. 2, this is beneficial to both of them. In general,

the introduction of deadlines belongs to a broader range of popular commitment devices that reduce the agent's set of choices, see e.g. [9, 10]. Note that the graph-based model lends itself to this approach as we can model any reduction of the agent's choices by simply removing the corresponding edges from  $G$  [5].

A second popular commitment device is to hand out rewards at intermediate states of the project [10]. In the graph-based model, we can do this by placing rewards at non-terminal nodes of  $G$ . We call such an assignment a *reward configuration*. This approach is especially interesting if the project designer's budget is only affected by rewards that are actually collected by the agent. As we will show in Sect. 2, this allows the construction of *exploitative* projects in which the agent is motivated by rewards it never claims. Considering the power and versatility of the two commitment devices mentioned above, Kleinberg and Oren pose the complexity of computing motivating subgraphs and reward configurations as two important open problems [5].

In an unpublished manuscript, Tang et al. address both of these problems [12]. First, they show that it is NP-hard to decide if  $G$  contains a motivating subgraph for a fixed reward placed at  $t$ . Secondly, they give NP-hardness results for three variations of the reward configuration problem: One in which the rewards must be positive, one in which rewards may also be negative and one in which every reward that is laid out must be collected. In each setting, the project designer is charged the absolute sum of the rewards placed on  $G$ .

**Our Contribution:** We will thoroughly analyze the complexity and approximability of computing motivating subgraphs as well as reward configurations. In Sect. 3, we will settle the complexity of finding a motivating subgraph for a fixed reward at  $t$ . First, we will show that the problem is polynomially solvable if  $\beta = 0$  or  $\beta = 1$ . We will then prove that it is NP-complete to decide the existence of a motivating subgraph for general  $\beta \in (0, 1)$ . Tang et al. showed NP-hardness via a reduction from 3-SAT [12]. In contrast, we use reduction from  $k$  DISJOINT CONNECTING PATHS. We believe this reduction to be simpler. More importantly, we will be able to generalize the reduction to obtain a hardness of approximation result at a later point.

Considering the hardness of the motivating subgraph problem, Sect. 4 will focus on an optimization version of the problem. More formally, we want to compute the minimum reward that must be placed at  $t$  such that  $G$  contains a motivating subgraph. We will propose a simple  $(1 + \sqrt{n})$ -approximation algorithm that outputs the reward and a corresponding motivating subgraph. As the main technical contribution of this paper, we will show that this approximation is asymptotically tight. In particular, we will prove that the problem cannot be approximated efficiently within a ratio less than  $\sqrt{n}/3$  unless  $P = NP$ . Thus, we resolve the approximability of the motivating subgraph problem.

Finally, Sect. 5 will explore the problem of finding reward configurations within a fixed total budget of at most  $b$ . We will examine a version of the problem that, in our view, is the most sensible one. First, only positive rewards may be laid out. This assumption is reasonable as it is not entirely clear how negative rewards should be implemented in practice and how they are accounted

for in the designer’s budget. Secondly, the designer must only pay for rewards that are actually collected by the agent. This setting is fundamentally different from the settings analyzed by Tang et al. as it allows exploitative solutions. We show that the problem can be solved in polynomial-time if  $\beta = 0$  or  $\beta = 1$ . Using a reduction from SET PACKING, we prove that deciding the existence of a motivating reward configuration is NP-complete for general  $\beta \in (0, 1)$ , even if  $b = 0$ . This immediately implies that the optimization problem of finding the minimum  $b$  for which a motivating reward configuration exists cannot be approximated efficiently within any ratio greater or equal to 1 unless  $P = NP$ .

## 2 The Formal Model

In the following, we will present Kleinberg and Oren’s graph-based model [5]. Let  $G = (V, E)$  be a finite directed acyclic graph. Associated with each edge  $(v, w)$  is a non-negative cost  $c_G(v, w)$ . Furthermore, the project designer may lay out positive rewards  $r_G(v)$  at arbitrary nodes  $v$ . We call  $r$  a *reward configuration*. An agent with bias factor  $\beta \in [0, 1]$  has to incrementally construct a path from a source  $s$  to a target  $t$ . Located at some node  $v$  different from  $t$ , the agent evaluates its *lowest perceived net cost*. For this purpose it considers all paths  $P$  from  $v$  to  $t$ . However, only the initial edge of  $P$  is accounted for by its actual value. All other costs and rewards along  $P$  are discounted by a factor of  $\beta$ . More precisely, let  $d_{G,r}(w)$  denote the cost of a cheapest path from some node  $w$  to  $t$  with respect to the actual costs and rewards. Note that although  $d_{G,r}(w)$  might be negative depending on  $r$ , no negative cycles can occur as  $G$  is acyclic. If no path exists, we assume that  $d_{G,r}(w) = \infty$ . The lowest perceived net cost is defined as  $\zeta_{G,r}(v) = \min\{c_G(v, w) + \beta d_{G,r}(w) \mid (v, w) \in E\}$  if  $v$  has at least one outgoing edge. Otherwise,  $\zeta_{G,r}(v) = \infty$ . If  $\zeta_{G,r}(v) > 0$ , then the agent has no motivation to continue the project and abandons. Conversely, if  $\zeta_{G,r}(v) \leq 0$ , the agent traverses an edge  $(v, w)$  for which  $c_{G,r}(v, w) + \beta d_{G,r}(w) = \zeta_{G,r}(v)$ . Ties are broken arbitrarily. Note that the agent could take more than one path from  $s$  to  $t$ . A project is called motivating if the agent successfully reaches  $t$  along all such paths. To simplify our notation, we will omit  $G$  and  $r$  in the index of  $c$ ,  $r$ ,  $d$  and  $\zeta$  whenever the graph and reward configuration is clear from context.

To illustrate the model, we consider the car wash problem from Sect. 1 once more. Assume that Alice’s car must be washed during the next  $m$  days with  $m > 50$ . The task graph  $G$  is depicted in Fig. 1. For each day  $i$  with  $1 \leq i \leq m$  there is a node  $v_i$ . Let  $v_1$  be the source. There is an edge  $(v_i, t)$  of cost  $i/50$  that




Fig. 1. Task graph of the car wash problem

represents the task of washing the car on day  $i$ . To keep the drawing simple, the edges  $(v_i, t)$  merge in Fig. 1. Furthermore, for every  $i < m$  there is an edge  $(v_i, v_{i+1})$  of cost 0 to postpone the task to the next day. Assume for now that Bob is located at some  $v_i$  with  $i < m$ . His perceived cost for procrastinating is at least  $\beta(i + 1)/50$ . This bound is tight if he plans to traverse  $(v_i, v_{i+1})$  and then  $(v_{i+1}, t)$ . Alternatively, his perceived cost for using  $(v_i, t)$  and washing the car on day  $i$  is  $i/50$ . Recall that Alice offers Bob a single reward  $r(t) = 1$  upon completing the car wash. Furthermore,  $\beta = 1/3$ . As a result, the minimum perceived net cost is  $\zeta(v_i) = \beta(i + 1)/50 - \beta$ . We conclude that Bob always prefers to wash the car on the next day instead of doing it right away. Moreover, for  $i < 50$  it holds true that  $\zeta(v_i) \leq 0$ . This means that during the first 49 days, Bob moves along  $(v_i, v_{i+1})$  believing that he will finish the project the next day. However, once Bob reaches  $v_{50}$  he suddenly realizes that  $\zeta(v_{50}) > 0$  and abandons. Therefore, the car wash problem in its current form is not motivating.

Next, assume that  $(v_{16}, v_{17})$  is deleted from  $G$ . This can be thought of as a deadline set by Alice at day  $i = 16$ . Let  $G'$  be the resulting subgraph. When Bob reaches  $v_{16}$ , he cannot procrastinate anymore but must wash the car to get a reward. The perceived net cost is  $\zeta_{G'}(v_{16}) = 16/50 - \beta = -1/75$ . Because this is less than 0, he washes the car. This makes  $G'$  a motivating subgraph. It is interesting to note that no reward configuration in  $G$  is motivating for a budget of  $b < (m/50)/\beta$ . This is because no matter how much reward is placed at any of the nodes, Bob prefers to procrastinate until the very last day.

To illustrate the power of reward configurations, we will consider a second scenario. Suppose that Alice offers Bob a new job. If he first washes her car, which by now incurs a cost of 1, and afterwards also mows her lawn, which has a cost of 6, he receives a reward of 10. What Bob is unaware of is that Alice does not care about the lawn. Instead, she tries to get Bob to wash the car for free. We model this project with a task graph  $G$  consisting of a path from  $s$  to  $t$  via the intermediate node  $v$  and another path from  $v$  to  $t$  via  $w$ . The edge  $(s, v)$  corresponds to the car wash and has a cost of 1. Furthermore,  $(v, w)$  corresponds to mowing of the lawn and has a cost of 6. The edges  $(v, t)$  and  $(w, t)$  are of cost 0. Assuming that  $\beta = 1/3$ , there is a reward configuration  $r$  for which Bob will wash the car but will not claim a reward. Suppose Alice sets  $r(w) = 10$ . In this case, Bob traverses  $(s, v)$  with a minimum perceived net cost of  $\zeta(s) = -1/3$  along the path  $s, v, w, t$ . When at  $v$ , Bob reevaluates the net cost for traversing  $(v, w)$  to  $8/3$ . In contrast, finishing the project right away along  $(v, t)$  has cost 0. As a result, Bob changes his plan and moves to  $t$  without collecting the reward, although he already washed the car.

### 3 Finding Motivating Subgraphs

In this Section, we assume that the project designer may only place a single reward at  $t$ . This way, no exploitative reward configurations are possible. We first argue that the problem of finding a motivating subgraph can be solved in polynomial-time if  $\beta = 0$  or  $\beta = 1$ . Although this claim might seem intuitive,



we will be able to generalize the idea to show the existence of an  $(1 + \sqrt{n})$ -approximation algorithm for general  $\beta \in (0, 1)$  in Sect. 4.

**Proposition 1.** *If  $\beta = 0$  or  $\beta = 1$ , it is possible to find a motivating subgraph in polynomial-time for arbitrary  $r(t) \geq 0$ .*

*Proof.* First, assume  $\beta = 0$ . Because the agent has no value for future rewards, it must walk along a path of cost 0. Otherwise, it would abandon once it encounters an edge of positive cost. If such a path exists, it itself is a motivating subgraph. Conversely, if no such path exists, no subgraph can be motivating. Next, assume  $\beta = 1$ . In this case, the agent behaves time-consistent and follows a cheapest path from  $s$  to  $t$ . Therefore,  $G$  contains a motivating subgraph if and only if there is a path from  $s$  to  $t$  with a total edge cost less or equal to  $r(t)$ . Any subgraph containing such a path is motivating. Clearly, if a motivating subgraph exists, it can be found efficiently in both scenarios, i.e.  $\beta = 0$  and  $\beta = 1$ .  $\square$

Unfortunately, computing motivating subgraphs for general  $\beta \in (0, 1)$  is more challenging. We will give evidence for this in Theorem 1 by showing that the corresponding decision problem, which we name MOTIVATING SUBGRAPH (MS), is NP-complete for general  $\beta \in (0, 1)$ .

**Definition 1.** *Given a task graph  $G$ , a reward  $r(t) \geq 0$  and a bias factor  $\beta \in [0, 1]$ , decide the existence of a motivating subgraph of  $G$ .*

To prove NP-completeness of MS, we must first show that MS is contained in NP. For this purpose we will argue that it can be decided in polynomial-time whether a task graph is motivating for a given reward configuration. Note that a naive approach that simply simulates the agents walk through  $G$  must fail as the agent might take more than one path whenever it is indifferent between two options. A possible solution that preserves polynomial-time bounds is presented in the following proposition.

**Proposition 2.** *For any task graph  $G$ , reward configuration  $r$  and bias factor  $\beta \in [0, 1]$ , it can be decided in polynomial-time if  $G$  is motivating.*

*Proof.* We modify  $G$  in the following way. For each node  $v$  we calculate the lowest perceived net cost  $\zeta_{G,r}(v)$ . Next, we take a copy of  $G$ , say  $G'$ , in which we remove all edges  $(v, w)$  for which  $\zeta_{G,r}(v) < c_G(v, w) + \beta d_{G,r}(w)$  or  $\zeta_{G,r}(v) > 0$ . In other words, we remove all edges from  $G'$  that do not minimize the agent's perceived net cost or are not motivating. Let  $V'$  be the set of all nodes that can be reached from  $s$  in  $G'$ . Observe that  $V'$  contains exactly those nodes that might be visited by the agent in  $G$ . Clearly,  $G$  is motivating if and only if the agent can reach  $t$  from all nodes of  $V'$  via some path in  $G'$ . This condition can be checked in polynomial-time.  $\square$

To show NP-hardness, we will use a reduction from  $k$  DISJOINT CONNECTING PATHS ( $k$ -DCP), which is defined as follows [3]:

**Definition 2.** Given a graph  $H$  and  $k$  disjoint node pairs  $(s_1, t_1), \dots, (s_k, t_k)$ , decide if  $H$  contains  $k$  mutually node-disjoint paths, one connecting every  $s_i$  to the corresponding  $t_i$ .

Lynch showed that  $k$ -DCP is NP-complete if  $H$  is undirected [8]. A simple modification of Lynch’s reduction, which can be found in the full version of this paper, implies that  $k$ -DCP is also NP-complete if  $H$  is directed and acyclic.

Before we finally tackle Theorem 1, we want to draw attention to a useful price structure that is common to all reductions presented in this paper. Imagine a directed path along  $k+1$  edges, such that the  $i$ -th edge has a cost of  $(1-\beta)^{k+1-i}$ . According to the following Lemma, the agent’s perceived cost for following the path to its end is 1 at every node except for the last.

**Lemma 1.** For every positive integer  $k$  and bias factor  $\beta \in [0, 1]$  it holds that:

$$(1-\beta)^k + \beta \left( \sum_{i=0}^{k-1} (1-\beta)^i \right) = 1.$$

The proof of Lemma 1 can be found in the full version of this paper. We are now ready to show NP-completeness of MS.

**Theorem 1.** MS is NP-complete for any bias factor  $\beta \in (0, 1)$ .

*Proof.* By Proposition 2, any motivating subgraph  $G'$  serves as a certificate for a “yes”-instance of MS. Consequently, MS is in NP. To complete the proof, we will establish NP-hardness via a polynomial reduction from  $k$ -DCP.

Consider an instance  $\mathcal{I}$  of  $k$ -DCP consisting of a directed acyclic graph  $H$  and  $k$  disjoint node pairs  $(s_1, t_1), \dots, (s_k, t_k)$ . We will embed  $H$  into the task graph  $G$  of an MS instance  $\mathcal{J}$  such that  $G$  has a motivating subgraph if and only if  $H$  has  $k$  disjoint connecting paths. For this purpose we assume that the encoding length of  $\beta \in (0, 1)$  is polynomial in that of  $\mathcal{I}$  and set  $r(t) = 1/\beta$ . The task graph  $G$ , which is illustrated in Fig. 2, is constructed as follows:

To get from  $s$  to  $t$ , the agent must follow the so called *main path* along intermediate nodes  $v_1, \dots, v_{k+3}$ . The first  $k+1$  edges of this main path each have a cost of  $(1-\beta)^3 - \varepsilon$ , with  $\varepsilon$  being a positive constant satisfying

$$\varepsilon < \min \left\{ \beta \frac{1-\beta}{k+1}, \beta \frac{(1-\beta)^3}{1+\beta} \right\}.$$

The last three edges have a cost of  $(1-\beta)^2$ ,  $1-\beta$  and 1, respectively. To keep the agent motivated, we introduce  $k$  *shortcuts* that connect every  $v_i$  with  $1 \leq i \leq k$  to  $t$  via the embedding of  $H$ . More formally, the  $i$ -th shortcut starts at  $v_i$  and is routed through a distinct node  $w_i$  via an edge of cost  $(1-\beta)^2$ . Node  $w_i$  is then connected to  $s_i$  via an edge of cost  $(k+1-i)(1-\beta)/(k+1)$ . Finally,  $t_i$  is connected to  $t$  via an edge of cost  $i(1-\beta)/(k+1) + 1$ . To keep Fig. 2 simple, the edges  $(t_i, t)$  are merged and their cost is depicted as two terms, namely  $i(1-\beta)/(k+1)$  and  $+1$ . Note that the prices of  $(w_i, s_i)$  and  $(t_i, t)$  complement




Fig. 2. Reduction from a general  $k$ -DCP instance:  $H$

each other, i.e. they sum to  $(1 - \beta) + 1$ . The edges of  $H$  all have a cost of 0. The resulting graph  $G$  is acyclic and its encoding length polynomial in  $\mathcal{I}$ .

It remains to show, that  $\mathcal{J}$  has a solution if and only if  $\mathcal{I}$  has one. ( $\Rightarrow$ ) First, suppose  $\mathcal{I}$  has a solution, i.e. there exist  $k$  node-disjoint connecting paths. Let  $G'$  be a subgraph of  $G$  obtained by deleting all edges of  $H$  that are not part of one of these paths. Furthermore, let  $s = v_0$  and assume the agent is located at  $v_i$  with  $0 \leq i \leq k$ . According to Lemma 1, the agent perceives a net cost of  $-\epsilon$  for taking the  $(i + 1)$ -st shortcut or if  $i = k$  for following the main path. In contrast, if  $0 < i \leq k$ , the perceived net cost of the  $i$ -th shortcut is 0. As a result, the agent follows the main path to  $v_{k+1}$  and then for lack of other options continues to  $t$ . We conclude that  $G'$  is a motivating subgraph of  $G$ .

( $\Leftarrow$ ) Due to space constraints, we only sketch this direction of the proof. A thorough analysis can be found in the full version of this paper. Suppose  $\mathcal{I}$  has no  $k$  node-disjoint connecting paths. Consequently, any subgraph of  $G$  must contain at least one shortcut  $i$  such that the cheapest path from  $v_i$  to  $t$  via  $w_i$  is different from  $(1 - \beta)^2 + (1 - \beta) + 1$ . We call such a shortcut *degenerate*. We distinguish between two scenarios: either the degenerate shortcuts are too expensive, and the agent loses motivation on the main path, or a degenerate shortcut is so cheap that the agent enters it. In the first case, the agent clearly abandons. In the latter case, the agent must traverse one of the edges  $(t_i, t)$  to reach  $t$ . However, since the price of  $(t_i, t)$  is greater than 1, a reward of  $r(t) = 1/\beta$  is not sufficiently motivating for the agent to take this step. Again, it abandons. Therefore, no subgraph can be motivating.  $\square$

### 4 Approximating Reward Optimal Subgraphs

Considering that the decision problem MS is NP-hard, the next and arguably natural question is whether good approximation algorithms exist. Therefore, we restate MS as an optimization problem that we call MS-OPT.

**Definition 3.** Given a task graph  $G$  and a bias factor  $\beta \in (0, 1)$ , determine the minimum reward  $r(t)$  such that  $G$  contains a motivating subgraph.

We will present two simple approximation algorithms: one that performs well for small values of  $\beta$  and one that leads to good solutions for large  $\beta$ . The algorithms return a reward  $r(t)$  as well as a corresponding motivating subgraph  $G'$ . Combining both algorithms eventually yields a general approximation algorithm with a ratio of  $(1 + \sqrt{n})$  for any  $\beta \in (0, 1)$ .

First, we assume that  $\beta$  is small. Because the agent is highly oblivious to the future, it is sensible to guide it along a path with minimal maximum edge cost. Paths with this property are called *minmax paths*. A minmax path can be computed easily in polynomial-time. For instance, starting with an empty subgraph, the edges of  $G$  can be inserted in non-decreasing order of cost until  $s$  and  $t$  become connected for the first time. Any path from  $s$  to  $t$  in the resulting subgraph is a minmax path. Our first algorithm, called `MINMAXPATHAPPROX`, computes a minmax path  $P$  from  $s$  to  $t$  and returns a subgraph  $G'$  whose edges are that of  $P$ . Furthermore,  $r(t)$  is chosen such that  $\max\{\zeta_{G',r}(v) \mid v \in P\} = 0$ . Clearly, this reward is sufficient to make  $G'$  motivating.

**Proposition 3.** `MINMAXPATHAPPROX` has an approximation ratio of  $1 + \beta n$ .

*Proof.* Let  $c$  denote the maximum cost among the edges of the minmax path  $P$  computed by `MINMAXPATHAPPROX`. By definition of  $P$ , the agent must encounter an edge of cost  $c$  or more in any subgraph that connects  $s$  with  $t$ . Thus the optimal reward is lower bounded by  $c/\beta$ . Conversely, the cost of every edge in  $P$ , of which there are at most  $n - 1$ , is  $c$  or less. This means that the reward returned by `MINMAXPATHAPPROX` is upper bounded by  $r(t) \leq c/\beta + (n - 2)c \leq c/\beta + nc$ . From this the desired approximation ratio of  $1 + \beta n$  follows immediately.  $\square$

Next, suppose that  $\beta$  is large and the agent is hardly present-biased at all. Our second algorithm, called `CHEAPESTPATHAPPROX`, simply computes a path  $P$  of minimum cost from  $s$  to  $t$  and returns a subgraph  $G'$  containing the edges of  $P$ . Again, the algorithm chooses  $r(t)$  such that  $\max\{\zeta_{G',r}(v) \mid v \in P\} = 0$ .

**Proposition 4.** `CHEAPESTPATHAPPROX` has an approximation ratio of  $1/\beta$ .

*Proof.* Let  $P$  be the path computed by `CHEAPESTPATHAPPROX` and  $c$  the total cost of  $P$ . At any node  $v$  of  $P$  the agent's perceived net cost is at most  $d_{G',r}(v) - \beta r(t)$ , which is less than  $c - \beta r(t)$ . The reward returned by `CHEAPESTPATHAPPROX` is therefore at most  $c/\beta$ . Conversely, when located at  $s$ , the agent perceives a cost of at least  $\beta c$  in any subgraph of  $G$ , including the optimal one. Consequently, a reward of  $c$  or more is required to motivate the agent. This establishes the approximation ratio of  $1/\beta$ .  $\square$

It is interesting to see how `MINMAXPATHAPPROX` and `CHEAPESTPATHAPPROX` generalize the algorithmic ideas of Proposition 1. If we combine the two and use `MINMAXPATHAPPROX` whenever  $\beta \leq 1/\sqrt{n}$  and `CHEAPESTPATHAPPROX` otherwise, we obtain a general approximation algorithm called `COMBINEDAPPROX`. Propositions 3 and 4 directly imply the following result.

**Theorem 2.** COMBINEDAPPROX has an approximation ratio of  $1 + \sqrt{n}$ .

Although the algorithmic techniques of COMBINEDAPPROX are simple, the following Theorem implies that asymptotically the approximation ratio is the best we can hope for in polynomial-time.

**Theorem 3.** MS-OPT is NP-hard to approximate within a ratio less than  $\sqrt{n}/3$ .

*Proof.* To establish hardness of approximation, we will use another reduction from  $k$ -DCP. Let  $\mathcal{I}$  be an instance of  $k$ -DCP that consists of a directed acyclic graph  $H$  and  $k$  disjoint node pairs  $(s_1, t_1), \dots, (s_k, t_k)$ . Furthermore, let  $\varrho$  be an arbitrary positive integer. The best choice of  $\varrho$  will be determined later. We will construct an instance  $\mathcal{J}$  of MS-OPT that consists of a task graph  $G$  and has the following two properties: (a) If  $\mathcal{I}$  has a solution, then  $G$  has a subgraph that is motivating for a reward of  $r(t) = 1/\beta$ . (b) If  $\mathcal{I}$  does not have a solution, then no subgraph of  $G$  is motivating for a reward of  $r(t) = \varrho/\beta$  or less. Consequently, any algorithm achieving an approximation ratio of  $\varrho$  or better must solve  $\mathcal{I}$ .

Unlike Theorem 1, the bias factor cannot be chosen arbitrarily anymore. Considering that Proposition 4 gives a  $(1/\beta)$ -approximation,  $\beta$  must be less than  $1/\varrho$ . For convenience, we set  $\beta = 1/(3\varrho + 3)$ . From a structural point of view, the task graph  $G$  consists of two units: the *embedding unit* and the *amplification unit*. The first unit contains an embedding of  $H$ , while the second unit amplifies approximation errors occurring in the embedding unit.


The overall structure of the embedding unit is similar to the task graph of Theorem 1. There exists a *main path* and  $k$  *shortcuts* that link to the embedding of  $H$ . However, there are some differences. First, the main path starts at the last node of the amplification unit  $u_{9\varrho^2}$  and passes  $k + 3\varrho + 3$  intermediate nodes  $v_1, \dots, v_{k+3\varrho+3}$  before it ends in  $t$ . The first  $k + 1$  edges of the main path each have a cost of  $(1 - \beta)^{3\varrho+3} - \varepsilon$ , where  $\varepsilon$  is a positive value satisfying

$$\varepsilon < \min \left\{ \beta \frac{(1 - \beta)^{3\varrho+1}}{k + 1}, \beta \frac{(1 - \beta)^{3\varrho+3}}{1 + \beta}, \frac{1}{1 + \varrho}, (1 - \beta)^{3\varrho+3} - \frac{1}{3} \right\}.$$

The remaining edges  $(v_i, v_{i+1})$  of the main path, with  $k < i \leq k + 3\varrho + 3$  and  $t = v_{k+3\varrho+3+1}$ , have an increasing cost of  $(1 - \beta)^{k+3\varrho+3-i}$ . Furthermore, the initial edge  $(v_i, w_i)$  of each shortcut has a cost of  $(1 - \beta)^{3\varrho+2}$ , while the edges  $(w_i, s_i)$  and  $(t_i, t)$  have complementing prices of  $(k + 1 - i)(1 - \beta)^{3\varrho+1}/(k + 1)$  and  $i(1 - \beta)^{3\varrho+1}/(k + 1) + \sum_{j=0}^{3\varrho} (1 - \beta)^j$ . All edges of  $H$  are again free of charge. As a result, the cost of each shortcut sums up to  $\sum_{j=0}^{3\varrho+2} (1 - \beta)^j$ .

The amplification unit, which is shown in Fig. 3, consists of an *amplification path* connecting  $s$  to  $u_{9\varrho^2}$  along the intermediate nodes  $u_1, \dots, u_{9\varrho^2-1}$ . Each edge of the amplification path has a cost of  $(1 - \beta)^{3\varrho+3} - \varepsilon$ . From every  $u_i$  there is also an edge of cost  $(1 - \beta)^{3\varrho+2}$  to a common node  $z$ . Node  $z$  is in turn connected to  $t$  via an edge of cost  $\sum_{j=0}^{3\varrho+1} (1 - \beta)^j$ .

To conclude the proof, we must show that our construction satisfies properties (a) and (b) stated above. We start with (a) and assume that  $k$  node-disjoint



**Fig. 3.** Amplification unit

paths exist in  $H$ . Let  $G'$  be a subgraph of  $G$  obtained by deleting all edges of  $H$  that are not part of such a path. Furthermore, we set  $r(t) = 1/\beta$  and  $s = u_0$ . When located at  $u_i$  with  $0 \leq i \leq 9\varrho^2$ , Lemma 1 suggests that the agent perceives a net cost of  $-\varepsilon$  for traversing  $(u_i, u_{i+1})$  and then following  $(u_{i+1}, z)$  or the first shortcut of the embedding unit if  $i = 9\varrho^2$ . Conversely, if  $i > 0$ , the agent evaluates the net cost of walking along  $(u_i, z)$  to 0. As a result, the agent follows the amplification path until it reaches  $v_1$ . From this point on it travels along the main path of the embedding unit until it eventually arrives at  $t$  for the same reasons given in Theorem 1. This means that  $G'$  is a motivating subgraph for a reward of  $r(t) = 1/\beta$ .

Due to space constraints, we refer to the full version of this paper for a thorough proof of statement (b). At this point we will confine ourselves to a brief sketch of the main ideas. Assume that  $\mathcal{I}$  has no  $k$  node-disjoint paths. As argued before in Theorem 1, at least one *degenerate shortcut* must exist in any subgraph of  $G$  whose minimum cost is different from the target value of  $\sum_{j=0}^{3\varrho+2} (1-\beta)^j$ . As a result, two scenarios are conceivable. First, the perceived cost of some shortcut is so low that the agent diverts from the main path. However, in this case the agent must pass one of the edges  $(t_i, t)$ , whose cost is greater than  $\varrho$ , to reach  $t$ . Consequently, no reward  $r(t) \leq \varrho/\beta$  can be motivating. If the first case does not apply, one can argue that the perceived net cost of the main path at  $u_{9\varrho^2}$  is greater than  $1 - \beta r(t)$ . To prevent the agent from moving to  $z$ , which results in cost greater than  $\varrho$  at  $(z, t)$ , all edges  $(u_i, z)$  must be removed from the amplification unit. However, in this case the agent perceives a net cost greater than  $\varrho - \beta r(t)$  at  $s$ . Again, no reward  $r(t) \leq \varrho/\beta$  is sufficiently motivating. If  $\varrho$  is the number of nodes in  $H$ , then our lower bound on the approximability of MS-OPT converges to  $\sqrt{n}/3$  as the size of  $H$  increases.  $\square$

## 5 Motivation Through Intermediate Rewards

In this section, we study the complexity of motivating agents through the strategic placement of rewards. In this scenario, the task graph must not be pruned. The goal is to minimize the total value of the rewards along the agent's walk from  $s$  to  $t$ . Similar to the previous setting of Sects. 3 and 4, a motivating reward configuration within a given budget  $b$  can be computed in polynomial-time if  $\beta = 0$  or  $\beta = 1$ .

**Proposition 5.** *A motivating reward configuration within budget  $b$  can be computed in polynomial-time for  $\beta = 0$  or  $\beta = 1$ .*

*Proof.* First, suppose that  $\beta = 0$ . In this case, the agent does not care for any future rewards and only traverses edges of cost 0. Let  $V'$  be the set of nodes that can be reached from  $s$  for cost 0. Note that  $V'$  contains exactly those nodes that might be visited by the agent independent of the specific reward configuration. As a result,  $G$  has a motivating reward configuration if and only if  $t$  can be reached from every node of  $V'$  for a cost of 0. Because no rewards need to be placed in this scenario, the budget constraint is always satisfied. Next, assume that  $\beta = 1$ . In this case the agent is time-consistent. Let  $c$  be the cost of a cheapest path from  $s$  to  $t$ . Setting  $r(t) = c$  yields a motivating and also optimal reward configuration. The required budget is  $c$ . Clearly both cases,  $\beta = 0$  and  $\beta = 1$  can be solved in polynomial time.  $\square$

As before, the problem becomes much harder for general  $\beta \in (0, 1)$ . In particular, the corresponding decision problem MOTIVATING REWARD CONFIGURATION (MRC), which we define below, is NP-hard.

**Definition 4.** *Given a task graph  $G$ , a budget  $b$  and a bias factor  $\beta \in [0, 1]$ , decide the existence of a motivating reward configuration  $r$  such that the total reward collected on any walk of the agent is at most  $b$ .*

The following proposition establishes membership of MRC in NP.

**Proposition 6.** *For any task graph  $G$ , reward configuration  $r$  and bias factor  $\beta \in [0, 1]$ , it is possible to decide in polynomial-time if  $r$  is motivating within a given budget  $b$ .*

A proof of Proposition 6 can be found in the full version of this paper. To show NP-hardness of MRC, we will use a reduction from SET PACKING (SP), cf. [3]. For convenience, the definition of SP is stated below.


**Definition 5.** *Given a collection of finite sets  $S_1, \dots, S_\ell$  and an integer  $k \leq \ell$ , decide if at least  $k$  of these sets are mutually disjoint.*

We are now ready to prove NP-completeness of MRC. Note that the problem remains hard even if the budget is 0.

**Theorem 4.** *MRC is NP-complete for any bias factor  $\beta \in (0, 1)$ , even if  $b = 0$ .*

*Proof.* By Proposition 6, we can use any motivating reward configuration  $r$  within budget  $b$  as certificate for a “yes”-instance of MRC. This establishes membership of MRC in NP. To prove NP-hardness we will present a polynomial-time reduction from SP to MRC. We focus on the case that  $b = 0$ . A modified reduction for budgets  $b > 0$  can be found in the full version of this paper.

Let  $\mathcal{I}$  be an instance of SP consisting of finite sets  $S_1, \dots, S_\ell$  and an integer  $k \leq \ell$ . We start by constructing an MRC instance  $\mathcal{J}$  that has a motivating reward configuration within a budget of  $b = 0$  if and only if  $\mathcal{I}$  has a solution.



**Fig. 4.** Reduction from the SP instance:  $S_1 = \{a, c, d\}$ ,  $S_2 = \{a, b\}$ ,  $S_3 = \{b, c, e\}$ ,  $S_4 = \{b, e\}$  and  $k = 2$

Figure 4 depicts the task graph  $G$  for a small sample instance of SP. In general,  $G$  consists of a source  $s$ , a target  $t$  and  $1 \leq i \leq k$  levels of nodes  $v_{i,j}$  with  $1 \leq j \leq \ell$ . For every  $v_{i,j}$  with  $i < k$  there is a so called *upward edge* to every node  $v_{i+1,j'}$  on the next level. To maintain readability, upward edges are omitted in Fig. 4. In addition to the upward edges, there is an edge from  $s$  to every node  $v_{1,j}$  on the bottom level and an edge towards  $t$  from every node  $v_{k,j}$  on the top level. The idea behind this construction is that the agent walks along the upward edges from  $s$  to  $t$  in such a way that the nodes  $v_{1,j}, \dots, v_{k,j'}$  on its path correspond to a collection of  $k$  mutually disjoint sets  $S_j, \dots, S_{j'}$ . The cost of the initial edges  $(s, v_{1,j})$  and all upward edges  $(v_{i,j}, v_{i+1,j'})$  is  $1 - \beta - \varepsilon$ . Note that  $\beta \in (0, 1)$  might be an arbitrary value with an encoding length that is polynomial in that of  $\mathcal{I}$ . Moreover  $\varepsilon$  is a positive value satisfying

$$\varepsilon < \min \left\{ \frac{(1 - \beta)^2}{k}, \frac{\beta - \beta^2}{k - 1 + \beta} \right\}.$$

The cost of the edges  $(v_{k,j}, t)$  is 0.

In order to motivate the agent, we add *shortcuts* to  $G$  that connect every  $v_{i,j}$  to  $t$  via an intermediate node  $w_{i,j}$ . The first edge  $(v_{i,j}, w_{i,j})$  has cost 1 and the second edge  $(w_{i,j}, t)$  has cost 0. In Fig. 4 the second edges are omitted for the sake of readability. Note that a reward of value less than  $1/\beta$  can be placed on  $w_{i,j}$  without the agent claiming it. Furthermore, if the reward is at least  $(1 - \varepsilon)/\beta$ , all edges  $(v_{i-1,j'}, v_{i,j})$ , or  $(s, v_{i,j})$  if  $i = 1$ , become motivating.

We finish our construction by connecting each node  $v_{i,j}$  with all nodes  $w_{i',j'}$  for which  $i' < i$  and  $S_j \cap S_{j'} \neq \emptyset$  via a *downward path*. Each downward path consists of two edges: the first one is of cost 0 and the second one is of cost  $(1 - \beta - k\varepsilon)/(\beta - \beta^2)$ . In Fig. 4, downward paths are drawn as single dashed edges. The idea behind these paths is to enforce the disjointness constraint of  $\mathcal{I}$ . In the next paragraph we will address this in more detail. But first note that  $G$  is an acyclic graph that is polynomial in the size of  $\mathcal{I}$ . It remains to show that  $\mathcal{J}$  has a solution if and only if  $\mathcal{I}$  has one.

Due to space constraints, we refer to the full version of this paper for a detailed proof. To offer some more insight, we briefly present a sketch of the main ideas. We first observe that the agent cannot enter any shortcut or downward



path on its way from  $s$  to  $t$ . The reason for this is that a positive reward must be placed onto such a path for the agent to enter it. However, once the agent enters a shortcut or downward path it either collects the reward or abandons. In both cases the given reward configuration is not motivating for a budget of 0. Consequently, the agent must climb from level to level until it reaches  $t$ . As a motivation, rewards need to be placed on selected nodes  $w_{i,j}$ . If the reward is chosen correctly, for instance  $r(w_{i,j}) = (1 - \varepsilon)/\beta$ , this is sufficiently motivating for the agent to move from any node on level  $i - 1$  to  $v_{i,j}$ , but not motivating enough for the agent to enter the shortcut from  $v_{i,j}$  to  $t$ . Next, assume that the agent is located at some node  $v_{i,j}$ . To prevent it from taking the downward path, no substantial rewards may be placed on any node  $w_{i',j'}$  with  $i' < i$  and  $S_j \cap S_{j'} \neq \emptyset$ . By construction of  $G$  such a reward configuration is possible if and only if  $\mathcal{I}$  has a feasible solution.  $\square$

Finally, we look at the optimization variant of MRC called MRC-OPT.

**Definition 6.** *Given a task graph  $G$  and a bias factor  $\beta \in (0, 1)$ , determine the infimum of all budgets  $b$  for which there exists a reward configuration  $r$  such that the total reward collected on any of the agent's walks is at most  $b$ .*

The fact that MRC is NP-complete for  $b = 0$  immediately implies that MRC-OPT does not permit any efficient approximation algorithm unless  $P = NP$ .

**Corollary 1.** *MRC-OPT is NP-hard to approximate within any ratio greater or equal to 1.*

## References

1. Akerlof, G.A.: Procrastination and obedience. *Am. Econ. Rev.* **81**, 1–19 (1991)
2. Bryan, G., Karlan, D., Nelson, S.: Commitment devices. *Ann. Rev. Econ.* **2**, 671–698 (2010)
3. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Co., New York (1979)
4. Gravin, N., Immorlica, N., Lucier, B., Pountourakis, E.: Procrastination with variable present bias. In: 17th ACM Conference on Economics and Computation, pp. 361–361. ACM, New York (2016)
5. Kleinberg, J., Oren, S.: Time-inconsistent planning: a computational problem in behavioral economics. In: 15th ACM Conference on Economics and Computation, pp. 547–564. ACM, New York (2014)
6. Kleinberg, J., Sigal, O., Manish, R.: Planning problems for sophisticated agents with present bias. In: 17th ACM Conference on Economics and Computation, pp. 343–360. ACM, New York (2016)
7. Laibson, D.: Golden eggs and hyperbolic discounting. *Q. J. Econ.* **112**, 443–477 (1997)
8. Lynch, J.F.: The equivalence of theorem proving and the interconnection problem. *SIGDA Newsl.* **5**, 31–36 (1975)
9. O'Donoghue, T., Rabin, M.: Doing it now or later. *Am. Econ. Rev.* **89**, 103–124 (1999)

10. O'Donoghue, T., Rabin, M.: Incentives and self control. In: *Advances in Economics and Econometrics: Theory and Application 9th World Congress*, vol. 2, pp. 215–245. Cambridge University Press, Cambridge (2006)
11. O'Donoghue, T., Rabin, M.: Procrastination on long-term projects. *J. Econ. Behav. Organ.* **66**, 161–175 (2008)
12. Tang, P., Teng, Y., Wang, Z., Xiao, S., Xu, Y.: *Computational issues in time-inconsistent planning* (2015)

# FPT Approximation Schemes for Maximizing Submodular Functions

Piotr Skowron<sup>(✉)</sup>

University of Oxford, Oxford, UK  
piotr.skowron@cs.ox.ac.uk

**Abstract.** We investigate the existence of approximation algorithms for maximization of submodular functions, that run in fixed parameter tractable (FPT) time. Given a non-decreasing submodular set function  $v : 2^X \rightarrow \mathbb{R}$  the goal is to select a subset  $S$  of  $K$  elements from  $X$  such that  $v(S)$  is maximized. We identify two properties of set functions, referred to as  $p$ -separability properties, and we argue that many real-life problems can be expressed as maximization of submodular,  $p$ -separable functions, with low values of the parameter  $p$ . We present FPT approximation schemes for the minimization and maximization variants of the problem, for several parameters that depend on characteristics of the optimized set function, such as  $p$  and  $K$ . We confirm that our algorithms are applicable to a broad class of problems, in particular to problems from computational social choice, such as item selection or winner determination under several multiwinner election systems.

## 1 Introduction

We study (exponential-time) approximation algorithms for maximizing non-decreasing submodular set functions. A set function  $v : 2^X \rightarrow \mathbb{R}$  is submodular if for each two subsets  $A \subseteq B \subset X$  and each element  $x \in X \setminus B$  it holds that  $v(A \cup \{x\}) - v(A) \geq v(B \cup \{x\}) - v(B)$ ;  $v$  is non-decreasing if for each two subsets  $A \subseteq B \subset X$  it holds that  $v(A) \leq v(B)$ . Our goal is to select a subset  $S$  of  $K$  elements from  $X$  such that the value  $v(S)$  is maximal.

Maximization of non-decreasing submodular functions is a very general problem that is extensively used in various research areas, from recommendation systems [21, 28], through voting theory [21, 29], image engineering [12, 13, 25], information retrieval [19, 34], network design [15, 16], clustering [22], speech recognition [20], to sparse methods [1, 6]. Algorithms for maximization of non-decreasing submodular functions are applicable to other general problems of fundamental significance, such as the MAXCOVER problem [4, 27]. The universal relevance of the problem implies that the existence of good (approximation) algorithms for it is highly desired.

Indeed, the problem has already received a considerable amount of attention in the scientific community. For instance, it is known that the greedy algorithm, i.e., the algorithm that starts with the empty set and in each of  $K$  consecutive steps adds to the partial solution such an element from  $X$  that increases the value

of the optimized function most, is an  $(1 - 1/e)$ -approximation algorithm for maximization of non-decreasing submodular functions [23]. The same approximation ratio can be achieved for the distributed [17] and online [30] variants of the problem. Algorithms for maximizing non-monotone submodular functions have been studied by Feige et al. [9], and the approximability of the problem with additional constraints has been investigated by Calinescu et al. [2], Sviridenko [31], Lee et al. [18], and Vondrák et al. [33]. Iwata et al. [11] have provided algorithmic view on minimizing submodular functions. For the survey on maximization of submodular functions we refer the reader to the work of Krause and Golovin [14].

Unfortunately, the approximation guarantees of the greedy algorithm cannot be improved without compromising the efficiency of computation. For example, the MAXCOVER problem can be expressed as maximization of a non-decreasing submodular function, yet it is known that under standard complexity assumptions no polynomial-time algorithm can approximate it better than with ratio  $(1 - 1/e)$  [8]. Motivated by this fact, and provoked by the desire to obtain better approximation guarantees, we turn our attention to algorithms that run in super-polynomial time. In our studies we follow the approach of parameterized complexity theory and look for algorithms that run in fixed parameter tractable time (in FPT time), for some natural parameters. To the best of our knowledge, FPT approximation of optimizing submodular functions has not been considered in the literature before.

Parameterized complexity theory aims at investigating how the complexity of a problem depends on the size of different parts of input instances, called parameters. An algorithm runs in FPT time for a parameter  $P$  if it solves each instance  $I$  of the problem in time  $O(f(|P|) \cdot \text{poly}(|I|))$ , where  $f$  is a computable function. This definition excludes a large class of algorithms, such as the ones with complexity  $O(|I|^{|P|})$ . From the point of view of parameterized complexity, FPT is seen as the class of easy problems. Intuitively, the complexity of an FPT algorithm consists of two parts:  $f(|P|)$ , which is relatively low for small values of the parameter, and  $\text{poly}(|I|)$  which is relatively low even for larger instances, because of polynomial relation between the computation time and the size of an instance. For details on parameterized complexity theory, we point the reader to appropriate overviews [5, 7, 10, 24].

We identify several parameters that we believe are suitable for a complexity analysis of maximization of non-decreasing submodular functions. Perhaps the most natural parameter to consider is the required size of solutions,  $K$ . Our other parameters depend on characteristics of the optimized set function. Specifically, we define a new property of set functions, called  $p$ -separability, and provide evidence that  $p$  is a natural parameter to consider. We do that in Sect. 4, by presenting several examples of real-life computational problems that can be expressed as maximization of submodular  $p$ -separable set functions, where the value of  $p$  is small.

Our main contribution is presentation and analysis of algorithms for the problem. We construct fixed parameter tractable approximation schemes, i.e., collections of algorithms that run in FPT time and that can achieve arbitrarily good approximation ratios. We provide algorithms for two variants of the

problem: in the first variant, referred to as the *maximization variant*, the goal is to maximize the value  $v(S)$ . In the second one, referred to as the *minimization variant*, the goal is to minimize  $(v(X) - v(S))$ . While these two variants of the problem have the same optimal solutions, they are not equivalent in terms of their approximability. Indeed, if there exists a solution  $S$  with objectively high value, i.e., if  $v(S)$  is close to  $v(X)$ , then approximation algorithm for the minimization variant of the problem will be usually superior. For instance, if there exists a solution  $S$  such that  $v(S) = 0.95 \cdot v(X)$ , then a 2-approximation algorithm for the minimization variant of the problem is guaranteed to return a solution with the value better than  $0.9 \cdot v(X)$ . On the other hand, a  $1/2$ -approximation algorithm for the maximization variant of the problem is allowed to return, in such a case, a solution with value  $0.475 \cdot v(X)$ . Conversely, if the value of an optimal solution is significantly lower than the value of the whole set  $X$ , then a good approximation algorithm for the maximization variant of the problem will produce solutions of a better quality.

Our algorithms run in FPT time for the parameter  $(K, p)$ , where  $K$  is the size of the solution, and  $p$  is the lowest value such that the set function is  $p$ -separable. To address the case of functions which are not  $p$ -separable for any reasonable values  $p$ , we define a weaker form of approximability, referred to as approximation of the *minimization-or-maximization variant*—here, the goal is to find a subset  $S$  that is good in one of the previous two metrics. Such algorithms are also desired as they are guaranteed to find good approximation solutions, provided high quality solutions exist (i.e., if values of the optimal solutions are close to  $v(X)$ ). We show that there exists a randomized FPT approximation scheme for minimization-or-maximization variant of the problem for the parameter  $(K, \sum_{x \in X} v(\{x\})/v(X))$ .

We believe that the consequences of our general results are quite significant. In particular, in Sect. 4, we prove the existence of FPT approximation schemes for some natural problems in the computational social choice, in the matching theory, and in the theoretical computer science.

## 2 Notation and Definitions

Let  $X$  denote the universe set. We consider a set function  $v : 2^X \rightarrow \mathbb{R}$  that is non-negative, i.e., such that for each  $S \subseteq X$  we have  $v(S) \geq 0$ . We say that a function  $v$  is non-decreasing if for each two subsets  $A \subseteq B \subseteq X$  it holds that  $v(B) \geq v(A)$ . A set function  $v$  is submodular if for each two subsets  $A \subseteq B \subset X$  and each element  $x \in X \setminus B$  it holds that  $v(A \cup \{x\}) - v(A) \geq v(B \cup \{x\}) - v(B)$ . There are numerous equivalent conditions characterizing submodular functions—for a survey we refer the reader to the seminal article of Nemhauser et. al. [23]. It is easy to see that if the set function  $v$  is non-decreasing and submodular, then for each two subsets  $A \subseteq B \subset X$  and each element  $x \in X$  it holds that  $v(A \cup \{x\}) - v(A) \geq v(B \cup \{x\}) - v(B)$  (here, we do not have to assume that  $x \in X \setminus B$ ).

Below, we define a new class of properties of set functions.

**Definition 1 (*p*-separable set function).** A submodular set function  $v : 2^X \rightarrow \mathbb{R}$  is:

1. *p*-superseparable, if for each  $S \subseteq X$  we have:

$$\sum_{x \in X} (v(S \cup \{x\}) - v(S)) \geq \left( \sum_{x \in X} v(\{x\}) \right) - p \cdot v(S), \tag{1}$$

2. *p*-subseparable, if for each  $S \subseteq X$  we have:

$$\sum_{x \in X} (v(S \cup \{x\}) - v(S)) \leq p \cdot v(X) - p \cdot v(S). \tag{2}$$

For better intuition on the above definitions, we refer the reader to Sect. 4 where we present several examples of natural problems which can be expressed as optimization of separable functions for low values of the parameter  $p$ . Indeed,  $p$  can be naturally bounded by  $|X|$ : it is easy to see that each monotone and submodular function is  $|X|$ -superseparable and  $|X|$ -subseparable. Yet, in Sect. 4 we show that the value of parameter  $p$  in many natural problems is significantly lower.

We observe that a linear combination of  $p$ -superseparable functions is  $p$ -superseparable. The same comment applies to  $p$ -subseparability. As we will see in Sect. 4, this observation is helpful in proving that certain set functions are  $p$ -separable.

In this paper we investigate the problem of selecting  $K$  elements from  $X$  that, altogether, maximize the value of the set function  $v$ .

**Definition 2 (BestKSubset).** For a set of elements  $X$ , a polynomially computable set function  $v : 2^X \rightarrow \mathbb{R}$ , and an integer  $K$ , the solution to the BESTK-SUBSET problem is such a set  $S \subseteq X$  that  $|S| \leq K$  and that  $v(S)$  is maximal.

We are specifically interested in finding approximation algorithms for the BESTKSUBSET problem. We focus on approximating two metrics: (i) the value  $v(S)$  in the maximization variant of the problem, and (ii) the value  $(v(X) - v(S))$  in its minimization variant.

**Definition 3 (Approximation algorithms).** Let  $S^*$  denote an optimal solution for BESTKSUBSET:

1. Fix  $\alpha$ ,  $0 < \alpha < 1$ .  $\mathcal{A}$  is an  $\alpha$ -approximation algorithm for the maximization variant of BESTKSUBSET, if for each instant  $I$  of BESTKSUBSET it returns a set  $S$  such that  $v(S) \geq \alpha v(S^*)$ .
2. Fix  $\alpha$ ,  $\alpha > 1$ .  $\mathcal{A}$  is an  $\alpha$ -approximation algorithm for the minimization variant of BESTKSUBSET, if for each instant  $I$  of BESTKSUBSET it returns a set  $S$  such that  $(v(X) - v(S)) \leq \alpha(v(X) - v(S^*))$ .
3. Fix  $\alpha$ ,  $\alpha > 1$ .  $\mathcal{A}$  is an  $\alpha$ -approximation algorithm for the minimization-or-maximization variant of BESTKSUBSET, if for each instant  $I$  of BESTKSUBSET it returns a set  $S$  such that  $v(S) \geq \frac{1}{\alpha} v(S^*)$  or  $(v(X) - v(S)) \leq \alpha(v(X) - v(S^*))$ .

The definition of an approximation algorithm for minimization-or-maximization variant of BESTKSUBSET requires some additional comment: this definition guarantees that the algorithm finds a good solution provided a high quality solution exists. In other words, if there exists an optimal solution  $S^*$  such that the value  $(v(X) - v(S^*))$  is low compared to  $v(S^*)$ , then the good approximation solution for the minimization variant of the problem is also a good solution for its maximization variant. For some parameters we present good approximation algorithms for the minimization-or-maximization variant of BESTKSUBSET, even though we do not have as good algorithms neither for the minimization nor maximization variants of the problem.

We are specifically interested in FPT approximation schemes. A collection of algorithms  $\mathcal{A}$  is an FPT approximation scheme for a parameter  $P$ , if for each constant  $\alpha$  there exists an  $\alpha$ -approximation algorithm in  $\mathcal{A}$  that runs in an FPT time for the parameter  $P$ .

### 3 Algorithms for Maximizing $p$ -Separable Submodular Functions

In this section we present our approximation algorithms for the two variants of the problem, formally stated in Definition 3, of the BESTKSUBSET problem. Our methods are inspired by the algorithms of Skowron and Faliszewski [27] for the MAXCOVER problem. We extend these algorithms to be applicable to the problem of maximizing more general submodular functions.

We start with presenting an FPT approximation scheme for BESTKSUBSET for submodular  $p$ -superseparable set functions. The algorithm, formally defined as Algorithm 1, gets as an input an instance of the problem and the required approximation ratio,  $\beta$ . It proceeds in two steps: first, it restricts the universe set by selecting a certain number of elements from  $X$  with the highest values of the set function  $v$ . Second, it takes the set  $\mathcal{A}$  of elements that were selected in the first step, computes the value of the set function for all  $K$ -element subsets of  $\mathcal{A}$ , and returns a subset with the highest value.

Algorithm 1 is an FPT approximation scheme for the maximization variant of the problem for the parameter  $(K, p)$ . Before we prove this fact, however, we note that under standard complexity theoretic assumptions, there exists no FPT

---

**Algorithm 1.** An algorithm for the BESTKSUBSET problem for non-negative, non-decreasing, submodular, and  $p$ -superseparable set functions.

---

**Parameters:**

$X$  — the set of elements.

$v$  — the submodular function  $v : 2^X \rightarrow \mathbb{R}$  that is  $p$ -superseparable.

$\beta$  — the required approximation ratio of the algorithm.

$\mathcal{A} \leftarrow \lceil \frac{pK}{(1-\beta)} + K \rceil$  elements  $x$  from  $X$  with highest values  $v(\{x\})$  ;

**return**  $K$ -element subset of  $\mathcal{A}$  with the highest value of  $v$  ;

---

exact algorithm for the problem. There even exists no FPT exact algorithm for the parameter  $K$  if  $p$  is a constant. This follows from our observation in Sect. 4.1, where we show that the MAXCOVER problem with frequencies bounded by  $p$  can be expressed as maximization of a non-negative, non-decreasing, submodular,  $p$ -superseparable set function, and from the fact that the MAXCOVER problem with frequencies bounded by a constant, for the parameter  $K$  belongs to the complexity class W[1] [27], and it is unlikely that  $W[1] \subseteq \text{FPT}$ .

**Theorem 1.** *For each non-negative, non-decreasing, submodular, and  $p$ -superseparable set function  $v : 2^X \rightarrow \mathbb{R}$  and for each  $0 \leq \beta < 1$ , Algorithm 1 outputs a  $\beta$ -approximate solution for the maximization variant of BESTKSUBSET, in time  $\text{poly}(n, m) \cdot \left(\frac{pK}{(1-\beta)^K} + K\right)$ .*

*Proof.* Consider an input instance  $I$  of the BESTKSUBSET problem. Let  $S$  and  $S^*$  be, respectively, the solution returned by Algorithm 1 and some optimal solution. We set  $\text{OPT} = v(S^*)$  as the value of an optimal solution.

We will show that  $v(S) \geq \beta \text{OPT}$ . Naturally, the value  $v(S)$  might be lower than  $v(S^*)$ . This might happen because  $\mathcal{A}$ , the set of the elements considered by the algorithm in its second step, might not contain some elements from  $S^*$ . We will show that  $\ell = |S^* \setminus \mathcal{A}|$  elements from  $S^* \setminus \mathcal{A}$  might be replaced by some elements from  $\mathcal{A}$  which are not present in  $S^*$ , in a way that decreases the value of  $S^*$  by at most a small fraction. After such replacement, we will end up with the set containing the elements from  $\mathcal{A}$  only. From this we will infer that the value of the best solution in  $\mathcal{A}$  is lower than the value of an optimal solution by at most a small factor.

Let us order the elements from  $S^* \setminus \mathcal{A}$  in some arbitrary way, and let us use the notation  $S^* \setminus \mathcal{A} = \{x_1, \dots, x_\ell\}$ . We will replace the elements  $\{x_1, \dots, x_\ell\}$  with the elements  $\{x'_1, \dots, x'_\ell\}$  (we will define these elements later), one by one, in  $\ell$  consecutive steps. Thus, in the  $i$ -th step we will replace  $x_i$  with  $x'_i$  in the set  $(S^* \setminus \{x_1, \dots, x_{i-1}\}) \cup \{x'_1, \dots, x'_{i-1}\}$ . The elements  $x'_1, \dots, x'_\ell$  are defined by induction, in the following way. Assume that we have already found elements  $x'_1, \dots, x'_{i-1}$  (for  $i = 1$  it means we have not yet found any element, i.e., that we are looking for the first element in the sequence). We define  $x'_i$  to be an element from  $\mathcal{A} \setminus (S^* \cup \{x'_1, \dots, x'_{i-1}\})$  that maximizes the value  $v\left((S^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_i\}\right)$ .

It may happen that after replacing  $x_i$  with  $x'_i$ , the value of the function  $v$  for the new set decreases. Let  $\Delta_i$  denote the value of such decrease (or increase if the algorithm were lucky—in such case  $\Delta_i$  would be negative):

$$\begin{aligned} \Delta_i &= v\left((C^* \setminus \{x_1, \dots, x_{i-1}\}) \cup \{x'_1, \dots, x'_{i-1}\}\right) \\ &\quad - v\left((C^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_i\}\right). \end{aligned}$$

By the construction of the set  $\mathcal{A}$  and the fact that  $x_i \notin \mathcal{A}$ , for every  $y \in \mathcal{A} \setminus (S^* \cup \{x'_1, \dots, x'_{i-1}\})$  we have that  $v(\{x_i\}) \leq v(\{y\})$ . By the way we choose the element  $x'_i$ , we know that for every  $y \in \mathcal{A} \setminus (S^* \cup \{x'_1, \dots, x'_{i-1}\})$ , we have:



$$\begin{aligned} \Delta_i &\leq v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_{i-1}\}) \cup \{x'_1, \dots, x'_{i-1}\}\right) \\ &\quad - v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}, y\}\right). \end{aligned}$$

Using submodularity and after reformulation we get:

$$\begin{aligned} \Delta_i &\leq v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}\}\right) + v(\{x_i\}) - v(\emptyset) \\ &\quad - v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}, y\}\right) \\ &\leq v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}\}\right) + v(\{y\}) - v(\emptyset) \\ &\quad - v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}, y\}\right). \end{aligned}$$

For any  $y \in X$  (in particular for  $y \notin \mathcal{A} \setminus (S^* \cup \{x'_1, \dots, x'_{i-1}\})$ ), by submodularity and monotonicity, we have that:

$$\begin{aligned} 0 &\leq v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}\}\right) + v(\{y\}) - v(\emptyset) \\ &\quad - v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}, y\}\right). \end{aligned}$$

Since the set function is non-negative, the inequalities above will still hold if we skip the fragment  $v(\emptyset)$ . Consequently, since the set function is  $p$ -superseparable, we get:

$$\begin{aligned} (|\mathcal{A}| - K)\Delta_i &\leq \sum_{y \in X} \left( v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}\}\right) + v(\{y\}) \right. \\ &\quad \left. - v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}, y\}\right) \right) \\ &\leq p \cdot v\left((\mathcal{C}^* \setminus \{x_1, \dots, x_i\}) \cup \{x'_1, \dots, x'_{i-1}\}\right) \leq p\text{OPT}. \end{aligned}$$

Which leads to:

$$\Delta_i \leq \frac{p\text{OPT}}{|\mathcal{A}| - K} = \frac{\text{OPT}p(1 - \beta)}{pK} = \frac{\text{OPT}(1 - \beta)}{K}.$$

Since  $\ell \leq K$ , we conclude that:

$$\sum_{i=1}^{\ell} \Delta_i \leq (1 - \beta)\text{OPT}.$$

That is, after replacing the elements from  $S^*$  that do not appear in  $\mathcal{A}$  with sets from  $\mathcal{A}$ , the optimal value is decreased by at most  $(1 - \beta)\text{OPT}$ . This means that there are  $K$  elements in  $\mathcal{A}$  for which the function  $v$  achieves the value equal to at least  $\beta\text{OPT}$ . Since the algorithm tries all size- $K$  subsets of  $\mathcal{A}$ , it finds a solution with such a value. □

**Algorithm 2.** An algorithm for the minimization variant of the BESTKSUBSET problem with a non-negative, non-decreasing, submodular, and  $p$ -subseparable set function.

**Parameters:**

- $X$  — the set of elements.
- $v$  — the submodular function  $v : 2^X \rightarrow \mathbb{R}$  that is  $p$ -subseparable.
- $\beta$  — the required approximation ratio of the algorithm
- $\epsilon$  — the allowed probability of achieving worse than  $\beta$  approximation ratio

**SingleRun():**

```

    S ← ∅;
    for i ← 0 to K do
        xr ← randomly select an element from X \ S
                with probability of selecting x proportional to v(S ∪ {x}) - v(S) ;
        S ← S ∪ {xr};
    return S;

```

**Main():** run **SingleRun()** for  $\lceil -\ln \epsilon / (\frac{\beta-1}{p\beta})^K \rceil$  times; return the best solution;

Next, we consider the minimization variant of BESTKSUBSET for the case of  $p$ -subseparable submodular set functions. In Algorithm 2 we present a randomized algorithm for the problem: the algorithm performs several independent runs. Each run, in Algorithm 2 described by the **SingleRun** procedure, builds the solution by selecting random elements in  $K$  consecutive steps. In each step, an element  $x$  is selected with the probability proportional to the marginal increase of the value of the set function caused by adding  $x$  to the partial solution. Theorem 2 below shows that if we repeat **SingleRun** a sufficient number of times, we are very likely to find a solution with the required approximation ratio.

**Theorem 2.** For each non-negative, non-decreasing, submodular,  $p$ -subseparable set function  $v : 2^X \rightarrow \mathbb{R}$  and for each  $0 \leq \beta < 1$ , Algorithm 2 outputs a  $\beta$ -approximate solution for the minimization variant of BESTKSUBSET, with probability  $(1 - \epsilon)$ . The time complexity of the algorithm is  $\text{poly}(n, m) \cdot \lceil -\ln \epsilon / (\frac{\beta-1}{p\beta})^K \rceil$ .

*Proof.* Let  $I$  be an instance of the BESTKSUBSET problem with  $v : 2^X \rightarrow \mathbb{R}$  being a non-negative, submodular,  $p$ -subseparable function. Let  $\beta, \beta > 1$ , and  $\epsilon, 0 < \epsilon < 1$  be the parameters of Algorithm 2. Let  $S^*$  be some optimal solution for  $I$ .

Let us consider a single call to **SingleRun** from the “for” loop within the function **Main**. Let  $p_s$  denote the probability that such a single invocation of the function **SingleRun** returns a  $\beta$ -approximate solution. We will prove the lower-bound of  $(\frac{\beta-1}{p\beta})^K$  for the value of  $p_s$ . Let  $Ev$  denote the event that during such an invocation, at the beginning of each iteration of the “for” loop within the function **SingleRun**, it holds that:

$$v(X) - v(S) > \beta(v(X) - v(S^*)). \tag{3}$$

Note that if the complementary event, denoted  $\overline{Ev}$ , occurs, then `SingleRun` definitely returns a  $\beta$ -approximate solution. The condition in Inequality 3 can be reformulated as follows:

$$\frac{v(S^*) - v(S)}{v(X) - v(S)} > \frac{\beta - 1}{\beta}. \tag{4}$$

Now, let us consider a single iteration of the “for” loop within the function `SingleRun`. Let  $S$  be the value of the partial solution at the beginning of this iteration and let  $p_{hit}$  denote the probability that in this iteration the element from  $S^*$  is added to the partial solution (thus, using notation from Algorithm 2,  $p_{hit}$  is the probability that  $x_r \in S^*$ ). Let us assess the conditional probability  $p_{hit|Ev}$ :

$$\begin{aligned} p_{hit|Ev} &= \frac{\sum_{x \in S^*} (v(S \cup \{x\}) - v(S))}{\sum_{x \in X} (v(S \cup \{x\}) - v(S))} \\ &\geq \frac{(v(S \cup \{x_1\}) - v(S)) + (v(S \cup \{x_1, x_2\}) - v(S \cup \{x_1\})) + \dots}{\sum_{x \in X} (v(S \cup \{x\}) - v(S))} \\ &= \frac{v(S \cup S^*) - v(S)}{\sum_{x \in X} (v(S \cup \{x\}) - v(S))} \\ &\geq \frac{v(S^*) - v(S)}{\sum_{x \in X} (v(S \cup \{x\}) - v(S))} && \text{non-decreasing} \\ &\geq \frac{v(S^*) - v(S)}{p(v(X) - v(S))} && p\text{-subseparability} \\ &\geq \frac{\beta - 1}{p\beta}. \tag{Eq. 4} \end{aligned}$$

Let  $p_{opt}$  denote the probability that the function `SingleRun` returns  $S^*$ , an optimal solution. We have that:

$$p_{opt|Ev} \geq (p_{hit|Ev})^K \geq \left(\frac{\beta - 1}{p\beta}\right)^K.$$

Altogether, combining all the above findings, we know that the probability that `SingleRun` returns a  $\beta$ -approximate solution is at least:

$$p_s \geq P(\overline{Ev}) + P(Ev)p_{opt|Ev} \geq p_{opt|Ev} \geq \left(\frac{\beta - 1}{p\beta}\right)^K. \tag{5}$$

The estimation in Inequality 5 can be obtained by observing that either the event  $Ev$  or  $\overline{Ev}$  must happen. If  $\overline{Ev}$  happens, then `SingleRun` definitely returns a  $\beta$ -approximate solution; if  $Ev$  happens, then we can lower-bound the probability of finding a  $\beta$ -approximate solution by the probability of finding an optimal one.

To conclude, we use the standard argument that if we make  $x = \lceil \frac{-\ln \epsilon}{p_s} \rceil$  independent calls to `SingleRun`, then the best output from these calls is a  $\beta$ -approximate solution with probability at least equal to:

$$1 - (1 - p_s)^x \geq 1 - e^{-\ln \epsilon} = 1 - \epsilon.$$

This completes the proof. □

Interestingly, we can slightly modify the proof of Theorem 2 so that it would apply with the more general parameter  $\frac{\sum_{x \in X} v(\{x\})}{v(X)}$ . On the other hand, for this parameter we give weaker approximation guarantees, by approximating the minimization-or-maximization instead of the minimization variant of the problem.

**Theorem 3.** *For each non-negative, non-decreasing and submodular set function  $v : 2^X \rightarrow \mathbb{R}$  there exists an FPT approximation scheme for the minimization-or-maximization variant of BESTKSUBSET problem with the parameter  $(K, \frac{\sum_{x \in X} v(\{x\})}{v(X)})$ .*

*Proof.* Let us fix  $\beta, \beta > 1$ , the required approximation ratio. Let  $p = \frac{\beta}{\beta-1} \cdot \frac{\sum_{x \in X} v(\{x\})}{v(X)}$ . We will show that Algorithm 2 with such value of the parameter  $p$  (this parameter is used to determine the number of iterations of the algorithm) is a  $\beta$ -approximation algorithm for the minimization-or-maximization variant of the problem. We repeat the reasoning from the proof of Theorem 2, with the following small modification. In the proof of Theorem 2 we defined  $Ev$  to denote the event that during a single invocation of the `SingleRun` function from Algorithm 2, at the beginning of each iteration of the “for” loop, it holds that:  $v(X) - v(S) > \beta(v(X) - v(S^*))$ . In this proof we modify this definition saying that  $Ev$  denotes the event when at the beginning of each iteration of the “for” loop within the function `SingleRun`, the following two conditions hold:

$$\begin{aligned} v(X) - v(S) &> \beta(v(X) - v(S^*)), \\ v(S) &< \frac{1}{\beta}v(S^*). \end{aligned}$$

Naturally, if the complementary event occurs, then `SingleRun` definitely returns a  $\beta$ -approximate solution for the minimization-or-maximization variant of the problem. In the proof of Theorem 2, we used at-most- $p$ -subseparability in the part that assumes that the event  $Ev$  happened, to show that:

$$\sum_{x \in X} (v(S \cup \{x\}) - v(S)) \leq p(v(X) - v(S)) \tag{6}$$

Here, we show that Inequality 6 also holds if we assume that the event  $Ev$  (using our redefinition of  $Ev$ ) happened:

$$\begin{aligned} \sum_{x \in X} \left( v(S \cup \{x\}) - v(S) \right) &\leq \sum_{x \in X} \left( v(\{x\}) - v(\emptyset) \right) \leq \sum_{x \in X} v(\{x\}) \\ &= p \cdot \frac{\beta - 1}{\beta} \cdot v(X) = p \cdot v(X) - p \cdot \frac{v(X)}{\beta} \\ &\leq p \cdot v(X) - p \cdot v(S). \end{aligned}$$

With these modifications the proof of Theorem 2 can be used in this case.  $\square$

Algorithm 2 can be applied to yet another variant of the problem. Let BESTSUBSET be defined similarly to BESTKSUBSET, with the following difference. In BESTSUBSET we are not putting any constraints on the size of the solution, but we rather look for the smallest possible set  $S$  such that  $v(S) = v(X)$ . Interestingly, Algorithm 2 can be used to find *exact* solutions to BESTSUBSET for non-negative, non-decreasing, submodular,  $p$ -subseparable set functions, and it will run in FPT time for the parameter  $(K, p)$ .

**Theorem 4.** *For each non-negative, non-decreasing, submodular,  $p$ -subseparable set function  $v : 2^X \rightarrow \mathbb{R}$ , the algorithm that runs Algorithm 2 for consecutive values of the parameter  $K$  until it finds a solution  $S$ , such that  $v(S) = v(X)$ , is a randomized FPT exact algorithm for the BESTSUBSET problem for the parameter  $(K, p)$ .*

*Proof.* The proof is provided in the full version of the paper [26].

## 4 Applications of the Algorithms

In this section we show that the assumption about  $p$ -separability of submodular set functions is plausible. We provide several examples of known computational problems that can be expressed as maximization of  $p$ -separable, submodular functions.

### 4.1 The MAXWEIGHTCOVER Problem

In this subsection we show that our algorithms are applicable to MAXWEIGHTCOVER, a generalized variant of the MAXCOVER problem.

In the MAXWEIGHTCOVER problem, we are given a universe set  $N = \{e_1, e_2, \dots, e_n\}$  of  $n$  elements and a collection  $X = \{S_1, \dots, S_m\}$  of  $m$  subsets of  $N$ . Each element  $e_i$  has its weight  $w_i$ . The goal is to find a subcollection  $\mathcal{C}$  of  $X$  of size at most  $K$  that maximizes the total weight of covered elements:

$$\sum_{i: i \in S \text{ for some } S \in \mathcal{C}} w_i.$$

A *frequency* of an element  $e_i$  is the number of sets that contain  $e_i$ . Frequency of elements is a natural parameter considered in the context of approximability of covering problems [32]. To the best of our knowledge, for polynomial-time algorithms, there exists no better guarantee for the MAXCOVER problem with bounded frequencies of elements than  $(1 - 1/e)$ . This is specifically interesting, since such an approximation algorithm exists for the very similar problem SETCOVER [32].

**Lemma 1.** *The MAXWEIGHTCOVER problem with the frequency of elements upper-bounded by  $p$  can be expressed as the maximization of a nonnegative, nondecreasing submodular function which is (i)  $p$ -superseparable, and (ii)  $p$ -subseparable.*

*Proof.* For each set  $\mathcal{C} \subseteq X$  we define  $v(\mathcal{C})$  as the total weight of elements covered by the sets from  $\mathcal{C}$ . Such defined  $v$  is nonnegative and submodular.

We observe that the weighted sum of  $p$ -superseparable set functions is also  $p$ -superseparable, and that the same argument applies to  $p$ -subseparability. Thus, it is sufficient to consider a function  $u_i$  which returns 1 for collections of sets that cover  $e_i$ , and 0 for the remaining ones. Observe that if the frequency of the elements is bounded by  $p$ , then  $\sum_{S \in X} u_i(\{S\})$ , the number of sets that cover  $e_i$ , is also bounded by  $p$ .

Let us fix a collection of sets  $\mathcal{C} \subseteq X$  and let us consider two cases. If  $e_i$  is covered by  $\mathcal{C}$ , then  $\sum_{S \in X} (u_i(\mathcal{C} \cup \{S\}) - u_i(\mathcal{C}))$  is equal to 0. But, in such case  $pu_i(\mathcal{C}) = p$  and the condition for  $p$ -superseparability holds. Naturally,  $u_i(X) = 1$ , thus the conditions for  $p$ -subseparability also holds.

If  $e_i$  is not covered by  $\mathcal{C}$ , then  $\sum_{S \in X} (u_i(\mathcal{C} \cup \{S\}) - u_i(\mathcal{C}))$  is equal to the number of sets that cover  $e_i$ , thus to  $\sum_{S \in X} u_i(\{S\})$ . This means that the condition for  $p$ -superseparability holds. If the frequency of the elements is upper-bounded by  $p$ , then  $\sum_{S \in X} (u_i(\mathcal{C} \cup \{S\}) - u_i(\mathcal{C}))$  is upper bounded by  $p$ , and since  $u_i(\mathcal{C}) = 0$ , the condition for  $p$ -subseparability holds. This proves the thesis. □

**Corollary 1.** *There exists an FPT approximation scheme for the maximization and minimization variant of the MAXWEIGHTCOVER for the parameter  $(K, p)$ , where  $p$  is the upper-bound on the frequency of the elements.*

Corollary 1 extends the recent results for the MAXCOVER problem [27]. Interestingly, Theorem 3 says that there exists a randomized FPT approximation scheme for the minimization-or-maximization variant of the MAXWEIGHTCOVER problem, for the parameter  $(K, p_{av})$ , where  $p_{av}$  is an average frequency of an element.

## 4.2 Other Applications

Due to space restrictions in this section we describe the other two applications of our results very briefly. For the thoughtful analysis of these two cases we refer the reader to the full version of the paper [26].

**Matching and assignment problems.** In the WEIGHTED-B-K-MATCHING problem we are given a set of vertices  $X \cup Y$ , a set of edges  $E$  (there are no edges neither between the vertices from  $X$  nor between the vertices from  $Y$ ), a weight function  $w : E \rightarrow \mathbb{R}$ , and a capacity function  $c : X \rightarrow \mathbb{Z}$ . The goal is to find a subset of edges with the maximal total weight, such that each vertex  $x \in X$  belongs to at most  $c(x)$  of the selected edges, each vertex  $y \in Y$

belongs to at most one of the selected edges, and altogether there are at most  $K$  vertices from  $X$  which belong to some of the selected edges.

Our results can be used to prove that there exists an FPT approximation scheme for the maximization variant of the WEIGHTED-B- $K$ -MATCHING for the parameter  $(K, p)$ , where  $p$  is a bound on the degree of vertices from  $Y$ .

**Item selection in multi-agent systems.** Our results can be also applied to the remarkably general model describing the problem of selecting a set of collective items for agents [28]. Let  $N = \{1, 2, \dots, n\}$  be the set of agents and let  $C = \{a_1, a_2, \dots, a_m\}$  be the set of *items*. Each agent  $i \in N$  is endowed with a *utility function*  $u_i : C \rightarrow \mathbb{R}$  that measures how much  $i$  desires each of the items. Our goal is to select  $K$  items, called *winners*, that in some sense would make the agents most satisfied. An OWA vector  $\alpha$  is a vector of  $K$  elements,  $\alpha = \langle \alpha_1, \dots, \alpha_K \rangle$ . Given an OWA vector  $\alpha$ , for each agent  $i$  and for each set of  $K$  items  $S$ , we define  $u_i(S)$ , the satisfaction of  $i$  from  $S$ , in the following way. Let  $u_1, u_2, \dots, u_K$  be the utilities from  $\{u_i(x) : x \in S\}$ , sorted in the descending order; then  $u_i(S) = \sum_{j=1}^K \alpha_j u_j$ . The satisfaction of all agents from  $S$  is defined as the sum of satisfactions of all the individuals from  $S$ .

This model captures various natural problems, from winner determination in multiwinner election systems, through recommendation systems, to location problems. For instance the problem of selecting  $K$  items under the OWA vector  $\alpha = \langle 1, 0, \dots, 0 \rangle$  boils down to the problem of winner determination under Chamberlin and Courant rule [3], or to the facility location problem. The problem for  $\alpha = \langle 1, 1/2, \dots, 1/K \rangle$  is equivalent to winner determination in the Proportional Approval Voting (PAV) system. For more examples of applications of this general model we refer the reader to the original work of Skowron et al. [28].

We say that the agents have  $k$ -approval utilities if each agent assigns utility equal to 1 to exactly  $k$  items, and utility equal to 0 to the remaining ones. Such  $k$ -approval utilities are very popular in the context of social choice, in particular in case of multi-winner election rules.

Our results can be used to prove that there exists an FPT approximation scheme for the maximization and minimization variants of the problem of selecting  $K$  items with  $k$ -approval utilities for the parameter  $(K, k)$ .

## 5 Conclusions

We have considered FPT approximation schemes for the problem of maximizing submodular set functions. There are many natural ways in which this research can be extended. We believe that one of the promising approaches is to consider the problem with additional constraints, such as knapsack constraints or matroid constraints.

**Acknowledgements.** The author thanks Piotr Faliszewski for his helpful comments. This research has been supported by Europe Research Grant ERC-StG 639945.

## References

1. Bach, F.: Structured sparsity-inducing norms through submodular functions. In: Proceedings of Advances in Neural Information Processing Systems 23 (NIPS 2010), pp. 118–126 (2010)
2. Calinescu, G., Chekuri, C., Pál, M., Vondrák, J.: Maximizing a submodular set function subject to a matroid constraint (extended abstract). In: Fischetti, M., Williamson, D.P. (eds.) IPCO 2007. LNCS, vol. 4513, pp. 182–196. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-72792-7\\_15](https://doi.org/10.1007/978-3-540-72792-7_15)
3. Chamberlin, B., Courant, P.: Representative deliberations and representative decisions: proportional representation and the borda rule. *Am. Polit. Sci. Rev.* **77**(3), 718–733 (1983)
4. Chvatal, V.: A greedy heuristic for the set-covering problem. *Math. Oper. Res.* **4**(3), 233–235 (1979)
5. Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: *Parameterized Algorithms*. Springer, Heidelberg (2015)
6. Das, A., Kempe, D.: Submodular meets spectral: greedy algorithms for subset selection, sparse approximation and dictionary selection. In: Getoor, L., Scheffer, T. (eds.) Proceedings of the 28th International Conference on Machine Learning (ICML 2011), New York, NY, USA, pp. 1057–1064 (2011)
7. Downey, R., Fellows, M.: *Parameterized Complexity*. Springer, New York (1999)
8. Feige, U.: A threshold of  $\ln n$  for approximating set cover. *J. ACM* **45**(4), 634–652 (1998)
9. Feige, U., Mirrokni, V.S., Vondrák, J.: Maximizing non-monotone submodular functions. *SIAM J. Comput.* **40**(4), 1133–1153 (2011)
10. Flum, J., Grohe, M.: *Parameterized Complexity Theory*. Springer, Heidelberg (2006)
11. Iwata, S., Fleischer, L., Fujishige, S.: A combinatorial strongly polynomial algorithm for minimizing submodular functions. *J. ACM* **48**(4), 761–777 (2001)
12. Jegelka, S., Bilmes, J.A.: Submodularity beyond submodular energies: coupling edges in graph cuts. In: Proceedings of the 24th IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011), pp. 1897–1904 (2011)
13. Kim, G., Xing, E.P., Fei-Fei, L., Kanade, T.: Distributed cosegmentation via submodular optimization on anisotropic diffusion. In: Proceedings of IEEE International Conference on Computer Vision (ICCV 2011), pp. 169–176 (2011)
14. Krause, A., Golovin, D.: Submodular function maximization. Technical report (2012)
15. Krause, A., Guestrin, C.: Near-optimal value of information in graphical models. In: Proceedings of Conference on Uncertainty in Artificial Intelligence (UAI 2005), July 2005
16. Krause, A., Singh, A., Guestrin, C.: Near-optimal sensor placements in Gaussian processes: theory, efficient algorithms and empirical studies. *J. Mach. Learn. Res.* **9**, 235–284 (2008)
17. Kumar, R., Moseley, B., Vassilvitskii, S., Vattani, A.: Fast greedy algorithms in MapReduce and streaming. In: Proceedings of the Twenty-Fifth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2013), pp. 1–10 (2013)
18. Lee, J., Mirrokni, V.S., Nagarajan, V., Sviridenko, M.: Non-monotone submodular maximization under matroid and knapsack constraints. In: Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing (STOC 2009), pp. 323–332 (2009)



19. Lin, H., Bilmes, J.: Multi-document summarization via budgeted maximization of submodular functions. In: Proceedings of Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics (HTL 2010), pp. 912–920. Association for Computational Linguistics, Stroudsburg (2010)
20. Lin, H., Bilmes, J.A.: How to select a good training-data subset for transcription: submodular active selection for sequences. In: Proceedings of 10th Annual Conference of the International Speech Communication Association (INTERSPEECH 2009), pp. 2859–2862 (2009)
21. Lu, T., Boutilier, C.: Budgeted social choice: from consensus to personalized decision making. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011), pp. 280–286 (2011)
22. Narasimhan, M., Nebojsa, J., Jeff, B.A.: Q-clustering. In: Weiss, Y., Schölkopf, B., Platt, J. (eds.) Proceedings of Advances in Neural Information Processing Systems 18 (NIPS 2010), pp. 979–986. MIT Press (2006)
23. Nemhauser, G., Wolsey, L., Fisher, M.: An analysis of approximations for maximizing submodular set functions. *Math. Program.* **14**(1), 265–294 (1978)
24. Niedermeier, R.: Invitation to Fixed-Parameter Algorithms. Oxford University Press, Oxford (2006)
25. Jegelka, S., Bach, F., Sra, S.: Reflection methods for user-friendly submodular optimization. In: Burges, C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K. (eds.) Proceedings of Advances in Neural Information Processing Systems 26 (NIPS 2013), pp. 1313–1321. Curran Associates Inc. (2013)
26. Skowron, P.: FPT approximation schemes for maximizing submodular functions. Technical report [arXiv:1510.00215](https://arxiv.org/abs/1510.00215) [cs.DS], January 2015
27. Skowron, P., Faliszewski, P.: Fully proportional representation with approval ballots: approximating the maxcover problem with bounded frequencies in FPT time. In: Proceedings of the 29th Conference on Artificial Intelligence (AAAI 2015), pp. 2124–2130 (2015)
28. Skowron, P., Faliszewski, P., Lang, J.: Finding a collective set of items: from proportional multirepresentation to group recommendation. *Artif. Intell.* **241**, 191–216 (2016)
29. Skowron, P., Faliszewski, P., Slinko, A.: Achieving fully proportional representation: approximability result. *Artif. Intell.* **222**, 67–103 (2015)
30. Streeter, M.J., Golovin, D.: An online algorithm for maximizing submodular functions. In: Proceedings of Advances in Neural Information Processing Systems 21 (NIPS 2008), pp. 1577–1584 (2008)
31. Sviridenko, M.: A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.* **32**(1), 41–43 (2004)
32. Vazirani, V.V.: Approximation Algorithms. Springer, New York (2001)
33. Vondrák, J., Chekuri, C., Zenklussen, R.: Submodular function maximization via the multilinear relaxation and contention resolution schemes. In: Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing (STOC 2011), pp. 783–792 (2011)
34. Yue, Y., Guestrin, C.: Linear submodular bandits and their application to diversified retrieval. In: Proceedings of Advances in Neural Information Processing Systems 24 (NIPS 2011), pp. 2483–2491 (2011)

# Bounds for the Convergence Time of Local Search in Scheduling Problems

Tobias Brunsch, Michael Etscheid<sup>(✉)</sup>, and Heiko Röglin

Department of Computer Science, University of Bonn, Bonn, Germany  
{brunsch,etscheid,roeglin}@cs.uni-bonn.de

**Abstract.** We study the convergence time of local search for a standard machine scheduling problem in which jobs are assigned to identical or related machines. Local search corresponds to the best response dynamics that arises when jobs selfishly try to minimize their costs. We assume that each machine runs a coordination mechanism that determines the order of execution of jobs assigned to it. We obtain various new polynomial and pseudo-polynomial bounds for the well-studied coordination mechanisms Makespan and Shortest-Job-First, using worst-case and smoothed analysis. We also introduce a natural coordination mechanism FIFO, which takes into account the order in which jobs arrive at a machine, and study both its impact on the convergence time and its price of anarchy.

## 1 Introduction

We analyze the following scheduling problem: Given  $m$  machines and  $n$  jobs, find an assignment of the jobs to the machines minimizing the maximum costs of a job, which are defined according to a coordination mechanism. The jobs may have different job sizes and the machines may have different machine speeds. A typical definition of the costs of a job is the sum of the job sizes assigned to the same machine divided by the machine speed, which is a natural choice when the makespan is to be minimized. In other contexts it might be more realistic to assume an order in which the jobs on a machine are executed and that a job only pays for the execution time of itself and all previous jobs.

Even in the case of identical machine speeds, the problem is known to be strongly NP-hard [10] and local search is a popular tool to approximate good solutions. Here, a job unilaterally changes its assignment and moves to another machine if it can reduce its costs this way. Throughout this paper, we assume a best response policy, i.e., a moving job selects a machine that minimizes its costs. If there is no job left that can improve its costs, we have attained a local optimum, which is guaranteed to be reached after a finite number of steps. Although the quality of the worst local optimum has been thoroughly analyzed [3, 5, 6, 9, 16], there is not much work about the convergence time needed to find one via local search.

---

This research was supported by ERC Starting Grant 306465 (BeyondWorstCase).

### 1.1 Terminology

Let us first describe the studied problem in detail. Consider an instance with  $m$  machines and  $n$  jobs. Each machine  $i$  has a *speed*  $s_i \in \mathbb{Q}_{>0}$  and each job  $j$  has a *job size*  $p_j \in \mathbb{Q}_{>0}$ . Let  $s_{\min}, s_{\max}, p_{\min}$ , and  $p_{\max}$  be the minimal and maximal speeds and job sizes. Let  $W = \sum_{j=1}^n p_j$  be the sum of the job sizes. For *identical machines*,  $s_{\max} = s_{\min} = 1$ , and for *unit-weight jobs*,  $p_{\max} = p_{\min} = 1$ .

For an assignment  $\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, m\}$  that maps the jobs to the machines, let  $L_i = \sum_{j \in \sigma^{-1}(i)} p_j / s_i$  be the *load* of machine  $i$ . The maximum load is called *makespan*. The *costs* of a job  $j$  are defined according to a *coordination mechanism*, which assigns costs to every job depending only on the set of jobs that have chosen the same machine, but not on the residual schedule.

1. In the *Makespan* model, all jobs assigned to the same machine are executed simultaneously such that the costs  $c_j^\sigma = L_{\sigma(j)}$  of a job  $j$  correspond to the load of its machine. This is the most common coordination mechanism and it corresponds to linear weighted congestion games on parallel links.
2. In the *FIFO* model, the jobs on each machine are executed one after another. Therefore, we need a permutation  $\pi$  on the jobs that determines the order in which the jobs on a machine get processed. The costs of a job  $j$  are then  $c_j^{(\sigma, \pi)} = \sum_{j' \in J_\sigma^\pi(j)} \frac{p_{j'}}{s_{\sigma(j)}}$ , where  $J_\sigma^\pi(j)$  is the set of jobs  $j'$  on the same machine with  $\pi(j') \leq \pi(j)$ . If a job  $j$  jumps to another machine, it is inserted as the last job, i.e.,  $\pi(j) = n$ .
3. In the *SJF* (shortest job first) model, the jobs are executed one after another, but the permutation of the jobs is at any time implicitly given by their job sizes where the smallest job on a machine is executed first. Ties for jobs of equal size are broken arbitrarily. This means that the costs of a job are defined as  $c_j^\sigma = \sum_{j': \sigma(j')=\sigma(j) \wedge \pi(j') \leq \pi(j)} \frac{p_{j'}}{s_{\sigma(j)}}$ , where  $\pi$  is a permutation of the jobs assigned to machine  $\sigma(j)$  such that  $\pi(j') < \pi(j)$  if  $p_{j'} < p_j$  and  $\pi(j') > \pi(j)$  if  $p_{j'} > p_j$ .

The FIFO model is not a coordination mechanism in the classical sense as the order in which the jobs are executed depends on previous iterations. Nevertheless, we believe that this model can easily be motivated by many real-world applications where the first-come, first-served principle is ubiquitous.

In the case of the Makespan and SJF models, we call  $\sigma$  a *schedule*. In the FIFO model, we call the tuple  $(\sigma, \pi)$  a *schedule*. Often, we omit the parameters  $\sigma$  and  $\pi$  if they are clear from the context, or we replace them by an iteration number  $t$ . Then we mean the schedule before the move of iteration  $t$  gets executed.

We say that a job is *unsatisfied* if it could improve its costs by jumping to a different machine. When an unsatisfied job jumps, it always jumps to a machine minimizing its costs, i.e., we consider best response dynamics. If there is no unsatisfied job, we call the current schedule a *local optimum*. The *convergence time* for an instance is the maximum number of jumps it can take starting from an arbitrary schedule until a local optimum is reached. The *price of anarchy* is the ratio of the makespans of the worst local optimum and the global optimum.

If there are several unsatisfied jobs, we choose the next job to jump according to a pivot rule:

- *Best Improvement*: Select a job for which the largest improvement of its costs is possible.
- *Random*: Select a job uniformly at random from the set of unsatisfied jobs.
- *Min Weight*: Select a smallest unsatisfied job.
- *Max Weight*: Select a largest unsatisfied job.
- *Fixed Priority*: Select the unsatisfied job with the largest priority according to a given order on the jobs. This pivot rule includes Min Weight and Max Weight as special cases.

## 1.2 Smoothed Analysis

Despite its bad running time, which can be exponential, and the large price of anarchy in theory, local search is a popular tool in practice as it typically delivers good local optima very quickly. At first glance, this seems like a contradiction, but the instances in the theoretical proofs are rather contrived and rarely observed in practice. To have a more realistic understanding of local search in theory, we use the framework of smoothed analysis introduced by Spielman and Teng [17] to explain the practical success of the simplex method. This model can be considered as a less pessimistic variant of worst-case analysis in which the adversarial input is subject to a small amount of random noise and it is by now a well-established alternative to worst-case analysis. This random noise can be motivated, for example, by measurement errors, numerical imprecision, and rounding errors, which often occur in practice. It can also model influences that cannot be quantified exactly but for which there is no reason to believe that they are adversarial.

We follow the more general model of smoothed analysis introduced by Beier and Vöcking [1]. In this model, the adversary is even allowed to specify the probability distribution of the random noise. The influence he can exert is described by a parameter  $\phi \geq 1$  denoting the maximum density of the noise. The model is formally defined as follows.

**Definition 1.** *In a  $\phi$ -smooth instance  $\mathcal{I}$ , the adversary chooses the following input data:*

- the number  $m$  of machines;
- arbitrary machine speeds  $s_1, \dots, s_m$  in the case of non-identical machines;
- the number  $n$  of jobs;
- for each  $p_j$ , a probability density  $f_j : [0, 1] \rightarrow [0, \phi]$  according to which  $p_j$  is chosen independently of the processing requirements of the other jobs.

*The smoothed convergence time is the worst expected convergence time of any  $\phi$ -smooth instance and the smoothed price of anarchy is the worst expected price of anarchy of any  $\phi$ -smooth instance.*

Note that the only perturbed part of the instance are the processing requirements. Formally, a  $\phi$ -smooth instance is not a single instance but a distribution over instances. The parameter  $\phi$  determines how powerful the adversary is. He can, for example, define an interval of length  $1/\phi$  for each job size from which it is drawn uniformly at random. Hence, for  $\phi = 1$  the model corresponds to an average-case analysis and for  $\phi \rightarrow \infty$  the adversary becomes as powerful as in a worst-case analysis.

## 2 Related Work and Results

Since its invention, smoothed analysis has been successfully applied in a variety of contexts. Two surveys [14, 18] summarize some of these results.

The notion of coordination mechanisms has been introduced by Christodoulou et al. [4] in the context of congestion games. There has been extensive research about the price of anarchy for the different coordination mechanisms. In the Makespan model it is constant for identical machines [9, 16] and  $\Theta\left(\min\left\{\frac{\log m}{\log \log m}, \log \frac{s_{\max}}{s_{\min}}\right\}\right)$  for related machines [5]. The smoothed price of anarchy for related machines is  $\Theta(\log \phi)$  regardless of whether the job sizes [3] or the machine speeds [6] are perturbed.

Immorlica et al. [13] showed a price of anarchy of  $2 - 1/m$  for identical and  $\Theta(\log m)$  for related machines for the SJF model, which is the same as for list schedules, i.e., schedules that are generated by a greedy assignment.

The FIFO model has been introduced implicitly by Brunsch et al. [3] through the equivalent concept of near list schedules which was used as a generalization of local optima w.r.t. the Makespan model and list schedules. They showed that the smoothed price of anarchy is  $\Theta(\log \phi)$ . We complement this by the corresponding worst-case results for identical and related machines to obtain the same tight bounds as in the SJF model.

There is less known about the convergence times in the different models. As we are up to our knowledge the first ones who consider the FIFO model, there are no previous results about convergence times. We show tight results for special cases like identical machines and several upper bounds depending on  $W/p_{\min}$  for different pivot rules in the general case. Although we conjecture polynomial bounds for all cases, we give the first non-trivial proofs for this natural problem. Immorlica et al. [13] showed for the SJF model that if the jobs are asked on a rotational basis if they want to jump, the convergence time is in  $O(n^2)$ . This is in sharp contrast to our result that for the Min Weight pivot rule it can take an exponential number of iterations even in the case of two identical machines.

Brucker et al. [2] considered the Makespan model with the difference that only jobs from a machine with maximum load—a so-called *critical* machine—are allowed to jump, i.e., a local optimum is reached as soon as every job on a critical machine is satisfied. They gave an algorithm that finds a local optimum after  $O(n^2)$  improving steps for identical machines. From this, one can easily derive an algorithm for identical machines in the Makespan model: Run Brucker's algorithm exhaustively until every job on a critical machine is satisfied. As on

identical machines the minimum load of a machine is monotonically increasing, these jobs cannot become unsatisfied again by any sequence of improving steps. Hence, the jobs on the critical machine are fixed and therefore we can remove the critical machine together with its assigned jobs from the instance. Repeating this argument yields a running time of  $O(n^2m)$  improving steps. As the monotonicity argument does not hold anymore in the case of related machines, we are not aware of a way to use similar results by Schuurman and Vredeveld [16] and Hurkens and Vredeveld [12] for Brucker’s model on related machines.

For the Makespan model and identical machines, Goldberg [11] considers randomized local search, where in each step a job and a machine are selected uniformly at random, and the job moves to that machine if it is an improving step. He shows that random local search converges in expected  $(m+n+\frac{p_{\max}}{p_{\min}})^{O(1)}$  time.

In the Makespan model, Feldmann et al. [8] provided an  $O(nm^2)$ -time Nashification algorithm, which, given an arbitrary schedule, computes a local optimum without increasing the social cost, i.e., the makespan in our case. They further showed that the convergence time on identical machines is bounded by  $\Omega(2^{\sqrt{n}})$  and  $O(2^n)$ . To be more precise, Even-Dar et al. [7] showed (again for identical machines) that the Max Weight and the Random pivot rule converge in  $n$  and  $O(n^2)$  steps, respectively, while the Min Weight pivot rule can take an exponential number of steps. We extend this result by showing that every pivot rule converges in  $O(n \cdot W/p_{\min})$  steps, which can be seen as a generalization of their result that every pivot rule converges in  $O(W + n)$  steps in the case of integer weights. For related machines and unit-weight jobs, Even-Dar et al. [7] showed that there is a pivot rule that converges in  $mn$  steps. We improve this by showing that the convergence time for any pivot rule with best response policy is exactly  $n$ . For the case of related machines and integral job sizes and machine speeds, they showed that any pivot rule converges in  $O(W^2 \cdot s_{\max}^2/s_{\min})$  steps. We prove a similar bound for the Best Improvement pivot rule on arbitrary weights. An overview of our results on convergence times is given in Tables 1, 2, and 3.

**Table 1.** FIFO convergence times

|                    |  |             |
|--------------------|--|-------------|
| Identical machines | $n - 1$                                  | (Theorem 1) |
| Unit-weight jobs   | $n$                                      | (Theorem 4) |
| Two machines       | $\Theta(n)$                              | (Theorem 5) |
| Best improvement   | $O(m^2n \cdot W/p_{\min})$               | (Theorem 6) |
| Random             | $O(m^2n^2 \cdot W/p_{\min})$             | (Theorem 7) |
| Fixed priority     | $O(n^2 \cdot W/p_{\min})$                | (Theorem 8) |
| lower bounds       | $\Omega(mn), \Omega(m^2)$ for Min Weight | (Theorem 9) |

### 2.1 Paper Organization

The remainder of this paper is organized as follows. In Sect. 3 we show how to convert superpolynomial deterministic convergence times to smoothed

**Table 2.** Makespan convergence times

|                    |                                 |              |
|--------------------|---------------------------------|--------------|
| Identical machines | $O(n \cdot W/p_{\min})$         | (Theorem 2)  |
| Unit-weight jobs   | $n$                             | (Theorem 4)  |
| Best improvement   | $O(m^2 n \cdot W^2/p_{\min}^2)$ | (Theorem 10) |

**Table 3.** SJF convergence times

|  |                        |              |
|--|------------------------|--------------|
| Max weight on two identical machines                     | $2^{\Omega(n)}$        | (Theorem 3)  |
| Max weight on two identical machines with random weights | $2^{\Omega(\sqrt{n})}$ | (Theorem 3)  |
| Min weight   | $n$                    | (Theorem 11) |
| Random   | $O(n^2)$               | (Theorem 11) |

polynomial convergence times. In Sects. 4 and 5 we deal with the special cases of identical machines and unit-weight jobs, respectively, before we turn to the more general case of related machines in Sect. 6. We conclude with the analysis of the price of anarchy in the FIFO model in Sect. 7 and some remarks in Sect. 8. Some of the proofs are deferred to a full version of this paper.

### 3 Smoothed Analysis

Some of our shown convergence times include the factor  $W/p_{\min}$ . While in the worst case this fraction can be exponentially large, in the smoothed setting they turn into expected polynomial convergence times.

**Lemma 1.** *If the convergence time is bounded by  $f(m, n) \cdot W/p_{\min}$  for some polynomial  $f$ , then the smoothed convergence time is bounded by  $O(f(m, n) \cdot n^3 \log(m) \cdot \phi)$ .*

Unfortunately, our result about the convergence time of the Best Improvement pivot rule in the Makespan model depends quadratically on  $p_{\min}$ . This does not allow us to derive an expected polynomial convergence time, but instead we can show that with high probability the convergence time is polynomially bounded.

**Lemma 2.** *The smoothed convergence time of the Best Improvement pivot rule in the Makespan model is in  $m^2 n^7 \phi^2$  with probability at least  $1 - 1/n$ .*

### 4 Identical Machines

In the FIFO model, the costs of a job decrease monotonically while the minimum load of a machine increases monotonically when considering identical machines. As a moving job always jumps to a machine with minimum load, every job can jump at most once. This leads to the following result.

**Theorem 1.** *In the FIFO model, for any pivot rule the worst-case running time is exactly  $n - 1$ .*

For the Makespan model, Even-Dar et al. [7] proved that the Min Weight pivot rule can take as many as  $\Omega((n/m^2)^{m-1})$  steps. They also showed that this is near to the worst case as every pivot rule terminates after  $O((\frac{n}{K} + 1)^K)$  steps, where  $K$  is the number of different job weights. We derive the bound  $O(n \cdot \frac{W}{p_{\min}})$  for arbitrary pivot rules, which is a significant improvement if  $\frac{W}{p_{\min}}$  is small. This bound is almost optimal as it is easy to see that the worst case instance used in [7] has  $\frac{p_{\max}}{p_{\min}} = (n/(m - 1))^{m-2}$ . It is also a generalization of the result that every pivot rule converges in  $O(W + n)$  steps in the case of integer weights.

**Theorem 2.** *In the Makespan model, every pivot rule terminates after  $O(n \cdot \frac{W}{p_{\min}})$  steps.*

*Proof.* As Even-Dar et al. [7] pointed out (without proof), after a job  $j$  moved to machine  $i$ , it can only be unsatisfied again after a strictly greater job moved to machine  $i$  in the meantime: A job always jumps to a machine with minimum load and the minimum load increases monotonically. Consider the last job  $j'$  entering machine  $i$  in iteration  $t'$  before job  $j$  jumps away from machine  $i$  in iteration  $t$ . Then machine  $i$  must be a machine with minimum load before iteration  $t'$ . Now if  $p_j > p_{j'}$ , then  $L_i^{t+1}$  would be strictly smaller than  $L_i^{t'}$ , which is a minimum load in a former iteration. If  $p_j = p_{j'}$ , then job  $j$  cannot be unsatisfied because job  $j'$  is not unsatisfied.

Based on their idea of push-out potentials, we define the potential  $\phi := \sum_{i=1}^m u_i^t \leq W$ , where  $u_i^t$  is the maximum total weight of jobs on machine  $i$  that could consecutively move away from  $i$ , starting in the schedule before iteration  $t$ . When a job  $j$  jumps from machine  $i$  to machine  $i'$ , then  $u_{i'}^t$  was 0 beforehand. As mentioned above, no job from any other machine than machine  $i'$  can become unsatisfied by the move of job  $j$  and thus the potential  $\phi$  decreases by at least  $u_i^t - u_i^{t+1} - u_{i'}^{t+1}$ .

If  $u_{i'}^{t+1}$  was larger than  $p_j$ , then there would be a sequence of moves from jobs away from machine  $i'$  such that the load of machine  $i'$  after these moves would be less than  $L_{i'}^t$ . But machine  $i'$  was a machine with minimum load before iteration  $t$ , a contradiction. Note also that  $u_i^t - u_i^{t+1} \geq p_j$ : Let  $J'$  be the jobs on machine  $i$  with total weight  $u_i^{t+1}$  which could consecutively jump away from machine  $i$  after iteration  $t$ . Then  $J' \cup \{j\}$  could consecutively jump away before iteration  $t$  and thus  $u_i^t \geq \sum_{j' \in J' \cup \{j\}} p_{j'} = u_i^{t+1} + p_{j'}$ . We can conclude that  $u_i^t - u_i^{t+1} - u_{i'}^{t+1} \geq 0$  and thus that  $\phi$  is actually a potential.

We call a jump of job  $j$  to machine  $i$  in iteration  $t$  *stable* if after that jump, another job moves to  $i$  before a job leaves  $i$ . As discussed above, through the stable jump the total potential of all machines except machine  $i$  decreases by at least  $p_j \geq p_{\min}$  and  $u_i^t = 0$  at time  $t'$  when the next job enters or leaves machine  $i$ . Hence, every stable jump induces a potential drop of at least  $p_{\min}$ . We maintain a set of indices: In the initial schedule, every job has an index attached to it. When a job  $j$  moves away from machine  $i$ , then the indices



attached to  $j$  get transferred to the job  $j'$  that moved last to machine  $i$ . If no such job exists, the indices get deleted. Afterwards, a new index gets attached to job  $j$  on its new machine if it was a stable jump.

When a job  $j$  moves to machine  $i$ , then no job on machine  $i$  was unsatisfied beforehand as  $j$  jumps to a machine with minimum load. Thus, when an index gets reattached from job  $j'$  to job  $j$ , then  $j$  made  $j'$  unsatisfied and thus  $p_j$  is strictly greater than  $p_{j'}$  because only larger jobs can make smaller jobs unsatisfied. Therefore, every index can be reattached at most  $n$  times. Furthermore, every time a job  $j$  jumps away from a machine  $i$ , it has at least one index attached to it: Assume to the contrary that it is the first jump without attached indices. If it is the first jump by job  $j$  or its last jump was stable, then there is by definition an attached index. Otherwise, there is a job  $j'$  that left machine  $i$  such that job  $j$  is the last job entering machine  $i$  beforehand and thus job  $j'$  transferred its indices to job  $j$ . Hence, the number of indices is at least one  $n$ th of the total number of jumps. There can only be  $W/p_{\min}$  many stable jumps as otherwise  $\phi$  would be negative. This yields the desired bound.  $\square$

Finally let us consider the SJF model. The Max Weight pivot rule in the SJF model can take an exponential number of steps even on two identical machines. Also an average-case analysis yields a superpolynomial convergence time. We do not consider the Random pivot rule and the Min Weight pivot rule in this section because for these rules we prove in Sect. 6.3 polynomial upper bounds even for the more general setting of related machines. We leave it as an open question whether the convergence time of the Best Improvement pivot rule is polynomial for identical machines.

**Theorem 3.** *In the SJF model, the convergence time of the Max Weight pivot rule is  $2^{\Omega(n)}$  even for two identical machines. The smoothed convergence time of the Max Weight pivot rule is  $2^{\Omega(\sqrt{n})}$  even for two identical machines and  $\phi = 1$ .*

## 5 Unit-Weight Jobs

In the case of unit-weight jobs, Even-Dar et al. [7] claimed that for Makespan, there exists a pivot rule which converges in  $mn$  steps and that there is a pivot rule with convergence time  $\Omega(mn)$  if jobs do not necessarily move to the machine yielding the biggest improvement but only have to improve their costs by jumping. We show that all pivot rules have linear convergence time if jobs have to jump to the best machine.

**Theorem 4.** *In both the FIFO and the Makespan model for unit-weight jobs, the convergence time for any pivot rule is  $n$  for any number  $m \geq 2$  of machines.*

## 6 Related Machines

For the most general case of related machines we use potential functions in order to show pseudo-polynomial convergence times for different pivot rules in both the FIFO and the Makespan model.

The potential  $\phi_{\text{FIFO}}$  used in the FIFO model is the Rosenthal potential introduced in [15], which is the sum of the execution times of the jobs. It is easy to see that  $\phi_{\text{FIFO}}$  decreases by at least  $\Delta$  when the jumping job improves its execution time by  $\Delta$ . It decreases even more if the jumping job was not on top of its original machine.

For the Makespan model, we use the potential

$$\phi_{\text{Makespan}} := \sum_{i=1}^m \frac{1}{s_i} \cdot \left( \left( \sum_{j \in \sigma^{-1}(i)} p_j \right)^2 + \sum_{j \in \sigma^{-1}(i)} p_j^2 \right),$$

defined by Even-Dar et al. [7].

The fastest machine has always load at most  $W/s_{\text{max}}$ . If there is a machine with load greater than  $2W/s_{\text{max}}$ , then a job from this machine can improve its costs by at least  $W/s_{\text{max}}$  by jumping to the fastest machine. This gives rise to the following lemma.

**Lemma 3.** *The following two statements hold:*

1. *If there is a machine with load greater than  $2W/s_{\text{max}}$ , the best improvement can be achieved by a jump from some job from a machine with load greater than  $W/s_{\text{max}}$  to a machine with load at most  $W/s_{\text{max}}$ .*
2. *If there is no machine with load greater than  $2W/s_{\text{max}}$ , then  $\phi_{\text{FIFO}} = O(n \cdot \frac{W}{s_{\text{max}}})$  and  $\phi_{\text{Makespan}} = O(\frac{W^2}{s_{\text{max}}})$ .*

**Corollary 1.** *For the Best Improvement pivot rule after  $n$  iterations and for the Random pivot rule after expected  $O(n \log n)$  iterations there is no machine left with load greater than  $2W/s_{\text{max}}$ .*

### 6.1 FIFO

Before we come to the general cases, let us first mention a linear-time result for the special case of  $m = 2$  machines.

**Theorem 5.** *In the FIFO model, the convergence time for any pivot rule on two related machines is at least  $n$  and at most  $2n - 2$ . There are pivot rules for which  $2n - 2$  is tight.*

The main idea of the following proofs is that if a job jumps that is not on top of its machine, the costs of all jobs above the moving job and thus the potential  $\phi_{\text{FIFO}}$  decrease by at least  $p_{\text{min}}/s_{\text{max}}$ . We are able to show that this must happen after a polynomial number of steps for the Best Improvement and for Fixed Priority pivot rules.

**Theorem 6.** *In the FIFO model, the convergence time of the Best Improvement pivot rule is in  $O(m^2 n \cdot W/p_{\text{min}})$ .*

*Proof.* According to Lemma 3 and Corollary 1, after  $O(n)$  iterations the potential  $\phi_{\text{FIFO}}$  is in  $O(n \cdot W/s_{\text{max}})$ . Hence, it suffices to show that in every sequence of  $m^2$  consecutive iterations,  $\phi_{\text{FIFO}}$  drops by at least  $p_{\text{min}}/s_{\text{max}}$ . Therefore, let us consider a sequence  $S$  of maximum length in which  $\phi_{\text{FIFO}}$  drops by strictly less than  $p_{\text{min}}/s_{\text{max}}$ . It is obvious that only jobs that are on top of some machine can jump as the running times of all the jobs above the moving job decrease by at least  $p_{\text{min}}/s_{\text{max}}$ .

For a given point in time, we call a job *active* if it jumps until the end of the sequence  $S$ . At any time, there can only be at most one active job on any machine. To see this, assume to the contrary that there are two active jobs  $j_1$  and  $j_2$  at the same time  $t_1$  on a machine  $i$ . Let job  $j_1$  w.l.o.g. be directly above job  $j_2$ , and let  $t_2 > t_1$  be the first iteration in which job  $j_2$  leaves machine  $i$  again. Define  $\alpha := c_{j_1}^{t_1} - c_{j_1}^{t_2}$  as the difference of  $j_1$ 's running times at time  $t_1$  and  $t_2$ . As job  $j_2$  was a top-most job in iteration  $t_2$  and no job below  $j_2$  could jump before  $j_2$  jumped, job  $j_1$  would have a running time of  $L_i^{t_1} - p_{j_2}/s_i$  if it jumped to machine  $i$  in the next step, yielding a total improvement of  $j_1$ 's running time of at least  $p_{\text{min}}/s_{\text{max}}$ . If  $j_1$  does not jump back to machine  $i$  in the next step, then either we have reached an equilibrium (then  $p_{j_2}/s_i \leq \alpha$ ) or there is a job (possibly also  $j_1$ ) that can improve by strictly more than  $p_{j_2}/s_i - \alpha$ . Hence, the potential drops by at least  $p_{j_2}/s_i \geq p_{\text{min}}/s_{\text{max}}$  during all the jumps of job  $j_1$  between  $t_1$  and  $t_2$  and the iteration following  $t_2 + 1$ .

Thus, we have shown that also at the beginning of the sequence  $S$  there are at most  $m$  active jobs as on each machine there is at most one active job. It also implies that no job  $j$  can jump back to a machine  $i$  it has already been onto as all jobs lying underneath  $j$  stay on machine  $i$  until the end of the sequence  $S$ . Hence, every job jumps at most  $m - 1$  times and the length of  $S$  is bounded from above by  $m(m - 1)$ . □

**Theorem 7.** *In the FIFO model, the expected convergence time of the Random pivot rule is in  $O(m^2 n^2 \cdot W/p_{\text{min}})$ .*

For Fixed Priority pivot rules, we cannot assume anymore that after a linear number of iterations there is no machine with load more than  $2W/s_{\text{max}}$  left and thus that  $\phi_{\text{FIFO}}$  is small. On the other hand, we know that the sum of the running times of all jobs that have already jumped is bounded by  $O(n \cdot W/s_{\text{max}})$  and we are able to show that during  $O(n)$  consecutive iterations, either a job jumps for the first time or the potential  $\phi_{\text{FIFO}}$  drops by at least  $p_{\text{min}}/s_{\text{max}}$ . In order to bound the potential by  $O(n \cdot W/s_{\text{max}})$ , we use the modified potential function

$$\phi'_{\text{FIFO}} := \sum_{j=1}^n \min \left\{ c_j, \frac{W}{s_{\text{max}}} \right\}.$$

**Theorem 8.** *In the FIFO model, the convergence time of any Fixed Priority pivot rule is in  $O(n^2 \cdot W/p_{\text{min}})$ .*

*Proof.* As  $0 \leq \phi'_{\text{FIFO}} \leq n \cdot W/s_{\text{max}}$ , we only have to show that during every sequence of  $n + 1$  steps, either  $\phi'_{\text{FIFO}}$  drops by at least  $p_{\text{min}}/s_{\text{max}}$  or a job must

jump for the very first time. In such a sequence, it must be the case that a job  $j_2$  jumps directly after a job  $j_1$ , where the priority of  $j_2$  is greater than the priority of  $j_1$ . This means that  $j_2$  jumps to the old machine  $i$  of job  $j_1$  as it could not jump before the move of  $j_1$ . If it was not  $j_1$ 's first jump, let  $t_2$  be the point in time between the two jumps by  $j_1$  and  $j_2$ , and let  $t_1$  be the point in time before  $j_1$  jumps the last time before  $t_2 - 1$ . As  $j_2$  does not want to jump to machine  $i$  at time  $t_1$ , but does this later at time  $t_2$ , it must be the case that  $L_i^{t_1} > L_i^{t_2}$ . Hence, between  $t_1 + 1$  and  $t_2 - 1$  a job  $j'$  assigned to machine  $i$  at time  $t_1$  must leave its machine. But during this time, job  $j_1$  lies above job  $j'$  yielding a running time improvement of  $p_{j'}/s_i \geq p_{\min}/s_{\max}$  for job  $j_1$  through the jump by  $j'$ . As  $j_1$  has jumped before, its running time before the jump by  $j'$  was already at most  $W/s_{\max}$ , meaning that also  $\phi_{\text{FIFO}}$  drops by at least  $p_{\min}/s_{\max}$ .  $\square$

The machine speeds do not occur in our bounds for the convergence times. Nevertheless, different machine speeds result in a higher convergence time than in the case of identical machines, as the following result shows. We believe that our proofs for the upper bounds on the convergence times are too pessimistic and thus we conjecture polynomial convergence times for all pivot rules. This is in contrast to the superpolynomial lower bounds in the Makespan and SJF model but a crucial difference is that the costs of a job can never increase in the FIFO model.

**Theorem 9.** *In the FIFO model, local search can take  $\Omega(mn)$  steps. The convergence time for the Min Weight pivot rule is in  $\Omega(m^2)$ .*

*Proof.* For the lower bound  $\Omega(mn)$ , let  $\ell \geq 1$  and  $k \geq 1$  be two integers. There are  $m = 2k + 1$  machines and  $n = k\ell + k + 1$  jobs split up in  $2k + 1$  job classes  $J_1, \dots, J_{2k+1}$ . The machine speeds are  $s_i = 2^{i-1}$  for  $1 \leq i \leq 2k$  and  $s_{2k+1} = 2^{2k+1}$ . The job classes  $J_1, \dots, J_k$  each contain  $\ell$  jobs with sizes  $2^0, \dots, 2^{\ell-1}$  and the job classes  $J_{k+1}, \dots, J_{2k+1}$  each contain a single job with size  $2^{\ell+j}$  for job class  $J_j$ .

Initially, each job class  $J_j$  is assigned to machine  $j$  and the jobs on a machine are processed in monotonically increasing order of the job sizes. We consider the following  $k$  rounds  $1, \dots, k$ . Before round  $i$  begins, the jobs from job class  $J_j$ ,  $j \leq k$ , are on machine  $j + i - 1$  such that they are processed in increasing order of the sizes, the jobs from job classes  $J_{k+1}, \dots, J_{k+i-1}$  are on machine  $2k + 1$  and the other jobs have not moved before. Then we let the single job from class  $J_{k+i}$  move from machine  $k + i$  to machine  $2k + 1$ . Thereupon, the jobs from class  $J_k$  move in ascending order of the sizes from machine  $k + i - 1$  to machine  $k + i$ , the jobs from class  $J_{k-1}$  move in ascending order of the sizes from machine  $k + i - 2$  to machine  $k + i - 1$  and so on. One can easily see that every job strictly decreases its costs while moving. All jobs from the job classes  $J_1, \dots, J_k$  move in every round. Hence, there are  $\Omega(k^2\ell) = \Omega(mn)$  iterations.

For the lower bound  $\Omega(m^2)$  for the Min Weight pivot rule, let again  $k$  be an integer and let  $\varepsilon > 0$  be appropriately small. There are  $m = n = 2k + 1$  machines and jobs. The machine speeds are  $s_i = 1 + i \cdot \varepsilon$  for  $1 \leq i \leq 2k$  and  $s_{2k+1} = 4k$ . The job sizes are  $p_j = 1 - j \cdot \varepsilon$  for  $1 \leq j \leq k$ ,  $p_j = 2 + 2j \cdot \varepsilon$  for  $k + 1 \leq j \leq 2k$ ,

and  $p_{2k+1} = 4k$ . Initially, every job  $j$  is assigned to machine  $j$  and the loads on the first  $k$  machines are less than 1,  $L_{k+1} = \dots = L_{2k} = 2$  and  $L_{2k+1} = 1$ . One can easily see that every job  $k + 1, \dots, 2k$  can move to machine  $2k + 1$  as  $L_{2k+1}$  remains to be less than 2 and that every such jump induces jumps from the jobs  $1, \dots, k$ . Hence, there are  $\Omega(k^2) = \Omega(m^2)$  iterations.  $\square$

### 6.2 Makespan

In this section, we consider the Best Improvement pivot rule in the Makespan model. We use the fact that the potential  $\phi_{\text{Makespan}}$  decreases by at least  $2p_{\min} \cdot p_{\min}/s_{\max}$  if a sequence of jobs decrease their running time by a total of  $p_{\min}/s_{\max}$  through jumping. This is due to a lemma by Even-Dar et al. [7] that if a jumping job  $j$  improves its execution time by  $\Delta$ , then  $\phi_{\text{Makespan}}$  drops exactly by  $2p_j\Delta$ .

Suppose that a job  $j$  wants to jump away from machine  $i$  to machine  $i'$  and there is a smaller job  $j'$  on machine  $i$ . At the current time, the costs of  $j$  and  $j'$  are the same as they are on the same machine. But the additional costs job  $j'$  would generate on any machine are strictly smaller than the additional costs job  $j$  would generate. Hence, job  $j'$  would have smaller costs on machine  $i'$  than job  $j$ . This leads to the following observation.

**Observation 1.** *When a job jumps away from a machine  $i$  according to the Best Improvement pivot rule, it was a smallest job on machine  $i$ .*

Let us now provide the main ideas of our proof. Imagine there are two jobs  $j_1, j_2$  on the same machine  $i$  and job  $j_1$  jumps away in iteration  $t_1$  making a small improvement directly before job  $j_2$  leaves machine  $i$  in iteration  $t_2 = t_1 + 1$ . Then job  $j_1$  could improve its running time by  $p_{j_2}/s_i$  by jumping back to machine  $i$  in iteration  $t_2 + 1$ . If, however,  $t_2 > t_1 + 1$ , it could happen that another job  $j_3$  from job  $j_1$ 's new machine leaves this machine leaving job  $j_1$  unable to jump back. But then job  $j_3$  is smaller than  $j_1$  according to Observation 1 and thus could jump to machine  $i$  in iteration  $t_2 + 1$  unless it already made a big improvement or another job from job  $j_3$ 's new machine jumped away in the meantime etc. Lemma 4 proves that the potential drops significantly during such a sequence.

**Lemma 4.** *If two jobs jump away from a machine  $i$  at iterations  $t < t'$  and no job enters machine  $i$  between  $t$  and  $t'$ , then the potential  $\phi_{\text{Makespan}}$  drops by at least  $p_{\min}^2/s_{\max}$  during the iterations  $t, \dots, t' + 1$  when using the Best Improvement pivot rule.*

Imagine now there are two jobs  $j_1, j_2$  entering the same machine  $i$  in two consecutive iterations  $t_1$  and  $t_2 = t_1 + 1$ , where job  $j_1$  moves first. Then job  $j_2$  would improve its running time by at least  $p_{j_1}/s_i$  if it jumped in iteration  $t_1$  as it also has the incentive to move to machine  $i$  after job  $j_1$ 's jump. But if  $t_2 > t_1 + 1$ , it could be that in iteration  $t_1$  job  $j_2$ 's running time is smaller than in iteration  $t_2$  and in the meantime another job  $j_3$  enters job  $j_2$ 's machine. If

job  $j_3$  is much larger than job  $j_2$ , then job  $j_2$  would improve much by jumping to job  $j_3$ 's old machine. Otherwise, job  $j_3$  could have moved to machine  $i$  in iteration  $t_1$  unless another job entered job  $j_3$ 's old machine in the meantime etc. Lemma 5 shows that also in this case the potential drops significantly.

**Lemma 5.** *If two jobs enter a machine  $i$  at iterations  $t' < t$  and no job leaves machine  $i$  between  $t'$  and  $t$ , then the potential  $\phi_{\text{Makespan}}$  drops by at least  $p_{\min}^2/(2 \cdot s_{\max})$  between  $t'$  and  $t + 1$  when using the Best Improvement pivot rule.*

Hence, we are able to show that if there is a machine to which two jobs migrate without a job leaving or from which two jobs leave without a job entering, the potential  $\phi_{\text{Makespan}}$  drops significantly. The proof then concludes with the observation that this must happen every  $O(m^2n)$  iterations.

**Theorem 10.** *In the Makespan model, the convergence time of the Best Improvement pivot rule is in  $O(m^2n \cdot W^2/p_{\min}^2)$ .*

*Proof.* According to Lemma 3, after  $O(n)$  iterations the potential  $\phi_{\text{Makespan}}$  is in  $O(W^2/s_{\max})$ . Hence, it suffices to show that in every sequence of  $m^2n$  consecutive iterations,  $\phi_{\text{Makespan}}$  drops by at least  $p_{\min}^2/(2s_{\max})$ .

Let  $S$  be a sequence of maximum length such that  $\phi_{\text{Makespan}}$  drops by less than  $p_{\min}^2/(2s_{\max})$ , lasting from iteration  $t_0$  to iteration  $t_\ell$ . We maintain a set of indices, which is empty at time  $t_0$ . When a job  $j$  jumps from a machine  $i_1$  to a machine  $i_2$  at iteration  $t \in \{t_0, \dots, t_\ell\}$  and if there has not been a job that jumped to machine  $i_1$  during the iterations  $t_0, \dots, t$ , generate a new index which gets attached to machine  $i_2$ . Otherwise, reattach the index previously attached to machine  $i_1$  to machine  $i_2$ . Lemma 4 shows that this is well-defined as there cannot be another job leaving machine  $i_1$  before another index gets attached to this machine.

At the end of the sequence, there can only be at most  $m$  indices. If an index gets reattached from machine  $i_1$  to machine  $i_2$  at iteration  $t$ , then  $L_{i_1}^t > L_{i_2}^{t+1}$ , i.e., the running time of the machine an index is attached to is strictly monotonically decreasing.

Consider an index that jumps with job  $j$  at iteration  $t$  and with job  $j'$  at iteration  $t'$  to the same machine  $i$ . Let  $j = j_1, j_2, \dots, j_\ell$  be the jobs that entered machine  $i$  and let  $j'_1, \dots, j'_\ell$  be the jobs that left machine  $i$  during the iterations  $t, \dots, t' - 1$  in this order. Lemmas 4 and 5 show that the order in which this happened must be  $j_1, j'_1, j_2, j'_2, \dots, j_\ell, j'_\ell$  and that the sequences have the same length, i.e., the sequences are well-defined. As always only a smallest job on a machine is able to achieve the best improvement and for every  $k$ , job  $j_k$  is on machine  $i$  when job  $j'_k$  leaves this machine, it must be the case that  $L_i^t \leq L_i^{t'}$ . But in the iterations  $t + 1$  and  $t' + 1$ , the same index is attached to machine  $i$ , meaning that  $L_i^t + p_j/s_i = L_i^{t+1} > L_i^{t'+1} = L_i^{t'} + p_{j'}/s_i$ , i.e.,  $p_j > p_{j'}$ . This means that an index cannot be attached twice to the same machine by a jump of the same job and thus an index gets reattached at most  $n \cdot m$  times. This concludes the proof.  $\square$

### 6.3 SJF

**Theorem 11.** *In the SJF model, the convergence time of the Min Weight pivot rule is exactly  $n$ , even on two machines. The expected convergence time of the Random pivot rule is less than  $n^2$ .*

## 7 Price of Anarchy for FIFO

Brunsch et al. [3] already showed that the smoothed price of anarchy for near list schedules in the Makespan model, which correspond to local optima in the FIFO model, is  $\Theta(\log \phi)$ . We give matching bounds for the deterministic case.

**Theorem 12.** *In the FIFO model, the price of anarchy for local search is  $\Theta(\log m)$  on related machines and  $2 - 1/m$  on identical machines.*

## 8 Concluding Remarks

We have shown several bounds for the convergence times of local search regarding three different coordination mechanisms on rational inputs. The choice of the right pivot rule decides in the Shortest Job First model between linear and exponential convergence times. The FIFO model is new but we believe that it is a realistic choice for many different real-life applications. We were able to show that every pivot rule converges in this model in linear time on identical machines and a large class of reasonable pivot rules converges in smoothed polynomial time on related machines. An interesting observation is that the machine speeds do not occur in any bound. We leave it as a conjecture that every pivot rule converges in polynomial time in the FIFO model. Another interesting open problem is whether the Best Improvement pivot rule in the Makespan model converges in smoothed or even deterministic polynomial time on related machines. We were only able to show that this happens with high probability when the input is perturbed.

**Acknowledgments.** We thank Clemens Rösner for helpful discussions about the lower bounds for the SJF model and the proof of Theorem 2.

## References

1. Beier, R., Vöcking, B.: Random knapsack in expected polynomial time. *J. Comput. Syst. Sci.* **69**(3), 306–329 (2004)
2. Brucker, P., Hurink, J., Werner, F.: Models and algorithms for planning and scheduling problems improving local search heuristics for some scheduling problems. Part II. *Discrete Appl. Math.* **72**(1), 47–69 (1997). doi:[http://dx.doi.org/10.1016/S0166-218X\(96\)00036-4](http://dx.doi.org/10.1016/S0166-218X(96)00036-4)
3. Brunsch, T., Röglin, H., Ruten, C., Vredeveld, T.: Smoothed performance guarantees for local search. *Math. Program.* **146**(1–2, Ser. A), 185–218 (2014). doi:[10.1007/s10107-013-0683-7](https://doi.org/10.1007/s10107-013-0683-7)

4. Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination mechanisms. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 345–357. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-27836-8\\_31](https://doi.org/10.1007/978-3-540-27836-8_31)
5. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. *Trans. Algorithms ACM* **3**(1) (2007)
6. Etscheid, M.: Performance guarantees for scheduling algorithms under perturbed machine speeds. In: Cai, L., Cheng, S.-W., Lam, T.-W. (eds.) ISAAC 2013. LNCS, vol. 8283, pp. 207–217. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-45030-3\\_20](https://doi.org/10.1007/978-3-642-45030-3_20)
7. Even-Dar, E., Kesselman, A., Mansour, Y.: Convergence time to Nash equilibria. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 502–513. Springer, Heidelberg (2003). doi:[10.1007/3-540-45061-0\\_41](https://doi.org/10.1007/3-540-45061-0_41)
8. Feldmann, R., Gairing, M., Lücking, T., Monien, B., Rode, M.: Nashification and the coordination ratio for a selfish routing game. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 514–526. Springer, Heidelberg (2003). doi:[10.1007/3-540-45061-0\\_42](https://doi.org/10.1007/3-540-45061-0_42)
9. Finn, G., Horowitz, E.: A linear time approximation algorithm for multiprocessor scheduling. *BIT* **19**, 312–320 (1979)
10. Garey, M.R., Johnson, D.S.: Complexity results for multiprocessor scheduling under resource constraints. *SIAM J. Comput.* **4**, 397–411 (1975)
11. Goldberg, P.W.: Bounds for the convergence rate of randomized local search in a multiplayer load-balancing game. In: Proceedings of the PODC 2004, pp. 131–140 (2004). doi:[10.1145/1011767.1011787](https://doi.org/10.1145/1011767.1011787)
12. Hurkens, C.A.J., Vredeveld, T.: Local search for multiprocessor scheduling: how many moves does it take to a local optimum? *Oper. Res. Lett.* **31**(2), 137–141 (2003)
13. Immorlica, N., Li, L., Mirrokni, V.S., Schulz, A.S.: Coordination mechanisms for selfish scheduling. *Theor. Comput. Sci.* **410**(17), 1589–1598 (2009). doi:[10.1016/j.tcs.2008.12.032](https://doi.org/10.1016/j.tcs.2008.12.032)
14. Manthey, B., Röglin, H.: Smoothed analysis: analysis of algorithms beyond worst case. *IT - Information Technology* **53**(6), 280–286 (2011)
15. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theor.* **2**, 65–67 (1973)
16. Schuurman, P., Vredeveld, T.: Performance guarantees of local search for multiprocessor scheduling. *Inform. J. Comput.* **19**(1), 52–63 (2007)
17. Spielman, D., Teng, S.-H.: Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time. *J. ACM* **51**(3), 385–463 (2004)
18. Spielman, D., Teng, S.-H.: Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *Commun. ACM* **52**(10), 76–84 (2009)



# On the Price of Stability of Undirected Multicast Games

Rupert Freeman, Samuel Haney<sup>(✉)</sup>, and Debmalya Panigrahi

Department of Computer Science, Duke University, Durham, NC 27708, USA  
{rupert, shaney, debmalya}@cs.duke.edu

**Abstract.** In multicast network design games, a set of agents choose paths from their source locations to a common sink with the goal of minimizing their individual costs, where the cost of an edge is divided equally among the agents using it. Since the work of Anshelevich et al. (FOCS 2004) that introduced network design games, the main open problem in this field has been the price of stability (PoS) of multicast games. For the special case of broadcast games (every vertex is a terminal, i.e., has an agent), a series of works has culminated in a constant upper bound on the PoS (Bilò et al., FOCS 2013). However, no significantly sub-logarithmic bound is known for multicast games. In this paper, we make progress toward resolving this question by showing a constant upper bound on the PoS of multicast games for quasi-bipartite graphs. These are graphs where all edges are between two terminals (as in broadcast games) or between a terminal and a nonterminal, but there is no edge between nonterminals. This represents a natural class of intermediate generality between broadcast and multicast games. In addition to the result itself, our techniques overcome some of the fundamental difficulties of analyzing the PoS of general multicast games, and are a promising step toward resolving this major open problem.

**Keywords:** Price of stability · Network design games · Cost sharing games

## 1 Introduction

In cost sharing network design games, we are given a graph/network  $G = (V, E)$  with edge costs and a set of users (agents/players) who want to send traffic from their respective source vertices to sink vertices. Every agent must choose a path along which to route traffic, and the cost of every edge is shared equally among all agents having the edge in their chosen path, i.e., using the edge to route traffic. This creates a *congestion game* since the players benefit from other players choosing the same resources. A Nash equilibrium is attained in this game

---

Rupert Freeman thanks NSF IIS-1527434 and ARO W911NF-12-1-0550 for support. Samuel Haney and Debmalya Panigrahi were supported in part by NSF Awards CCF-1527084 and CCF-1535972l.

when no agent has incentive to unilaterally deviate from her current routing path. The social cost of such a game is the sum of costs of edges being used in at least one routing path, and efficiency of the game is measured by the ratio of the social cost in an equilibrium state to that in an optimal state. (The optimal state is defined as one where the social cost is minimized, but the agents need not be in equilibrium.) The maximum value of this ratio (i.e., for the most expensive equilibrium state) is called the *price of anarchy* of the game, while the minimum value (i.e., for the least expensive equilibrium state) is called its *price of stability*. It is well known that even for the most restricted settings, the price of anarchy can be  $\Omega(n)$  for  $n$  agents. Therefore, the main question of research interest has been to bound the price of stability (POS) of this class of congestion games.

Anshelevich *et al.* [2] introduced network design games and obtained a bound of  $O(\log n)$  on the POS in directed networks with arbitrary source-sink pairs. While this is tight for directed networks, they left determining tighter bounds on the POS in undirected networks as an open question. Subsequent work has focused on the case of all agents sharing a common sink (called *multicast games*) and its restricted subclass where every vertex has an agent residing at it (called *broadcast games*). These problems are natural analogs of the Steiner tree and minimum spanning tree (MST) problems in a game-theoretic setting. For broadcast games, Fiat *et al.* [13] improved the POS bound to  $O(\log \log n)$ , which was subsequently improved to  $O(\log \log \log n)$  by Lee and Ligett [15], and ultimately to  $O(1)$  by Bilò *et al.* [5]. For multicast games, however, progress has been much slower, and the only improvement over the  $O(\log n)$  result of Anshelevich *et al.* is a bound of  $O(\log n / \log \log n)$  due to Li [16]. In contrast, the best known lower bounds on the PoS of both broadcast and multicast games are small constants [4]. As a result, determining the POS of multicast games has become a compelling open question in the area of network games.

In this paper, we achieve progress toward answering this question. In the multicast setting, a vertex is said to be a *terminal* if it has an agent on it, else it is called a *nonterminal*. Note that in the broadcast problem, there are no nonterminals and all the edges are between terminal vertices. In this paper, we consider multicast games in *quasi-bipartite* graphs: all edges are either between two terminals, or between a nonterminal and a terminal. (That is, there is no edge with both nonterminal endpoints.) This is a natural setting of intermediate generality between broadcast and multicast games. Moreover, quasi-bipartite graphs have been widely studied for the Steiner tree problem (see, e.g., [6, 7, 17, 18]) and has provided insights for the problem on general graphs. Our main result is an  $O(1)$  bound on the POS of multicast games in quasi-bipartite graphs.

**Theorem 1.** *The price of stability of multicast games in quasi-bipartite graphs is a constant.*

In addition to the result itself, our techniques overcome some of the fundamental difficulties of analyzing the PoS of general multicast games, and therefore represent a promising step toward resolving this important open problem. To illustrate this point, we outline the salient features of our analysis below.

The previous POS bounds for multicast games [2, 16] are based on analyzing a potential function  $\phi_e$  defined on each edge  $e$  as its cost scaled by the harmonic of the number of agents using the edge, i.e.,  $\phi_e = \text{cost}(e) \cdot (1 + 1/2 + 1/3 + \dots + 1/j)$  where  $j$  is the number of terminals using  $e$ . The overall potential is  $\phi = \sum_e \phi_e$ . When an agent changes her routing path (called a *move*), this potential exactly tracks the change in her shared cost. If the move is an *improving* one, then the shared cost of the agent decreases and so too does the potential. As a consequence, for an arbitrary sequence of improving moves starting with the optimal Steiner tree, the potential decreases in each move until a Nash Equilibrium (NE) is reached. This immediately yields a POS bound of  $H(n) = O(\log n)$  [2]. To see this, note that the potential of any configuration is bounded below by its cost, and above by its cost times  $H(n)$ . Then, letting  $S_{NE}$  be the Nash equilibrium reached, and  $T^*$  be the optimal routing tree, we have

$$c(S_{NE}) \leq \phi(S_{NE}) \leq \phi(T^*) \leq H(n)c(T^*).$$

This bound was later improved to  $O(\log n / \log \log n)$  by Li [16] with a similar but more careful accounting argument.

The previous POS bounds for broadcast games [5, 13, 15] use a different strategy. As in the case of multicast games, these results analyze a game dynamics that starts with an optimal solution (MST) and ends in an NE. However, the sequence of moves is carefully constructed — the moves are not arbitrary improving moves. At a high level, the sequence follows the same pattern in all the previous results for broadcast games:

1. Perform a *critical* move: Allow some terminal  $v$  to switch its path to introduce a single new edge into the solution, that is not in the optimal routing tree and is adjacent to  $v$ . This edge is associated with  $v$  and denoted  $e_v$ . Any edge introduced by the algorithm in any move other than a critical move uses only edges in the current routing tree, and edges in the optimal routing tree. Therefore, we only need to account for edges added by critical moves.
2. Perform a sequence of moves to ensure that the routing tree is *homogenous*. That is, the difference in costs of a pair of terminals is bounded by a function of the length of the path between them on the optimal routing tree. For example, suppose two terminals  $w$  and  $w'$  differ in cost by more than the length of the path between them in the optimal routing tree. Then the terminal with larger cost has an improving move that uses this path, and then the other terminal's path to the root. Such a move introduces only edges in the optimal routing tree.
3. *Absorb* a set of terminals around  $v$  in the shortest path metric defined on the optimal tree: terminals  $w$  replace their current strategy with the path in the optimal routing tree to  $v$ , and then  $v$ 's path to the root. If  $w$  had an associated edge  $e_w$ , introduced via a previous critical move, it is removed from the solution in this step.

The absorbing step allows us to account for the cost of edges added via critical moves, by arguing that vertices associated with critical edges of similar

length must be well-separated on the optimal routing tree. If edges  $e_u$  and  $e_v$  are not far apart, the second edge to be added would be removed from the solution via the absorbing step.

Homogeneity facilitates absorption: Suppose  $v$  has performed a critical move adding edge  $e_v$ , and let  $w$  be some other terminal. While  $v$  pays  $c(e_v)$  to use edge  $e_v$ ,  $w$  would only pay  $c(e_v)/2$  to use  $e_v$ , since it would split the cost with  $v$ . That is, if  $w$  bought a path to  $v$  and then used  $v$ 's path to the root, it would save at least  $c(e_v)/2$  over  $v$ 's current cost. If the current costs paid by  $v$  and  $w$  are not too different, and the distance between  $v$  and  $w$  not too large, then such a move is improving for  $w$ .

The previous results differ in how well they can homogenize: the tighter the bound on the difference in costs of a pair of terminals as a function of the length of the path between them in the optimal routing tree, the larger the radius in the absorb step. In turn, a larger radius of absorption establishes a larger separation between edges with similar cost, which yields a tighter bound on the PoS.

This homogenization-absorption framework has not previously been extended to multicast games. The main difficulty is that there can be nonterminals that are in the routing tree at equilibrium but are not in the optimal tree. No edge incident on these vertices is in the optimal tree metric, and therefore these vertices cannot be included in the homogenization process. So, any critical edge incident on such a vertex cannot be charged via absorption. This creates the following basic problem: what metric can we use for the homogenization-absorption framework that will satisfy the following two properties?

1. The metric is feasible – the sum of all edge costs in (a spanning tree of) the metric is bounded by the cost of the optimal routing tree. These edges can therefore be added or removed at will, without need to perform another set of moves to pay for them (in contrast to critical edges). This allows us to homogenize using these edges.
2. The metric either includes all vertices (as is the case with the optimal tree metric for broadcast games), or if there are vertices not included in the metric, critical edges adjacent to these vertices can be accounted for separately, outside the homogenization-absorption framework.

We create such a metric for quasi-bipartite graphs, allowing us to extend the homogenization-absorption framework to multicast games. Our metric is based on a dynamic tree containing all the terminals and a dynamic set of nonterminals. We show that under certain conditions, we can include the shortest edge incident on a nonterminal vertex, even if it is not in the optimal routing tree, in this dynamic tree. These edges are added and removed throughout the course of the algorithm. Our new metric is now defined by shortest path distances on this dynamic tree: the optimal routing tree extended with these special edges. We ensure homogeneity not on the optimal routing tree, but on this dynamic metric. Likewise, absorption happens on this new metric. We define the metric in such a way that the following hold:

1. The metric is feasible. That is, the total cost of all edges in the dynamic tree is within a constant factor of the cost of the optimal tree.

2. Consider some critical edge  $e_v$  such that the corresponding vertex  $v$  is not in the metric. That is, it was not possible to add the shortest edge adjacent to  $v$  to the dynamic tree while keeping it feasible. Therefore,  $v$  is at infinite distance from every other vertex in this metric, ruling out homogenization. Then,  $e_v$  can be accounted for separately, outside the homogenization-absorption framework.

For the remaining edges  $e_v$  such that  $v$  is in the metric, we account for them by using the homogenization-absorption framework. Our main technical contribution is in creating this feasible dynamic metric, going beyond the use of static optimal metrics in broadcast games. While the proof of feasibility currently relies on the quasi-bipartiteness of the underlying graph, we believe that this new idea of a feasible dynamic metric is a promising ingredient for multicast games in general graphs.

In the rest of the paper, we present the algorithm in detail, and provide an outline of its analysis. Details of the analysis are deferred to the full version of the paper due to space constraints.

### 1.1 Related Work

Recall that the upper bounds for PoS are a (large) constant and  $O\left(\frac{\log n}{\log \log n}\right)$  for broadcast and multicast games, respectively. The corresponding best known lower bounds are 1.818 and 1.862 respectively by Bilò *et al.* [4], leaving a significant gap, even for broadcast games. Moreover, Lee and Ligett [15] show that obtaining superconstant lower bounds, even for multicast games where they might exist, is beyond current techniques. While this lends credence to the belief that the PoS of multicast games is  $O(1)$ , Kawase and Makino [14] have shown that the potential function approach of Anshelevich *et al.* [2] cannot yield a constant bound on the PoS, even for broadcast games. In fact, Bilò *et al.* [5] used a different approach for broadcast games, as do we for multicast games on quasi-bipartite graphs.

Various special cases of network design games have also been considered. For small instances ( $n = 2, 3, 4$ ), both upper [10] and lower [3] bounds have been studied. [10] show upper bounds of 1.65 and  $4/3$  for two and three players respectively. For weighted players, Anshelevich *et al.* [2] showed that pure Nash equilibria exist for  $n = 2$ , but the possibility of a corresponding result for  $n \geq 3$  was refuted by Chen and Roughgarden [9], who also provided a logarithmic upper bound on the PoS. An almost matching lower bound was later given by Albers [1]. Recently, Fanelli *et al.* [12], showed that the PoS of network design games on undirected rings is  $3/2$ .

Network design games have also been studied for specific dynamics. In particular, starting with an empty graph, suppose agents arrive online and choose their best response paths. After all arrivals, agents make improving moves until an NE is reached. The worst-case inefficiency of this process was determined to be poly-logarithmic by Charikar *et al.* [8], who also posed the question of bounding the inefficiency if the arrivals and moves are arbitrarily interleaved.

This question remains open. Upper and lower bounds for the strong POA of undirected network design games have also been investigated [1, 11]. They show that the price of anarchy in this setting is  $\Theta(\log n)$ .

## 2 Preliminaries

Let  $G = (V, E)$  be an undirected edge-weighted graph and let  $c(e)$  denote the cost of edge  $e$ . Let  $U \subseteq V$  be a set of *terminals* and  $r \in U$ . In an instance of a *network design game*, each terminal  $u$  is associated with a *player*, or *agent*, that must select a path from  $u$  to  $r$ . We consider instances in which  $G$  is *quasi-bipartite*, that is no edge  $e$  has two nonterminal end points.

A *solution*, or *state*, is a set of paths connecting each player to the root. Let  $\mathcal{S}$  be the set of all possible solutions. For a solution  $S$ , a terminal  $u$ , and some subset  $E'$  of the edges in the graph, let  $c_u^{E'}(S) = \sum_{e \in E'} c(e)/n_e(S)$  be the cost paid by  $u$  for using edges in  $E'$ , where  $n_e(S)$  is the number of players using edge  $e$  in state  $S$ . Let  $p_u(S)$  be the set of edges used by  $u$  to connect to the root in  $S$  and let  $c_u(S) = c_u^{p_u(S)}(S)$  be the total cost paid by  $u$  to use those edges. For a nonterminal  $v$ , if every terminal  $u$  with  $v \in p_u(S)$  uses the same path from  $v$  to the root then define  $p_v(S)$  to be this path from  $v$  to  $r$ , and  $c_v(S) = c_v^{p_v(S)}(S)$ . Additionally, we will sometimes refer to the cost a vertex  $v$  pays, even if  $v$  is a nonterminal. By this we mean  $c_v(S)$ . For any vertex  $v \in \mathcal{S}$ , let  $e_v$  be the edge in  $p_v(S)$  with  $v$  as an endpoint.

Let  $\Phi : \mathcal{S} \rightarrow \mathbb{R}_+$  be the potential function introduced by Rosenthal [19], defined by

$$\Phi(S) = \sum_{e \in E} c(e)H_{n_e(S)} = c(e) \left( 1 + \frac{1}{2} + \dots + \frac{1}{n_e(S)} \right).$$

Let  $u \in U$  and suppose  $S$  and  $S'$  are states for which  $p_v(S) = p_v(S')$  for all players  $v \neq u$ . Then  $\Phi(S') - \Phi(S) = c_u(S') - c_u(S)$ . In particular, if a single player changes their path to a path of lower cost, the potential decreases.

The goal of each player is to find a path of minimum cost. A solution where no player can benefit by unilaterally changing their path is called a *Nash Equilibrium*. Let  $T^*$  be a solution that minimizes the total cost paid. Note that  $T^*$  is a minimum Steiner tree for  $G$ . The *price of stability* (PoS) is the ratio between the minimum cost of a Nash equilibrium and the cost of  $T^*$ .

Let  $p_{T^*}(u, v)$  be the path in  $T^*$  between  $u$  and  $v$ . Let  $v_1, \dots, v_n$  be the vertices of  $T^*$  in the order they appear in a depth first search of  $T^*$ . Let  $MC$ , the “main cycle”, be the concatenation of  $p_{T^*}(v_1, v_2), p_{T^*}(v_2, v_3), \dots, p_{T^*}(v_{n-1}, v_n), p_{T^*}(v_n, v_1)$ . Note that each edge in  $T^*$  appears exactly twice in  $MC$ . The following property will be helpful:

**Fact 2.** *Any  $x$  to  $y$  path in  $MC$  completely contains  $p_{T^*}(x, y)$ .*

Define the class of edge  $e$ ,  $class(e)$ , as  $\alpha$  if  $256^\alpha \leq c(e) < 256^{\alpha+1}$ . Without loss of generality, we assume that  $c(e) \geq 1$  for all  $e \in E$ , so the minimum possible

edge class is 0. For simplicity, define  $\lfloor c(e) \rfloor = 256^{class(e)}$ , a lower bound for  $c(e)$ , and  $\lceil c(e) \rceil = 256^{class(e)+1}$ , an upper bound for  $c(e)$ .

For each nonterminal  $v$ , let  $\sigma_v$  be the minimum cost edge adjacent to  $v$  in  $G$ . Let  $t_v$  be the terminal adjacent to  $\sigma_v$ . Let  $T^+$  be the extended optimal metric:  $T^* \cup \{\sigma_v\}_{v \in V}$ . We maintain a dynamic set of nonterminals  $\mathcal{Z}_S = \{w \notin T^* : c(\sigma_w) \leq \lfloor c(e_w) \rfloor / 64\}$ . That is,  $\mathcal{Z}_S$  are those nonterminals  $w$  in solution  $S$  whose first edge  $e_w$  has cost within a constant factor of the cost of  $\sigma_w$ . For any  $w \in S$ , if  $\sigma_w$  is added to  $S$  while  $w \in \mathcal{Z}_S$ , then we show that we will be able to pay for  $\sigma_w$  if it remains in the final solution. In the algorithm, we denote the current state by  $S_{curr}$ . For ease of notation, we define  $\mathcal{Z} = \mathcal{Z}_{S_{curr}}$ .

The remaining definitions are modifications of key definitions from [5]. The interval around vertex  $v \in T^*$  with budget  $y$ ,  $I_{v,y}$ , is the concatenation of its right and left intervals,  $I_{v,y}^+$  and  $I_{v,y}^-$ , where  $I_{v,y}^+$  is the maximal contiguous interval in  $MC$  with  $v$  a left endpoint such that

$$2 \sum_{\alpha \geq 0} 256^{\alpha+1} H_{n_{I^+, \alpha}}^2 \leq y,$$

where  $n_{I^+, \alpha}$  is the number of edges of class  $\alpha$  in  $I_{v,y}^+$  (repeated edges are counted every time they appear). We define  $I_{v,y}^-$  similarly.

The neighborhood of  $v$  in state  $S$ ,  $N_S(v)$  is an interval around  $v$  as well as certain  $w \notin T^*$  with  $t_w$  in the interval. Formally,

$$N_S(v) = \begin{cases} I_{v, \lfloor \frac{c(e_v)}{56} \rfloor} \cup \left\{ w \in \mathcal{Z}_S \mid t_w \in I_{v, \lfloor \frac{c(e_v)}{56} \rfloor} \text{ and } c(\sigma_w) \leq \frac{\lfloor c(e_v) \rfloor}{64} \right\} & \text{if } v \in T^*, \\ I_{t_v, \lfloor \frac{c(e_v)}{56} \rfloor} \cup \left\{ w \in \mathcal{Z}_S \mid t_w \in I_{t_v, \lfloor \frac{c(e_v)}{56} \rfloor} \text{ and } c(\sigma_w) \leq \frac{\lfloor c(e_v) \rfloor}{64} \right\} & \text{otherwise.} \end{cases}$$

$N_S^+(v)$  and  $N_S^-(v)$  are the right and left intervals of the neighborhood respectively (that is, the portions of  $N_S(v)$  to the right and left of  $v$  or  $t_v$  respectively). We denote  $N_{S_{curr}}(v)$  as  $N(v)$ . Roughly speaking, we are going to charge the cost of edges in the final solution not in  $T^*$  to the interval portions of non-overlapping right neighborhoods. A path  $X = p_{T^*}(x, y)$  is *homogenous* if

$$|c_x(S) - c_y(S)| \leq 4 \sum_{\alpha \geq 0} 256^{\alpha+1} H_{n_{X, \alpha}}^2.$$


If  $X = p_{T^*}(x, y) \subseteq N(v) \cap T^*$  is a homogenous path then

$$|c_x(S) - c_y(S)| \leq 4 \sum_{\alpha \geq 0} 256^{\alpha+1} H_{n_{X, \alpha}}^2 \leq 8 \sum_{\alpha \geq 0} 256^{\alpha+1} H_{n_{N^+(v), \alpha}}^2 \leq \lfloor c(e) \rfloor / 14.$$

$N(v)$  is homogenous if the following holds: For all  $x, y \in N(v)$  with  $x, y \neq u_v$ , a special vertex to be defined later, such that the path in  $T^+$  from  $x$  to  $y$  does not contain  $v$ ,  $|c_x(S_{curr}) - c_y(S_{curr})| \leq \frac{23 \lfloor c(e_v) \rfloor}{112}$ . Homogenous neighborhoods allow us to bound the difference in cost between any two vertices in  $N(v)$  which will be useful when arguing that players have improving strategy changes.

### 3 Algorithm

The initial state of the algorithm is the minimum cost tree  $T^*$  connecting all the terminals to the root. The algorithm carefully schedules a series of potential-reducing moves. (Recall the potential function  $\Phi(S) = \sum_{e \in E} c(e)H_{n_e(S)}$  introduced in Sect. 2). Since there are finitely many states possible, such a series of moves must always be finite. Since any improving move reduces potential, we must be at a Nash equilibrium if there is no potential reducing move. These moves are scheduled such that if any edge outside of  $T^*$  is introduced, it is subsequently accounted for by charging to some part of  $T^*$ . In particular, we will show that at any point in the process, and therefore in the equilibrium state at the end, the total cost of these edges is bounded by  $O(1) \cdot c(T^*)$ .



**Fig. 1.** Types of critical improving moves. Dotted edges represent the new edges being added.

The algorithm is a series of loops, which we run repeatedly until we reach a Nash equilibrium. Each loop begins with a terminal,  $a$ , performing either a *safe* improving move, or a *critical* improving move. In both cases,  $a$  switches strategy to follow a new path to the root. Let  $S$  be the state before the start of the loop. A safe improving move is one which results in some state  $S' \subseteq T^* \cup S$ , i.e., the new path of  $a$  contains edges currently in  $S$  and edges in the optimal tree  $T^*$ . A safe improving move requires no additional accounting on our part. A critical improving move on the other hand introduces one or two new edges that must be accounted for (see Fig. 1). We will show later that in any non-equilibrium state, a safe or critical improving move always exists (see Lemma 3).

The algorithm will use a sequence of (potential-reducing) moves to account for the new edges introduced by a critical move. At a high level, each of these edges is accounted for in the following way. Let  $e_v$  be the edge in question, and  $v$  be the first vertex using  $e_v$  on its path to the root.

1. In some neighborhood around  $v$ , perform a sequence of moves to ensure that for every pair of vertices (excluding  $v$  and at most one other special vertex), the difference in shared costs of these vertices is not too large. (Recall that the while nonterminals do not pay anything, the shared cost of a nonterminal  $u$  is defined to be  $c_u(S)$ , the cost that a terminal using  $u$  pays on its subpath from  $u$  to the root). This sequence of moves must be potential-reducing, and cannot add any edges outside of  $T^* \cup S$  to the solution.



2. For every vertex  $y$  in the neighborhood around  $v$ ,  $v$  has an alternative path to the root consisting of the path in  $T^+$  to  $y$ , and  $y$ 's path to the root. (Recall from Sect. 2 that  $T^+$  is the optimal tree,  $T^*$ , augmented with minimum cost edges incident on nonterminals  $\{\sigma_w : w \text{ is a nonterminal}\}$ ).
  - (a) If a  $y$  exists for which this alternative path is improving for  $v$ , then  $v$  can switch to this new path and  $e_v$  will be removed from the solution.
  - (b) If every path is *not* improving for  $v$ , then we show that every vertex in the neighborhood of  $v$  has an improving move that uses  $e_v$ .

These steps ensure that we either remove  $e_v$  from the solution, or else for any vertex  $y$  in the neighborhood we remove edge  $e_y \notin T^*$  from the solution. We elaborate on the steps above, referencing the subroutines described in Algorithm 2 – HOMOGENIZE, ABSORB, and MAKETREE:

**Step 1:** This is accomplished in two ways. For any path in  $T^*$ , the HOMOGENIZE subroutine ensures that a path in  $T^*$  is homogenous. Recall that this gives a bound (relative to the cost of  $e_v$ ) on the difference in shared costs of the endpoints of the path. Additionally, for any pair of adjacent vertices, if the difference in the shared costs is more than the cost of the edge between them, then one vertex must have an improving move through this edge. This move adds no edges outside of  $T^*$ . The second way of bounding differences in shared cost is much weaker, but we will use it only a small number of times. Overall, the path between any two vertices in the neighborhood will comprise homogenous segments connected by edges whose cost is bounded by the second method above. Adding up the cost bounds for these segments gives us the total bound.

**Step 2(a):** The purpose of this step is to establish that either the shared cost of  $v$  is not much larger than the shared cost of every other vertex in its neighborhood, or that we can otherwise remove  $e_v$  from the solution. If the shared cost of  $v$  is much larger than some other vertex in the neighborhood, then it is also much larger than the shared cost of an adjacent vertex (call it  $q$ ) in  $T^+$ . This is because every pair of vertices in the neighborhood have a similar shared cost (by Step 1). Then,  $v$  has a lower cost path to the root consisting of the  $(v, q)$  edge, combined with  $q$ 's current path to the root. Such a move would remove  $e_v$  from the solution.

**Step 2(b):** If we reach this step, we need to account for the cost of  $e_v$  by making every other vertex in the neighborhood give up its first edge, if that edge is not in  $T^+$ . This ensures that at the end, the edges in the solution that are not in  $T^+$  will be very far apart. This is accomplished via the ABSORB function:  $v$  is currently paying the entire cost of  $e_v$ , while any vertex that would switch to using  $v$ 's path to the root would only pay at most half the cost of  $e_v$ . Furthermore, if vertices close to  $v$  in  $T^+$  switch first, vertices farther from  $v$  (who must pay a higher cost to buy a path to  $v$ ) will reap the benefits of more sharing, and therefore a further reduction in shared cost. This is formalized in the definition of ABSORB.

There are some other details which we mention here before moving on to a more formal description of the algorithm:

- If  $v$  is a nonterminal, let  $u_v$  be the terminal that added  $v$  as part the critical move. We avoid including  $u_v$  in any path provided to the HOMOGENIZE subroutine. This is because HOMOGENIZE switches the strategies of terminals to follow the strategy of some terminal on input path. If terminals were switched to follow  $u_v$ 's path, this would increase the sharing on  $e_v$ , when it is required at the beginning of Step (2b) that only one terminal is using  $e_v$ . When  $v$  is a terminal, then  $u_v$  is undefined and this problem does not exist. We define two versions of a loop of the algorithm, defined as MAINLOOP in Algorithm 1, to account for this difference.
- We have only described how to account for a single edge, but sometimes a critical move adds two new edges that must be accounted for. Suppose  $e_a$  and  $e_b$  are the new edges added by  $a$  ( $a$  is a terminal and  $b$  is a nonterminal). Then we run MAINLOOP( $e_b$ ) first, and then MAINLOOP( $e_a$ ). The first loop does not increase sharing on  $e_a$ , so the second loop is still valid.
- We assume the existence of a function MAKE TREE. This function takes as input a set of strategies. Its output is a new set of strategies such that (1) the new set of strategies has lower potential than the old set, (2) the edge set of the new strategies is a subset of the old edge set, and (3) the edge set of the new strategies is a tree. In particular, MAKE TREE( $S_{curr} \setminus \{p_{u_v}(S_{curr}), p_v(S_{curr})\}$ ), used on line 9 does not increase sharing on  $e_v$ , since  $v$  and  $u_v$  are the only two vertices using  $e_v$  on their path to the root. MAKE TREE( $S_{curr} \setminus \{p_{u_v}(S_{curr}), p_v(S_{curr})\}$ ) will also not increase sharing on  $e_{u_v}$  if this edge has just been added (and therefore  $u_v$  is the only vertex using the edge). We will not go into more detail about this function, since an identical function was used in both [5,13].
- We assume that all edges in  $E$  with  $c(e) > c(T^*)$  have been removed from the graph. This is without loss of generality: if the final state  $S_f$  is a Nash equilibrium, then  $S_f$  is still an equilibrium after reintroducing  $e$  with  $c(e) > c(T^*)$ . This is because any vertex with an improving move that adds such an edge  $e$  also has a path to the root (in  $T^*$ ) with total cost less than  $c(e)$ .

We walk through the pseudocode next: We execute the MAINLOOP function given in Algorithm 1 either once or twice, once for each edge not in  $T^* \cup S$  that is added by a critical move. If two edges have been added, we execute in the order MAINLOOP( $e_b$ ) then MAINLOOP( $e_a$ ) (where  $a$  is the terminal and  $b$  the nonterminal). We define two versions of MAINLOOP( $e_v$ ), one when  $v$  is a terminal, and one when  $v$  is a nonterminal, appearing on lines 17 and 1 respectively. When  $v$  is a nonterminal, we denote the terminal which added  $e_v$  to the solution as part of the initial improving move as  $u_v$ . For brevity, we define  $u_v$  as “empty” when  $v$  is a terminal. Thus if  $v$  is a terminal, define  $N(v) \setminus \{u_v\} = N(v)$ .

The **while** loops at lines 2 and 18 terminate with  $N(v)$  being homogenous. For any violated **if** statement within the **while** loop, we perform a move that reduces potential, and does not increase sharing on  $e_v$ , or on  $e_{u_v}$  if it was added along with  $e_v$  as part of  $u_v$ 's critical move. If none of these **if** conditions

```

1: function MAINLOOP( $e_v$ )  $\triangleright v$  is a nonterminal and  $u_v$  the terminal which added
    $e_v$  as part of a critical move.
2:   while any of the following if conditions are true do
3:     if  $\exists X = p_{T^*}(x, y) \in N(v) \cap T^*$  with  $u_v, v \notin X$  and  $X$  not homogenous then
       HOMOGENIZE( $X$ )
4:     if  $\exists x, y \in N(v) \setminus \{v\}$  adjacent to  $u_v$  with  $c_x(S_{curr}) - c_y(S_{curr}) > c(x, u_v) +$ 
        $c(u_v, y)$  then
5:       Replace  $x$ 's strategy with  $(x, u_v) \cup (u_v, y) \cup p_y(S_{curr})$ .
6:     if  $\exists w \in N(v) \setminus T^*$  such that  $t_w \neq v, u_v$  with  $|c_w(S_{curr}) - c_{t_w}(S_{curr})| > c(\sigma_w)$ 
       then
7:       Assuming WLOG  $c_{t_w}(S_{curr}) > c_w(S_{curr})$ , replace  $t_w$ 's strategy with
        $\sigma_w \cup p_w(S_{curr})$ .
8:     if  $S_{curr} \setminus \{p_{u_v}(S_{curr}), p_v(S_{curr})\}$  is not a tree then
9:       MAKETREE( $S_{curr} \setminus \{p_{u_v}(S_{curr}), p_v(S_{curr})\}$ )
10:    for  $q \in N(v) \setminus \{v, u_v\}$  adjacent in  $T^+$  to either  $v$  or  $u_v$  do
11:      if  $c(v, q) + c_q(S_{curr}) < c_v(S_{curr})$  then
12:         $v$  changes strategy to  $(v, q) \cup p_q(S_{curr})$ .
13:      return
14:      Repeat the previous 3 lines substituting  $u_v$  for  $v$ .
15:       $\triangleright$  Note that  $u_v$  changing strategy will remove  $v$  from the solution.
16:    ABSORB( $v$ )

17: function MAINLOOP( $e_v$ )  $\triangleright v$  a terminal.
18:   while any of the following if conditions are true do
19:     if  $\exists X = p_{T^*}(x, y) \in N(v) \cap T^*$  with  $v \notin X$  and  $X$  is not homogenous then
       HOMOGENIZE( $X$ )
20:     if  $\exists w \in N(v) \setminus T^*$  such that  $t_w \neq v$  with  $|c_w(S_{curr}) - c_{t_w}(S_{curr})| > c(\sigma_w)$ 
       then
21:       Assuming WLOG  $c_{t_w}(S_{curr}) > c_w(S_{curr})$ , replace  $t_w$ 's strategy with
        $\sigma_w \cup p_w(S_{curr})$ .
22:     if  $S_{curr} \setminus \{p_v(S_{curr})\}$  is not a tree then MAKETREE( $S_{curr} \setminus \{p_v(S_{curr})\}$ )
23:     for  $q \in N(v)$  adjacent in  $T^+$  to  $v$  do
24:       if  $c(v, q) + c_q(S_{curr}) < c_v(S_{curr})$  then
25:          $v$  changes strategy to  $(v, q) \cup p_q(S_{curr})$ .
26:       return
27:     ABSORB( $v$ )

```

**Algorithm 1.** Main loop to be executed for each edge added to the solution as part of a critical move.

hold,  $N(v)$  is homogenous. Therefore, this **while** loop eventually terminates in a homogenous state.

We next ensure that the cost that  $v$  pays is similar to the cost every other vertex in  $N(v)$  pays. If these costs are not close, we can show that the condition at line 11/24 will be true, and  $e_v$  will be deleted from the solution.

```

25: function HOMOGENIZE( $X = p_{T^*}(x, y)$ )
26:   Let  $X = (x = x_1, x_2, \dots, x_k, x_{k+1} = y)$ 
27:   Let  $S'$  be the current state.
28:   for  $i \leftarrow 1$  to  $k$  do
29:     for  $j \leftarrow i$  down to 1 do
30:       Change  $x_j$ 's strategy to  $p_{T^*}(x_j, x_{i+1}) \cup p_{x_i}(S)$ .
31:       if  $\Phi(S_{curr}) < \Phi(S')$  then return
32:       else Reset state to  $S'$ 

Require:  $c_q(S) \geq c_v(S) - \frac{2 \cdot |c(e_v)|}{7} \quad \forall q \in N(v) \setminus \{u_v\}$ 
33: function ABSORB( $v$ )  $\triangleright v$  absorbs  $N(v) \setminus \{u_v\}$ 
34:   for  $q \in N(v) \cap T^* \setminus \{u_v\}$  in breadth-first order from  $r$  according to  $T^*$  do
35:     if  $v \notin T^*$  then Change  $q$ 's strategy along with its descendants to
36:        $p_{T^*}(q, t_v) \cup \sigma_v \cup p_v(S)$ .
37:     else Change  $q$ 's strategy along with its descendants to  $p_{T^*}(q, v) \cup p_v(S)$ .
38:   Let  $S'$  be the current state.
39:   for  $q \in N(v) \setminus T^*$ , in reverse breadth-first order from  $r$  according to  $S'$  do
40:     Change  $q$ 's strategy along with its descendants to  $\sigma_q \cup p_{t_q}(S')$ .

```

**Algorithm 2.** Helper functions for Algorithm 1.

If  $e_v$  is still present at this point, we finally call the ABSORB function. We use the precondition of the ABSORB function to show that the switches made by all the vertices in  $N(v)$  are improving, and therefore reduce potential.

Note that although we do not make this explicit, if at any point  $S_{curr}$  contains edges that are not part of  $p_u(S_{curr})$  for any terminal  $u$ , these edges are deleted immediately. This ensures that any nonterminal in  $S_{curr}$  is always used as part of some terminal's path to  $r$ .

**Outline of Analysis.** We first show that all parts of the algorithm reduce potential, guaranteeing that the algorithm terminates (by the definition of the potential function, the minimum decrease in potential is bounded away from 0). Most steps in the algorithm involve single terminals making improving moves, and therefore these steps reduce potential. There are two parts of the algorithm for which it is not immediately obvious that potential is reduced: the HOMOGENIZE function and the ABSORB function. The lemma below states that HOMOGENIZE reduces potential, and we give its proof in the full paper.

**Lemma 1.** *Suppose there is a path  $X = p_{T^*}(x, y) \in N(v)$  which is not homogeneous. Let  $(x = x_1, x_2, \dots, x_k, x_{k+1} = y)$  be the sequence of vertices in  $X$ . Then there exists a prefix of  $X$ ,  $(x_1, \dots, x_i)$ , such that the sequence of moves in which each  $x_j$ ,  $j \in \{1, \dots, i\}$ , switches its strategy to  $p_{T^*}(x_j, x_{i+1}) \cup p_{x_{i+1}}(S)$  reduces potential.*

Proof that the precondition for the ABSORB function is satisfied (homogeneity is required here) is deferred to the full version. If it is satisfied, we can show that the ABSORB function reduces potential.

**Lemma 2.** *If  $c_q(S_{curr}) \geq c_v(S_{curr}) - \frac{2 \cdot \lfloor c(e_v) \rfloor}{7}$  for all  $q \neq u_v \in N(v)$ , then every strategy change in ABSORB reduces potential.*

Lemmas 1 and 2 imply that the entire main loop is potential reducing. Since the minimum decrease in potential is bounded away from zero, and the potential is always at least zero, the algorithm necessarily terminates. However, termination alone does not guarantee that the final state is a Nash equilibrium. Since we have restricted the set of moves that the algorithm can perform, we must show that whenever an improving move is available to some terminal, there is also an improving move that is either a safe or critical move (proof in full version).

**Lemma 3.** *The final state reached by the algorithm,  $S_f$ , is a Nash equilibrium.*

Finally, we show our main result, i.e., that  $c(S_f) = O(c(T^*))$ . To establish the theorem, it is sufficient to show that  $c(S_f \setminus T^*) = O(c(T^*))$ . We devise a charging scheme that distributes the cost of edges in  $S_f \setminus T^*$  among edges in  $T^*$ . Each  $e \in S_f \setminus T^*$  must be an  $e_v$  edge for some vertex  $v$ . Furthermore, these  $e_v$  edges were not later removed as the result of an absorbing process initiated from another  $e_{v'}$ . At a high level, this allows us to distribute the cost of each  $e_v$  to the edges in the neighborhood  $N(v) \cap T^*$ , since the  $\text{ABSORB}(v)$  function removes many other  $e_{v'}$  edges where  $v' \in N(v)$  from the solution.

We first consider a set of edges that we will not charge to their neighborhood. Define  $E_\sigma = \{e_v \in S_f \mid v \text{ is a nonterminal, } \frac{\lfloor c(e_v) \rfloor}{64} \leq c(\sigma_v)\}$ . We bound the cost of  $E_\sigma$  by the cost of edges in  $S_f \setminus E_\sigma$  (proof in full version).

**Lemma 4.**  $c(E_\sigma) = O(c(S_f \setminus E_\sigma))$ .

Our goal now is to find a set of edges  $e_v$  such that the right neighborhoods associated with edges of the same class are not overlapping. In the absence of nonterminals, this is simple: For every edge in  $S_f \setminus T^*$ , the right neighborhoods of vertices corresponding to edges of the same class being overlapping implies that each edge is contained in the other's neighborhood. Therefore, we argue that the second edge to arrive would have deleted the first through the ABSORB function, which gives a contradiction. With nonterminals, the same property does not hold. When edge  $e_v$  is added for some nonterminal  $v$ ,  $e_{u_v}$  will not be deleted from the solution, even if  $u_v$  falls in  $v$ 's neighborhood. The presence of  $\sigma_v$  for which no  $\text{MAINLOOP}(\sigma_v)$  was run (added, e.g., in line 35) further complicates things. To show that no right neighborhoods overlap, we will therefore remove some edges from  $S_f \setminus (T^* \cup E_\sigma)$ .

For nonterminal  $v$ , if  $v$  is adjacent to at least two edges in  $S_f \setminus (T^* \cup E_\sigma)$  and  $\sigma_v$  is one such edge, remove  $\sigma_v$  and charge it to one of the remaining edges adjacent to  $v$ . Next, for any pair of edges  $e_u$  and  $e_v$  in  $S_f \setminus (T^* \cup E_\sigma)$  such that  $u$  was the terminal which added  $e_v$ , we delete the smaller of  $e_u$  and  $e_v$  and charge it to the remaining edge. We are left with a set of edges which we denote  $E^*$ , each of which has been charged by at most two edges that were removed (and each edge removed is charged to some edge in  $E^*$ ).

Our argument will charge to each edge in  $T^*$  at most one edge in  $E^*$  of each class. To make the argument simpler, it is desirable to charge those  $\sigma_v$ 's for

which  $\text{MAINLOOP}(\sigma_v)$  was never run to higher classes than their actual classes. To this end, we increase the cost of each such  $\sigma_v$  to  $c(e_{\sigma_v})$ , the cost of the first edge on  $v$ 's path in the state just before  $\sigma_v$  was added.

**Lemma 5.** *For edges  $e_u, e_v \in E^*$ , if  $\text{class}(e_v) = \text{class}(e_u)$ , then  $N^+(v)$  and  $N^+(u)$  are disjoint.*

Given Lemma 5, the scheme from [5] for distributing the cost of each  $e_v$  to its neighborhood can be applied directly. This leads to Theorem 1. For the details of this analysis, the reader is referred to the full version of the paper.

## References

1. Albers, S.: On the value of coordination in network design. *SIAM J. Comput.* **38**(6), 2273–2302 (2009)
2. 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 J. Comput.* **38**(4), 1602–1623 (2008)
3. Bilò, V., Bove, R.: Bounds on the price of stability of undirected network design games with three players. *J. Interconnect. Netw.* **12**(1–2), 1–17 (2011)
4. Bilò, V., Caragiannis, I., Fanelli, A., Monaco, G.: Improved lower bounds on the price of stability of undirected network design games. *Theory Comput. Syst.* **52**(4), 668–686 (2013)
5. Bilò, V., Flammini, M., Moscardelli, L.: The price of stability for undirected broadcast network design with fair cost allocation is constant. In: *FOCS*, pp. 638–647 (2013)
6. Byrka, J., Grandoni, F., Rothvoß, T., Sanità, L.: Steiner tree approximation via iterative randomized rounding. *J. ACM* **60**(1), 6 (2013)
7. Chakrabarty, D., Devanur, N.R., Vazirani, V.V.: New geometry-inspired relaxations and algorithms for the metric steiner tree problem. *Math. Program.* **130**(1), 1–32 (2011)
8. Charikar, M., Karlo, H.J., Mathieu, C., Naor, J., Saks, M.E.: Online multicast with egalitarian cost sharing. In: *Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2008, Munich, Germany, 14–16 June 2008*, pp. 70–76 (2008)
9. Chen, H.-L., Roughgarden, T.: Network design with weighted players. *Theory Comput. Syst.* **45**(2), 302–324 (2009)
10. Christodoulou, G., Chung, C., Ligett, K., Pyrga, E., Stee, R.: On the price of stability for undirected network design. In: Bampis, E., Jansen, K. (eds.) *WAOA 2009*. LNCS, vol. 5893, pp. 86–97. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-12450-1\\_8](https://doi.org/10.1007/978-3-642-12450-1_8)
11. Epstein, A., Feldman, M., Mansour, Y.: Strong equilibrium in cost sharing connection games. *Games Econ. Behav.* **67**(1), 51–68 (2009)
12. Fanelli, A., Leniowski, D., Monaco, G., Sankowski, P.: The ring design game with fair cost allocation. *Theor. Comput. Sci.* **562**, 90–100 (2015)
13. 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). doi:[10.1007/11786986\\_53](https://doi.org/10.1007/11786986_53)

14. Kawase, Y., Makino, K.: Nash equilibria with minimum potential in undirected broadcast games. *Theor. Comput. Sci.* **482**, 33–47 (2013)
15. Lee, E., Ligett, K.: Improved bounds on the price of stability in network cost sharing games. In: *EC*, pp. 607–620 (2013)
16. Li, J.: An  $o(\log(n)/\log(\log(n)))$  upper bound on the price of stability for undirected shapley network design games. *Inf. Process. Lett.* **109**(15), 876–878 (2009)
17. Rajagopalan, S., Vazirani, V.V.: On the bidirected cut relaxation for the metric steiner tree problem. In: *SODA*, pp. 742–751 (1999)
18. Robins, G., Zelikovsky, A.: Tighter bounds for graph steiner tree approximation. *SIAM J. Discret. Math.* **19**(1), 122–134 (2005)
19. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **2**(1), 65–67 (1973)

# Efficiency and Budget Balance

Swaprava Nath<sup>(✉)</sup> and Tuomas Sandholm

Carnegie Mellon University, Pittsburgh, USA  
{swapravn,sandholm}@cs.cmu.edu

**Abstract.** We study *efficiency* and *budget balance* for designing mechanisms in general quasi-linear domains. Green and Laffont [13] proved that one cannot generically achieve both. We consider strategyproof budget-balanced mechanisms that are approximately efficient. For deterministic mechanisms, we show that a strategyproof and budget-balanced mechanism must have a *sink* agent whose valuation function is ignored in selecting an alternative, and she is compensated with the payments made by the other agents. We assume the valuations of the agents come from a bounded open interval. This result strengthens Green and Laffont’s impossibility result by showing that even in a restricted domain of valuations, there does not exist a mechanism that is strategyproof, budget balanced, and takes every agent’s valuation into consideration—a corollary of which is that it cannot be efficient. Using this result, we find a tight lower bound on the inefficiencies of strategyproof, budget-balanced mechanisms in this domain. The bound shows that the inefficiency asymptotically disappears when the number of agents is large—a result close in spirit to Green and Laffont [13, Theorem 9.4]. However, our results provide worst-case bounds and the best possible rate of convergence. Next, we consider minimizing any convex combination of inefficiency and budget imbalance. We show that if the valuations are unrestricted, no deterministic mechanism can do asymptotically better than minimizing inefficiency alone. Finally, we investigate randomized mechanisms and provide improved lower bounds on expected inefficiency. We give a tight lower bound for an interesting class of strategyproof, budget-balanced, randomized mechanisms. We also use an optimization-based approach—in the spirit of *automated mechanism design*—to provide a lower bound on the minimum achievable inefficiency of any randomized mechanism. Experiments with real data from two applications show that the inefficiency for a simple randomized mechanism is 5–100 times smaller than the worst case. This relative difference increases with the number of agents.

## 1 Introduction

Consider a group of friends deciding which movie to watch together. The movie can be watched in someone’s home by renting it or at any of a number of

---

We are grateful to Ioannis Caragiannis, Vincent Conitzer, Debasis Mishra, Hervé Moulin, Ariel Procaccia, and Arunava Sen for useful discussions. Nath is funded by the Fulbright-Nehru postdoctoral fellowship. Sandholm is funded by the National Science Foundation under grants 1320620, 1546752, and 1617590, and by the ARO under award W911NF-16-1-0061.



movie theaters. Each of these choices incurs a cost. Since individual preferences are different and sometimes conflicting, the final choice may not make everybody maximally satisfied. This may cause some of the agents to misreport their preferences or drop out of the plan. To alleviate this problem, one can think of monetary transfers so that the friends who get their more-preferred choice pay more than the friends that get their less-preferred choice. Desirable properties of such a choice and payment rule are that (1) the total side payments (transfers among the friends) should sum to zero, so there is no surplus or deficit, and (2) the choice is efficient, that is, the movie that is selected maximizes the sum of all the friends' valuations. Since the valuations are private information of the friends, an efficient decision requires the valuations to be revealed truthfully. This simple example is representative of many joint decision-making problems that often involve monetary transfers. Consider, for example, a group of firms sharing time on a jointly-owned supercomputer, city dwellers deciding on the location and choice of a public project (e.g., stadium, subway, or library), mobile service providers dividing spectrum among themselves, or a student body deciding which musician or art performer to invite to entertain at their annual function. These problems all call for efficient joint decision making and involve—or could involve depending on the application—monetary transfers.

This is a ubiquitous problem in practice and a classic problem in the academic literature. We study the standard model of this problem where the agents' utilities are *quasi-linear*: each agent's utility is her valuation for the selected alternative (e.g., the choice of movie) minus the money she has to pay. A classic goal is to select an *efficient* alternative, that is, the one that maximizes the sum of the agents' valuations (also known as *social welfare*). We will study the problem of designing strategyproof mechanisms, that is, mechanisms where each agent is best off revealing the truth regardless of what other agents reveal.

Even though there are mechanisms that select efficient alternatives in a truthful manner (e.g., the Vickrey-Clarke-Groves (VCG) mechanism [5, 14, 34]), the transfers by the individuals do not sum to zero (in public goods settings, the VCG mechanism leads to too much money being collected from the agents). The execution of such a mechanism needs an external mediator who consumes the surplus (or may need to pay the deficit), to keep the mechanism truthful and efficient—a phenomenon known as 'money burning' in literature. In our movie selection example, this implies that we need a third party who will collect the additional money paid by the individuals, which is highly impractical in many settings. This has attracted significant criticism of the VCG mechanism [30]. Ideally, one would like to design strategyproof mechanisms that are efficient and *budget balanced*, that is, they do not have any surplus or deficit. Green and Laffont [13] proved a seminal impossibility for this setting: in the general quasi-linear domain, strategyproof, efficient mechanisms cannot be budget balanced.

In this paper, we primarily focus on the problem of minimizing inefficiency subject to budget balance in the general setting of quasi-linear utilities. This is because, in the applications of interest to this paper (e.g., movie selection), budget balance is more critical than efficiency. However, we show that for a large set of agents, the per-agent inefficiency vanishes. We also show that

for deterministic settings, optimizing the sum (or any convex combination) of efficiency and budget balance—which seems to be the most sensible objective—does not provide any asymptotic benefit over maximizing efficiency subject to budget balance.

## 1.1 Contributions of this Paper

In this paper, we assume that the agents' valuations are picked from a bounded open interval. In Sect. 3, we characterize the structure of truthful, budget balanced, *deterministic* mechanisms in this restricted domain, and show that any such mechanism must have a *sink* agent,<sup>1</sup> whose reported valuation function does not impact the choice of alternative and she gets the payments made by the other agents (Theorem 1). This result strengthens the Green and Laffont impossibility by showing that even in a restricted domain of bounded valuations, there does not exist a mechanism that is strategyproof, budget balanced, and takes every agent's valuation into consideration—a corollary of which is that it cannot be efficient. With the help of this characterization, we find the optimal deterministic mechanism that minimizes the inefficiency. This provides a tight lower bound on the inefficiency of deterministic, strategyproof, budget-balanced mechanisms. By inefficiency of a mechanism in this paper, we mean the worst-case inefficiency over all valuation profiles. We provide a precise rate of decay ( $\frac{1}{n}$ ) of the inefficiency with the increase in the number of agents (Theorem 2). This implies that the inefficiency vanishes for large number of agents. To contrast this mechanism with the class of mechanisms that minimize budget imbalance subject to efficiency, we considered the joint minimization problem of a convex combination of *inefficiency* and budget imbalance, and observed that it does *not* provide any asymptotic benefit over the previous problem. Due to limited space, we discuss this only in the full version of this paper [28].

We investigate the advantages of randomized mechanisms in Sect. 4. We first consider the class of *generalized sink* mechanisms. These mechanisms have, for every possible valuation profile, a probability distribution over the agents that determines each agent's chance of becoming the sink. This class of mechanisms is budget balanced by design. We show examples where mechanisms from this class are not strategyproof (Algorithm 2), and then isolate an interesting subclass whose mechanisms are strategyproof, the *modified irrelevant sink mechanisms* (Algorithm 3). We show that no mechanism from this class can perform better than the deterministic mechanisms if the number of alternatives is greater than the number of agents (Theorem 3). Since inefficiency (weakly) increases with the

---

<sup>1</sup> Mechanisms using this idea have been presented with different names in the literature. The original paper by Green and Laffont [13] refers to this kind of agents as a *sample* of the population. Later Gary-Bobo and Jaaidane [11] formalized the randomized version of this mechanism which is known as *polling* mechanism. Faltings [9] refers to this as an *excluded coalition* (when there are multiple such agents) and Moulin [25] mentions this as *residual claimants*. However, we use the term 'sink' for brevity and convenience, and our paper considers a different setup and optimization objective.

number of alternatives (Theorem 4), we consider the extreme case of two alternatives and compare the performances of different mechanisms. We show that a naïve uniform random sink mechanism and the modified irrelevant sink mechanism (Algorithm 3) perform equally well (Theorems 5 and 6) and reduce the inefficiency by a constant factor of 2 from that of the deterministic mechanisms. However, the optimal, strategyproof, budget-balanced, randomized mechanism performs better than these mechanisms. Since the structure of strategyproof randomized mechanisms for general quasi-linear utilities is unknown,<sup>2</sup> we take an optimization-based approach to find the best mechanism for the special case of two agents. This approach is known in the literature as *automated mechanism design* [6]. For an overview, see [32]. We discretize the range of the valuations into finite levels and show that when the number of levels increases—thereby making the lower bound tighter to the actual open-interval problem—the improvement factor reduces to less than 5 (Fig. 1). This is a significant improvement over the class of randomized sink mechanisms, which only improve over the best deterministic mechanism by a factor of 2.

We present experiments using real data from two applications. They show that in practice the inefficiency is significantly smaller and has a faster rate of decay than the worst case bounds (Sect. 5). We conclude the paper in Sect. 6 and present future research directions. Owing to the page limitation, the complete details of the results and the proofs are available in the full version of this paper [28].

## 1.2 Relationships to Prior Literature

The Green-Laffont impossibility result motivated the research direction of designing efficient mechanisms that are minimally budget imbalanced. The approach is to redistribute the surplus money in a way that satisfies truthfulness of the mechanism [3, 4]. The *worst case optimal* and *optimal in expectation* guarantees have been given for this class of mechanisms in restricted settings [16, 17, 25]. The performance of this class of *redistribution* mechanisms has been evaluated in interesting special domains such as allocating single or multiple (identical or heterogeneous) objects [15]. Also, mechanisms have been developed and analyzed that are budget balanced (or no deficit) and minimize the inefficiency in special settings [18, 22, 24]. Characterization of strategyproof budget-balanced mechanisms in the setting of cost-sharing is explored by Moulin and Shenker [26] and its quantitative guarantees are presented by Roughgarden and Sundararajan [31]. If the distribution of the agents' valuations is known and we assume common knowledge among the agents over those priors, the strategyproofness requirement can be weakened to Bayesian incentive compatibility. In that weaker framework, mechanisms can extract full expected efficiency and achieve budget balance [1, 7]. However, those mechanisms use knowledge of the priors. Therefore, in the general quasi-linear setting, for mechanisms without priors, it is an important open

<sup>2</sup> For randomized mechanisms, results involving special domains are known, e.g., facility location [10, 29, 33], auctions [8], kidney exchange [2], and most of these mechanisms aim for specific objectives.

question to characterize the class of strategyproof budget-balanced mechanisms, to find such mechanisms that minimize inefficiency, and to find strategyproof mechanisms that minimize the sum (or other convex combination) of inefficiency and budget imbalance. This paper addresses this important research gap in the general quasi-linear setting, for both deterministic and randomized settings. Our approach is also prior-free—the strategyproofness guarantees consider the worst-case scenarios. We show that the answers are asymptotically positive: even in such a general setup, the Green-Laffont impossibility is not too restrictive when the number of agents is large, and our mechanisms seem to work well on real-world datasets.

## 2 Model and Definitions

We denote the set of agents by  $N = \{1, 2, \dots, n\}$  and the set of alternatives by  $A = \{a_1, a_2, \dots, a_m\}$ . We assume that each agent’s valuation is drawn from an open interval  $(-\frac{M}{2}, \frac{M}{2}) \subset \mathbb{R}$ , that is, the valuation of agent  $i$  is a mapping  $v_i : A \rightarrow (-\frac{M}{2}, \frac{M}{2}), \forall i \in N$  and is a private information. Denote the set of all such valuations of agent  $i$  as  $V_i$  and the set of valuation profiles by  $V = \times_{i \in N} V_i$ .

A *mechanism* is a tuple of two functions  $\langle f, \mathbf{p} \rangle$ , where  $f$  is called the *social choice function* (SCF) that selects the *allocation* and  $\mathbf{p} = (p_1, p_2, \dots, p_n)$  is the vector of *payments*,  $p_i : V \rightarrow \mathbb{R}, \forall i \in N$ . The utility of agent  $i$  for an alternative  $a$  and valuation profile  $v \equiv (v_i, v_{-i})$  is given by the *quasi-linear* function:  $v_i(a) - p_i(v_i, v_{-i})$ . For *deterministic* mechanisms,  $f : V \rightarrow A$  is a deterministic mapping, while for *randomized* mechanisms, the allocation function  $f$  is a lottery over the alternatives, that is,  $f : V \rightarrow \Delta A$ . With a slight abuse of notation, we denote  $v_i(f(v_i, v_{-i})) \equiv \mathbb{E}_{a \sim f(v_i, v_{-i})} v_i(a) = \int v_i(a) \cdot f(v_i, v_{-i})$  to be the expected valuation of agent  $i$  for a randomized mechanism. The following definitions are standard in the mechanism design literature.

**Definition 1 (Strategyproofness).** A mechanism  $\langle f, \mathbf{p} \rangle$  is strategyproof if for all  $v \equiv (v_i, v_{-i}) \in V$ ,

$$v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i}), \quad \forall v'_i \in V_i, i \in N.$$

**Definition 2 (Efficiency).** An allocation  $f$  is efficient if it maximizes social welfare, that is,  $f(v) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} v_i(a), \forall v \in V$ .

**Definition 3 (Budget Balance).** A payment function  $p_i : V \rightarrow \mathbb{R}, i \in N$  is budget balanced if  $\sum_{i \in N} p_i(v) = 0, \forall v \in V$ .

In addition, in parts of this paper we will consider mechanisms that are oblivious to the alternatives—a property known as *neutrality*. To define this, we consider a permutation  $\pi : A \rightarrow A$  of the alternatives. Therefore,  $\pi$  over a randomized mechanism and over a valuation profile will imply that the probability masses and the valuations of the agents are permuted over the alternatives according to  $\pi$ , respectively.<sup>3</sup>

<sup>3</sup> We have overloaded the notation of  $\pi$  following the convention in social choice literature (see, e.g., Myerson [27]). The notation  $\pi(v)$  denotes the valuation profile where the alternatives are permuted according to  $\pi$ .

**Definition 4 (Neutrality).** *A mechanism  $\langle f, \mathbf{p} \rangle$  is neutral if for every permutation of the alternatives  $\pi$  (where  $\pi(v) \neq v$ ) we have*

$$\pi(f(v)) = f(\pi(v)) \quad \text{and} \quad p_i(\pi(v)) = p_i(v), \quad \forall v \in V, \forall i \in N.$$

Note that efficient social choice functions are neutral and the Green-Laffont result implicitly assumes this property.

The most important class of allocation functions in the context of deterministic mechanisms are *affine maximizers*, defined as follows.

**Definition 5 (Affine Maximizers).** *An allocation function  $f$  is an affine maximizer if there exist real numbers  $w_i \geq 0, i \in N$ , not all zeros, and a function  $\kappa : A \rightarrow \mathbb{R}$  such that  $f(v) \in \operatorname{argmax}_{a \in A} (\sum_{i \in N} w_i v_i(a) + \kappa(a))$ .*

As we will explain in the body of this paper, we will focus on *neutral* affine maximizers [23], where the function  $\kappa$  is zero.

$$f(v) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} w_i v_i(a) \quad \text{neutral affine maximizer} \quad (1)$$

The following property of the mechanism ensures that two different payment functions of an agent, say  $i$ , that implement the same social choice function differ from each other by a function that does not depend on the valuation of agent  $i$ .<sup>4</sup>

**Definition 6 (Revenue Equivalence).** *An allocation  $f$  satisfies revenue equivalence if for any two payment rules  $p$  and  $p'$  that make  $f$  strategyproof, there exist functions  $h_i : V_{-i} \rightarrow \mathbb{R}$ , such that*

$$p_i(v_i, v_{-i}) = p'_i(v_i, v_{-i}) + h_i(v_{-i}), \quad \forall v_i \in V_i, \forall v_{-i} \in V_{-i}, \forall i \in N.$$

The metrics of inefficiency we consider in this paper are defined as follows.

**Definition 7 (Sample Inefficiency).** *The sample inefficiency for a deterministic mechanism  $\langle f, \mathbf{p} \rangle$  is:*

$$r_n^M(f) := \frac{1}{nM} \sup_{v \in V} \left[ \max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right]. \quad (2)$$

The metric is adapted to expected sample inefficiency for randomized mechanisms:

$$r_n^M(f) := \frac{1}{nM} \sup_{v \in V} \left\{ \mathbb{E}_{f(v)} \left[ \max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right] \right\}. \quad (3)$$

The majority of this paper is devoted to finding strategyproof and budget balanced mechanisms  $\langle f, \mathbf{p} \rangle$  that minimize the sample inefficiency.

---

<sup>4</sup> This definition is a generalization of auction revenue equivalence and is commonly used in the social choice literature (see, e.g., Heydenreich et al. [21]).

A different, but commonly used, metric of inefficiency in the literature is the worst-case ratio of the social welfare of the mechanism and the maximum social welfare:  $\inf_{v \in V} \frac{\sum_{i \in N} v_i(f(v))}{\max_{a \in A} \sum_{i \in N} v_i(a)}$ . A conclusion similar to what we prove in this paper: “*inefficiency vanishes when  $n \rightarrow \infty$* ”, holds in that metric as well, but unlike our metric, that metric would require an additional assumption that the valuations are positive, which is not always the case in a quasi-linear domain.

We are now ready to start presenting our results. We begin with deterministic mechanisms that are strategyproof and budget balanced.

### 3 Deterministic, Strategyproof, Budget-Balanced Mechanisms

Before presenting the main result of this section, we formally define a class of mechanisms we call *sink* mechanisms. A sink mechanism has one or more *sink* agents, given by the set  $S \subset N$ , picked a priori, whose valuations are not used when computing the allocation (i.e.,  $f(v) = f(v_{-S})$ ) and the sink agents do not pay anything and together they receive the payments made by the other agents. The advantage of a sink mechanism is that it is strategyproof if it is strategyproof for the agents other than the sink agents and the surplus is divided among the sink agents in some reasonable manner, and sink mechanisms are budget balanced by design. An example of a sink mechanism is where  $S = \{i_s\}$  (only one sink agent) and  $f(v_{-i_s})$  chooses an alternative that would be efficient had agent  $i_s$  not exist, that is,  $f(v_{-i_s}) = \operatorname{argmax}_{a \in A} \sum_{i \in N \setminus \{i_s\}} v_i(a)$ . The Clarke [5] payment rule can be applied here to make the mechanism strategyproof for the rest of the agents—that is, for agents other than  $i_s$ ,  $p_i(v_{-i_s}) = \max_{a \in A} \sum_{j \in N \setminus \{i_s, i\}} v_j(a) - \sum_{j \in N \setminus \{i_s, i\}} v_j(f(v_{-i_s}))$ ,  $\forall i \in N \setminus \{i_s\}$ . Paying agent  $i_s$  the ‘leftover’ money (that is,  $p_{i_s}(v_{-i_s}) = -\sum_{j \in N \setminus \{i_s\}} p_j(v_{-i_s})$ ) makes the mechanism budget balanced. Our first result establishes that the existence of a sink agent is not only sufficient but also *necessary* for deterministic mechanisms.

**Theorem 1.** *Any deterministic, strategyproof, budget-balanced, neutral mechanism  $\langle f, \mathbf{p} \rangle$  in the domain  $V$  has at least one sink agent.*<sup>5</sup>

All proofs are provided in the full version of this paper [28]. This proof involves two steps. First, we leverage the fact that a mechanism that satisfies the stated axioms must necessarily be a neutral affine maximizer (Eq. 1) and has a specific structure for payments. The characterization of the payment structure comes

<sup>5</sup> Green and Laffont’s impossibility result holds for efficient mechanisms, and all efficient mechanisms are neutral. However, there could be instances where multiple alternatives are efficient, i.e., there is a tie. The neutrality of an efficient rule is up to tie-breaking, and Green-Laffont applies no matter how the tie is broken. Similarly, our result also holds irrespective of how the tie is broken. Therefore, this theorem covers and generalizes that result since having at least one sink agent implies that the outcome cannot be efficient.

from the revenue equivalence result. The second part of the proof shows that for such payment functions, it is impossible to have no sink agents (identified as agents that have zero weights,  $w_i = 0$ , in the affine maximizer). This is shown in a contrapositive manner—assuming that there is no sink agent, we construct valuation profiles that lead to a contradiction to budget balance.

The next goal is to find the mechanism in this class that gives the *lowest* sample inefficiency (Eq.2). In the proof of the next theorem (presented in [28]) we show that this is achieved when there is exactly one sink and the neutral affine maximizer weights are equal for all other agents. This, in turn, yields the following lower bound on inefficiency.

**Theorem 2.** *For every deterministic, strategyproof, budget-balanced, neutral mechanism  $\langle f, \mathbf{p} \rangle$  over  $V$ ,  $r_n^M(f) \geq \frac{1}{n}$ . This bound is tight.*

## 4 Randomized, Strategyproof, Budget-Balanced Mechanisms

In Sect. 3, we saw that the best sample inefficiency achieved by a deterministic budget balanced mechanism is  $\frac{1}{n}$ . In this section, we discuss how the inefficiency can be reduced by considering randomized mechanisms. An intuitive approach is to consider a mechanism where each agent is picked as a sink with probability  $\frac{1}{n}$ .

**Definition 8 (Naïve Randomized Sink).** *A naïve randomized sink (NRS) mechanism picks every agent as a sink w.p.  $\frac{1}{n}$  and takes the efficient allocation without that agent. The payments of the non-sink agents are VCG payments without the sink. The surplus is transferred to the sink.*

Clearly, this mechanism is strategyproof, budget balanced, and neutral by design. One can anticipate that this may not yield the best achievable inefficiency bound. Unlike deterministic mechanisms, very little is known about the structure of randomized strategyproof mechanisms in the general quasi-linear setting. Furthermore, we consider mechanisms that are budget-balanced in addition. Hence, even though we can obtain an upper bound on the expected sample inefficiency ( $r_n^M(f)$ ) by considering specific mechanisms like the NRS mechanism described above, the problem of providing a lower bound (i.e., no randomized mechanism can achieve a smaller  $r_n^M(f)$  than a given number), seems elusive in the general quasi-linear setting.

Therefore, in the following two subsections, we consider two approaches, respectively. First, we show lower bounds in a special class of strategyproof, budget-balanced, randomized mechanisms. Second, we provide a lower bound of the optimal, strategyproof, budget-balanced, randomized mechanism for two agents and two alternatives, using a discrete relaxation of the original problem (in the spirit of *automated mechanism design* [6, 32]). However, the problems of finding a mechanism that matches this lower bound and extending the lower bound to any number of agents and alternatives are left as future work.

## 4.1 Generalized Sink Mechanisms

In the first approach, we consider a broad class of randomized, budget-balanced mechanisms, which we coin *generalized sink mechanisms*. In this class, the probability of an agent  $i$  to become a sink is dependent on the valuation profile  $v \in V$ , and we consider mechanisms with only *one* sink, i.e., if the probability vector returned by a generalized sink mechanism is  $g(v)$ , then w.p.  $g_i(v)$ , agent  $i$  is treated as the *only* sink agent.<sup>6</sup> Clearly, the naïve randomized sink mechanism belongs to this class. Once agent  $i$  is picked as a sink, the alternative chosen is the *efficient* one *without* agent  $i$ . All agents  $j \neq i$  are charged a Clarke tax payment in the world without  $i$ , and the surplus amount of money is transferred to the sink agent  $i$ . Algorithm 1 shows the steps of a generic mechanism in this class.

---

### ALGORITHM 1. Generalized Sink Mechanisms, $\mathcal{G}$

---

- 1: **Input:** a valuation profile  $v \in V$
  - 2: A generic mechanism in this class is characterized by a probability distribution over the agents  $N$  (which may depend on the valuation profile),  $g : V \rightarrow \Delta N$
  - 3: The mechanism randomly picks one agent  $i$  in  $N$  with probability  $g_i(v)$
  - 4: Treat agent  $i$  as the sink
- 

Clearly, not every mechanism in this class is strategyproof. The crucial aspect is how the probabilities of choosing the sink are decided. If the probability  $g_i(v)$  depends on the valuation of agent  $i$ , that is,  $v_i$ , then there is a chance for agent  $i$  to misreport  $v_i$  to have higher (or lower) probability of being a sink (being a sink could be beneficial since she gets all the surplus). For example, the *irrelevant sink* mechanism given in Algorithm 2 is *not* strategyproof.

---

### ALGORITHM 2. Irrelevant Sink Mechanism (not strategyproof)

---

- 1: **Input:** a valuation profile  $v \in V$
  - 2: **for** agent  $i$  in  $N$  **do**
  - 3:   Define:  $a^*(v_{-i}) = \operatorname{argmax}_{a \in A} \sum_{j \neq i} v_j(a)$
  - 4:   **if**  $\sum_{j \neq i} v_j(a^*(v_{-i})) - \sum_{j \neq i} v_j(a) > M$  for all  $a \in A \setminus \{a^*(v_{-i})\}$  **then**
  - 5:     Call  $i$  an irrelevant agent
  - 6:   **if** irrelevant agent is found **then**
  - 7:     Arbitrarily pick one of them as a sink with probability 1
  - 8:   **else**
  - 9:     Pick an agent  $i$  with probability  $\frac{1}{n}$  and treat as sink
- 

In the full version of this paper [28], we provide an counterexample to strategyproofness of this mechanism. However, a small modification of the previous mechanism leads to a strategyproof generalized sink mechanism. This shows

<sup>6</sup> One can think of a more general class of sink mechanisms where multiple agents are treated as sink agents simultaneously. However, it is easy to see—by a similar argument to that in the context of deterministic mechanisms—that using multiple sinks cannot decrease inefficiency.



that the class of generalized sink mechanisms is indeed richer than the constant probability sink mechanisms. In the modified version, we pick a default sink with a certain probability, which will be the sink if there exists no irrelevant agent among the rest of the agents. The change here is that when an agent is picked as a default sink, her valuation has no effect in deciding the sink. See Algorithm 3.

---

**ALGORITHM 3.** Modified Irrelevant Sink Mechanism (strategyproof)

---

- 1: **Input:** a valuation profile  $v \in V$
  - 2: Pick agent  $i$  as a *default sink* with probability  $p_i$
  - 3: **for** agent  $j$  in  $N \setminus \{i\}$  **do**
  - 4:     **if** irrelevant agent(s) found within  $N \setminus \{i\}$  **then**
  - 5:         Arbitrarily pick one of them as a sink
  - 6:         Irrelevant agent is found
  - 7: **if** no irrelevant agent is found within  $N \setminus \{i\}$  **then**
  - 8:     Treat agent  $i$  as sink
- 

It is easy to verify that this mechanism is strategyproof. Interestingly, no generalized sink mechanism can improve the expected sample inefficiency over deterministic mechanisms if there are more alternatives than agents ( $m > n$ ).

**Theorem 3 (Generalized Sink for  $m > n$ ).** *If  $m > n$ , every generalized sink mechanism has expected sample inefficiency  $\geq \frac{1}{n}$ .*

The proof is critically dependent on  $m > n$ . However, we can hope for a smaller inefficiency if the number of alternatives is small. We state this intuition formally as follows.

**Theorem 4 (Increasing Inefficiency with  $m$ ).** *For every mechanism  $f$  and for a fixed number of agents  $n$ , the expected sample inefficiency is non-decreasing in  $m$ , i.e.,  $r_{n,m_1}^M(f) \geq r_{n,m_2}^M(f), \forall m_1 > m_2$ .<sup>7</sup>*

Theorems 3 and 4 suggest that in order to minimize inefficiency, one must have a small number of alternatives. So from now on, we consider the extreme case with  $m = 2$ , where we investigate the advantages of randomization.

For two alternatives, the following theorem shows that the naïve randomized sink (NRS) mechanism reduces the inefficiency by a factor of two.

**Theorem 5 (Naïve Randomized Sink).** *For  $m = 2$ , the expected sample inefficiency of the NRS mechanism is  $\frac{1}{n^2} \lceil \frac{n}{2} \rceil \sim \frac{1}{2n}$ .*

Even though the modified irrelevant sink (MIS) mechanism (Algorithm 3) is more sophisticated than NRS, it turns out that both of them have the same inefficiency on every valuation profile. Both mechanisms choose a single agent as a sink. The default sink for MIS is chosen uniformly at random, identical to the choice of the sink for NRS. If there does not exist an irrelevant sink in the rest of the agents, the inefficiency remains the same as that for the default sink,

---

<sup>7</sup> We overload the notation for the expected sample inefficiency  $r_n$  with  $r_{n,m}$  to make the number of alternatives explicit.

which is identical to the inefficiency of NRS for that choice of sink. But even if an irrelevant sink exists, by the construction of the irrelevant sink, the resulting alternative is the efficient alternative for the agents except the default sink. This outcome would have resulted even if the default sink were chosen as the sink. Therefore, the inefficiencies in MIS and NRS mechanisms are the same. Hence, we get the following theorem.

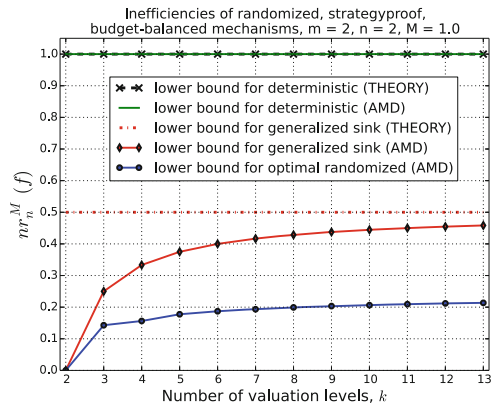
**Theorem 6 (Modified Irrelevant Sink).** *For  $m = 2$ , the expected sample inefficiency of the MIS mechanism (Algorithm 3) is at least  $\frac{1}{n^2} \lceil \frac{n}{2} \rceil \sim \frac{1}{2n}$ .*

### 4.2 Unrestricted Randomized Mechanisms

We now move on to study optimal randomized mechanisms without restricting attention necessarily to generalized sink mechanisms. For a fixed number of agents, minimizing the expected sample inefficiency is equivalent to minimizing the expected absolute inefficiency given by  $nr_n^M(f)$ . Finding a mechanism that achieves the minimum absolute inefficiency can be posed as the following optimization problem.

$$\begin{aligned}
 \min_{f, \mathbf{p}} \quad & \sup_{v \in V} \left[ \max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right] \\
 \text{s.t.} \quad & v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \\
 & \geq v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i}), \forall v_i, v'_i, v_{-i}, \forall i \in N \\
 & \sum_{a \in A} f_a(v) = 1, \forall v \in V, \\
 & \sum_{i \in N} p_i(v) = 0, \forall v \in V, \\
 & f_a(v) \geq 0, \forall v \in V, a \in A.
 \end{aligned} \tag{4}$$

The objective function denotes the absolute inefficiency. The first set of inequalities in the constraints denote the strategyproofness requirement, where the term  $v_i(f(v)) = v_i \cdot f(v)$  denotes the expected valuation of agent  $i$  due to the randomized mechanism  $f$ . The second and last set of inequalities ensure that the  $f_a(v)$ 's are valid probability distributions, and the third set of inequalities ensure that the budget is balanced. The optimization is over the social choice functions  $f$  and the payments  $\mathbf{p}$ , where the  $f$  variables are non-negative but the  $p$  variables are unrestricted. Clearly, this is a linear program (LP), which has an uncountable number of constraints (because the equalities



**Fig. 1.** Lower bound for the discrete relaxation of the inefficiency minimization LP.

clearly, this is a linear program (LP), which has an uncountable number of constraints (because the equalities

and inequalities have to be satisfied at all  $v \in V$ , which are the profiles of valuation functions mapping alternatives to an open interval). We address this optimization problem using finite constrained optimization techniques by discretizing the valuation levels. We assume that each agent’s valuations are uniformly discretized with  $k$  levels in  $[-M/2, M/2]$ , which makes the set of valuation profiles  $V$  finite. The optimal value of such a discretized relaxation of the constraints provides a lower bound on the optimal value of the original problem. This is because the discretized relaxation of the valuations only increases the feasible set since some of the constraints are removed, that is, more  $f$ ’s and  $p$ ’s satisfy the constraints, allowing a potentially lower value to be achieved for the minimization objective.

We conducted a form of automated mechanism design [6, 32] by solving this LP using Gurobi [19] for increasing values of  $k$ . We apply the same optimization-based approach for generalized sink and the deterministic cases as well, even though for these cases we have theoretical bounds. The solid lines in Fig. 1 show the optimization-based results (denoted as AMD) and the dotted lines show the theoretical bounds. Note that for deterministic case, the theoretical and optimization-based approaches overlap since the inefficiency is unity even with two valuation levels. The convergence of the optimization-based approach for generalized sink mechanism shows the efficacy of the approach and helps to predict the convergence point for the optimal randomized mechanism. One can see that the lower bound is greater than 0.2 for the optimal mechanism, but it seems to converge to a value much lower than 0.5.

### 5 Experiments with Real Data

In this section we investigate the average and worst-sink performances of the the naïve randomized sink (NRS) (Definition 8) mechanism on real datasets of user preferences. Going back to the example of movie selection by a group of friends (Sect. 1), we consider several sizes of the group. A small group consists of tens of friends, while if the decision involves screening a movie at a school auditorium, the group size could easily be in the hundreds. This is why we consider group sizes spanning from 10 to 210 in steps of 50.

A similar situation occurs when a group of people decides which comedian/musician to invite in a social gathering, where they need to pay the cost of bringing the performer. Keeping these motivating situations in mind, we used two datasets that closely represent the scenarios discussed. We used the MovieLens 20M dataset [20] and the Jester dataset [12] to compare the average and worst-case performances of NRS. The first dataset contains preferences for movies,




Fig. 2. Naïve random sink mechanism

while the second contains preferences for online jokes. The MovieLens 20M dataset (ml-20m) describes users' ratings between 1 and 5 stars from MovieLens, a movie recommendation service. It contains 20,000,263 ratings across 27,278 movies. These data were created from the ratings of 138,493 users between January 09, 1995 and March 31, 2015. For our experiment, we sampled the preferences of a specific number of users (shown as agents on the x-axis of Figs. 2 and 3) 100 times uniformly at random from the whole set of users that rated a particular genre of movies, and computed the sample inefficiency on this sampled set and plotted it along with the standard deviation.

The Jester dataset (`jester-data-1`) used in our experiment contains data from 24,983 users who have rated 36 or more jokes, a matrix with dimensions 24983 X 100, and is obtained from Jester, an online joke recommendation system.<sup>8</sup>

Figure 2 shows that the real preferences of users yield much lower expected sample inefficiencies for the naïve randomized sink (NRS) mechanism than the theoretical worst-case guarantee. The improvement ranges from roughly a factor of 5 (for a group size of 10) to almost 100 (for a group size of 210). This also indicates that the rate of decay of the inefficiency with the size of the group is faster than the theoretical guarantee. The bars in Figs. 2 and 3 show the *average* (w.r.t. the randomly selected users) expected sample inefficiency (Eq. 3) and the inefficiency of the worst sink of the NRS mechanism respectively with the standard deviations around them.

By the arguments preceding Theorem 6 and since MIS (Algorithm 3) also picks exactly one sink, it is easy to see that the average inefficiency and inefficiency of the worst sink of MIS will be same as NRS.

## 6 Conclusions and Future Research

In this paper, we considered the classic question of the interplay between efficiency and budget balance, properties that are incompatible with strategyproofness due to the Green-Laffont impossibility result, in the general quasi-linear setting. We proved the limits of possibility in the context of deterministic




Fig. 3. Worst-sink behavior

<sup>8</sup> In both datasets there are missing values because a user has typically not rated all movies/jokes. Before our experiment, we filled the missing values with a random realization of ratings drawn from the empirical distribution for that alternative (movie or joke). The empirical distribution of an alternative is created from the histogram of the available ratings of the users. We cleaned the dataset by keeping only those alternatives that have at least 10 or more available ratings and filled the rest using their empirical distributions.

mechanisms for both efficiency and budget balance. For randomized mechanisms, we identified a class of mechanisms that perform better than deterministic ones. We used an optimization-based scheme to find the optimal randomized mechanism. Experiments with real datasets showed that the values (rate of decay) of inefficiency are significantly smaller (faster) than those of the theoretical worst case. Future research includes studying the structure of the optimal randomized mechanisms that achieve the (theoretical) improved efficiency. Future work also includes investigating the rate of improvement of the optimal bound for a general number of agents.

## References

1. Arrow, K.: The property rights doctrine and demand revelation under incomplete information. In: *Economics and Human Welfare*. New York Academic Press (1979)
2. Ashlagi, I., Fischer, F., Kash, I.A., Procaccia, A.D.: Mix and match: a strategyproof mechanism for multi-hospital kidney exchange. *Games Econ. Behav.* **91**, 284–296 (2015)
3. Bailey, M.J.: The demand revealing process: to distribute the surplus. *Public Choice* **91**(2), 107–126 (1997)
4. Cavallo, R.: Optimal decision-making with minimal waste: strategyproof redistribution of VCG payments. In: *Proceedings of the Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 882–889 (2006)
5. Clarke, E.: Multipart pricing of public goods. *Public Choice* **8**, 19–33 (1971)
6. Conitzer, V., Sandholm, T.: Complexity of mechanism design. In: *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 103–110. Morgan Kaufmann Publishers Inc. (2002)
7. d’Aspremont, C., Gérard-Varet, L.A.: Incentives and incomplete information. *J. Public Econ.* **11**(1), 25–45 (1979)
8. Dobzinski, S., Nisan, N., Schapira, M.: Truthful randomized mechanisms for combinatorial auctions. In: *Proceedings of the Annual ACM Symposium on Theory of Computing*, pp. 644–652. ACM (2006)
9. Faltings, B.: A budget-balanced, incentive-compatible scheme for social choice. In: Faratin, P., Rodríguez-Aguilar, J.A. (eds.) *AMEC 2004. LNCS (LNAI)*, vol. 3435, pp. 30–43. Springer, Heidelberg (2006). doi:[10.1007/11575726\\_3](https://doi.org/10.1007/11575726_3)
10. Feldman, M., Wilf, Y.: Randomized strategyproof mechanisms for facility location, the mini-sum-of-squares objective [arXiv:1108.1762](https://arxiv.org/abs/1108.1762) (2011)
11. Robert, R.J., Jaaidane, T.: Polling mechanisms and the demand revelation problem. *J. Public Econ.* **76**(2), 203–238 (2000)
12. Goldberg, K., Roeder, T., Gupta, D., Perkins, C.: Eigentaste: a constant time collaborative filtering algorithm. *Inf. Retr.* **4**(2), 133–151 (2001)
13. Green, J.R., Laffont, J.-J.: *Incentives in Public Decision Making*. North-Holland, Amsterdam (1979)
14. Groves, T.: Incentives in teams. *Econometrica* **41**(4), 617–631 (1973)
15. Gujar, S.P., Narahari, Y.: Redistribution mechanisms for assignment of heterogeneous objects. *J. Artif. Intell. Res.* **41**, 131–154 (2011)
16. Guo, M., Conitzer, V.: Optimal-in-expectation redistribution mechanisms. In: *Proceedings of the Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1047–1054 (2008)

17. Guo, M., Conitzer, V.: Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games Econ. Behav.* **67**(1), 69–98 (2009)
18. Guo, M., Conitzer, V.: Better redistribution with inefficient allocation in multi-unit auctions. *Artif. Intell.* **216**, 287–308 (2014)
19. Gurobi. Gurobi optimizer reference manual (2015)
20. Harper, F.M., Konstan, J.A.: The movielens datasets: history and context. *ACM Trans. Interact. Intell. Syst. (TiiS)* **5**(4), 19 (2016)
21. Heydenreich, B., Müller, R., Uetz, M., Vohra, R.V.: Characterization of revenue equivalence. *Econometrica* **77**(1), 307–316 (2009)
22. Massó, J., Nicolò, A., Sen, A., Sharma, T., Ülkü, L.: On cost sharing in the provision of a binary and excludable public good. *J. Econ. Theory* **155**, 30–49 (2015)
23. Mishra, D., Sen, A.: Roberts’ theorem with neutrality: a social welfare ordering approach. *Games Econ. Behav.* **75**(1), 283–298 (2012)
24. Mishra, D., Sharma, T.: On the optimality of the green-laffont mechanism. Preliminary Draft (2016)
25. Moulin, H.: Almost budget-balanced VCG mechanisms to assign multiple objects. *J. Econ. Theory* **144**(1), 96–119 (2009)
26. Moulin, H., Shenker, S.: Strategyproof sharing of submodular costs: budget balance versus efficiency. *Econ. Theor.* **18**(3), 511–533 (2001)
27. Myerson, R.B.: Fundamentals of social choice theory. *Q. J. Polit. Sci.* **8**(3), 305–337 (2013)
28. Nath, S., Sandholm, T.: Efficiency and budget balance in general quasi-linear domains. [arXiv:1610.01443](https://arxiv.org/abs/1610.01443) [cs.GT] (2016)
29. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. In: Proceedings of the 10th ACM Conference on Electronic Commerce (EC), pp. 177–186. ACM (2009)
30. Rothkopf, M.H.: Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Oper. Res.* **55**(2), 191–197 (2007)
31. Roughgarden, T., Sundararajan, M.: Quantifying inefficiency in cost-sharing mechanisms. *J. ACM (JACM)* **56**(4), 23 (2009)
32. Sandholm, T.: Automated mechanism design: a new application area for search algorithms. In: Rossi, F. (ed.) CP 2003. LNCS, vol. 2833, pp. 19–36. Springer, Heidelberg (2003). doi:[10.1007/978-3-540-45193-8\\_2](https://doi.org/10.1007/978-3-540-45193-8_2)
33. Thang, N.K.: On randomized strategy-proof mechanisms without payment for facility location games. In: Proceedings of the Web and Internet Economics (WINE), Stanford, USA, pp. 13–16 (2010)
34. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Financ.* **16**(1), 8–37 (1961)

# The Core of the Participatory Budgeting Problem

Brandon Fain<sup>1</sup>(✉), Ashish Goel<sup>2</sup>, and Kamesh Munagala<sup>1</sup>

<sup>1</sup> Department of Computer Science, Duke University, 308 Research Drive,  
Durham, NC 27708, USA

{btfain,kamesh}@cs.duke.edu

<sup>2</sup> Management Science and Engineering Department, Stanford University, Stanford,  
CA 94305, USA

ashishg@stanford.edu

**Abstract.** In *participatory budgeting*, communities collectively decide on the allocation of public tax dollars for local public projects. In this work, we consider the question of fairly aggregating preferences to determine an allocation of funds to projects. We argue that the classic game theoretic notion of core captures fairness in the setting. To compute the core, we first develop a novel characterization of a public goods market equilibrium called the *Lindahl equilibrium*. We then provide the first polynomial time algorithm for computing such an equilibrium for a broad set of utility functions. We empirically show that the core can be efficiently computed for utility functions that naturally model data from real participatory budgeting instances, and examine the relation of the core with the welfare objective. Finally, we address concerns of incentives and mechanism design by developing a randomized approximately dominant-strategy truthful mechanism building on the Exponential Mechanism from differential privacy.

## 1 Introduction

Transparency and citizen involvement are fundamental goals for a healthy democracy. Participatory Budgeting (PB) [4, 22] is a process by which a municipal organization (eg. a city or a district) puts a small amount of its budget to direct vote by its residents. PB is growing in popularity, with over 30 such elections conducted in 2015. Implementing participatory budgeting requires careful consideration of how to aggregate the preferences of community members into an actionable project funding plan. In this work, we model participatory budgeting as a fair resource allocation problem. We note that this problem is different

---

B. Fain—Supported by NSF grants CCF-1637397 and IIS-1447554.

A. Goel—Supported by the Army Research Office Grant No. 116388, the Office of Naval Research Grant No. 11904718, by NSF grant CCF-1637418, and by the Stanford Cyber Initiative.

K. Munagala—Supported by NSF grants CCF-1408784, CCF-1637397, and IIS-1447554.

from standard fair resource allocation because of *public goods*: The allocated goods benefit all users simultaneously. We model this problem as a central body fairly allocating public goods according to preferences reported by the community members (or users), subject to a budget constraint. It is important to note that in participatory democracy, equitable and fair outcomes are an important systemic goal.

*Model of Fairness:* In a participatory budgeting setting, there are  $k$  projects (or items) and  $n$  voters (or agents) who participate. Unlike in a private good economy, it is usually the case that  $k \ll n$ . There is an overall budget  $B$  available for funding projects. An *allocation* is a  $k$ -dimensional vector  $\mathbf{x} \in \mathbb{R}^k$  with  $\mathbf{x} \geq 0$  and  $\sum_{j=1}^k x_j \leq B$ . The quantity  $x_j$  denotes the funding for project  $j$ . We assume voters report a cardinal *utility function*. We denote the utility of an agent  $i$  given an allocation  $\mathbf{x}$  as  $U_i(\mathbf{x})$ , and we assume this function is continuous, non-decreasing, and concave.

In this model, we study *fair allocations*. In this paper, the concept of fairness with which we work is the core. This notion is borrowed from cooperative game theory and was first phrased in game theoretic terms in [25]. It has been studied extensively even in public goods settings [8, 19].

**Definition 1.** *An allocation  $\mathbf{x}$  is a core solution if there is no subset  $S$  of agents who, given a budget of  $(|S|/n)B$ , could compute an allocation  $\mathbf{y}$  where every user in  $S$  receives strictly more utility in  $\mathbf{y}$  than  $\mathbf{x}$ , i.e.,  $\forall i \in S, U_i(\mathbf{y}) > U_i(\mathbf{x})$ .*

**Definition 2.** *For  $\alpha > 1$ , an allocation  $\mathbf{x}$  lies in the  $\alpha$ -approximate multiplicative (resp. additive) core if for any subset  $S$  of agents, there is no allocation  $\mathbf{y}$  using a budget of  $(|S|/n)B$ , s.t.  $U_i(\mathbf{y}) > \alpha U_i(\mathbf{x})$  (resp.  $U_i(\mathbf{y}) > U_i(\mathbf{x}) + \alpha$ ) for all  $i \in S$ .*

Note that when  $S = \{1, 2, \dots, n\}$ , the above constraints encode (weak) Pareto-Efficiency. Further, when  $S$  is a singleton voter, the core captures *Sharing Incentive*, meaning that the voter gets at least as much utility as she would have obtained with budget  $B/n$  dedicated to just her. In general, the core captures a *group sharing incentive*: No community of users suffers envy with respect to its share of the overall budget.

*Some Clarifying Examples:* We briefly consider some examples to clarify the concept of the core and compare it with other definitions of fairness. For simplicity in these examples, assume the utility function of the agents is linear, so  $U_i(\mathbf{x}) = \sum_{j=1}^k u_{ij}x_j$ . Also, assume that there is a unit size budget and all projects are of unit size. First note that the core will produce a very different outcome from approval voting. In the example in Fig. 1(a), the majority has one more vote than the minority, yet they are exclusively privileged by approval voting (funding projects in order of number of votes). The remaining examples compare the core with other fair allocations. Figure 1(b) shows that the naive fair allocation to allow every agent to determine  $1/n$  of the overall allocation is not Pareto-efficient. In Fig. 1(c), while a max-min fair allocation favors one voter



at the expense of all others, the core solution funds items *in proportion* to the number of voters preferring them.

*High-Level Goals:* At a high level, we explore three related questions in Sects. 2, 3, and 4 respectively: (i) Can we efficiently compute core allocations for reasonably general utility functions? (ii) What do these allocations look like for data generated by real participatory budgeting instances under utility functions motivated by that data? (iii) For simple utility functions, can we develop a truthful mechanism for computing core allocations without payments?

We positively answer the first and third question using techniques from optimization and differential privacy to develop the algorithmic understanding of the Lindahl Equilibrium; a market based notion we will define shortly. For the second question, we use our theoretical results to develop principled heuristics that we validate using real voting data. Before proceeding however, we turn to consider utility functions more precisely.

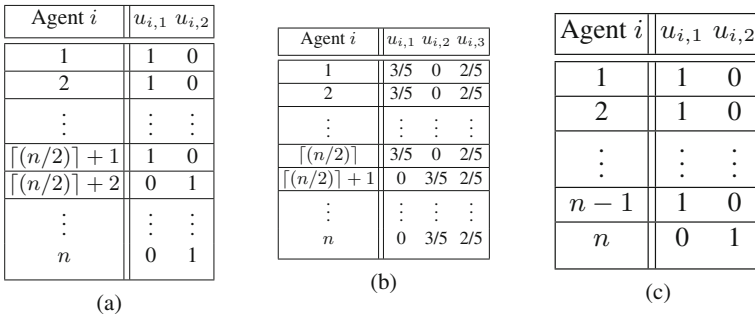


Fig. 1. Some examples clarifying the concept of core.

### 1.1 Utility Functions

We consider utility functions generalizing the linear utility functions used in previous examples. These utility functions, which we term SCALAR SEPARABLE, have the form

$$U_i(\mathbf{x}) = \sum_j u_{ij} f_j(x_j)$$

for every agent  $i$  where  $\{f_j\}$  are smooth, non-decreasing, and concave, and  $\mathbf{u}_i \geq 0$ . By  $\sum_j$  we always mean the sum over the  $k$  projects. SCALAR SEPARABLE utilities are fairly general and well-motivated. First, this concept encompasses linear utilities and several other canonical utility functions (see below). Secondly, if voters express scalar valued preferences (such as up/down approval voting), SCALAR SEPARABLE utilities provide a natural way of converting these votes into cardinal utility functions. In fact, as we discuss below, we will do precisely this when handling real data. We consider two subclasses that we term NON-SATIATING and SATURATING utilities respectively. Each arises naturally in settings related to participatory budgeting.

*Non-satiating Functions:* For our main computational result in Sect. 2, we consider a subclass of utility functions that we term NON-SATIATING.

**Definition 3.** A differentiable, strictly increasing, concave function  $f$  is called NON-SATIATING if  $xf'_j(x)$  is monotonically increasing and equal to 0 when  $x = 0$ .

This is effectively a condition that the functions grow at least as fast as  $\ln x$ . Several utility functions used for modeling substitutes and complements fall in this class. For instance, constant elasticity of substitution (CES) utility functions where

$$U_i(\mathbf{x}) = \left( \sum_j u_{ij} x_j^\rho \right)^{\frac{1}{\rho}} \quad \text{for } \rho \in (0, 1]$$

can be monotonically transformed into NON-SATIATING utilities.<sup>1</sup> CES functions are also *homogeneous of degree 1*, meaning that  $U_i(\alpha\mathbf{x}) = \alpha U_i(\mathbf{x})$  for any scalar  $\alpha \geq 0$ . When  $\rho = 1$ , this captures linear utilities. When  $\rho \rightarrow 0$ , these are COBB-DOUGLAS utilities which for  $\alpha_{ij} > 0$  such that  $\sum_j \alpha_{ij} = 1$ , can be written as  $U_i(\mathbf{x}) = \prod_j x_j^{\alpha_{ij}}$ .

*Saturating Functions:* Note that SCALAR SEPARABLE utilities assume projects are divisible. Fractional allocations make sense in their own right in several scenarios: Budget allocations between goals such as defense and education at a state or national level are typically fractional, and so are allocations to improve utilities such as libraries, parks, gyms, roads, etc. However, in the settings for which we have real data, the projects are indivisible and have a monetary cost  $s_j$ , so that we have the additional constraint  $x_j \in \{0, s_j\}$  on the allocations. We describe such data from the Stanford Participatory Budgeting Platform [21] in greater detail in Sect. 3. We therefore need utility functions that model budgets in individual projects. These utility functions must also be simple to account for the limited information elicited in practice. For example, in the voting data that we use in our experiments, each voter receives an upper bound on how many projects she can select, and the ballot cast by a voter is simply the subset of projects she selects. A related voting scheme implemented in practice, called Knapsack Voting [10], has similar elicitation properties. For modeling these two considerations, we consider SATURATING utilities.

**Definition 4.** SATURATING utility functions  $U_i(\mathbf{x})$  have the form  $\sum_j u_{ij} \min\left(\frac{x_j}{s_j}, 1\right)$ .

For converting our voting data into a SATURATING utility, we set  $s_j$  to be the budget of project  $j$ , and set  $u_{ij}$  to 1 if agent  $i$  votes for project  $j$  and 0 otherwise. Note that if  $x_j = s_j$ , then the utility of any agent who voted for this item is 1. This implies the total utility of an agent is the number (or total fraction) of items that he voted for that are present in the final allocation. Clearly, SATURATING

---

<sup>1</sup> Note that the core remains unchanged if utilities undergo a monotone transform.

utilities do not satisfy Definition 3. However, we will connect NON-SATIATING and SATURATING utilities by developing an approximation algorithm and heuristic for computing core allocations in the saturating model using results developed for the NON-SATIATING model.

## 1.2 Computing Core Solutions via the Lindahl Equilibrium

In a fairly general public goods setting, there is a market based notion of fairness due to Lindahl [15] and Samuelson [24] termed the *Lindahl equilibrium*, which is based on setting different prices for the public goods for different agents. The market on which the Lindahl equilibrium is defined is a mixed market of public and private goods. We present a definition below that is specialized to just a public goods market relevant for participatory budgeting.

**Definition 5.** *In a public goods market with budget  $B$ , per-voter prices  $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n$  each in  $\mathbb{R}_+^k$  and allocation  $\mathbf{x} \in \mathbb{R}_+^k$  constitute a **Lindahl equilibrium** if the following two conditions hold:*

1. *For every agent  $i$ , the utility  $U_i(\mathbf{y}_i)$  is maximized subject to  $\mathbf{p}_i \cdot \mathbf{y}_i \leq B/n$  when  $\mathbf{y}_i = \mathbf{x}$ ; and*
2. *The **profit** defined as  $(\sum_i \mathbf{p}_i) \cdot \mathbf{z} - \|\mathbf{z}\|_1$ , subject to  $\mathbf{z} \geq 0$  is maximized when  $\mathbf{z} = \mathbf{x}$ .*

The price vector for every agent is traditionally interpreted as a tax. However, unlike in private goods markets, in our case these prices (or taxes) are purely hypothetical; we are only interested in the allocation that results at equilibrium (in fact, we eliminate the prices from our characterization of the equilibrium). Under innocuous conditions for the mixed public and private goods market, Foley proved that the Lindahl equilibrium exists and lies in the core [8]. This remains true in our specialized instance of the problem; the omitted proof is a trivial adaption from [8]. Thus, computing a Lindahl equilibrium is sufficient for the purpose of computing a core allocation. However, Foley only proves existence of the equilibrium via a fixed point argument that does not lend itself to efficient computation.

## 1.3 Our Results

In Sect. 2, we present a simple characterization of the Lindahl equilibrium in terms of the allocation variables and a means of efficient computation for NON-SATIATING utilities. Together, this results in an efficient algorithm for computing the core exactly for NON-SATIATING utilities via convex programming. As far as we are aware, this is the *first* non-trivial computational result for the Lindahl equilibrium.

As a consequence of our characterization, if the utility functions are homogeneous of degree 1 and concave (or any monotone transform thereof), then the *proportionally fair* allocation, the extension of the Nash Bargaining solution [20])

that maximizes  $\sum_i \log U_i(\mathbf{x})$ , computes the Lindahl equilibrium. This mirrors similar results for computing a Fisher equilibrium in private good markets [12]. In addition, we show that for homogeneous functions, quadratic voting [14] can be used to elicit the gradient of the proportional fairness objective, pointing to practical implementations in the field. For more general utility functions, our potential function can be viewed as a regularized version of the proportional fairness objective written on a non-linear transform of the utility function – a result that is new to the best of our knowledge. We also note that the class of NON-SATIATING utilities includes many functions that are not monotone transforms of homogeneous functions of degree 1, and for some of these functions, computing a Fisher equilibrium is intractable [27].

In Sect. 3, we consider the question of computing core solutions for real world data sets from the Stanford Participatory Budgeting Platform [21] that we model using SATURATING utility functions as discussed in Sect. 1.1. We present an approximation algorithm as well as a heuristic implementation inspired by our characterization. On real data, we find that this heuristic efficiently computes the exact core. Surprisingly, the resulting outcomes match the welfare optimal solutions on the same utility functions.

In Sect. 4, we address incentive concerns. Truthfulness has long been considered a serious problem for the allocation of public goods [11, 19]. We study *asymptotic approximate truthfulness* [16]. Truthfulness means that for any agent  $i$ , reporting the true utility function  $U_i(\mathbf{x})$  maximizes the expected utility of agent  $i$ , subject to all agents  $i' \neq i$  reporting true utility functions and agent  $i$  knowing these. Our notion is asymptotic in the sense that  $n \gg k$ , which is reasonable in practice. We show that when agents' utilities are linear (and more generally, homogeneous of degree 1), there is an efficient randomized mechanism that implements an  $\epsilon$ -approximate core solution as a dominant strategy for large  $n$ . We use the Exponential Mechanism [18] from differential privacy to achieve this. The application of the Exponential Mechanism is not straightforward since the proportional fairness objective (that computes the Lindahl equilibrium) is not separable when used as a scoring function; the allocation variables are common to all agents. Furthermore, this objective varies widely when one agent misreports utility. We define a scoring function directly based on the *gradient condition* of proportional fairness to circumvent this hurdle.

Most proofs are omitted throughout and can be found in the full version of the paper [7]. We also relegate many of the details in Sect. 4 to the full version.

## 1.4 Related Work

The general literature characterizing private good market equilibrium and computation is extensive [12, 13, 26, 27]; however, there is relatively little known about computational results for public goods. The proportional fairness algorithm, which has been extensively studied in private good markets [12, 20], need not find solutions in the core for SCALAR SEPARABLE utilities, and we can view our computational results as providing a non-trivial generalization of the proportional fairness concept to Lindahl equilibria.

Our work is related to designing truthful mechanisms for combinatorial public projects [6]. However, these works focused on the social welfare objective and utilized payments as does the well known VCG mechanism [5, 28], which is impractical for the application of participatory budgeting. Though public good markets are truthful in the Bayesian sense in the large market limit [1–3] because they are envy-free by definition, we seek dominant strategy truthful mechanisms. These are non-trivial to design for public good markets even in the large market limit. The problem of truthful allocation of public goods without payments is considered in the context of the facility location problem in [23]; however, the setting is unrelated to ours and the authors are concerned with the social welfare or the total dis-utility, not the core.

## 2 Non-Satiating Utilities: Characterization and Computation

Recall that in the participatory budgeting problem, there are  $k$  items (or projects) and  $n$  agents (or voters). It is typically the case that  $k \ll n$ . We will denote a generic voter by  $i$  and a generic item by  $j$ . There is an overall budget of  $B$ . An *allocation* is a  $k$ -dimensional vector  $\mathbf{x} \in \mathbb{R}^k$  with  $\mathbf{x} \geq 0$  and  $\sum_{j=1}^k x_j \leq B$ . We consider scalar separable utility, where the utility of an agent  $i$  given an allocation  $\mathbf{x}$  is  $U_i(\mathbf{x}) = \sum_j u_{ij} f_j(x_j)$ , where  $\{f_j\}$  are smooth, non-decreasing, and concave, and  $\mathbf{u}_i \geq 0$ .

### 2.1 Characterization

Recall that in order to compute a core allocation it is sufficient to compute a Lindahl equilibrium (Definition 5). Our first result is a characterization of the Lindahl equilibrium that uses the optimality conditions to eliminates the price variables entirely.

**Theorem 1.** *An allocation  $\mathbf{x} \geq 0$  corresponds to a Lindahl equilibrium if and only if*

$$\sum_i \left( \frac{u_{ij} f'_j(x_j)}{\sum_m u_{im} x_m f'_m(x_m)} \right) \leq \frac{n}{B} \tag{1}$$

for all items  $j$ , where this inequality is tight when  $x_j > 0$ .

### 2.2 Efficient Computation

We now present our main computational result that builds on the characterization above to give the first non-trivial poly-time method for computing the Lindahl equilibrium. We need the non-satiation assumption on the functions  $\{f_j\}$  given in Definition 3.

**Theorem 2.** *When  $U_i(\mathbf{x}) = \sum_j u_{ij} f_j(x_j)$  where  $\{f_j\}$  satisfy Definition 3, the Lindahl equilibrium (and therefore a core solution) is the solution to a convex program.*

*Proof.* Recall the characterization of the Lindahl equilibrium from Theorem 1. Define  $z_j = x_j f'_j(x_j)$ . Note that  $x_j = 0$  iff  $z_j = 0$ . Since  $f_j$  satisfies non-satiation, this function is continuous and monotonically increasing, and hence invertible. Let  $h_j$  be this inverse such that  $h_j(z_j) = x_j$ . Let  $r_j(z_j) = h_j(z_j)/z_j = 1/f'_j(x_j)$ . The Lindahl equilibrium characterization therefore simplifies to:

$$\sum_i \left( \frac{u_{ij}}{\sum_m u_{im} z_m} \right) \leq \frac{n}{B} r_j(z_j)$$

with the inequality being tight when  $z_j > 0$ . Let  $R_j(z_j)$  be the indefinite integral of  $r_j$  (with respect to  $z_j$ ). Define the following potential function

$$\Phi(\mathbf{z}) = \sum_i \log \left( \sum_j u_{ij} z_j \right) - \left( \frac{n}{B} \right) \sum_j R_j(z_j) \tag{2}$$

We claim that  $\Phi(\mathbf{z})$  is concave in  $\mathbf{z}$ . The first term in the summation is trivially concave. Also, since  $f'_j(x_j)$  is a decreasing function,  $1/f'_j(x_j)$  is increasing in  $x_j$ . Since  $r_j(z_j) = 1/f'_j(x_j)$ , this is increasing in  $x_j$  and hence in  $z_j$ . This implies  $R_j(z_j)$  is convex, showing the second term in the summation is concave as well. It is easy to check that the optimality conditions of maximizing  $\Phi(\mathbf{z})$  subject to  $\mathbf{z} \geq 0$  are exactly the conditions for the Lindahl equilibrium. This shows that the Lindahl equilibrium corresponds to the solution to the convex program maximizing  $\Phi(\mathbf{z})$ .  $\square$

As we show in the full paper [7], an approximately optimal solution to the convex program gives an approximate core solution, which implies polynomial time computation to arbitrary accuracy. We note that the non-satiation condition essentially implies that  $f_j(x_j)$  should grow faster than  $\ln x_j$ . In combination with the assumption that  $f_j(x_j)$  is concave, this leaves us with a broad class of concave functions for which the Lindahl equilibrium and hence the core can be efficiently computed.

### 2.3 Connection to Proportional Fairness

The following is now a simple corollary of Theorem 2.

**Corollary 1.** *If  $U_i(\mathbf{x})$  is linear, i.e.,  $U_i(\mathbf{x}) = \sum_j u_{ij} x_j$ , or more generally, if it is homogeneous of degree 1, then the Lindahl equilibrium coincides with the proportionally fair allocation that maximizes  $\sum_i \log U_i(\mathbf{x})$  subject to  $\|\mathbf{x}\|_1 \leq B$  and  $\mathbf{x} \geq 0$ .*

The proof for the linear case is direct, and that for homogeneous functions uses a standard change of variables [13] and is omitted. As mentioned in Sect. 1.1, an interesting special case of homogeneous functions concerns Cobb-Douglas utilities, where  $U_i(\mathbf{x}) = \prod_j x_j^{\alpha_{ij}}$  where  $\sum_j \alpha_{ij} = 1$  and  $\alpha_{ij} > 0$ . In this case, if a single agent could choose the whole allocation, the optimal choice would be

$x_j = \alpha_j B$ . Suppose every agent  $i$  reveals these optimal allocations for themselves for every item  $j$ ; call this  $x_{ij}$ . Then it is easy to check that the Lindahl equilibrium sets  $x_j = \frac{1}{n} \sum_i x_{ij}$ , which is simply the average of the individual monetary allocations.

*Elicitation via Quadratic Voting.* For homogeneous functions, it is easy to show that the gradient of the proportional fairness objective in direction  $x_j$  is given by:

$$\sum_i \frac{\frac{\partial}{\partial x_j} U_i(\mathbf{x})}{\sum_m x_m \frac{\partial}{\partial x_m} U_i(\mathbf{x})} - \frac{n}{B}$$

Suppose users  $i$  are drawn from some large population, and we were to perform stochastic gradient descent by sampling a random user  $i$ , and estimating the gradient. The above expression shows that this needs estimating the relative magnitudes of  $\left\{ \frac{\partial}{\partial x_m} U_i(\mathbf{x}) \right\}$ , since these terms are present both in the numerator and denominator. The relative magnitudes at any point  $\mathbf{x}_t$  can be estimated by presenting user  $i$  with a ball of radius  $\epsilon$  around  $\mathbf{x}_t$  and asking the user to maximize her utility,  $U_i(\mathbf{x})$ . This is termed *quadratic voting* [14], and gives a way to *elicit* enough information from individual voters in order to perform stochastic gradient descent and compute the proportionally fair allocation.

*Beyond Proportional Fairness.* When the utility functions are not homogeneous, it is not clear how to express the potential function in Eq. (2) as running proportional fairness on a transformed space of allocations. For instance, if  $U_i(\mathbf{x}) = \sum_j u_{ij} x_j^{\alpha_j}$ ,

$$\Phi(\mathbf{x}) = \sum_i \log \left( \sum_j \alpha_j u_{ij} x_j^{\alpha_j} \right) - \frac{n}{B} \sum_j \alpha_j x_j$$

This involves a non-linear transform of the utility function and a regularization term, which proportional fairness on any transformed input space does not capture. We also observe that running proportional fairness directly can be far away from the core. Consider an instance where agents are partitioned into groups  $G_j$  where all agents in a group have non-zero utility for only item  $j$ , with utility function  $u_{ij} f_j(x_j) = x_j^{\alpha_j}$  for some  $\alpha_j \in (0, 1)$ . Since all groups have disjoint preferences, the core solution allocates  $x_j$  in proportion to  $|G_j|$ . However, proportional fairness maximizes  $\sum_j |G_j| \log(x_j^{\alpha_j}) = \sum_j \alpha_j |G_j| \log x_j$ , which allocates  $x_j$  in proportion to  $\alpha_j |G_j|$ .

### 3 Saturating Utilities: Approximation and Experiments

We now move to the question of modeling and analyzing real participatory budgeting data. We use data from seven different elections that used the Stanford

Participatory Budgeting Platform (SPBP). This platform (<http://pbstanford.org>) [9, 10] has been used by over 25 PB elections for digital voting and incorporates multiple voting mechanisms including  $K$ -approval, knapsack, ranking, and comparisons.

Voters are presented with a ballot containing descriptions of the candidate public projects with associated budgets as well as an overall budget. They can vote for at most a certain number of these projects, typically 4 or 5 (this voting method is called  $K$ -approval). Note that the projects chosen by a voter can exceed the total budget. The data set is therefore a 0/1 matrix on projects and voters, where a 1 denotes a vote by the voter for the project. The number of voters,  $n$ , ranges between 200 and 3000 in our datasets, and the number of items  $k$  is at most 30. A typical example is presented in Table 1.

For modeling such data, we need utility functions that respect the budget constraints of individual projects. It is natural to use the SATURATING utility model (see Sect. 1.1), where the utility of user  $i$  is

$$U_i(\mathbf{x}) = \sum_j u_{ij} \min(x_j/s_j, 1)$$

where  $s_j$  is the budget of project  $j$ , and  $u_{ij}$  is 1 if  $i$  votes for  $j$  and 0 otherwise. Therefore, the utility for  $i$  if  $j$  is chosen in the final allocation is  $u_{ij} \in \{0, 1\}$ . Clearly, this function does not satisfy Definition 3. We first show that an approximation to the core (see Definition 2) can be efficiently computed using a NON-SATIATING relaxation of the utility model. The proof is relegated to the full version [7].

**Theorem 3.** *Given a collection of SATURATING utility functions, let  $s = \min_j s_j$ . Then, for any  $\epsilon > 0$ , an  $\alpha$ -approximate multiplicative core can be efficiently computed, where  $\alpha = (1/\epsilon)(B/s)^\epsilon + 1 - 1/\epsilon$ . For  $\epsilon = \log(B/s)$ , we have a  $O(\log \frac{B}{s})$  approximation to the core.*

### 3.1 Heuristically Computing the Exact Core

We now show an even stronger result empirically: We can efficiently compute the *exact* core solutions under the SATURATING utility model on our real-world data sets. We conclude this section with some observations on the relationship between welfare maximizing and core allocations in the saturating model.

Our heuristic crucially uses the characterization in Theorem 1. Let  $x_j \in [0, s_j]$  denote the current allocation to item  $j$ , and let  $y_j = f'_j(x_j)$ . The following complementarity condition relates  $x_j$  and  $y_j$ :

$$\forall j, \quad y_j \leq \frac{1}{s_j} \quad \text{and} \quad x_j < s_j \quad \Rightarrow \quad y_j = \frac{1}{s_j}$$

The Lindahl equilibrium condition in Theorem 1 can be written as:

$$\forall j, \quad \frac{B}{n} \sum_i \frac{u_{ij} y_j}{\sum_m u_{im} x_m y_m} \leq 1$$



with equality when  $x_j > 0$ . Given  $\mathbf{x}_{-j}$  and  $\mathbf{y}_{-j}$ , we perform binary search on  $x_j, y_j$  to satisfy the above non-linear equation subject to complementarity on  $x_j, y_j$ . We repeat this process, at each step choosing that item  $j$  with the largest additive violation in the above inequality. We iterate until the Lindahl conditions for all items are satisfied to accuracy  $\epsilon$ . (e.g.,  $\epsilon = 1/n$ ). As we show in the full version of the paper, if this process converges, the result is an  $\epsilon$ -approximate additive core solution.

One issue is that these dynamics are not theoretically guaranteed to converge. Even empirically, there are instances where we observe cycling. To address this issue, we perturb the vote matrix by small additive noise, so that  $u_{ij} \leftarrow u_{ij} + \text{Uniform}(0, \alpha)$ , where  $\alpha$  is a small constant like  $1/k^2$ . We empirically observe that the process now converges. In Fig. 2, we show this behavior for three datasets with at least 2000 voters and 10 items each. The convergence is comparable for all seven of our data sets; only three are shown for the sake of readability.

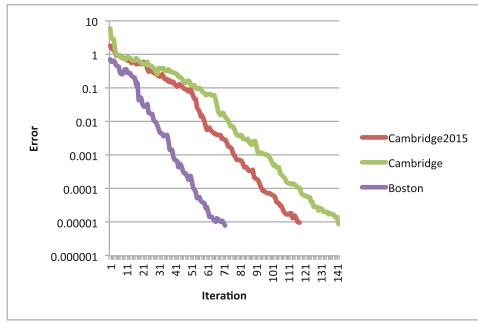


Fig. 2. Plot of error  $\epsilon$  in Lindahl conditions as a function of number of iterations.

**Observation 1.** *Despite lacking a theoretical guarantee of convergence for SATURATING utilities, we are able to consistently compute near-exact core solutions for our data sets using binary search on the complementarity conditions.*

### 3.2 Comparing the Core with WELFARE

Given that we can compute the core exactly, we investigate its structure on our datasets. For item  $j$ , let  $n_j$  denote the number of votes received. Recall that these votes come from simple approval voting and that  $s_j$  is the budget (cost) of the project. We define the following vote aggregation schemes that we will use for comparison. We can compute fractional welfare and core allocations straightforwardly, but the final allocation needs to be integral. To do this, the items are sorted in a certain order. Consider items in this order and add the item if its budget is less than the remaining total budget, stopping when all items are exhausted. Importantly, both aggregation schemes use the same utility model.

- CORE: Compute a fractional core allocation as described in Sect. 3.1. Let  $x_j$  denote the fractional allocation of item  $j$ . Sort the items in descending order of order  $\frac{x_j}{s_j}$ , which is the fraction to which item  $j$  is funded in the fractional allocation.
- WELFARE: Sort the items in descending order of  $\frac{n_j}{s_j}$ . This is the allocation that maximizes total (fractional) utility in the SATURATING utility model (Definition 4).

We compare the outcomes of these algorithms for data sets from seven different real world instances of participatory budgeting. We consider two measures of the similarities of outcomes: the Jaccard index and Budget similarity. The Jaccard index for two integral allocations is the ratio of the size of their intersection to the size of their union. The Budget similarity for two fractional allocations  $\mathbf{x}$  and  $\mathbf{z}$  is defined as  $\frac{\sum_j \min(x_j, z_j)}{B}$ . Here,  $\mathbf{x}$  is the actual monetary amount allocated to the project in the fractional allocation. Our results are shown in Fig. 3, and one example is presented in Table 1.

**Observation 2.** *CORE and WELFARE compute the same integer allocations on almost all of our data sets, showing WELFARE produces fair allocations in practice. Furthermore, since the fractional allocation produced by WELFARE is an integer allocation except for one item, the high Budget similarity between WELFARE and CORE implies that the fractional core produces almost integer allocations.*




Fig. 3. Similarity Scores for CORE vs WELFARE.

The above observation that the CORE empirically coincides with WELFARE is quite surprising. It is easy to construct examples where the core allocation will be very different from welfare maximization. This is particularly pronounced when there is a significant minority of voters who have orthogonal preferences from the majority. One possible explanation for our observation is that users might have approximately independent random preferences over the projects. We explore this possibility more formally in the full version [7], where we show that this is not the complete explanation.

**Table 1.** Aggregation results for Boston. The Budget column lists the project’s budget in dollars. The final two columns list the allocation of the project as a fraction of its budget, so that an integral allocation corresponds to 1.

| Project                               | Budget    | Votes | CORE | WELFARE |
|---------------------------------------|-----------|-------|------|---------|
| Wicked Free Wifi 2.0                  | \$119,000 | 2,054 | 1.00 | 1.00    |
| Water Bottle Refill Stations at Parks | \$260,000 | 1,794 | 1.00 | 1.00    |
| Hubway Extensions                     | \$101,600 | 737   | 1.00 | 1.00    |
| Bowdoin St. Roadway Resurfacing       | \$100,000 | 611   | 1.00 | 1.00    |
| Bike Lane Installation                | \$200,000 | 771   | 0.74 | 1.00    |
| Track at Walker Park                  | \$240,000 | 672   | 0.33 | 0.91    |
| BCYF HP Dance Studio Renovation       | \$286,000 | 759   | 0.31 | 0.00    |
| BLA Gym Renovations                   | \$475,000 | 1,044 | 0.20 | 0.00    |
| Ringer Park Renovation                | \$280,000 | 546   | 0.02 | 0.00    |
| Green Renovation for BCYF Pino        | \$250,000 | 452   | 0.01 | 0.00    |

### 4 Homogeneous Utilities: Mechanism Design

In this section, we develop a randomized mechanism that finds an approximately core solution with high probability while ensuring approximate dominant-strategy truthfulness for all agents. In the spirit of [16], we assume the large market limit so that  $n \gg k$ ; in particular, we assume  $k = o(\sqrt{n})$ . We present the mechanism for linear utility functions where  $U_i(\mathbf{x}) = \sum_{j=1}^k u_{ij}x_j$ , noting that it easily generalizes to degree one homogeneous functions. The values of  $u_{ij}$  are reported by the agents. Without loss of generality, these are normalized so that  $\|\mathbf{u}_i\|_1 = 1$ . Also without loss of generality, let  $B$  be normalized to 1. Recall from Corollary 1 that for linear utility functions, the proportional fairness algorithm that maximizes  $\sum_i \log U_i(\mathbf{x})$  subject to  $\|\mathbf{x}\|_1 \leq 1$  and  $\mathbf{x} \geq 0$  computes the Lindahl equilibrium.

We will design additive approximations to the core (see Definition 2) that achieve approximate truthfulness in an additive sense. We use the Exponential Mechanism [18] to achieve approximate truthfulness.

Fix a constant  $\gamma \in (0, 1)$  to be chosen later. We first define the convex set of feasible allocations as  $\mathcal{P} := \{\mathbf{x} : \mathbf{x} \geq n^{-\gamma}, \|\mathbf{x}\|_1 \leq 1\}$ . Note that all such allocations are restricted to allocating at least  $n^{-\gamma}$  to each project. Since the utility vector of any agent is normalized so  $\|\mathbf{u}_i\|_1 = 1$ , this implies that every agent gets a baseline utility of at least  $n^{-\gamma}$ , a fact we use frequently. We define the following scoring function, which is based on the gradient optimality condition of Proportional Fairness:

$$q(\mathbf{x}) := n - n^{-\gamma} \max_{\mathbf{y} \in \mathcal{P}} \left( \sum_i \frac{U_i(\mathbf{y})}{U_i(\mathbf{x})} \right)$$

We will approximately and truthfully maximize this scoring function. The trade off in defining the scoring function is between reducing the sensitivity of the function to the report of an individual agent and thus improving the approximation to truthfulness, and having just enough sensitivity so that the mechanism defined in terms of the scoring function provides a good approximation to the core. This scoring function, derived from the gradient condition of the proportional fairness program, provides this balance. There are several details, which we present in [7]. The formal mechanism is defined below.

**Definition 6.** *Define  $\mu$  to be a uniform probability distribution over all feasible allocations  $\mathbf{x} \in \mathcal{P}$ . For a given set of utilities, let the mechanism  $\zeta_q^\epsilon$  be given by the rule:*

$$\zeta_q^\epsilon := \text{choose } \mathbf{x} \text{ with probability proportional to } e^{\epsilon q(\mathbf{x})} \mu(\mathbf{x})$$

The primary result of this section demonstrates that  $\zeta_q^\epsilon$  can find an approximate core solution while providing approximate truthfulness.

**Theorem 4.**  *$\zeta_q^\epsilon$  is  $(e^{2\epsilon} - 1)$ -approximately truthful. Furthermore, if  $k$  is  $o(\sqrt{n})$  and  $\frac{1}{\epsilon} > \frac{kn}{(n-k^2) \ln n}$  then  $\zeta_q^\epsilon$  can be used to choose an allocation  $\mathbf{x}$  that is an  $O\left(\frac{k \ln n}{\epsilon \sqrt{n}}\right)$ -approximate additive core solution w.p.  $1 - \frac{1}{n}$ .*

It can be shown that  $e^{\epsilon q(\mathbf{x})} \mu(\mathbf{x})$  is log-concave, so that  $\zeta_q^\epsilon$  can be sampled in polynomial time [17] with small additive error in truthfulness.

## 5 Conclusion

In this paper, we have initiated the computational study of the Lindahl equilibrium in order to address fair resource allocation in the context of participatory budgeting. Our work is just the first step towards understanding participatory budgeting specifically and the fair allocation of public goods more generally. We do not yet understand the computational complexity for more general utility functions. Is computing the Lindahl equilibrium for public goods computationally hard or is there a polynomial time algorithm even without the non-satiating assumption? Our experimental results leave open intriguing questions about modeling real voting data. In particular, is there a more formal explanation of why welfare appears fair in practice? Also, is there a different way to elicit more information from voters for a more precise modeling of their utility than just approval voting?

**Acknowledgements.** We thank Anilesh Krishnaswamy for useful discussions, and the Stanford Crowdsourced Democracy Team for the use of their data.

## References

1. Azevedo, E.M., Budish, E.B.: Strategy-proofness in the large. In: Chicago Booth Research Paper, pp. 13–35 (2013)
2. Azevedo, E.M., Budish, E.: Strategyproofness in the large as a desideratum for market design. In: EC, pp. 55–55 (2012)
3. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. *J. Polit. Econ.* **119**(6), 1061–1103 (2011)
4. Cabannes, Y.: Participatory budgeting: a significant contribution to participatory democracy. *Environ. Urbanization* **16**(1), 27–46 (2004)
5. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* **11**, 17–33 (1971)
6. Dughmi, S.: A truthful randomized mechanism for combinatorial public projects via convex optimization. CoRR abs/1103.0041 (2011)
7. Fain, B., Goel, A., Munagala, K.: The core of the participatory budgeting problem. arXiv e-prints CoRR:1610.03474 (2016). <https://arxiv.org/pdf/1610.03474v2.pdf>
8. Foley, D.K.: Lindahl’s solution and the core of an economy with public goods. *Econometrica* **38**(1), 66–72 (1970)
9. Goel, A., Krishnaswamy, A.K., Sakshuwong, S.: Budget aggregation via knapsack voting: welfare-maximization and strategy-proofness. In: Collective Intelligence (2016). [http://web.stanford.edu/~anilesh/publications/GoKrSu\\_budget.pdf](http://web.stanford.edu/~anilesh/publications/GoKrSu_budget.pdf)
10. Goel, A., Krishnaswamy, A.K., Sakshuwong, S., Aitamurto, T.: Knapsack voting. In: Collective Intelligence (2015). [http://web.stanford.edu/~anilesh/publications/knapsack\\_voting.pdf](http://web.stanford.edu/~anilesh/publications/knapsack_voting.pdf)
11. Groves, T., Ledyard, J.: Optimal allocation of public goods: a solution to the “free rider” problem. *Econometrica* **45**(4), 783–809 (1977)
12. Jain, K., Vazirani, V.V.: Eisenberg-gale markets: algorithms and structural properties. In: Proceedings of the ACM Symposium Theory of Computing, STOC 2007, pp. 364–373 (2007)
13. Jain, K., Vazirani, V.V., Ye, Y.: Market equilibria for homothetic, quasi-concave utilities and economies of scale in production. In: Proceedings of the ACM-SIAM Symposium Discrete Algorithms, SODA 2005, pp. 63–71 (2005)
14. Lalley, S.P., Weyl, E.G.: Quadratic voting. Available at SSRN 2003531 (2015)
15. Lindahl, E.: Just taxation: a positive solution. In: Musgrave, R.A., Peacock, A.T. (eds.) *Classics in the Theory of Public Finance*. Palgrave Macmillan, Basingstoke (1958)
16. Liu, Q., Pycia, M.: Ordinal efficiency, fairness, and incentives in large markets (2012)
17. Lovász, L., Vempala, S.: The geometry of logconcave functions and sampling algorithms. *Random Struct. Algorithms* **30**(3), 307–358 (2007). <http://dl.acm.org/citation.cfm?id=1234701>
18. McSherry, F., Talwar, K.: Mechanism design via differential privacy. In: Annual IEEE Symposium on Foundations of Computer Science (FOCS) (2007)
19. Muench, T.J.: The core and the lindahl equilibrium of an economy with a public good: an example. *J. Econ. Theory* **4**(2), 241–255 (1972)
20. Nash, J.F.: The bargaining problem. *Econometrica* **18**(2), 155–162 (1950)
21. PB-Stanford: Stanford participatory budgeting platform. Website (2015). <https://pbstanford.org>
22. PBP: Where has it worked? - the participatory budgeting project (2016). <http://www.participatorybudgeting.org/about-participatory-budgeting/where-has-it-worked/>

23. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. In: Proceedings of the 10th ACM EC, EC 2009, pp. 177–186 (2009)
24. Samuelson, P.A.: The pure theory of public expenditure. *Rev. Econ. Stat.* **36**(4), 387–389 (1954). <http://www.jstor.org/stable/1925895>
25. Scarf, H.E.: The core of an  $n$  person game. *Econometrica* **35**(1), 50–69 (1967)
26. Vazirani, V.V.: Nash bargaining via flexible budget markets. In: Fleischer, R., Xu, J. (eds.) AAIM 2008. LNCS, vol. 5034, p. 2. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-68880-8\\_2](https://doi.org/10.1007/978-3-540-68880-8_2)
27. Vazirani, V.V., Yannakakis, M.: Market equilibrium under separable, piecewise-linear, concave utilities. *J. ACM* **58**(3), 10:1–10:10 (2011)
28. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Finan.* **16**(1), 8–37 (1961)

# Approximating Gains-from-Trade in Bilateral Trading

Liad Blumrosen<sup>1</sup>(✉) and Yehonatan Mizrahi<sup>2</sup>

<sup>1</sup> School of Business, The Hebrew University, Jerusalem, Israel  
Blumrosen@gmail.com

<sup>2</sup> The Hebrew University, Jerusalem, Israel

**Abstract.** We consider the design of platforms that facilitate trade between a single seller and a single buyer. The most efficient mechanisms for such settings are complex and sometimes even intractable, and we therefore aim to design simple mechanisms that perform approximately well. We devise a mechanism that always guarantees at least  $1/e$  of the optimal expected gain-from-trade for every set of distributions (assuming monotone hazard rate of the buyer’s distribution). Our main mechanism is extremely simple, and achieves this approximation in Bayes-Nash equilibrium. Moreover, our mechanism approximates the optimal gain-from-trade, which is a strictly harder task than approximating efficiency. Our main impossibility result shows that no Bayes-Nash incentive compatible mechanism can achieve better approximation than  $2/e$  to the optimal gain from trade. We also bound the power of Bayes-Nash incentive compatible mechanisms for approximating the expected efficiency.

## 1 Introduction

When we look at the global commerce landscape in the Internet era, we can see that most of the products and services are sold on platforms that involve users of different roles, usually sellers and buyers. In such environments, the “auctioneer” or the “social planner” is the platform designer and not any one of the sellers (as in classic auction settings). For example, online ads are sold via exchange markets where advertisers bid for ad slots and content providers seek to maximize profit. Another example is the recent Incentive Auctions run by the US FCC [1], where spectrum is traded between TV stations and wireless communication companies. Internet commerce giants like Amazon and eBay are essentially large-scale platforms that mitigate trade between sellers and buyers for a myriad of products, and Airbnb is a marketplace where travelers seek to purchase accommodation from various vendors. The design of such two-sided markets brings in major challenges for mechanism designers, and it has been the focus of a series of recent papers (e.g., [13, 14, 23, 24, 26]).

In this paper we study the simplest two-sided market, known as the *Bilateral Trade* setting. In this setting, a single seller owns an item, and can consume it and gain a value  $s$ ; a single buyer is interested in purchasing the item that

can give him a value  $b$ . Since both values are private, agreeing on a price in an incentive-compatible mechanism may be hard. Indeed, the celebrated impossibility result by Myerson and Satterthwaite [21] claims that no Bayes-Nash incentive compatible mechanism can simultaneously achieve full efficiency (that is, perform a trade when  $b > s$ ) and be *budget balanced* (BB) and *individually rational* (IR).<sup>1</sup> In situations where budget balance and individual rationality are hard constraints, one thus has to compromise and design mechanisms with approximate expected efficiency. In their original paper, Myerson and Satterthwaite [21] characterized the “second-best” mechanism, that is, the mechanism that maximizes efficiency subject to the BB and IR constraints. However, this second-best mechanism is often too complex to implement, as it involves solving a set of differential equations which is a challenging task in the bilateral-trade setting, and seems to be completely intractable when the setting is even slightly generalized. Moreover, even if one is able to implement it, determining how well this second-best mechanism performs, compared to the optimal (“first-best”) efficiency, is not a trivial task.

There are two standard measures that quantify the efficiency of allocations in mechanisms. The first one is the expected *efficiency* (or social welfare), that is, the expected value of the player that obtains the item. The second measure is the expected *gain from trade* (GFT), which is the expected value of:  $b - s$  when a trade happens, and 0 otherwise. While the maximal efficiency and the maximal gain-from-trade are achieved by the same allocation rule, it is clear that from an approximation perspective approximating the GFT is a harder task. Every  $c$  approximation to the gain-from-trade implies a  $c$  approximation to the expected efficiency, but the opposite does not hold (this easy observation will be discussed in the sequel of the paper). For example, think about an instance where both  $s$  and  $b$  are distributed over the support  $[1, 2]$ . Every mechanism clearly gains efficiency of at least 1 and of at most 2, and thus every mechanism guarantees 1/2 approximation to the efficiency. However, designing a mechanism that attains 1/2 of the expected GFT is completely non trivial. Approximating the GFT is a notoriously hard analytical problem, and in this paper we devise simple mechanisms that approximate this objective function.

## 1.1 Our Results

A series of recent works compared the power of simple mechanisms and optimal (yet complex) mechanisms (e.g., [2, 6–8, 11, 16, 18]). Most of these results consider simple mechanisms that are *dominant-strategy* incentive compatible (DSIC). For the bilateral-trade problem, however, it was shown by Blumrosen and Dobzinski [5] that no DSIC mechanism can guarantee any constant approximation to the expected GFT. The weakness of DSIC mechanisms relates to the fact that they

---

<sup>1</sup> A mechanism is *budget balanced* if the mechanism does not gain any profit nor requires any subsidies. A mechanism is *individually rational* if the utility of each player cannot decrease by participating in the mechanism. Formal definitions will be given later in the paper.



are essentially restricted to posting a single price to the agents, where this price cannot depend on the actual bids of the agents. In this paper, we devise a mechanism that achieves approximate efficiency in Bayes-Nash incentive compatibility (BNIC). This follows a recent line of research, mostly for combinatorial auction settings, that compared the power of simple BNIC mechanism to optimal outcomes (see, e.g., [3,9,25]). Our main result in this paper is a mechanism with extremely simple rules in which simple Bayes-Nash equilibrium strategies obtain a constant approximation ratio. This mechanism circumvents the DSIC limitations, and the final price may depend on the seller's value. More precisely, this mechanism admits a unique Bayes-Nash equilibrium with at least  $1/e$  of the optimal ("first-best") GFT whenever the distribution of the buyer's value satisfies the monotone hazard rate (MHR) property (with no restrictions on the seller's distribution). We stress that, as we observe later in the paper, no DSIC mechanism can approximate the GFT even for distributions that satisfy the MHR condition.

**Theorem 1.** *When the distribution of the buyer's value satisfies the monotone hazard rate condition, there is a "simple" Bayes-Nash incentive-compatible, individually-rational and budget-balanced mechanism which always achieves at least a  $\frac{1}{e}$ -fraction of the optimal expected gain from trade.*

In this mechanism, the seller offers a take-it-or-leave-it price to the buyer, who then decides whether to accept it or not.<sup>2</sup> This mechanism is simple in several ways: first, the mechanism designer needs no distributional knowledge. The seller does need to know the distribution of the buyer in order to compute his optimal offer, but the buyer's strategy does not involve any distributional considerations. The computation required from the seller for computing her optimal offer is as complex as determining the monopoly price in the presence of a single buyer, which is known to have a simple closed-form formula and can be computed easily even in practical settings (e.g., [22]).

We note that this approximation result also implies the same approximation factor for the "second-best" mechanism.<sup>3</sup> That is, it follows that the expected gain-from-trade in the optimal BNIC mechanism cannot fall below a  $1/e$  fraction of the optimal (first-best) gain-from-trade. Furthermore, the theorem demonstrates how this bound can be achieved even by simple, more practical, mechanisms.

We strengthen this approximation result in two respects. We first prove that the approximation ratio achieved by the mechanism is actually  $\frac{1+c}{e}$ , where  $c \in [0, 1]$  is a constant that depends on the buyer's distribution (and more specifically, on the steepness of the virtual valuation function); for example, for the

<sup>2</sup> We note that our mechanism satisfies two stronger and desired versions of the above economic properties: it is *strongly* budget balanced, i.e., the sum of payments is always exactly zero; it is also *ex-post* individually rational, i.e., agents cannot lose in every instance and not only in expectation.

<sup>3</sup> Note that the characterization of the "second-best" mechanism by [21] requires that both agents have Myerson-regular distributions, while we require the stronger MHR assumption for the buyer and require nothing for the seller.

uniform distribution  $c = 0.5$ , so the approximation bound in this case is actually  $\frac{1.5}{e} \cong 0.55$ . We then prove that given a stronger condition on the buyer's distribution, namely, that the *hazard-rate is concave*, we can significantly improve the approximation bound for the GFT to  $2/e \cong 0.74$ .<sup>4</sup> We give an example for an MHR distribution with a non-concave hazard rate, for which the approximation achieved by our mechanism is strictly worse than  $2/e$ ; therefore, the concavity assumption is necessary for the analysis of our mechanism.

Our main impossibility result in this paper shows that no BNIC mechanism can guarantee an approximation ratio better than  $2/e$ .

**Theorem 2.** *There is no Bayes-Nash incentive compatible, individually rational<sup>5</sup> and budget balanced mechanism that guarantees a  $\frac{2}{e}$ -fraction of the optimal expected gain from trade. Moreover, this holds even when both distributions admit the MHR property.*

Unlike the impossibility results for DSIC mechanisms [4,10], there is no simple characterization for BNIC mechanisms; therefore, our proof relies on solving the complex “second-best” mechanism by [21] for carefully chosen distributions, and analyze its equilibrium properties. The buyer's distribution for which the bound is proven admits concave hazard rate, so this bound matches the above  $2/e$  bound for this family of distributions.

Our final impossibility result bounds the power of BNIC mechanisms for approximating the expected efficiency (all the results described so far concerned approximating gains-from-trade). We show that no BNIC mechanism can guarantee better than a 0.93-approximation to the optimal efficiency. Although this bound appears to be weak compared to the other impossibility results, this is the strongest impossibility result for BNIC mechanisms that we are aware of. We know [5] that there are BNIC mechanisms (actually, even DSIC mechanisms) that achieve a  $1 - 1/e$  approximation to the optimal efficiency. This leaves a considerable gap for BNIC mechanisms between 0.63 and 0.93.

## 1.2 More Related Work

McAfee [18] studied a similar problem to ours, i.e., how simple mechanisms can approximate the gain-from-trade in bilateral-trade settings. He proved that half of the expected gain from trade can be achieved via a DSIC mechanism for settings where the median of the buyer distribution is greater than the median of the seller's distribution. The mechanism simply posts any price between the medians as a take-it-or-leave-it offer to both agents. As mentioned, this bound cannot be generalized with DSIC mechanisms for general distributions [5], or even to MHR distributions. We overcome this impossibility by relaxing the incentive constraints from DSIC to BNIC. The Bilateral Trade problem for non quasi-linear settings was recently studied in [15].

<sup>4</sup> Concavity of the hazard rate is satisfied by some standard distributions (e.g., exponential, Weibull(2,1), etc.), and does not hold for some other distributions (e.g., uniform on  $[0, 1]$ ).

<sup>5</sup> We consider the weaker version of interim IR, which makes the proof only harder.

Blumrosen and Dobzinski [4, 5] designed simple DSIC mechanisms that approximate the expected efficiency for Bilateral trade and more complex settings. [4, 5] were inspired by McAfee's work and used the medians of the distributions as a major design tool. [4, 5] showed how features that are used in mechanisms for Bilateral Trade can be used in more general exchange frameworks, and even constructed black-box reductions from other settings to Bilateral Trade. This highlights the importance of understanding the basic bilateral-trade problem for the design of more complex markets. Colini-Baldeschi et al. [10] further studied approximation mechanisms in exchange settings under strong budget balance, and proved, among other results, an impossibility result of 0.749 for the efficiency approximation obtained by DSIC mechanisms in the bilateral trade problem.

Two-sided markets have been extensively studied in the last three decades. McAfee [17] designed an elegant DSIC, BB and IR mechanism for two sided markets with homogenous goods, which is nearly efficient in large markets. Other work about asymptotic efficiency of two-sided markets include [14, 23, 24]. Dutting et al. [13] developed a modular approach for the design of two-sided markets, based on the deferred-acceptance heuristics from [19].

We continue as follows: We present the model and a brief survey of some relevant existing results in Sect. 2. Our main positive results are given in Sect. 3, and our negative results appear in Sect. 4.

## 2 Model

The bilateral trade problem involves two agents, a seller and a buyer. The seller owns one indivisible item from which he gains a value  $s$ . The buyer gains a value  $b$  from the same item after purchasing it. In fact,  $s$  and  $b$  are drawn from two independent distributions  $F_s$  and  $F_b$  which correspond to the two random variables  $S$  and  $B$  respectively. Each of the two agents does not know the realization of the other agent's value, but the distributions are public knowledge. In our analysis we shall assume the existence of the density functions  $f_s$  and  $f_b$  for the seller and the buyer respectively. Furthermore, we assume that both agents are risk neutral and that the prices and values are commensurable.

Based on their values, the seller and the buyer simultaneously report their bids, denoted by  $\sigma(s)$  and  $\beta(b)$  respectively, to the trading mechanism. The mechanism is defined by the two functions  $t(\beta, \sigma)$  and  $p(\beta, \sigma)$ , both known to the agents, such that the item is transferred from the seller to the buyer at price  $t(\beta, \sigma)$  with probability  $p(\beta, \sigma)$ . We will be focusing on deterministic mechanisms, such that the item is transferred iff  $p(\beta, \sigma) = 1$ .

As previously mentioned, the two main measures that will be analyzed throughout this paper are the expected *gains from trade* and the expected *efficiency*. Given a mechanism  $M = \langle t, p \rangle$  and two agents with distributions  $F_b$  and  $F_s$ , these two measures, denoted by  $GFT_M^{F_b, F_s}$  and  $EFF_M^{F_b, F_s}$  respectively, are defined as follows (when they are clear, the notations  $M$ ,  $F_s$  or  $F_b$  are omitted):

$$GFT_M^{F_b, F_s} = E[(B - S) \cdot p(\beta(B), \sigma(S))]$$

$$EFF_M^{F_b, F_s} = E[B \cdot p(\beta(B), \sigma(S)) + S \cdot (1 - p(\beta(B), \sigma(S)))]$$

From these definitions it becomes clear that  $EFF_M^{F_b, F_s} = GFT_M^{F_b, F_s} + E[S]$ . In the fully efficient case (i.e., when  $p(\beta(b), \sigma(s)) = 1$  iff  $b \geq s$ ), the measures are  $GFT_{OPT}^{F_b, F_s} = E[\max\{B - S, 0\}]$  and  $EFF_{OPT}^{F_b, F_s} = E[\max\{B, S\}]$ . We note that, by definition, maximizing GFT also implies maximizing efficiency. The fully efficient allocation is our benchmark for our approximation results; we say that for a pair of such distributions, a mechanism  $M$  achieves a  $k$ -approximation to the optimal GFT if  $\frac{GFT_M^{F_b, F_s}}{GFT_{OPT}^{F_b, F_s}} \geq k$  and similarly for  $EFF$ , and we note that it always holds that  $\frac{EFF_M^{F_b, F_s}}{EFF_{OPT}^{F_b, F_s}} \geq \frac{GFT_M^{F_b, F_s}}{GFT_{OPT}^{F_b, F_s}}$ .<sup>6</sup>

### 2.1 The Hazard Rate of a Distribution

We now present some definitions, properties and notations regarding the *Hazard Rate* of a general distribution  $F$  with density  $f$  that has a non-negative support. These are used in our main approximation results in the next section.

We begin by defining the *Hazard Rate* of such distribution by  $h(x) = \frac{f(x)}{1-F(x)}$ . The *Cumulative Hazard Function* of  $F$  is defined by  $H(x) = -\ln(1 - F(x))$  for every  $x \geq 0$  (which is not to the right of  $F$ 's support). We note that  $e^{-H(x)} = 1 - F(x)$ , and that  $H(0) = 0$ . Differentiating yields  $H'(x) = h(x)$ , and we get that  $H(x) = \int_0^x h(t) dt + k$  for some  $k$ . Placing  $x = 0$  shows that  $k = 0$ .

We continue by defining the *Virtual Valuation Function* of an agent with such distribution by  $\varphi(x) = x - \frac{1-F(x)}{f(x)}$ .

Moreover, we also define the *Monotone Hazard Rate (MHR)* property of a distribution, which simply states that  $h$  is monotone non-decreasing. This property also implies that  $\varphi$  is monotone increasing, a state in which we often call  $F$  a *regular distribution*<sup>7</sup>. We note that in this case, since  $\varphi$  is strictly monotone, its inverse function exists.

In this paper, we only require such hazard rate assumptions for the buyer's distribution; thus, when we use these notations they shall be associated with  $F_b$ .

### 2.2 Bayes-Nash IC: The Second-Best Mechanism

While it was proved in [21] that no IR and BB mechanism is fully efficient in BNIC, the same paper present a characterization of the mechanisms that

<sup>6</sup> This follows from  $EFF_M \cdot GFT_{OPT} \geq GFT_M \cdot EFF_{OPT}$  which is equivalent by definition to the inequality  $(GFT_M + E[S]) \cdot GFT_{OPT} \geq GFT_M \cdot (GFT_{OPT} + E[S])$  that holds by  $GFT_{OPT} \geq GFT_M$ .

<sup>7</sup> Most of the literature assumes a weaker condition, that the  $\varphi$  is non-decreasing. In our paper we often use the inverse function of  $\varphi$ , and the notations become much simpler when  $\varphi$  is strictly increasing. Moreover, our main results consider MHR distributions that imply that  $\varphi$  is always strictly increasing.

maximize GFT subject to the IR and BB constraints. We will now describe this “second-best” mechanism for bilateral trade from [21], which is used later in our inapproximability results. We will denote this mechanism by *MS*.

As stated [21], in order to derive the correct approximation results using this mechanism, we need to assume that the support of  $F_b$  is  $[\underline{b}, \bar{b}]$  or  $[\underline{b}, \infty)$  for some  $\bar{b} \geq \underline{b} \geq 0$  and that the support of  $F_s$  is  $[\underline{s}, \bar{s}]$  or  $[\underline{s}, \infty)$  for some  $\bar{s} \geq \underline{s} \geq 0$ . As in [21], we assume regularity of the distributions, i.e., that the functions  $b - \frac{1-F_b(b)}{f_b(b)}$  and  $s + \frac{F_s(s)}{f_s(s)}$  are monotone increasing. Using the fact that this mechanism is truthful, i.e., in a Bayes-Nash equilibrium  $\beta(b) = b$  and  $\sigma(s) = s$ , the mechanism is defined by:

$$p^\alpha(\beta(b), \sigma(s)) = \begin{cases} 1 & \text{if } s + \alpha \cdot \frac{F_s(s)}{f_s(s)} \leq b - \alpha \cdot \frac{1-F_b(b)}{f_b(b)} \\ 0 & \text{otherwise.} \end{cases}$$

The appropriate parameter is the unique (as proved in [21])  $\alpha \in (0, 1]$  that solves the following equation, presented for the bounded supports case (and similar for the unbounded case):

$$\int_{\underline{b}}^{\bar{b}} \int_{\underline{s}}^{\bar{s}} \left( \left( b - \frac{1-F_b(b)}{f_b(b)} \right) - \left( s + \frac{F_s(s)}{f_s(s)} \right) \right) \cdot p^\alpha(b, s) f_b(b) f_s(s) dsdb = 0$$

The appropriate payment function can be determined ad hoc, given the distributions. Nonetheless, we note that it is not necessary in order to analyze the *GFT* and *EFF* measures.

### 3 A Constant Approximation for the Gains from Trade

In this section we present a simple mechanism that approximates the optimal gains from trade for bilateral trade settings. The mechanism has no dominant-strategy equilibrium, and the results are achieved in Bayes-Nash equilibrium. We call this mechanism the *Seller-Offering Mechanism* (abbreviated as *SO*).

#### The Seller-Offering (SO) Mechanism:

- The seller offers a take-it-or-leave-it price  $t$  to the buyer, who chooses whether to accept it or not.
- If the buyer accepts the price, a trade occurs at price  $t$ . Otherwise, no trade occurs and no payments are transferred.

We note that at first glance, it seems as if this mechanism does not fall into formal model of bilateral trade mechanisms we defined earlier, since it is two-staged and not simultaneous. However, using  $p(\beta, \sigma) = 1_{\{\beta \geq \sigma\}}(\beta, \sigma)$  and  $t(\beta, \sigma) = \sigma$  in the original scheme yields the same results.

### 3.1 Some Technical Definitions and Observations

For our results in this section, it suffices to assume that the support of  $F_b$  is  $[\underline{b}, \bar{b}]$  or  $[\underline{b}, \infty)$  for some  $\underline{b} \geq 0$ , the support of  $F_s$  is contained in  $[0, \infty)$ ,  $f_b$  is differentiable and  $F_b$  adheres to the MHR assumption.

The inverse virtual valuation  $\varphi^{-1}(\cdot)$  turns out to be very useful in our analysis. This inverse function is not well defined for all possible values, therefore we frequently use its extension denoted by  $\overline{\varphi^{-1}}(\cdot)$ .

**Definition 1.** Under the aforementioned assumptions, we define the Extended Inverse Virtual Valuation Function,  $\overline{\varphi^{-1}}(x)$ , to be the continuous extension of  $\varphi^{-1}(x)$ :

Since  $\varphi(x)$  is increasing,  $\varphi^{-1}(x)$  is undefined for  $x \leq \varphi(\underline{b})$ , and in case  $F_b$ 's support is  $[\underline{b}, \bar{b}]$ , it is also undefined for  $x \geq \varphi(\bar{b}) = \bar{b}$ . The left part is extended using  $\overline{\varphi^{-1}}(x) = \underline{b}$  and the right part using  $\overline{\varphi^{-1}}(x) = x$ .<sup>8</sup>

We continue by showing some useful technical observations regarding these functions, used later in our proofs:

**Observation 2.** For every  $x$  in their domain,  $\varphi(x) \leq x$  and  $\overline{\varphi^{-1}}(x) \geq x$ .

*Proof.* The first inequality follows from  $\varphi(x) = x - \frac{1}{h(x)} \leq x$  since  $h$  is positive. The second follows from the fact that  $\varphi^{-1}$  is the reflection of  $\varphi$  with respect to the line  $y=x$ , and since the extension of it preserves the inequality.

**Observation 3.** For every  $x \geq \varphi(\underline{b})$ :

1. If  $\bar{b} \geq x$  then  $\overline{\varphi^{-1}}(x) - x = \frac{1 - F_b(\varphi^{-1}(x))}{f_b(\varphi^{-1}(x))} = \frac{1}{h(\varphi^{-1}(x))}$ .
2. If  $\bar{b} \leq x$  then  $\overline{\varphi^{-1}}(x) - x = 0$ .

*Proof.* For the first case, it holds that  $\overline{\varphi^{-1}}(x) - x = \varphi^{-1}(x) - \varphi(\varphi^{-1}(x)) = \varphi^{-1}(x) - \varphi^{-1}(x) + \frac{1 - F_b(\varphi^{-1}(x))}{f_b(\varphi^{-1}(x))} = \frac{1 - F_b(\varphi^{-1}(x))}{f_b(\varphi^{-1}(x))} = \frac{1}{h(\varphi^{-1}(x))}$  by the definitions of  $\varphi$  and  $h$ . For the second case, by the definition of the right extension of  $\varphi^{-1}$ , it holds that  $\overline{\varphi^{-1}}(x) - x = x - x = 0$ .

**Observation 4.** For every  $\bar{b} \geq x \geq \varphi(\underline{b})$  it holds that  $\frac{d\varphi^{-1}(x)}{dx} = \frac{1}{1 + \frac{h'(\varphi^{-1}(x))}{(h(\varphi^{-1}(x)))^2}} \in$

$[0, 1]$  under the MHR assumption.

*Proof.* We remind that  $\varphi(x) = x - \frac{1}{h(x)}$ . By the reciprocal rule, differentiating yields  $\varphi'(x) = 1 - \frac{0 - h'(x)}{(h(x))^2} = 1 + \frac{h'(x)}{(h(x))^2}$ . Furthermore,  $\frac{d\varphi^{-1}(x)}{dx} = \frac{1}{\varphi'(\varphi^{-1}(x))}$  by the derivative of an inverse function. Thus, the identity follows by plugging  $\varphi^{-1}(x)$  in the derivative. We also note that by the MHR assumption,  $\frac{h'(\varphi^{-1}(x))}{(h(\varphi^{-1}(x)))^2} \geq 0$ , hence  $\frac{d\varphi^{-1}(x)}{dx} \in [0, 1]$ .

<sup>8</sup> We note that  $\overline{\varphi^{-1}}(x)$  is defined for every  $x$ , even when  $F_b$ 's support is  $[\underline{b}, \infty)$ , since  $\varphi(x)$  is unbounded from above in that case. This can be seen by noticing that for every  $y \in \mathbb{R}$ , choosing  $x > \max\{\frac{1}{h(\underline{b})} + y + 1, \underline{b}\}$  yields  $\varphi(x) = x - \frac{1}{h(x)} \geq \frac{1}{h(\underline{b})} + y + 1 - \frac{1}{h(\underline{b})} > y$  by the MHR assumption.

### 3.2 Analysis of the Seller-Offering Mechanism

Although not admitting a dominant-strategy equilibrium, the above Seller- Offering mechanism induces quite straightforward Bayes-Nash equilibrium strategies for the agents. In equilibrium, the seller offers the monopoly price given his own value for the item, that is,  $\varphi^{-1}(s)$  (as in [20]), and the bidder will simply bid truthfully to accept the deal if its value exceeds the offered price.<sup>9</sup> This is an immediate application of Myerson’s theory [20], but for completeness, a proof is given in the full version of the paper.

**Proposition 5.** *For every MHR distribution  $F_b$  for the buyer and every distribution  $F_s$  for the seller, the bids  $\beta(b) = b$  and  $\sigma(s) = \varphi^{-1}(s)$  form a Bayes-Nash equilibrium in the Seller-Offering Mechanism.*

In the following lemma we present a convenient representation of  $GFT_{OPT}$  and  $GFT_{SO}$  which we heavily use. The representation of  $GFT_{OPT}$  is also shown in [18]. We prove this lemma in the full version of the paper.

**Lemma 6.** *For every MHR distribution  $F_b$  for the buyer and every distribution  $F_s$  for the seller, the following equalities hold:*

$$GFT_{OPT} = \int_0^\infty F_s(s) \cdot (1 - F_b(s)) ds$$

$$GFT_{SO} = \int_0^\infty F_s(s) \cdot \left(1 + \frac{d\overline{\varphi^{-1}}(s)}{ds}\right) \cdot (1 - F_b(\overline{\varphi^{-1}}(s))) ds$$

We now turn to proving the main result of the paper, concerning the constant approximation guarantee obtained using the Seller-Offering mechanism. This approximation result is parameterized by a parameter  $c$  that describes the steepness of the buyer’s virtual function.

**Definition 7.** *We define the Virtual Steepness Constant of an MHR distribution  $F$  with a differentiable density  $f$  by  $c = \min_s \frac{d\varphi^{-1}(s)}{ds}$ . We note that  $c$  is in fact the reciprocal of the virtual valuation function’s Lipschitz constant, since  $\min_s \frac{d\varphi^{-1}(s)}{ds} = \min_s \frac{1}{\varphi'(\varphi^{-1}(s))} = \frac{1}{\max_s \varphi'(s)}$ .*

Our theorem shows that given the MHR condition on the buyer’s valuation, our mechanism attains a  $\frac{1+c}{e}$  fraction of the optimal gains-from-trade. Since by Observation 4 we have that  $c \in [0, 1]$ , this approximation is at least  $\frac{1}{e}$  for all possible distributions.

**Theorem 8.** *For every MHR distribution  $F_b$  for the buyer and every distribution  $F_s$  for the seller, the Seller-Offering Mechanism obtains a  $\frac{1+c}{e}$ -approximation to the optimal gains from trade.*

<sup>9</sup> Recall that  $\varphi$  denotes the virtual valuation of the buyer, and the seller use the details of this distribution to determine what price to post.

*Proof.* We remind that in Lemma 6, we concluded that  $GFT_{OPT} = \int_0^\infty F_s(s) \cdot (1 - F_b(s)) ds$  and that  $GFT_{SO} = \int_0^\infty F_s(s) \cdot (1 + \frac{d\varphi^{-1}(s)}{ds}) \cdot (1 - F_b(\varphi^{-1}(s))) ds$ . We therefore analyze the relation between  $(1 + \frac{d\varphi^{-1}(s)}{ds}) \cdot (1 - F_b(\varphi^{-1}(s)))$  and  $(1 - F_b(s))$  for every  $s \geq 0$ .

If  $s \geq \bar{b}$ , both terms are 0 (we use Observation 2 for the first term). If  $s \leq \varphi(\bar{b})$  then  $(1 + \frac{d\varphi^{-1}(s)}{ds}) \cdot (1 - F_b(\varphi^{-1}(s))) = (1 + 0) \cdot (1 - F_b(\bar{b})) = 1 = (1 - F_b(s))$ . The first equality follows from  $\varphi^{-1}(s) = \bar{b}$  for such  $s$ , and the last equality follows from  $\bar{b} \geq \varphi(\bar{b})$  as noted in Observation 2. The ratio between these two terms is 1, which is greater than  $\frac{1+c}{e}$ .

We now focus on the case where  $\bar{b} \geq s \geq \max\{0, \varphi(\bar{b})\}$ , such that  $\varphi^{-1}(s) = \varphi^{-1}(s)$ , and we show that  $e \cdot (1 - F_b(\varphi^{-1}(s))) \geq 1 - F_b(s)$ . The Cumulative Hazard Function  $H$  of the buyer is the integral of the monotone increasing function  $h$ , hence  $H$  is convex. Therefore, the line tangent to  $H$  at any point is below the function. In other words, fixing  $x_0 \in [0, \bar{b}]$ , for every  $x \in [0, \bar{b}]$  it holds that  $H(x) \geq H(x_0) + h(x_0) \cdot (x - x_0)$ . By Observation 3.3, choosing  $x_0 = \varphi^{-1}(s)$  we get that for every  $x \in [0, \bar{b}]$ , and specifically  $x = s$ , it holds that:

$$\begin{aligned} H(x) &\geq H(\varphi^{-1}(x)) + h(\varphi^{-1}(x)) \cdot (x - \varphi^{-1}(x)) = \\ &= H(\varphi^{-1}(x)) + h(\varphi^{-1}(x)) \cdot \left( -\frac{1}{h(\varphi^{-1}(x))} \right) = H(\varphi^{-1}(x)) - 1 \end{aligned}$$

Hence:

$$1 - F_b(s) = e^{-H(s)} \leq e^{-H(\varphi^{-1}(s))+1} = e \cdot e^{-H(\varphi^{-1}(s))} = e \cdot (1 - F_b(\varphi^{-1}(s)))$$

The first and the last equalities in the last line follow from the definition of  $H$  as described in Sect. 2.1.

Concluding, we get that for every  $s \geq 0$  it holds that

$$F_s(s) \cdot (1 + \frac{d\varphi^{-1}(s)}{ds}) \cdot (1 - F_b(\varphi^{-1}(s))) \geq F_s(s) \cdot \frac{1+c}{e} \cdot (1 - F_b(s))$$

Integrating both parts and by the monotonicity of the integral, we get that by Lemma 6 we have that  $GFT_{SO} \geq \frac{1+c}{e} \cdot GFT_{OPT}$ .

We now proceed to proving an amplified version of this theorem. In the proof of Theorem 8 we relied on a linear approximation of  $H$ . The next theorem utilizes a quadratic approximation of  $H$  (via the Taylor expansion) to improve the bound, but requires an additional assumption, the concavity of  $h$ . With this additional assumption, the approximation can be improved to  $2/e$ .

The following technical lemma, which is proved in the full version, manifests the importance of the concavity assumption and is used for proving the theorem below.

**Lemma 9.** *Let  $f$  be a twice differentiable function that has a concave derivative, and let  $T(x)$  be a second degree Taylor polynomial at  $x_0$ , i.e., a quadratic approximation at this point. Then for every  $x \leq x_0$  it holds that  $T(x) \leq f(x)$ .*



**Theorem 10.** *For every MHR distribution  $F_b$  with a concave hazard rate for the buyer and every distribution  $F_s$  for the seller, the Seller-Offering Mechanism obtains a  $\frac{2}{e}$ -approximation to the optimal gains from trade.*

A proof can be found in the full version of the paper.

We can now use Theorem 10 to separate the power of DSIC and BNIC mechanisms in terms of approximating the gains from trade. The following proposition shows that there are instances where no DSIC mechanism can obtain a constant approximation to the gains from trade, but as the relevant distributions satisfy MHR and admit a concave hazard function, Theorem 10 implies the existence of BNIC mechanisms with  $\frac{2}{e}$  approximation.

**Proposition 11.** *There exists a pair of distributions  $F_b$  for the buyer and  $F_s$  for the seller, for which no DSIC mechanism that is IR and BB can achieve a constant approximation to the optimal gains from trade, while there exists a BNIC mechanism that is IR and BB that does achieve a  $2/e$ -approximation to the optimal gains from trade for them.*

*Proof.* Consider the two distributions  $F_b \sim \text{Exponential}(1)$  and  $F_s(x) = \lambda(e^{x-t} - e^{-t})$  with  $\lambda = \frac{1}{1-e^{-t}}$  on the support  $[0, t]$ . In [5], Blumrosen and Dobzinski analyze the scenario in which  $F_b(x) = \lambda(1 - e^{-x})$  on the support  $[0, t]$  and  $F_s$  is the same as above. They show that every fixed price mechanism achieves at most  $O(1/t)$ -approximation to the optimal gains from trade in this case. By taking  $t$  that tends to infinity, this buyer’s distribution converges to Exponential(1) while  $1/t$  converges to 0. Alternatively, a direct calculation using the original distributions yields these results. Since it is well known that every DSIC mechanism that is IR and BB is a fixed price mechanism (see, e.g., [10, 12] and the references therein), the first part follows.

We note that in this case,  $h(x) = 1$  which is a constant function and therefore the MHR and concavity assumptions hold. Thus, by Theorem 3.10 the SO Mechanism indeed obtains a  $2/e$ -approximation to the optimal gains from trade.

Lastly, the following proposition signifies the necessity of  $h$ ’s concavity assumption for Theorem 10. We also show that the analysis of Theorem 10 is tight, and for some distributions (that satisfy MHR and concave hazard rate) our mechanism achieves exactly  $2/e$  approximation. A proof can be found in the full version of the paper.

**Proposition 12.** *Using the Seller-Offering Mechanism:*

1. *There exists an MHR distribution  $F_b$  with a non-concave hazard rate, and a distribution  $F_s$  for the seller, such that the mechanism achieves an approximation to the optimal GFT which is strictly worse than  $2/e$ .*
2. *There exists an MHR distribution  $F_b$  with a concave hazard rate  $h$  and a distribution  $F_s$  for the seller, such that  $h$  is concave and  $\frac{GFT_{SO}}{GFT_{OPT}} = \frac{2}{e}$ .*

## 4 Inapproximability Results

In this section, we present impossibility results for approximating the gains from trade and efficiency using BNIC mechanism.

In the previous section, we presented an IR, BB and BNIC mechanism that guarantees a  $1/e$ -approximation to the optimal gains from trade for any pair of distributions under standard MHR assumptions. A question that naturally arises concerns the limitations of BNIC mechanisms in this our setting. The following theorem addresses that question and shows that no BNIC mechanism can maintain IR and BB and guarantee more than  $2/e$  approximation. Moreover, this holds even when the distributions satisfy the MHR condition.<sup>10</sup> We also note that this result is proven for the case where the buyer’s distribution has concave hazard rate, and thus it matches the positive result in Theorem 10 when this condition is satisfied.

**Theorem 13.** *No BNIC mechanism which is IR and BB can guarantee an approximation to the optimal gains from trade which is better than  $2/e$ . This holds even if both distributions satisfy the MHR condition.*

The proof relies on the Second-Best mechanism devised by Myerson and Satterthwaite in [21]. We show that for every  $\epsilon > 0$  there exists a pair of distributions such that  $\frac{GFT_{MS}}{GFT_{OPT}} < \frac{2}{e} + \epsilon$ . Since by its definition, no BNIC mechanism which is IR and BB can achieve a better approximation than this mechanism for these distributions, the claim follows. In fact, the relevant distributions are exactly the ones used in Proposition 11, i.e.,  $F_b \sim Exponential(1)$  and  $F_s(x) = \lambda(e^{x-t} - e^{-t})$  with  $\lambda = \frac{1}{1-e^{-t}}$  on the support  $[0, t]$ . The distributions satisfy the MHR property. We remind that the second-best solution requires that  $b - \frac{1-F_b(b)}{f_b(b)}$  and  $s + \frac{F_s(s)}{f_s(s)}$  are monotone increasing, and indeed this property holds for  $b - \frac{1-F_b(b)}{f_b(b)} = b - 1$  and  $s + \frac{F_s(s)}{f_s(s)} = s + 1 - e^{-s}$ . The full proof can be found in the full version of the paper.

We conclude by showing a similar result for the expected efficiency in the bilateral trade setting. As the previous proof illustrates, and as supported by simulations using various distributions, the Second-Best mechanism achieves a relatively low approximation to the optimal GFT when the buyer’s values tend to be low and the seller’s values tend to be high. Since this is normally associated with low expected gains from trade, and since  $EFF = GFT + E[S]$ , these scenarios often produce high approximation to the optimal efficiency. Thus, it seems that tackling the question of finding an approximation to that measure that cannot be guaranteed requires observing somewhat more balanced scenarios.

We remark that [10] studied this question for the DSIC case, and showed that no DSIC mechanism which is IR and BB can guarantee a 0.749-approximation to the optimal efficiency. The following theorem shows a similar result for BNIC mechanisms, and is proved in the full version of the paper. While this bound

---

<sup>10</sup> This theorem holds for a weaker notion of interim individual rationality (as in [21]); This clearly strengthens the result.

appears to be weak compared to the bound on the GFT in Theorem 13, we are not aware of any stronger bound for this problem. The best positive result to date for this problem is by [5], who showed a DSIC (and thus also BNIC) mechanism that guarantees about 0.63 fraction of the optimal efficiency.

**Theorem 14.** *No BNIC mechanism which is IR and BB can guarantee an approximation to the optimal efficiency which is better than 0.934.*

## 5 Conclusion

This paper considers the bilateral-trade problem, which is a fundamental problem in economics for more than three decades and it demonstrates the simplest form of two sided markets. We hope that developing understanding of this fundamental problem may also be helpful in the design of more general two sided markets.

Our main result is a mechanism that achieves at least  $1/e$  fraction of the optimal gain from trade, assuming that the distribution of the buyer satisfies MHR. The mechanism is simple, Bayes-Nash incentive compatible, strongly budget balanced and ex-post individually rational. The bound also implies that the most efficient mechanism subject to the IR and BB, which was characterized in the seminal paper of [21], must also achieve at least the same fraction of the optimal gain-from-trade. Our main impossibility result shows that no BNIC mechanism can guarantee an approximation which is better than  $2/e$ .

The main open question that is raised in this paper is whether the MHR assumption (on the buyer's side) is really required for achieving a constant approximation to the gain from trade via BNIC mechanisms. In other words, is there a BNIC, IR and BB mechanism that guarantees a constant approximation to the gain from trade for all distributions? We note that Myerson and Satterthwaite's [21] characterization of the "second-best" mechanism was not general, and assumed that the distributions are regular (a slightly weaker assumption than MHR).

A second interesting open question concerns closing the relatively-wide gap between the lower and the upper bound for the efficiency-maximizing problem by DSIC mechanisms. The best currently known approximation for this problem is 0.63 [5], while our impossibility result gives a bound of 0.93. As these results are given for Bayes-Nash incentive compatible mechanisms, the analysis can be challenging.

## References

1. Ausubel, L., Levin, J., Milgrom, P., Segal, I.: Incentive auction rules option and discussion (2012). Appendix to the FCCs NPRM on Incentive Auctions
2. Babaioff, M., Immorlica, N., Lucier, B., Weinberg, S.M.: A simple and approximately optimal mechanism for an additive buyer. In: 2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS), pp. 21–30 (2014)
3. Bhawalkar, K., Roughgarden, T.: Welfare guarantees for combinatorial auctions with item bidding. In: Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, pp. 700–709 (2011)

4. Blumrosen, L., Dobzinski, S.: Reallocation mechanisms. In: The 15th ACM Conference on Economics and Computation, EC 2014, p. 617 (2014)
5. Blumrosen, L., Dobzinski, S.: (Almost) efficient mechanisms for bilateral trading (2016). Working paper
6. Blumrosen, L., Holenstein, T.: Posted prices vs. negotiations: an asymptotic analysis. In: ACM Conference on Electronic Commerce (2008)
7. Blumrosen, L., Nisan, N., Segal, I.: Auctions with severely bounded communications. *J. Artif. Intell. Res.* **28**, 233–266 (2006)
8. Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, pp. 311–320 (2010)
9. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian combinatorial auctions. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008. LNCS, vol. 5125, pp. 820–832. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-70575-8\\_67](https://doi.org/10.1007/978-3-540-70575-8_67)
10. Colini-Baldeschi, R., de Keijzer, B., Leonardi, S., Turchetta, S.: Approximately efficient double auctions with strong budget balance. In: SODA 2016 (2016)
11. Dhangwatnotai, P., Roughgarden, T., Yan, Q.: Revenue maximization with a single sample. In: The 11th ACM Conference on Electronic Commerce, pp. 129–138 (2010)
12. Drexler, M., Kleiner, A.: Optimal private good allocation: the case for a balanced budget. *Games Econ. Behav.* **94**, 169–181 (2015)
13. Dutting, P., Roughgarden, T., Talgam-Cohen, I.: Modularity and greed in double auctions. In: EC 2014 (2014)
14. Fudenberg, D., Mobius, M., Szeidl, A.: Existence of equilibrium in large double auctions. *J. Econ. Theory* **133**(1), 550–567 (2007)
15. Garratt, R., Pycia, M.: Efficient bilateral trade (2015). Working paper
16. Hartline, J.D., Roughgarden, T.: Simple versus optimal mechanisms. In: The 10th ACM Conference on Electronic Commerce, EC 2009, pp. 225–234 (2009)
17. McAfee, P.R.: Amicable divorce: dissolving a partnership with simple mechanisms. *J. Econ. Theory* **56**(2), 266–293 (1992)
18. McAfee, P.R.: The gains from trade under fixed price mechanisms. *Appl. Econ. Res. Bull.* **1**, 1–10 (2008)
19. Milgrom, P., Segal, I.: Deferred-acceptance heuristic auctions (2013). Working paper, Stanford University
20. Myerson, R.B.: Optimal auction design. *Math. Oper. Res.* **6**(1), 58–73 (1981)
21. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. *J. Econ. Theory* **28**, 265–281 (1983)
22. Ostrovsky, M., Schwarz, M.: Reserve prices in internet advertising auctions: a field experiment. In: The 12th ACM Conference on Electronic Commerce, pp. 59–60 (2011)
23. Rustichini, A., Satterthwaite, M.A., Williams, S.R.: Convergence to efficiency in a simple market with incomplete information. *Econometrica* **62**(5), 1041–63 (1994)
24. Satterthwaite, M.A., Williams, S.R.: The optimality of a simple market mechanism. *Econometrica* **70**(5), 1841–1863 (2002)
25. Syrgkanis, V., Tardos, E.: Composable and efficient mechanisms. In: The 43th Annual ACM Symposium on Theory of Computing, STOC 2013, pp. 211–220 (2013)
26. Wilson, R.: Incentive efficiency of double auctions. *Econometrica* **53**(5), 1101–1115 (1985)

# Coverage, Matching, and Beyond: New Results on Budgeted Mechanism Design

Georgios Amanatidis<sup>(✉)</sup>, Georgios Birmpas, and Evangelos Markakis

Department of Informatics, Athens University of Economics and Business,  
Athens, Greece

{gamana, gebirbas, markakis}@aub.gr

**Abstract.** We study a type of reverse (procurement) auction problems in the presence of budget constraints. The general algorithmic problem is to purchase a set of resources, which come at a cost, so as not to exceed a given budget and at the same time maximize a given valuation function. This framework captures the budgeted version of several well known optimization problems, and when the resources are owned by strategic agents the goal is to design truthful and budget feasible mechanisms. We first obtain mechanisms with an improved approximation ratio for weighted *coverage valuations*, a special class of submodular functions. We then provide a general scheme for designing randomized and deterministic polynomial time mechanisms for a class of XOS problems. This class contains problems whose feasible set forms an *independence system* (a more general structure than matroids), and some representative problems include, among others, finding maximum weighted matchings and maximum weighted matroid members. For most of these problems, only randomized mechanisms with very high approximation ratios were known prior to our results.

## 1 Introduction

In this work, we study a class of mechanism design problems under a budget constraint. Consider a reverse auction setting, where a single buyer wants to select a subset, among a set  $A$  of agents, for performing some tasks. Each agent  $i$  comes at a cost  $c_i$ , in the case that he is chosen. The buyer has a budget  $B$  and a valuation function  $v(\cdot)$ , so that  $v(S)$  is the derived value if  $S \subseteq A$  is the chosen set. The purely algorithmic version then asks to maximize the generated value subject to the constraint that the total cost of the selected agents should not exceed  $B$ . Some of these problems are motivated by crowdsourcing scenarios and related applications, where agents can be viewed as workers, e.g., [4]. Apart from that, they form natural budgeted versions of well known optimization problems.

In the mechanism design version that we focus on, the cost  $c_i$  is private information for each agent  $i$ . Hence, we want to design mechanisms that are incentive

---

The full version of this paper can be found in [3]. Research supported by an internal research funding program of the Athens University of Economics and Business.

compatible, individually rational, and budget feasible, i.e. the sum of the payments to the agents does not exceed  $B$ . Note that the payments here can be higher than the actual costs in order to induce truthfulness. Budget feasibility is a tricky property that makes the problem more challenging, as it already rules out well known mechanisms such as VCG. Although the algorithmic versions of such problems often admit constant factor approximation algorithms, it is not clear how to appropriately convert them into truthful budget feasible mechanisms. Therefore, the question of interest is to find mechanisms that achieve the best possible approximation for the optimal value of  $v(\cdot)$  under these constraints. We stress that the question is nontrivial even if we allow exponential time algorithms, since computational power does not necessarily make the problem easier (see also the discussion in [9]).

Budgeted mechanism design was first studied by Singer [14] when  $v(\cdot)$  is an additive or a nondecreasing submodular function. Later on, follow up works have also provided more results for XOS and subadditive functions. Although these results shed more light on our understanding of the problem, there are still several interesting issues that remain unresolved. First, the current results on submodular valuations are not known to be tight. Further, and most importantly, when going beyond submodularity, to XOS functions, we are not even aware of general mechanisms with small approximation guarantees, let alone deterministic polynomial time mechanisms.

**Contribution:** We first demonstrate (Sect. 3) how to obtain improved deterministic budget feasible mechanisms for weighted *coverage valuations*, a notable subclass of submodular functions. This class has already received attention in previous works [14, 15], motivated by problems related to influence maximization in social networks. Our mechanism reduces roughly by half (from 31.03 to 15.45) the known approximation of [15] and also generalizes it to the weighted version of coverage functions. We then move to our main result (Sect. 4), which is a general scheme for obtaining randomized and deterministic polynomial time approximations for a subclass of XOS problems, that contains the budgeted versions of several well known optimization problems. We first illustrate our ideas in Sect. 4.1 on the budgeted matching problem, where  $v(S)$  is defined as the maximum weight matching that can be derived from the edges of  $S$ . For this problem only a randomized 768-approximation was known [5]. Our approach yields a randomized 3-approximation and a deterministic 4-approximation. Then in Sect. 4.2 we show how to generalize these results to problems with a similar combinatorial structure, where the set of feasible solutions forms an *independence system*. These structures are more general than matroids (they do not always satisfy the exchange property) and some representative problems that are captured include finding maximum weighted matroid members, maximum weighted  $k$ -D-matchings, and maximum weighted independent sets.

**Related Work:** The study of budget feasible mechanisms, as considered here, was initiated by Singer [14], who gave a randomized constant factor approximation mechanism for nondecreasing submodular functions. Later, Chen et al. [7]

significantly improved these approximation ratios, obtaining a randomized, polynomial time mechanism achieving a 7.91-approximation and a deterministic one with a 8.34-approximation. Their deterministic mechanism does not run in polynomial time in general, but it can be modified to do so for special cases (see Sect. 3). As an example, Singer [15] followed a similar approach to obtain a deterministic, polynomial time, 31.03-approximation mechanism for the unweighted version of Budgeted Max Coverage, a class that we also consider in Sect. 3. Along these lines, Horel et al. [12] consider another family of submodular functions and give a deterministic, polynomial time, constant approximation for the so-called Experimental Design Problem, under a mild relaxation on truthfulness. For subadditive functions Dobzinski et al. [9] suggested a randomized  $O(\log^2 n)$ -approximation mechanism. This was later improved to  $O(\log n / \log \log n)$  by Bei et al. [5], who also gave a randomized  $O(1)$ -approximation mechanism for XOS functions, albeit in exponential time. Recently, there is also a line of works under the *large market* assumption (where no participant can affect significantly the market outcome). Under this assumption, Anari et al. [4] resolved the additive case, giving a  $\frac{e}{e-1}$ -approximation mechanism and a matching lower bound. Further results for large markets were obtained by Goel et al. [11] for a crowdsourcing problem with matching constraints (and hence a non submodular objective).

## 2 Definitions and Notation

We use  $A = [n] = \{1, 2, \dots, n\}$  to denote a set of  $n$  agents. Each agent  $i$  is associated with a private cost  $c_i$ , denoting the cost for participating in the solution. We consider a procurement auction, where the auctioneer is equipped with a valuation function  $v : 2^A \rightarrow \mathbb{Q}^+$  and a positive budget  $B$ . For  $S \subseteq A$ ,  $v(S)$  is the value derived by the auctioneer if the set  $S$  is selected. Therefore, the goal is to select a set  $S$  that maximizes  $v(S)$  subject to the constraint  $\sum_{i \in S} c_i \leq B$ .

We consider valuation functions that are non-decreasing, i.e.  $v(S) \leq v(T)$  for any  $S \subseteq T \subseteq A$ . Throughout our work, we will focus on valuations that come from two natural classes of functions, namely submodular and XOS functions defined below.

**Definition 1.** *A valuation function, defined on  $2^A$  for some set  $A$ , is*

(i) *submodular, if  $v(S \cup \{i\}) - v(S) \geq v(T \cup \{i\}) - v(T)$  for any  $S \subset T \subset A$ , and  $i \notin T$ .*

(ii) *XOS or fractionally subadditive, if there exist additive functions  $\alpha_1, \dots, \alpha_r$ , for some finite  $r$ , such that  $v(S) = \max\{\alpha_1(S), \alpha_2(S), \dots, \alpha_r(S)\}$ .*

We note that the class XOS is a strict superclass of submodular valuations.

**Mechanism Design.** Each agent here only has his cost as private information, hence we are in the domain of single-parameter problems. A mechanism  $\mathcal{M} = (f, p)$  in our context consists of an outcome rule  $f$  and a payment rule  $p$ . Given a vector of cost declarations,  $b = (b_i)_{i \in A}$ , where  $b_i$  denotes the cost reported

by agent  $i$ , the mechanism selects the set  $f(b)$ . At the same time, it computes payments  $p(b) = (p_i(b))_{i \in N}$  where  $p_i(b)$  denotes the payment issued to agent  $i$ .

The main properties we want to ensure for our mechanisms in this work are *truthfulness* (reporting  $c_i$  is a dominant strategy for every agent  $i$ ), *individual rationality* ( $p_i(b) \geq 0$  for every  $i \in A$ , and  $p_i(b) \geq c_i$ , for every  $i \in f(b)$ ), and *budget feasibility* ( $\sum_{i \in A} p_i(b) \leq B$  for every  $b$ ).

When referring to randomized mechanisms, the notion of truthfulness we use is *universal truthfulness*, which means that the mechanism is a probability distribution over deterministic truthful mechanisms.

For single-parameter problems we use the characterization by Myerson [13] for deriving truthful mechanisms. In particular, we say that an outcome rule  $f$  is *monotone*, if for every agent  $i \in A$ , and any vector of cost declarations  $b$ , if  $i \in f(b)$ , then  $i \in f(b'_i, b_{-i})$  for  $b'_i \leq b_i$ . This simply means that if an agent is selected in the outcome by declaring a cost  $b_i$ , then he should also be selected if he declares a lower cost.

**Lemma 1.** *Given a monotone algorithm  $f$ , there is a unique payment scheme  $p$  such that  $(f, p)$  is a truthful and individually rational mechanism, given by  $p_i(b) = \sup_{b_i \in [c_i, \infty)} \{b_i : i \in f(b_i, b_{-i})\}$  when  $i \in f(b)$ , and  $p_i(b) = 0$  otherwise.*

Lemma 1 is known as Myerson's lemma, and the payments are often referred to as *threshold payments*. Myerson's lemma simplifies the design of truthful mechanisms by focusing only on constructing monotone algorithms and not having to worry about the payment scheme. In this work, we always assume that the underlying payment scheme is given by Myerson's lemma.

### 3 Deterministic Mechanisms for Submodular Objectives

We begin our exposition with submodular valuations, and show in Sect. 3.1 how to obtain an improved approximation for a subclass of such functions. To do this, we exploit the approach by Chen et al. [7], starting with their mechanism:

---

MECHANISM-SM( $A, B$ ) [7]

---

- 1 Set  $A = \{i \mid c_i \leq B\}$  and  $i^* \in \arg \max_{i \in A} v(i)$
  - 2 **if**  $\frac{1+4e+\sqrt{1+24e^2}}{2(e-1)} \cdot v(i^*) \geq \text{OPT}(A \setminus \{i^*\}, B)$  **then**
  - 3    | **return**  $i^*$
  - 4 **else**
  - 5    | **return** GREEDY-SM( $A, B/2$ )
-



In MECHANISM-SM, an agent  $i^*$  of maximum value is compared with an optimal solution at the instance  $A \setminus \{i^*\}$  with budget  $B$ . Then, either  $i^*$  or GREEDY-SM( $A, B/2$ ) is returned.

---

GREEDY-SM( $A, B/2$ ) [7]

---

```

1 Let  $k = 1$  and  $S = \emptyset$ 
2 while  $k \leq |A|$  and  $c_k \leq \frac{B}{2} \cdot \frac{v(S \cup \{k\}) - v(S)}{v(S \cup \{k\})}$  do
3    $S = S \cup \{k\}$ 
4    $k = k + 1$ 
5 return  $S$ 

```

---

GREEDY-SM is a greedy algorithm that picks agents according to their ratio of marginal value over cost, given that this cost is not too large. For the sake of presentation, we assume the agents are sorted in descending order with respect to this ratio. The marginal value of each agent is calculated with respect to the previous agents in the ordering, i.e.  $1 = \arg \max_{j \in A} \frac{v(j)}{c_j}$  and  $i = \arg \max_{j \in A \setminus [i-1]} \frac{v([j]) - v([j-1])}{c_j}$  for  $i \geq 2$ .

**Lemma 2** [7]. GREEDY-SM( $A, B/2$ ) is monotone and outputs a set  $S$  such that  $v(S) \geq \frac{e-1}{3e} \cdot \text{OPT}(A, B) - \frac{2}{3} \cdot v(i^*)$ . Using the payments of Myerson’s lemma, the mechanism is truthful, individually rational, and budget feasible.

MECHANISM-SM is deterministic and by using Lemma 2, it can be shown that it achieves an approximation factor of 8.34 for any nondecreasing submodular objective. However, it is not guaranteed to run in polynomial time, since we need to compute  $\text{OPT}(A \setminus \{i^*\}, B)$ , and more often than not, submodular maximization problems turn out to be NP-hard. An obvious question here is whether we can use an approximate solution instead, but it is not hard to see that by doing so we might sacrifice truthfulness. As a way out, Chen et al. [7] mention that instead of  $\text{OPT}(A \setminus \{i^*\}, B)$ , an optimal solution to a fractional relaxation of the problem can be used. Although this does not always make the mechanism run in polynomial time, it helps in some cases.

Suppose that for a specific submodular objective, the budgeted maximization problem can be expressed as an ILP, the corresponding LP relaxation of which can be solved in polynomial time. Further, suppose that for any instance  $I$  and any budget  $B$ , the optimal fractional solution  $\text{OPT}_f(I, B)$  is within a constant factor of the optimal integral solution  $\text{OPT}(I, B)$ . Then replacing  $\text{OPT}(A \setminus \{i^*\}, B)$  by  $\text{OPT}_f(A \setminus \{i^*\}, B)$  in MECHANISM-SM still gives a truthful, constant approximation. In fact, we give a variant of MECHANISM-SM below, where the constants have been appropriately tuned, so as to optimize the achieved approximation ratio. Specifically, suppose that the valuation function is such that  $\text{OPT}_f(I, B) \leq \rho \cdot \text{OPT}(I, B)$ , for any  $I$  and any  $B$ . Let  $\gamma = \sqrt{1 + 4(\rho - 1)e + 4(\rho^2 + 4\rho + 1)e^2}$  and  $\alpha = \frac{1 + 2(\rho + 1)e + \gamma}{2(e - 1)}$ .

---

MECHANISM-SM-FRAC( $A, B$ )

---

```

1 Set  $A = \{i \mid c_i \leq B\}$  and  $i^* \in \arg \max_{i \in A} v(i)$ 
2 if  $\alpha \cdot v(i^*) \geq \text{OPT}_f(A \setminus \{i^*\}, B)$  then
3   return  $i^*$ 
4 else
5   return GREEDY-SM( $A, B/2$ )

```

---

**Theorem 1.** MECHANISM-SM-FRAC is truthful, individually rational, and budget feasible with approximation ratio  $\frac{2(\rho+2)e-1+\gamma}{2(e-1)}$ . Moreover, it is deterministic and runs in polynomial time given a polynomial time exact algorithm for computing  $\text{OPT}_f(A \setminus \{i^*\}, B)$ .

Due to space constraints, the proof of Theorem 1 (as well as all the missing proofs in the subsequent sections) is deferred to the full version of the paper.

### 3.1 Budgeted Max Weighted Coverage

We consider the class of *weighted coverage valuations*, a special class of submodular functions. Their unweighted version was studied by Singer in [14] and [15], motivated by the problem of influence maximization over social networks. On a different note, the problem can also be thought of as a crowdsourcing problem, where each (single-minded) worker  $i$  is able to execute only the set of tasks  $S_i$ .

*Budgeted Max Weighted Coverage.* Given a set of subsets  $\{S_i \mid i \in [m]\}$  of a ground set  $[n]$ , along with costs  $c_1, c_2, \dots, c_m$ , on the subsets, weights  $w_1, \dots, w_n$ , on the ground elements, and a positive budget  $B$ , find  $X \subseteq [m]$  so that  $v(X) = \sum_{j \in T} w_j$ , where  $T = \bigcup_{i \in X} S_i$ , is maximized subject to  $\sum_{i \in X} c_i \leq B$ .

In [15], Singer takes an approach similar to what led to MECHANISM-SM-FRAC, but suggests a different polynomial time mechanism for Budgeted Max Coverage that is deterministic, truthful, budget feasible, and achieves approximation ratio 31.03. Here we generalize and improve this result by showing that there is a deterministic, truthful, budget feasible, polynomial time 15.45-approximate mechanism for the Budgeted *Weighted* Max Coverage problem.

For all  $j \in [n]$  define  $T_j = \{i \mid j \in S_i\}$ . We begin with a LP formulation of this problem, where without loss of generality we assume that  $c_i \leq B, \forall i \in [m]$  (otherwise we could just discard any subsets with cost greater than  $B$ ).

$$\text{maximize: } \sum_{j \in [n]} w_j z_j \tag{1}$$

$$\text{subject to: } \sum_{i \in T_j} x_i \geq z_j, \quad \forall j \in [n] \tag{2}$$

$$\sum_{i \in [m]} c_i x_i \leq B \tag{3}$$

$$0 \leq x_i, z_j \leq 1, \quad \forall i \in [m], \forall j \in [n] \tag{4}$$

$$x_i \in \{0, 1\}, \quad \forall i \in [m] \tag{5}$$

It is not hard to see that (1)–(5) is a natural ILP formulation for Budgeted Max Weighted Coverage and (1)–(4) is its linear relaxation. For the rest of this subsection, let  $\text{OPT}(I, B)$  and  $\text{OPT}_f(I, B)$  denote the optimal solutions to (1)–(5) and (1)–(4) respectively for instance  $I$  and budget  $B$ .

To show how these two are related we use the technique of pipage rounding [1, 2]. We should note here that Ageev and Sviridenko [2] use the above linear programs to obtain a (non-truthful)  $\frac{e}{e-1}$ -approximation LP-based algorithm that uses pipage rounding on a number of different instances of the problem. However, in their algorithm  $\text{OPT}(I, B)$  is never compared directly to  $\text{OPT}_f(I, B)$ , and therefore we cannot get the desired bound from there.

**Lemma 3.** *Given the fractional relaxation (1)–(4) for Budgeted Max Weighted Coverage, we have that for any instance  $I$  and any budget  $B$ ,  $\text{OPT}_f(I, B) \leq \frac{2e}{e-1} \cdot \text{OPT}(I, B)$ .*

Combining Theorem 1 and Lemma 3 we get the following result.

**Corollary 1.** *There exists a deterministic, truthful, individually rational, budget feasible 15.45-approximate mechanism for Budgeted Max Weighted Coverage that runs in polynomial time.*

## 4 Going Beyond Submodularity

Going beyond submodular valuations is even more challenging. The first attempt with a non-submodular objective was due to Chen et al. [7], who gave a  $(2 + \sqrt{2})$ -approximation mechanism for a non-submodular variation of Knapsack. For the more general class of subadditive functions Dobzinski et al. [9] suggested a randomized  $O(\log^2 n)$ -approximation mechanism, and later, Bei et al. [5] provided randomized, truthful, budget feasible mechanisms with approximation ratio 768 for XOS objectives and  $O\left(\frac{\log n}{\log \log n}\right)$  for subadditive objectives.

More recently, Goel et al. [11] study a budgeted maximization problem with matching constraints, which is not submodular, and they achieve an approximation ratio of  $3 + o(1)$  with a deterministic mechanism, but under the large market assumption<sup>1</sup> (their mechanism has an unbounded ratio in general). Essentially, they use the same greedy approach with Singer [14] and Chen et al. [7] but seen as a descending price auction. A very similar mechanism was also briefly discussed in Anari et al. [4] for Knapsack under the large market assumption.

We are building on this idea of gradually decreasing a global upper bound on the payment per value ratio to get all the results of this section. We first use Budgeted Max Weighted Matching in Subsect. 4.1, as an illustrative example of how this approach works, but the exact same approach gives the same approximation guarantees for a number of different XOS problems that can be seen as appropriately restricted generalizations of Knapsack. We elaborate further on this in Subsect. 4.2, and we even extend these ideas to problems where the unbudgeted versions are not easy.

<sup>1</sup> A market is said to be large if the number of participants is large enough that no single person can affect significantly the market outcome, i.e.  $\max_i c_i/B = o(1)$ .

## 4.1 Budgeted Max Weighted Matching

We revisit the following budgeted matching problem.

*Budgeted Max Weighted Matching.* Given a budget  $B$ , and a graph  $G = (V, E)$ , where each edge  $e_i \in E$  has a cost  $c_i$  and a value  $v_i$ , find a matching  $M$  of maximum value subject to  $\sum_{i \in M} c_i \leq B$ .

Here we study the mechanism design version of the problem, where the values are known to the mechanism and the edges are viewed as single-parameter strategic agents whose cost is private information.<sup>2</sup> Note that in order to formulate the problem to fit the general description given in the beginning of Sect. 2, we can define the valuation function as follows (as also mentioned in [5]): for any subset of edges  $S \subseteq E$ ,  $v(S)$  is taken to be the value of the maximum weighted matching of  $G$  that only uses edges in  $S$ . This function turns out to be XOS, but not submodular. Hence, by [5], there exists a randomized, 768-approximation, that is truthful and budget feasible.

We provide both deterministic and randomized polynomial time mechanisms with a much improved approximation ratio, based on selecting an outcome among two candidate solutions. The first solution comes from the greedy mechanism GREEDY-ISK described below. The main idea behind the mechanism is that in each iteration there is an implicit common upper bound on the rate that determines the payment of each winner in the candidate outcome of that iteration. More specifically, if the  $i$ th iteration is the final iteration (i.e. the condition in line 5 is true), the common payment per value for each of the winners is upper bounded by  $\min\{B/v(M), c_{i-1}/v_{i-1}\}$ . This upper bound decreases with each iteration, while the set of active agents is shrinking, until budget feasibility is achieved. At the same time we ensure the mechanism is monotone and returns enough value.

---

GREEDY-ISK( $A, v, c, B, f$ )

---

```

1 Set  $A = \{i \mid c_i \leq B\}$ 
2 Possibly rename elements of  $A$  so that  $\frac{c_1}{v_1} \geq \frac{c_2}{v_2} \geq \dots \geq \frac{c_m}{v_m}$ 
3 for  $i = 1$  to  $m$  do
4    $M = f(A, v)$ 
5   if  $v(M) \cdot \frac{c_i}{v_i} \leq B$  then
6      $\lfloor$  return  $M$ 
7   else
8      $\lfloor A = A \setminus \{i\}$ 

```

---

We assume that the mechanism also takes as input a deterministic exact algorithm  $f$  for the unbudgeted Max Weighted Matching, e.g., Edmond's algorithm [10]. Later, in Subsect. 4.2 the choice of  $f$  will depend on the underlying

<sup>2</sup> The work of Singer [14] also studies a type of a budgeted matching problem. That objective, however, is OXS (a subclass of submodular objectives), and differs significantly from ours, which is not submodular [16].

unbudgeted problem. Finally, note that our mechanisms are named after the generalization we study in Subsect. 4.2, namely Independence System Knapsack problems.

We now exhibit some desirable properties of GREEDY-ISK, starting with truthfulness.

**Lemma 4.** *Mechanism GREEDY-ISK is monotone, and hence truthful and individually rational.*

*Proof.* By Lemma 1, we just need to show that the allocation rule is monotone, i.e. a winning agent remains a winner if he decreases his cost. Initially note that in line 4 the mechanism computes an optimal matching  $M$  (without a budget constraint) using only the values of the edges, thus it cannot be manipulated given the set of active edges  $A$ .

Fix a vector  $c_{-j}$  for the costs of the other agents, and suppose that when agent  $j$  declares  $c_j$ , he is in the matching  $M$  returned in the final iteration, say  $k$ , of GREEDY-ISK. Let agent  $j$  now report  $c'_j < c_j$  to the mechanism. This makes him agent  $j' \geq j$  in the new instance, but does not affect the relative ordering of the other agents (although a few of them may move down one position). Therefore, GREEDY-ISK will run exactly as before for each iteration  $i < k$  and in the beginning of the  $k$ th iteration, it will produce the exact same matching  $M$ . Then in line 5, there are 2 cases to examine. If in the initial instance  $j > k$ , then we have the exact same ratio  $\frac{c_k}{v_k}$  to consider, and the algorithm will terminate with  $M$  (since it did so in the initial instance). In the second case,  $j = k$  in the initial instance. This means that now at the  $k$ th iteration, we either have the same agent with the reduced ratio  $\frac{c'_k}{v_k}$  (since now  $c'_k = c'_j$ ) or we have the agent who was in position  $k + 1$  in the initial instance with ratio equal to the original  $\frac{c_{k+1}}{v_{k+1}}$ . Therefore, the new ratio  $\frac{c_k}{v_k}$  that the algorithm considers in this iteration is at most equal to the original ratio  $\frac{c_k}{v_k}$ . Thus, the condition in line 5 is satisfied, and the mechanism will return  $M$ . We conclude that an agent who is in the matching, remains in the matching by decreasing his cost.  $\square$

We also make the following remark, which can be derived by the same arguments used in the proof of Lemma 4. This property is crucial for derandomizing our mechanisms both here and in the next subsection.

*Remark 1.* There is no agent  $i$  that can manipulate the output set of GREEDY-ISK given that  $i$  is guaranteed to be a winner. That is, fix  $c_{-i}$  and let  $M$  and  $M'$  be the winning sets when  $i$  bids  $c_i$  and  $c'_i$  respectively; if  $i \in M \cap M'$ , then  $M = M'$ .

We move on to prove that the mechanism will never exceed the budget  $B$ , by establishing an appropriate upper bound on every winning bid.

**Lemma 5.** *Mechanism GREEDY-ISK is budget feasible.*

*Proof.* We will show that the threshold payment of Lemma 1 cannot be higher than  $\frac{v_i B}{v(M)}$  for any winning agent  $i$ . Fix a vector  $c_{-i}$  for all agents other than  $i$  and recall that the threshold payment, given  $c_{-i}$ , is the maximum cost that  $i$  can declare and still be included in the solution. So, towards a contradiction, suppose that agent  $i$  declares a cost  $c_i > \frac{v_i B}{v(M)}$  and he is a winner. Let  $j$  denote the iteration where the mechanism GREEDY-ISK terminates and the matching  $M$  is returned. By the construction of the mechanism, and since  $i \in M$ , we have that  $\frac{c_j}{v_j} \geq \frac{c_i}{v_i}$ . Since  $j$  is the last iteration, we also have by line 5 that  $v(M) \frac{c_j}{v_j} \leq B$ . Hence  $v(M) \frac{c_i}{v_i} \leq v(M) \frac{c_j}{v_j} \leq B$  that leads to the contradiction  $c_i \leq \frac{v_i B}{v(M)}$ . Therefore, the payment of each winning agent  $i$  is bounded by  $\frac{v_i B}{v(M)}$ , and the total payment of the mechanism is  $\sum_{i \in M} p_i \leq \sum_{i \in M} \frac{v_i B}{v(M)} = B$ .  $\square$

Finally, we analyze the quality of the solution produced by GREEDY-ISK.

**Lemma 6.** *Mechanism GREEDY-ISK produces a matching with value at least  $\frac{1}{2}(v(M^*) - v_{i^*})$ , where  $M^*$  is an optimal solution to the given instance of Budgeted Max Weighted Matching, and  $i^*$  has maximum value among the budget feasible edges of  $G$ , i.e.  $i^* \in \arg \max_{i \in F} v(i)$  where  $F = \{i \in E(G) \mid c_i \leq B\}$ .*

We can now state our randomized mechanism for the problem (where the constants below have been optimized to get the best ratio).

---

|   |
|---|
| <b>RAND-ISK</b>   |
| <ol style="list-style-type: none"> <li>1 Set <math>A = \{i \mid c_i \leq B\}</math> and <math>i^* \in \arg \max_{i \in A} v(i)</math></li> <li>2 With probability 1/3 return <math>i^*</math> and with probability 2/3 return GREEDY-ISK(<math>A, v, c, B, f</math>)</li> </ol> |

---

**Theorem 2.** *RAND-ISK is a universally truthful, individually rational, budget feasible, polynomial time randomized mechanism, achieving a 3-approximation in expectation, for the Budgeted Max Weighted Matching problem.*

*Proof.* Universal truthfulness and individual rationality follow from Lemma 4 and the fact that the simple mechanism that returns  $i^*$  and pays him  $B$  is truthful and individually rational. Regarding budget feasibility, just notice that if  $i^*$  is returned then the threshold payment is exactly  $B$ , otherwise the payments of GREEDY-ISK are used, so budget feasibility follows from Lemma 5. Finally, if  $M$  is the outcome of RAND-ISK, then directly by Lemma 6 we have  $E(M) \geq \frac{2}{3} \cdot \frac{1}{2}(v(M^*) - v_{i^*}) + \frac{1}{3}v_{i^*} = \frac{1}{3}v(M^*)$ , thus proving the approximation ratio.  $\square$

**Derandomization.** We close this subsection by providing a deterministic polynomial time mechanism with a slightly worse approximation ratio. Note that in contrast to MECHANISM-SM or MECHANISM-SM-FRAC, here  $i^*$  is directly compared to its alternative, which is just an approximate solution, without sacrificing truthfulness. This is due to Remark 1. Moreover, although taking the maximum of two truthful algorithms does not always yield a truthful mechanism, this is the case for the mechanism below.

---

DET-ISK

---

```

1 Set  $A = \{i \mid c_i \leq B\}$  and  $i^* \in \arg \max_{i \in A} v(i)$ 
2 if  $v_{i^*} \geq \text{GREEDY-ISK}(A \setminus \{i^*\}, v, c_{-i^*}, B, f)$  then
3   return  $i^*$ 
4 else
5   return  $\text{GREEDY-ISK}(A \setminus \{i^*\}, v, c_{-i^*}, B, f)$ 

```

---

**Theorem 3.** *DET-ISK is a truthful, individually rational, budget feasible, polynomial time deterministic mechanism, achieving a 4-approximation ratio for the Budgeted Max Weighted Matching problem.*

Although the proof is omitted, we should mention that the analysis of DET-ISK is tight, i.e. there exist instances where the value of the optimal solution is four times the value of the mechanism’s output.

*Remark 2.* Chen et al. [7] prove lower bounds for Knapsack, namely there is no deterministic (resp. randomized) truthful, budget feasible mechanism for Knapsack that achieves an approximation ratio better than  $1 + \sqrt{2}$  (resp. 2). These lower bounds hold here as well, because when the given graph  $G$  is a matching to begin with, Budgeted Max Weighted Matching reduces to Knapsack.

### 4.2 A Generalization to Other Objectives

Our approach can tackle a number of different problems that have certain structural similarities with Budgeted Max Weighted Matching. Here, we define a class of such problems for which GREEDY-ISK—given an appropriate subroutine  $f$ —produces truthful, individually rational, budget feasible mechanisms with good approximation guarantees. Two crucial properties of the matching problem were used in the previous subsection: (i) every subset of a matching is itself a matching, and (ii) the objective function becomes additive when restricted to matchings. These two properties is all we need, and note that (i) and (ii) are exactly what makes the set of matchings of a graph an *independence system*.

**Definition 2.** *An independence system is a pair  $(U, I)$ , where  $U$  is an arbitrary finite set and  $I \subseteq 2^U$  is a family of subsets, whose members are called the independent sets of  $U$  and satisfy:*

- (i)  $\emptyset \in I$
- (ii) If  $B \in I$  and  $A \subseteq B$ , then  $A \in I$ .

Below we define a variant of Knapsack where the feasible solutions are constrained to an independence system. This is a generalization of knapsack problems subject to matroid constraints, which are more common in the literature.

*Independence System Knapsack.* Given an independence system  $(U, I)$  with costs  $c_i$  and values  $v_i$  on the elements of  $U$ , as well as a budget  $B$ , find  $M \in I$  that maximizes  $\sum_{i \in M} v_i$  subject to  $\sum_{i \in M} c_i \leq B$ .

Note that for plain Knapsack  $U = [n], I = 2^{[n]}$ , while for Budgeted Max Weighted Matching  $U$  is the set of edges of a given graph  $G$  and  $I$  is the set of all matchings of  $G$ . There exist several other problems that are special cases of Independence System Knapsack, like

- *Budgeted Max Weighted Forest* where  $U$  is the set of edges of a given graph  $G$  and  $I$  is the set of acyclic subgraphs of  $G$ ,
- *Budgeted Max Weighted Matroid Member* where  $(U, I)$  is a matroid<sup>3</sup> (Budgeted Max Weighted Forest is a special case of this problem),
- *Budgeted Max Independent Set* where  $U$  is the set of vertices of a given graph  $G$  and  $I$  is the set of independence sets of  $G$ , and
- *Budgeted Max Weighted  $k$ -D-Matching* where  $U$  is the set of hyperedges of a  $k$ -uniform  $k$ -partite hypergraph  $H$  and  $I$  is the set of all  $k$ -dimensional matchings of  $H$ .

The following can be easily derived as in the case of Budgeted Max Weighted Matching.

**Lemma 7.** *Every problem that can be formulated as an Independence System Knapsack problem belongs to the class XOS.*

Clearly it is not always possible to find an optimal solution to Independence System Knapsack in polynomial time, even if we remove the budget constraint. Putting the running time aside, however, GREEDY-ISK combined with an exact algorithm  $f$  for the problem makes RAND-ISK (resp. DET-ISK) a 3-approximate randomized (resp. 4-approximate deterministic) truthful, individually rational, budget feasible mechanism.

Moreover, when the unbudgeted underlying problem is easy—as is the case for Max Weighted Matching, Max Weighted Forest, and Max Weighted Matroid Member—the mechanisms run in polynomial time. Even if the unbudgeted underlying problem is *NP*-hard, as long as there is a polynomial time  $\rho(n)$ -approximation we get  $O(\rho(n))$ -approximate, truthful, individually rational, budget feasible mechanisms, e.g., for Budgeted Max Weighted  $k$ -D-Matching this translates to a  $O(k)$ -approximation mechanism. Here,  $n$  is the size of the input, and we should mention that the independent sets of  $U$  may not be explicitly given. Typically we assume an *independence oracle* that decides for any  $X \subseteq U$  whether  $X \in I$ . However, note that in most of the cases of Independence System Knapsack mentioned above (with the exception of Budgeted Max Weighted Matroid Member) we are given a combinatorial, succinct representation of  $I$  and therefore there is no need to assume access to an oracle.

When using a  $\rho(n)$ -approximation algorithm we should adjust the probabilities in RAND-ISK, namely we should use  $\frac{2\rho(n)}{2\rho(n)+1}$  instead of  $2/3$  and  $\frac{1}{2\rho(n)+1}$  instead of  $1/3$ . Moreover, for both mechanisms and without loss of generality, we assume that for every  $i \in U$  we have  $\{i\} \in I$ , or else  $i$  can be excluded from the initial set  $A$  of active elements that is given as input to the mechanisms.

<sup>3</sup> A *matroid*  $(U, I)$  is an independence system that also has the *exchange property*: If  $A, B \in I$  and  $|A| < |B|$ , then there exists  $x \in B \setminus A$  such that  $A \cup \{x\} \in I$ .



Following closely the analysis of Subject. 4.1, we get the next theorem.

**Theorem 4.** *If a deterministic  $\rho(n)$ -approximation algorithm  $f$  for the unbudgeted version of Independence System Knapsack is given as an auxiliary input to GREEDY-ISK, then RAND-ISK (resp. DET-ISK) becomes a  $(2\rho(n) + 1)$ -approximate randomized (resp.  $(2\rho(n) + 2)$ -approximate deterministic) truthful, individually rational, budget feasible mechanism. Moreover, if  $f$  runs in polynomial time so do the mechanisms.*

Combining Theorem 4 with the polynomial time  $(k - 1)$ -approximation algorithm of Chan and Lau [6] for Max Weighted  $k$ -D-Matching, and the fact that Max Weighted Forest and Max Weighted Matroid Member (given a polynomial time independence oracle) can be solved in polynomial time (see, e.g., [8]), we get the following corollary.

**Corollary 2.** *We can obtain*

(i) *randomized 3-approximation mechanisms and deterministic 4-approximation mechanisms for Budgeted Max Weighted Forest and Budgeted Max Weighted Matroid Member that run in polynomial time.*

(ii) *randomized 3-approximation mechanisms and deterministic 4-approximation mechanisms for Budgeted Max Weighted Independent Set and Budgeted Max Weighted  $k$ -D-Matching.*

(iii) *for any  $k \geq 3$ , a randomized  $(2k - 1)$ -approximation mechanism and a deterministic  $2k$ -approximation mechanism for Budgeted Max Weighted  $k$ -D-Matching that run in polynomial time.*

*Remark 3.* Max Weighted Independent Set and Max Weighted  $k$ -D-Matching are not submodular, as was the case for Max Weighted Matching. Max Weighted Matroid Member (and thus Max Weighted Forest), on the other hand, is submodular and therefore the results of [7] apply. However, our approach significantly improves both the approximation ratio and the running time.

Naturally, Remark 2 applies here as well. For every problem stated in this section there is no deterministic (resp. randomized) truthful, budget feasible mechanism with better approximation ratio than  $1 + \sqrt{2}$  (resp. 2). These lower bounds are independent of any complexity assumption.

## 5 Conclusions

We have studied further the problem of designing truthful and budget feasible mechanisms for budgeted versions of well known optimization problems. Especially for the XOS problems we considered, only randomized mechanisms with very high approximation ratios were known prior to our result. There are still many interesting open problems that are worth further exploration in the context of budgeted mechanism design. First, for the case of submodular functions, even though we do have a better understanding for designing mechanisms given all the previous works, the current results are still not known to be tight. We also

want to stress that the literature has mostly considered nondecreasing submodular functions. Dobzinski et al. [9] gave a constant approximation mechanism for the Budgeted Max Cut problem, however it remains a very interesting problem for future work to obtain mechanisms for general non-monotone submodular valuations. Furthermore, for the XOS class, the picture is way more challenging. We would like to identify more problems that admit better approximation guarantees, even with exponential time mechanisms. A component that seems to be missing at the moment is a characterization of truthful and budget feasible mechanisms. We believe that obtaining characterization results would be crucial in resolving the above questions.

## References

1. Ageev, A.A., Sviridenko, M.I.: Approximation algorithms for maximum coverage and max cut with given sizes of parts. In: Cornuéjols, G., Burkard, R.E., Woeginger, G.J. (eds.) IPCO 1999. LNCS, vol. 1610, pp. 17–30. Springer, Heidelberg (1999). doi:[10.1007/3-540-48777-8\\_2](https://doi.org/10.1007/3-540-48777-8_2)
2. Ageev, A.A., Sviridenko, M.: Pipe rounding: a new method of constructing algorithms with proven performance guarantee. *J. Comb. Optim.* **8**(3), 307–328 (2004)
3. Amanatidis, G., Birmpas, G., Markakis, E.: Coverage, matching, and beyond: new results on budgeted mechanism design. CoRR abs/1610.00901 (2016), <https://arxiv.org/abs/1610.00901>
4. Anari, N., Goel, G., Nikzad, A.: Mechanism design for crowdsourcing: an optimal  $1-1/e$  competitive budget-feasible mechanism for large markets. In: 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18–21, 2014, pp. 266–275 (2014)
5. Bei, X., Chen, N., Gravin, N., Lu, P.: Budget feasible mechanism design: from prior-free to Bayesian. In: Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19–22, 2012, pp. 449–458 (2012)
6. Chan, Y.H., Lau, L.C.: On linear and semidefinite programming relaxations for hypergraph matching. *Math. Program.* **135**(1–2), 123–148 (2012)
7. Chen, N., Gravin, N., Lu, P.: On the approximability of budget feasible mechanisms. In: Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23–25, 2011, pp. 685–699 (2011)
8. Cook, W.J., Cunningham, W.H., Pulleyblank, W.R., Schrijver, A.: *Combinatorial Optimization*. John Wiley & Sons Inc., New York (1998)
9. Dobzinski, S., Papadimitriou, C.H., Singer, Y.: Mechanisms for complement-free procurement. In: Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5–9, 2011, pp. 273–282 (2011)
10. Edmonds, J.: Maximum matching and a polyhedron with 0, 1 vertices. *J. Res. Natl. Bur. Stan.* **69**(B), 125–130 (1965)
11. Goel, G., Nikzad, A., Singla, A.: Mechanism design for crowdsourcing markets with heterogeneous tasks. In: Proceedings of the Second AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2014, Pittsburgh, Pennsylvania, USA, November 2–4, 2014 (2014)

12. Horel, T., Ioannidis, S., Muthukrishnan, S.: Budget feasible mechanisms for experimental design. In: Pardo, A., Viola, A. (eds.) LATIN 2014. LNCS, vol. 8392, pp. 719–730. Springer, Heidelberg (2014). doi:[10.1007/978-3-642-54423-1\\_62](https://doi.org/10.1007/978-3-642-54423-1_62)
13. Myerson, R.: Optimal auction design. *Math. Oper. Res.* **6**(1), 58–73 (1981)
14. Singer, Y.: Budget feasible mechanisms. In: 51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, Las Vegas, Nevada, USA, October 23–26, 2010, pp. 765–774 (2010)
15. Singer, Y.: How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In: Proceedings of the Fifth International Conference on Web Search and Web Data Mining, WSDM 2012, Seattle, WA, USA, February 8–12, 2012, pp. 733–742 (2012)
16. Singer, Y.: Personal communication (2016)

# Strategic Network Formation with Attack and Immunization

Sanjeev Goyal<sup>1</sup>, Shahin Jabbari<sup>2</sup>(✉), Michael Kearns<sup>2</sup>, Sanjeev Khanna<sup>2</sup>,  
and Jamie Morgenstern<sup>2</sup>

<sup>1</sup> Faculty of Economics, University of Cambridge, Cambridge, UK  
sg472@cam.ac.uk

<sup>2</sup> Department of Computer and Information Sciences,  
University of Pennsylvania, Philadelphia, USA  
{jabbari,mkearns,sanjeev,jamiemor}@cis.upenn.edu

**Abstract.** Strategic network formation arises in settings where agents receive some benefit from their connectedness to other agents, but also incur costs for forming these links. We consider a new network formation game that incorporates an adversarial attack, as well as *immunization* or protection against the attack. An agent’s network benefit is the expected size of her connected component post-attack, and agents may also choose to immunize themselves from attack at some additional cost. Our framework can be viewed as a stylized model of settings where *reachability* rather than centrality is the primary interest (as in many technological networks such as the Internet), and vertices may be vulnerable to attacks (such as viruses), but may also reduce risk via potentially costly measures (such as an anti-virus software).

Our main theoretical contributions include a strong bound on the edge density at equilibrium. In particular, we show that under a very mild assumption on the adversary’s attack model, every equilibrium network contains at most only  $2n-4$  edges for  $n \geq 4$ , where  $n$  denotes the number of agents and this upper bound is tight. We also show that social welfare does not significantly erode: every non-trivial equilibrium with respect to several adversarial attack models asymptotically has social welfare at least as that of any equilibrium in the original attack-free model.

We complement our sharp theoretical results by a behavioral experiment on our game with over 100 participants, where despite the complexity of the game, the resulting network was surprisingly close to equilibrium.

## 1 Introduction

In network formation games, distributed and strategic agents receive benefit from their connectedness to others, but also incur some cost for forming these links. Much research in this area [4, 6, 9] studies the structure of equilibrium networks

---

The full version of this paper with all the omitted details is available at <https://arxiv.org/abs/1511.05196>.

formed as the result of various choices for the network benefit function, as well as the social welfare in equilibria. In many such games, the costs incurred from forming links are direct: each edge costs  $C_E > 0$  for an agent to purchase. Recently, motivated by scenarios as diverse as financial crises, terrorism and technological vulnerability, games with indirect connectivity costs have been considered: an agent's connections expose her to negative, contagious shocks.

We begin with the well-studied *reachability* network formation game [4], in which players purchase links to each other, and enjoy a network benefit equal to the size of their connected component in the formed graph. We modify this model by introducing an adversary who is allowed to examine the network, and choose a single vertex or player to attack. This attack then spreads throughout the entire connected component of the originally attacked vertex, destroying all of these vertices. Crucially however, players also have the option of purchasing *immunization* against attack. Thus the attack spreads only to those non-immunized (or *vulnerable*) vertices reachable from the originally attacked vertex. We examine several natural adversarial attacks such as an adversary that seeks to maximize destruction, an adversary that randomly selects a vertex for the start of infection and an adversary that seeks to minimize the social welfare of the network post-attack to name a few. A player's overall payoff is thus the expected size of her post-attack component, minus her edge and immunization expenditures.<sup>1</sup>

Our game can be viewed as a stylized model for settings where reachability rather than centrality is the primary interest in joining a network vulnerable to adversarial attack. Examples include technological networks such as the Internet, where packet transmission times are sufficiently low that being "central" [9] or a "hub" [6] is less of a concern, but in the presence of attacks such as viruses or DDoS, mere reachability may be compromised. Parties may reduce risks via costly measures such as anti-virus. In a financial setting, vertices might represent banks and edges credit/debt agreements. The introduction of an attractive but extremely risky asset is a threat or attack on the network that naturally seeks its largest accessible market, but can be mitigated by individual institutions adopting balance sheet requirements or leverage restrictions. In a biological setting, vertices could represent humans, and edges physical proximity or contact. The attack could be an actual biological virus that randomly infects an individual and spreads by physical contact through the network; again, individuals may have the option of immunization. While our simplified model is obviously not directly applicable to any of these examples in detail, we do believe our results provide some high-level insights about the strategic tensions in such scenarios.

---

<sup>1</sup> The spread of the initial attack to reachable non-immunized vertices is deterministic in our model, and the protection of immunized vertices is absolute. It is also natural to consider relaxations such as probabilistic attack spreading and imperfect immunization, as well as generalizations such as multiple initial attack vertices. However, as we shall see, even the basic model we study here exhibits substantial complexity. We refer the reader to the full version for a discussion on possible extensions/relaxations.

Immunization against attack has recently been studied in games played on a network where risk of contagious shocks are present [7] but only in the setting in which the network is first designed by a centralized party, after which agents make individual immunization decisions. We endogenize both these aspects, which leads to a model incomparable to this earlier work.

The original reachability game [4] permitted a sharp and simple characterization of the equilibria: any tree as well as the empty graph. We demonstrate that once attack and immunization are introduced, the set of possible equilibria becomes considerably more complex, including networks that contain multiple cycles, as well as others which are disconnected but nonempty. This diversity leads to our primary questions of interest: How dense can equilibria become? In particular, does the presence of the attacker encourage the creation of massive redundancy of connectivity? Also does the introduction of attack and immunization result in dramatically lower social welfare compared to the original game?

**Our Results and Techniques.** The main theoretical contributions of this work are to show that our game still exhibits edge sparsity at equilibrium, and has high social welfare properties despite the presence of attacks. First we show that under a mild assumption on the adversary's attack model, the equilibrium networks with  $n \geq 4$  players have at most  $2n - 4$  edges, fewer than twice as many edges as any nonempty equilibria of the original game without attack. We prove this by introducing an abstract representation of the network and use tools from graph theory to upper bound the resources globally invested by the players to mitigate connectivity disruptions due to any attack.

We then show that with respect to several attack models, in any equilibrium with at least one edge and one immunized vertex, the resulting network is connected. This implies that any *new* equilibrium network (i.e. one which was not an equilibrium of the original reachability game) is either a sparse but connected graph, or is a forest of unimmunized vertices. The latter occurs only in the rather unnatural case where the cost of immunization or edges grows with the population size, and in the former case we further show the social welfare is at least  $n^2 - O(n^{5/3})$  – asymptotically the maximum possible with a polynomial rate of convergence. These results provide us with a complete picture of welfare in our model. We prove the welfare lower bound by showing that there cannot be many targeted vertices who are *critical* for global connectivity, where critical is defined formally in terms of both the vertex's probability of attack and the size of the components remaining after the attack. Thus players myopically optimizing their own utility create highly resilient networks in presence of attack.

We conclude by reporting on a behavioral experiment on our network formation game with over 100 participants, where despite the complexity of the game, the resulting network was surprisingly close to equilibrium.

**Organization.** We formally present our model and review some related work in Sect. 2. In Sect. 3 we briefly describe some interesting topologies that arise as equilibria and then prove our sparsity result. We present our lower bound on welfare in Sect. 4. Section 5 describe our behavioral experiment.

In the full version, we provide simulations demonstrating fast and general convergence of *swapstable* best response, a type of limited best response which generalizes linkstable best response but is more powerful in our game. The computational complexity of full best response dynamics was unknown to us at the time of conducting our simulations but this question has been recently studied by Ihde et al. [13]. The simulations illustrate a number of interesting further features of equilibria e.g. heavy-tailed degree distributions. Whether swapstable best response provably converges (as seen empirically) is an open question.

## 2 Model

We assume the  $n$  vertices of a graph (network) correspond to individual players. Each player has the choice to purchase edges to other players at a cost of  $C_E > 0$  per edge. Each player additionally decides whether to immunize herself at a cost of  $C_I > 0$  or remain *vulnerable*.

A (pure) *strategy* for player  $i$  (denoted by  $s_i$ ) is a pair consisting of the subset of players  $i$  purchased an edge to and her immunization choice. Formally, we denote the subset of edges which  $i$  buys an edge to as  $x_i \subseteq \{1, \dots, n\}$ , and the binary variable  $y_i \in \{0, 1\}$  as her immunization choice ( $y_i = 1$  when  $i$  immunizes). Then  $s_i = (x_i, y_i)$ . We assume that edge purchases are unilateral i.e. players do not need approval in order to purchase an edge to another but that the connectivity benefits and risks are bilateral. We restrict our attention to pure strategy equilibria and our results show they exist and are structurally diverse.

Let  $\mathbf{s} = (s_1, \dots, s_n)$  denote the strategy profile for all the players. Fixing  $\mathbf{s}$ , the set of edges purchased by all the players induces an undirected graph and the set of immunization decisions forms a bipartition of the vertices. We denote a game *state* as a pair  $(G, \mathcal{I})$ , where  $G = (V, E)$  is the undirected graph induced by the edges purchased by the players and  $\mathcal{I} \subseteq V$  is the set of players who decide to immunize. We use  $\mathcal{U} = V \setminus \mathcal{I}$  to denote the vulnerable vertices i.e. the players who decide not to immunize. We refer to a subset of vertices of  $\mathcal{U}$  as a *vulnerable region* if they form a maximally connected component. We denote the set of vulnerable regions by  $\mathcal{V} = \{\mathcal{V}_1, \dots, \mathcal{V}_k\}$  where each  $\mathcal{V}_i$  is a vulnerable region.

Fixing a game state  $(G, \mathcal{I})$ , the adversary inspects the formed network and the immunization pattern and chooses to attack some vertex. If the adversary attacks a vulnerable vertex  $v \in \mathcal{U}$ , then the attack starts at  $v$  and spreads, killing  $v$  and any other vulnerable vertices reachable from  $v$ . Immunized vertices act as “firewalls” through which the attack cannot spread. We point out that in this work we restrict the adversary to only pick one seed to start the attack.

More precisely, the adversary is specified by a function that defines a probability distribution over vulnerable regions. We refer to a vulnerable region with non-zero probability of attack as a *targeted region* and the vulnerable vertices inside of a targeted region as *targeted vertices*. We denote the targeted regions by  $\mathcal{T} = \{\mathcal{T}_1, \dots, \mathcal{T}_{k'}\}$  where each  $\mathcal{T}' \in \mathcal{T}$  denotes a targeted region.<sup>2</sup>

<sup>2</sup> The index  $k'$  in the definition of  $\mathcal{T}$  satisfies  $k' \leq k$  (see  $k$  in the definition of  $\mathcal{V}$ ).

$\mathcal{T} = \emptyset$  corresponds to the adversary making no attack, so player  $i$ 's *utility* (or *payoff*) is equal to the size of her connected component minus her expenses (edge purchases and immunization). When  $|\mathcal{T}| > 0$ , player's  $i$  expected utility (fixing a game state) is equal to the expected size of her connected component<sup>3</sup> less her expenditures, where the expectation is taken over the adversary's choice of attack (a distribution on  $\mathcal{T}$ ). Formally, let  $\Pr[\mathcal{T}']$  denote the probability of attack to targeted region  $\mathcal{T}'$  and  $CC_i(\mathcal{T}')$  the size of the connected component of player  $i$  post-attack to  $\mathcal{T}'$ . Then the expected utility of  $i$  in strategy profile  $\mathbf{s}$  denoted by  $u_i(\mathbf{s})$  is precisely

$$u_i(\mathbf{s}) = \sum_{\mathcal{T}' \in \mathcal{T}} \left( \Pr[\mathcal{T}'] CC_i(\mathcal{T}') \right) - |x_i|C_E - y_i C_I.$$

We refer to the sum of expected utilities of all the players playing  $\mathbf{s}$  as the (*social welfare*) of  $\mathbf{s}$ .

**Examples of Adversaries.** We highlight several natural adversaries that fit into our framework. We begin with a natural adversary whose goal is to maximize the number of agents killed.

**Definition 1.** *The maximum carnage adversary attacks the vulnerable region of maximum size. If there are multiple such regions, the adversary picks one of them uniformly at random. Once a targeted region is selected, the adversary selects a vertex inside of that region uniformly at random to start the attack.*

So a targeted region with respect to a maximum carnage adversary is a vulnerable region of maximum size and the adversary defines a uniform distribution over such regions (see Fig. 1). Another natural but less sophisticated adversary starts an attack by picking a vulnerable vertex at random.



**Fig. 1.** Blue and red vertices denote  $\mathcal{I}$  and  $\mathcal{U}$ , respectively. The probability of attack to the vulnerable regions denoted by  $\mathcal{V}_1, \mathcal{V}_2$  and  $\mathcal{V}_3$  (in that order) for each adversary are as follows. maximum carnage: 0.5, 0, 0.5; random attack: 0.4, 0.2, 0.4; maximum disruption: 0, 1, 0. (Color figure online)

**Definition 2.** *The random attack adversary attacks a vulnerable vertex uniformly at random.*

So every vulnerable vertex is targeted with respect to the random attack adversary and the adversary induces a distribution over targeted regions such that the probability of attack to a targeted region is proportional to its size (see Fig. 1). Lastly, we define another natural adversary whose goal is to minimize the post-attack welfare.

<sup>3</sup> If a vertex is killed, the size of her connected component is defined to be 0.



**Definition 3.** *The maximum disruption adversary attacks the vulnerable region which minimizes the post-attack social welfare. If there are multiple such regions, the adversary picks one of them uniformly at random. Once a targeted region is selected for the attack, the adversary selects a vertex inside of that region uniformly at random to start the attack.*

Thus the maximum disruption adversary only attacks those vulnerable regions which minimize the post-attack welfare and the adversary defines a uniform distribution over such regions (see Fig. 1).

**Equilibrium Concepts.** We analyze the networks formed in our game under two types of equilibria. We model each of the  $n$  players as strategic agents who choose deterministically which edges to purchase and whether or not to immunize, knowing the exogenous behavior of the adversary defined as above. We say a strategy profile  $\mathbf{s}$  is a *pure strategy Nash equilibrium* (Nash equilibrium for short) if, for any player  $i$ , fixing the behavior of the other players to be  $\mathbf{s}_{-i}$ , the expected utility for  $i$  cannot strictly increase playing any action  $\mathbf{s}'_i$  over  $\mathbf{s}_i$ .

In addition to Nash, we study another equilibrium concept that is closely related to linkstable equilibrium [5], a bounded-rationality generalization of Nash. We call this concept *swapstable equilibrium*.<sup>4</sup> A strategy profile is a swapstable equilibrium if no agent's expected utility (fixing other agents' strategies) can strictly improve under any of the following *swap deviations*: (1) dropping any single purchased edge, (2) purchasing any single unpurchased edge, (3) dropping any single purchased edge and purchasing any single unpurchased edge, (4) any one of the deviations above and also changing the immunization status.

The first two deviations correspond to the standard linkstability. The third permits the more powerful *swapping* of one purchased edge for another. The last additionally allows reversing immunization status. Our interest in swapstable networks derives from the fact that while they only consider "simple" or "local" deviation rules, they share several properties with Nash networks that linkstable networks do not. Hence, swapstability is a bounded rationality concept that moves us closer to full Nash. Intuitively, in our game (and in many of our proofs), we exploit the fact that if a player is connected to some other set of vertices via an edge to a targeted vertex, and that set also contains an immune vertex, the player would prefer to connect to the immune vertex instead. This deviation involves a swap not just a single addition or deletion. It is worth mentioning explicitly that by definition every Nash equilibrium is a swapstable equilibrium and every swapstable equilibrium is a linkstable equilibrium. The reverse of none of these statements are true in our game. Also the set of equilibrium networks with respect to adversaries defined in Definitions 1, 2 and 3 are disjoint.

## 2.1 Related Work

Our paper is a contribution to the study of strategic network design and defense. The problem has been extensively studied in economics, electrical engineering,

<sup>4</sup> Lenzner [17] introduced this equilibrium concept under the name *greedy equilibrium*.

and computer science (see e.g. [1, 2, 11, 18]). Most of the existing work takes the network as given and examines optimal security choices (see e.g. [3, 8, 12, 14, 16]). To the best of our knowledge, our paper offers the first model in which both links and defense (immunization) are chosen by the players.

Combining linking and immunization within a common framework yields new insights. We start with a discussion of the network formation literature. In a setting with no attack, our model with respect to the maximum carnage adversary reduces to the original model of one-sided reachability network formation of Goyal [4]. They showed that a Nash equilibrium network is either a tree or an empty network. By contrast, we show that in the presence of a security threat, Nash networks exhibit very different properties: both networks containing cycles and partially connected networks can emerge in equilibrium. We also show that while networks may contain cycles, they are sparse (we provide a tight upper bound on the number of links in any equilibrium network of our game).

Regarding security, a recent paper by Cerdeiro et al. [7] studies optimal design of networks in a setting where players make immunization choices against a maximum carnage adversary but the network design is given. They show that an optimal network is either a hub-spoke or a network containing  $k$ -critical vertices<sup>5</sup> or a partially connected network (a  $k$ -critical vertex can secure  $n - k$  vertices by immunization). We extend this work by showing that there is a pressure toward the emergence of  $k$ -critical vertices even when linking is decentralized. We also contribute to the study of welfare costs of decentralization. Cerdeiro et al. [7] show that the Price of Anarchy (PoA) is bounded, when the network is centrally designed while immunization is decentralized (their welfare measure includes the edge expenditures of the planner). By contrast, we show that the PoA is unbounded when both decisions are decentralized. Although we also show that non-trivial equilibrium networks with respect to various adversaries have a PoA very near 1. This highlights the key role of linking and resonates with the original results on the PoA of pure network formation games [10].

Recently Blume et al. [6] study network formation where new links generate direct (but not reachability) benefits, infection can flow through paths of connections and immunization is not a choice. They demonstrate a fundamental tension between socially optimal and stable networks: the former lie just below a linking threshold that keeps contagion under check, while the latter admit linking just above this threshold, leading to extensive contagion and very low payoffs.

Finally, Kliemann [15] introduced a reachability game with attacks but without defense. In their model, the attack happens after the network is formed and the adversary destroys exactly one *link* (with no spread) according to a probability distribution over links that can depend on the structure of the network. They show equilibrium networks are chord-free and hence sparse. We also show an abstract representation of equilibrium networks in our model corresponds to chord-free graphs and then use this observation to prove sparsity. While both models lead to chord-free graphs in equilibria, the analysis of *why* these graphs


---

<sup>5</sup> Vertex  $v$  is  $k$ -critical in a connected network if the size of the largest connected component after removing  $v$  is  $k$ .

are chord-free is quite different. In their model, the deletion of a single link destroys at most one path between any pair of vertices. So if there were two edge-disjoint paths between any pairs of vertices, they will certainly remain connected after any attack. In our model the adversary attacks a vertex and the attack can spread and delete many links. This leads to a more delicate analysis. The welfare analysis is also quite different, since the deletion of an edge can cause a network to have at most two connected components, while the deletion of vertices might lead to many connected components.

### 3 Sparsity

In contrast to the original game [4], our game exhibits equilibrium networks with cycles, as well as disconnected but non-empty graphs. Figure 2 gives several examples of Nash networks with respect to the maximum carnage adversary for small populations, each of which is representative of a broad family of equilibria for large populations and a range of values for  $C_E$  and  $C_I$ .<sup>6</sup> So the tight characterization of the original game, where equilibrium networks are either empty graph or trees, fails to hold for our game. However, we show that an approximate version of this characterization continues to hold for several adversaries.



**Fig. 2.** Examples of equilibria with respect to the maximum carnage adversary: (a) Forest equilibrium,  $C_E = 1$  and  $C_I = 9$ ; (b) cycle equilibrium,  $C_E = 1.5$  and  $C_I = 3$ ; (c) 4-petal flower equilibrium,  $C_E = 0.1$  and  $C_I = 3$ , (d) Complete bipartite equilibrium,  $C_E = 0.1$  and  $C_I = 4$ . (Color figure online)

We show that despite the existence of equilibria containing cycles as shown in Fig. 2, under a very mild restriction on the adversary, *any* Nash, swapstable or linkstable equilibrium network of our game has at most  $2n - 4$  edges and is thus quite sparse. Moreover, this upper bound is tight as the generalized complete bipartite graph in Fig. 2d has exactly  $2n - 4$  edges.

The rest of this section is organized as follows. We start by defining a natural restriction on the adversary. We then propose an abstract view of the networks in our game and proceed to show that the abstract network is chord-free in equilibria with respect to the restricted adversary. We finally derive the edge density of the original network by connecting its edge density to the density of the abstract network. We start by defining equivalence classes for networks.

<sup>6</sup> We represent immunized and vulnerable vertices as blue and red, respectively. Although we treat the networks as undirected (the benefits and risks are bilateral), we use directed edges in some figures to denote which player purchased the edge.


**Definition 4.** Let  $G_1 = (V, E_1)$  and  $G_2 = (V, E_2)$  be two networks.  $G_1$  and  $G_2$  are equivalent if for all vertices  $v \in V$ , the connected component of  $v$  is the same in both  $G_1$  and  $G_2$  for every possible choice of initial attack vertex in  $V$ .

Based on equivalence, we make the following natural restriction on the adversary.

**Assumption 1.** An adversary is well-behaved if on any pair of equivalent networks  $G_1 = (V, E_1)$  and  $G_2 = (V, E_2)$ , the probability that a vertex  $v \in V$  is chosen for attack, is the same.

The adversaries in Definitions 1–3 are all well-behaved. Next, we abstract the network formed by the agents and analyze the edge density in the abstraction.

Let  $G = (V, E)$  be any network,  $\mathcal{I} \subseteq V$  the immunized vertices and  $\mathcal{V}_1, \dots, \mathcal{V}_k$  the vulnerable regions in  $G$ . In the abstract network every vulnerable region in  $G$  is contracted to a single vertex. Formally, let  $G' = (V', E')$  be the abstract network. Define  $V' = \mathcal{I} \cup \{u_1, \dots, u_k\}$  where each  $u_i$  represents a contracted vulnerable region of  $G$ .  $E'$  is constructed as follows. For any edge  $(v_1, v_2) \in E$  such that  $v_1, v_2 \in \mathcal{I}$  there is an edge  $(v_1, v_2) \in E'$ . For any edge  $(v_1, v_2) \in E$  such that  $v_1 \in \mathcal{V}_i$  for some  $i$  and  $v_2 \in \mathcal{I}$  there is an edge  $(u_i, v_2) \in E'$  where  $u_i$  denotes the contracted vulnerable region of  $G$  that  $v_1$  belongs to. For any edge  $(v_1, v_2)$  such that  $v_1, v_2 \in \mathcal{V}_i$  for some  $i$  there is no edge in  $G'$  (see Fig. 3).



**Fig. 3.** Example of original and abstract network. Blue: immunized vertices in both networks. Red: the vulnerable vertices and regions in the original and abstract network, respectively. (Color figure online)

We next show that if  $G$  is an equilibrium network then  $G'$  is a chord-free graph. We defer all the omitted proofs to the full version.

**Lemma 1.** Let  $G = (V, E)$  be a Nash, swapstable or linkstable equilibrium network and  $G' = (V', E')$  the abstraction of  $G$ . Then  $G'$  is a chord-free graph if the adversary is well-behaved.

As the next step we bound the edge density of chord-free networks in Theorem 1 using tools from the graph theory literature.

**Theorem 1.** Let  $G = (V, E)$  be a chord-free graph on  $n \geq 4$  vertices. Then  $|E| \leq 2n - 4$ .<sup>7</sup>

<sup>7</sup> Kliemann [15] proved Theorem 1 with a different technique for a density bound of  $2n - 1$  for all  $n$ .

Theorem 1 implies the edge density of the abstract network  $G' = (V', E')$  is at most  $2|V'| - 4$ . To derive the edge density of the original network, we first show that any vulnerable region in  $G$  is a tree when  $G$  is an equilibrium network.

**Lemma 2.** *Let  $G = (V, E)$  be a Nash, swapstable or linkstable equilibrium network. Then any vulnerable region in  $G$  is a tree if the adversary is well-behaved.*

We use Lemmas 1, 2 and Theorem 1 to prove our sparsity result.

**Theorem 2.** *Let  $G = (V, E)$  be a Nash, swapstable or linkstable equilibrium network on  $n \geq 4$  vertices. Then  $|E| \leq 2n - 4$  for any well-behaved adversary.*

## 4 Connectivity and Social Welfare in Equilibria

The results of Sect. 3 show that despite the potential presence of cycles at equilibrium, there are still sharp limits on collective expenditure on edges. However, they do not directly lower bound the welfare, due to connectivity concerns: if the graph could become highly fragmented after the attack, or is sufficiently fragmented prior to the attack, the reachability benefits to players could be sharply lower than in the attack-free reachability game. We now show that when  $C_I$  and  $C_E > 1$  are both constants with respect to  $n$ ,<sup>8</sup> none of these concerns are realized in any “interesting” equilibrium network, described precisely below.

In the original reachability game [4], the *maximum* welfare achievable in any equilibrium is  $n^2 - O(n)$ . Here we will show that the welfare achievable in any “non-trivial” equilibrium is  $n^2 - O(n^{5/3})$ . Obviously with no restrictions on the adversary and the parameters this cannot be true. Just as in the original game, for  $C_E > 1$ , the empty graph with a social welfare of only  $O(n)$  remains an equilibrium in our game with respect to all the natural adversaries in Sect. 2. We thus assume the equilibrium network contains at least *one* edge and at least *one* immunized vertex. We refer to all equilibrium networks that satisfy the above assumption as *non-trivial* equilibria. They capture the equilibria that are new to our game compared to the original attack-free setting — the network is not empty, and at least one player has chosen immunization.

Limiting attention to non-trivial equilibria is *necessary* if we hope to guarantee that the welfare at equilibrium is  $\Omega(n^2)$  when  $C_E > 1$ . As already noted, without the edge assumption, the empty graph is an equilibrium with respect to several natural adversaries. Furthermore, without the immunization assumption,  $n/3$  disjoint components where each component consists of 3 vulnerable vertices is an equilibrium (for carefully chosen  $C_E$  and  $C_I$ ) with respect to e.g. the maximum carnage adversary. In both cases, the social welfare is only  $O(n)$ .

Similar to Sect. 3, to get any meaningful results for the welfare we need to restrict the adversary. To simplify presentation, we state and analyze our results for the maximum carnage adversary. We later show how these results (or their slight modifications) can be extended to several other adversaries.

<sup>8</sup> We view this condition as the most interesting regime, since in natural circumstances we do not expect the edge or immunization costs to grow with the population size.

Consider any connected component that contains an immunized vertex and an edge in a non-trivial equilibrium network with respect to the maximum carnage adversary. We first show that any targeted region in such component (if exists) has size 1 when  $C_E > 1$ .

**Lemma 3.** *Let  $G$  be a non-trivial Nash or swapstable equilibrium network with respect to the maximum carnage adversary. Then in any component of  $G$  with at least one immunized vertex and at least one edge, the targeted regions (if they exist) are singletons when  $C_E > 1$ .*

We then show that non-trivial equilibrium networks with respect to the maximum carnage adversary are connected when  $C_E > 1$ .

**Theorem 3.** *Let  $G$  be a non-trivial Nash, swapstable or linkstable equilibrium network with respect to the maximum carnage adversary. Then,  $G$  is a connected graph when  $C_E > 1$ .*

So any non-trivial equilibrium network with respect to maximum carnage adversary is a connected network with targeted regions of size 1. Finally, we state our main result regarding the welfare in such non-trivial equilibria.

**Theorem 4.** *Let  $G$  be a non-trivial Nash or swapstable equilibrium network on  $n$  vertices with respect to the maximum carnage adversary. If  $C_E$  and  $C_I$  are constants (independent of  $n$ ) and  $C_E > 1$  then the welfare of  $G$  is  $n^2 - O(n^{5/3})$ .*

Our proof techniques for Theorem 4 might not extend to non-trivial linkstable equilibrium networks with respect to the maximum carnage adversary since such networks can have targeted regions of size bigger than 1 when  $C_E > 1$ .

**Remarks.** We proved our sparsity result with a rather mild restriction on the adversary. However, we presented our welfare results only with respect to the maximum carnage adversary. Our proofs in this section only rely on the following two properties of the maximum carnage adversary: (1) Adding an edge between any 2 vertices (at least 1 of which is immunized) does not change the distribution of the attack and (2) Breaking a link inside of a targeted region does not increase the probability of attack to the targeted region while at the same time does not decrease the probability of attack to any other vulnerable region. The same properties hold for the random attack adversary and other adversaries that set the probability of attack to a vulnerable region directly proportional to an increasing function of the size of the region. Thus our welfare results extend to random attack adversary and other such adversaries without any modifications.

However, some natural adversaries (e.g. the maximum disruption adversary) might not satisfy these properties. While the same techniques might not be directly applicable to such adversaries, it is still possible to reason about the welfare using different methods e.g. we can still show that in any non-trivial and *connected* equilibrium with respect to the maximum disruption adversary, when  $C_E$  and  $C_I$  are constants and  $C_E > 1$ , then the welfare is  $n^2 - O(n^{5/3})$ . See the full version for more details.

## 5 A Behavioral Experiment

To complement our theory, we conducted a behavioral experiment on our game with 118 participants. The participants were students in an undergraduate survey course on network science at the University of Pennsylvania. As training, participants were given a detailed document and lecture on the game, with simple examples of payoffs for players on small graphs under various edge purchase and immunization decisions. (See <http://www.cis.upenn.edu/~mkearns/teaching/NetworkedLife/NetworkFormationExperiment2015.pdf> for the training document provided to participants.) Participation was a course requirement, and students were instructed that their grade on the assignment would be exactly equal to their payoffs according to the rules of the game.

The payoffs used the maximum carnage adversary, with costs of  $C_E = 5$  and  $C_I = 20$ . With  $n = 118$  participants (so a maximum connectivity benefit of 118 points), it felt that these values made edge purchases and immunization significant expenses and thus worth careful deliberation. Second, running swapstable best response simulations using these values generally resulted in non-trivial equilibria with high welfare, whereas raising  $C_E$  and  $C_I$  significantly generally resulted in empty or fragmented graphs with low welfare.


In a game of such complexity, with so many participants, it is unreasonable and uninteresting to formulate the experiment as a one-shot simultaneous move game. Rather, some form of communication must be allowed. We chose to conduct the experiment in a courtyard with the single ground rule that *all conversations be quiet and local* i.e. in order to hear what a participant was saying to others, one should have to stand next to them.

Other than the quiet rule, there were no restrictions on the nature of conversations: participants were free to enter agreements, make promises or threats and move freely. However, it was made clear that any agreements or bargains struck would *not* be enforced by the rules of the experiment (thus were non-binding). Each subject was given a handout that required them to indicate which other subjects they chose to purchase edges to (if any), and whether or not they chose to purchase immunization. The handout contained a list of subject names, along with a unique identification number for each subject used to indicate edge purchases. Thus subjects knew the names of the others as well as their assigned ID numbers. An entire class session was devoted to the experiment, but subjects were free to (irrevocably) turn in their handout at any time and leave. Subjects committed and exited sequentially, and the entire duration was approximately 30 min. During the experiment, subjects tended to gather quickly in small discussion groups that reformed frequently, with subjects moving freely from group to group. It is clear from the outcome that despite adherence to the quiet rule, the subjects engaged in widespread coordination via this rapid mixing.

In the left panel of Fig. 4, we show the final undirected network formed by the edge purchases and immunization decisions. The graph is clearly anchored by two main immunized hub vertices, each with many spokes who purchased their single edge to the respective hub. These two large hubs are both directly connected, as well as by a longer “bridge” of three vulnerable vertices. There is

also a smaller hub with just a handful of spokes, again connected to one of the larger hubs via a chain of two vulnerable vertices.

For the payoffs, inspection of the network reveals that there are 2 groups of 3 vertices that are the largest vulnerable connected components, and thus are the targets of the attack. These 6 players are each killed with probability  $1/2$  for a payoff that is only half that of the wealthiest players (the vulnerable spokes of degree 1). In between are the players who purchased immunization including the 3 hubs and 2 immunized spokes. The immunized spoke of the upper hub is unnecessarily so, while the immunized spoke in the lower hub is best responding — had they not purchased immunization, they would have formed a unique largest vulnerable component of size 4 and thus been killed with certainty.



**Fig. 4.** Left: the final undirected network formed by the edge purchases and immunization decisions (blue for immunized, red for vulnerable). Right: a “nearby” Nash network. (Color figure online)

It is striking how many properties the behavioral network shares with the theory: multiple hub-spoke structures with sparse connecting bridges, resulting in high welfare and a heavy-tailed degree distribution; a couple of cycles. To quantify how far the behavioral network is from equilibrium we use it as the starting point for swapstable best response dynamics and run it until convergence. In the right panel of Fig. 4, we show the resulting Nash network reached from the behavioral network in only 4 rounds of swapstable dynamics, and with only 15 of 118 vertices updating their choices. The dynamics simply *clean up* some suboptimal behavioral decisions — the vulnerable bridges between hubs are replaced by direct edges, the other targeted group of three spokes drops their fatal edges, and immunizing spokes no longer do so.

Participants were required to complete a survey after the experiment: they were asked to comment on any strategies they contemplated prior to the experiment; whether and how those strategies changed during the experiment; and what strategies or behaviors they observed in other participants.

Many subjects reported entering the experiment with not just a strategy for themselves, but also a “master plan” they hoped to convince others to join. One frequently reported plan involved variations on cycles. Though little thought seems to have been given to how to coordinate a global ordering in a cycle via only the quiet rule. Another frequently cited plan involved the hub-spoke. Although most strategies are based on abstractions, others reported planning to use real-world social relationships e.g. connecting to students they know.

Of course, of particular interest are the surveys of the hubs. One seems to report an altruistic motivation for purchasing immunization, hoping to maximize



welfare. In contrast, the other displays a more Machiavellian attitude and was willing to immunize in the hopes of creating 3 distinct groups of participants: the “winners” who would connect to the hub; the hub with slightly lower payoff; a large group of “losers” deliberately left out of the hub-spoke structure.

It is clear from the surveys that the word quickly spread during the experiment to connect to hubs and that many participants joined though not without some reported mistrust and hesitation.


**Acknowledgments.** We thank Chandra Chekuri, Yang Li and Aaron Roth for useful suggestions. We also thank anonymous reviewers for detailed suggestions regarding some of the proofs. Sanjeev Khanna is supported in part by National Science Foundation grants CCF-1552909, CCF-1617851, and IIS-1447470.

## References

1. Alpcan, T., Baar, T.: *Network security: a decision and game-theoretic approach*, 1st edn. Cambridge University Press, Cambridge (2010)
2. Anderson, R.: *Security engineering: a guide to building dependable distributed systems*, 2nd edn. Wiley Publishing (2008)
3. Aspnes, J., Chang, K., Yampolskiy, A.: Inoculation strategies for victims of viruses and the sum of squares partition problem. *J. Comput. Syst. Sci.* **72**(6), 1077–1093 (2006)
4. Bala, V., Goyal, S.: A noncooperative model of network formation. *Econometrica* **68**(5), 1181–1230 (2000)
5. Bloch, F., Jackson, M.: Definitions of equilibrium in network formation games. *Int. J. Game Theo.* **34**(3), 305–318 (2006)
6. Blume, L., Easley, D., Kleinberg, J., Kleinberg, R., Tardos, É.: Network formation in the presence of contagious risk. In: *EC*, pp. 1–10 (2011)
7. Cerdeiro, D., Dziubinski, M., Goyal, S.: Contagion risk and network design. Working Paper (2014)
8. Cunningham, W.: Optimal attack and reinforcement of a network. *J. ACM* **32**(3), 549–561 (1985)
9. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C., Shenker, S.: On a network creation game. In: *PODC*, pp. 347–351 (2003)
10. Goyal, S.: *Connections: an introduction to the economics of networks*. Princeton University Press, Princeton (2007)
11. Goyal, S.: *Conflicts and Networks*. The Oxford Handbook on the Economics of Networks (2015)
12. Gueye, A., Walrand, J., Anantharam, V.: A network topology design game, how to choose communication links in an adversarial environment. In: *GameNets* (2011)
13. Ihde, S., Keßler, C., Neubert, S., Schumann, D., Lenzner, P., Friedrich, T.: Efficient best-response computation for strategic network formation under attack. *CoRR* abs/1610.01861 (2016)
14. Kearns, M., Ortiz, L.: Algorithms for interdependent security games. In: *NIPS*, pp. 561–568 (2003)
15. Kliemann, L.: The price of anarchy for network formation in an adversary model. *Games* **2**(3), 302–332 (2011)

16. Laszka, A., Szeszlér, D., Buttyán, L.: Linear loss function for the network blocking game: an efficient model for measuring network robustness and link criticality. In: Grossklags, J., Walrand, J. (eds.) *GameSec 2012*. LNCS, vol. 7638, pp. 152–170. Springer Berlin Heidelberg, Berlin, Heidelberg (2012). doi:[10.1007/978-3-642-34266-0\\_9](https://doi.org/10.1007/978-3-642-34266-0_9)
17. Lenzner, P.: Greedy selfish network creation. In: Goldberg, P.W. (ed.) *WINE 2012*. LNCS, vol. 7695, pp. 142–155. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-35311-6\\_11](https://doi.org/10.1007/978-3-642-35311-6_11)
18. Roy, S., Ellis, C., Shiva, S., Dasgupta, D., Shandilya, V., Wu, Q.: A survey of game theory as applied to network security. In: *HICSS*, pp. 1–10 (2010)

# Opinion Formation Games with Dynamic Social Influences

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

<sup>1</sup> Department of Mathematics and Physics, University of Salento, Lecce, Italy  
vittorio.bilo@unisalento.it

<sup>2</sup> CNRS, (UMR-6211), Caen, France  
angelo.fanelli@unicaen.fr

<sup>3</sup> Department of Economic Studies, University of Chieti-Pescara, Pescara, Italy  
luca.moscardelli@unich.it

**Abstract.** We introduce opinion formation games with dynamic social influences, where opinion formation and social relationships co-evolve in a cross-influencing manner. We show that these games always admit an ordinal potential, and so, pure Nash equilibria, and we design a polynomial time algorithm for computing the set of all pure Nash equilibria and the set of all social optima of a given game. We also derive non-tight upper and lower bounds on the price of anarchy and stability which only depend on the players' stubbornness, that is, on the scaling factor used to counterbalance the cost that a player incurs for disagreeing with the society and the cost she incurs for disagreeing with her innate beliefs.

## 1 Introduction

*Opinion formation* is a sociological process by which an individual, possibly starting from her innate viewpoint, shapes her belief on a certain subject as a result of the interaction with others (*social influence*).

Several interesting models have been proposed in the literature to assess this phenomenon. In the seminal *DeGroot model* [4], each individual  $i$  has an opinion  $z_i$ , lying on a real line, which is iteratively updated to the average of the opinions expressed by her acquaintances, e.g., neighbors in a social network. Subsequent models, as the *HK model* by Hegselmann and Krause [8] and the *DW model* by Weisbuch et al. [13], restrict the social influence to only those individuals whose expressed opinion is within a certain distance to  $z_i$  (the *confidence region* of individual  $i$ ). The *FJ model* by Friedkin and Johnsen [7] assumes that individual  $i$  also has an innate opinion  $s_i$  and  $i$ 's expressed opinion is then updated by counterbalancing the effects of the social influence with the disagreement between  $z_i$  and  $s_i$ .

All the above models share the common assumption that the social influence each individual has to undergo remains fixed during the whole duration of the

---

This work was partially supported by the project ANR-14-CE24-0007-01 “CoCoRiCo-CoDec”.

process, e.g., the social network is a static graph. This assumption has been relaxed in some recent works [1, 5, 9, 10] which are based on the evidence that opinion formation and friends selection are often **co-evolving processes** in real life. In particular, Holme and Newman [10] consider the *DeGroot model* (and its generalizations) in which at each step a certain individual  $i$  is selected and ( $i$ ) with probability  $\alpha$ ,  $i$  replaces a random individual from her set of acquaintances with a random individual from the set of people whose expressed opinion coincides with  $z_i$ ; ( $ii$ ) with probability  $1 - \alpha$ , a random individual in the set of  $i$ 's acquaintances changes her opinion to  $z_i$ . Bhawalkar et al. [1], instead, consider the *FJ model* in which the disagreement with the innate opinion and the social influence are both expressed as individual's specific functions; moreover, for a given positive integer  $K$ , they also investigate the variant in which, for each individual  $i$ , the set of acquaintances is formed by the  $K$  individuals whose expressed opinion is at minimum distance from  $s_i$ .

The co-evolutionary opinion formation models of Holme and Newman [10] and Bhawalkar et al. [1] still assume that the underlying social relationships are not completely dynamic, as they only allow for an individual's set of acquaintances to vary over time. Quantitatively speaking, this means that the social influence that an individual exercises on somebody else can only have a dichotomic behavior: it may appear or disappear, but, whenever present, its magnitude remains fixed.

Since real-life social relationships may either strengthen or weaken over time, it is natural to assume that so will also evolve the attitude that an individual may have on influencing a friend's opinion. Moreover, due to homophily, i.e., the tendency of individuals to associate and bond with similar others, it also happens that an individual's expressed opinion influences in turn the strength of her social relationships. Based on these evidences, Bhawalkar et al. [1] conclude their paper by proposing a general co-evolutionary opinion formation game with dynamic (i.e., opinion-dependent) social relationships.

## 1.1 Our Contribution

Bhawalkar et al. [1] only show that their proposed co-evolutionary opinion formation games with dynamic social relationships always admit pure Nash equilibria. To the best of our knowledge, despite the relevance of their paper, no progresses have been done so far on (specializations of) this model. In this work, we try to fill this gap by embarking on the study of a basic, yet interesting class of opinion formation games with dynamic social relationships.

Let  $\mathbf{z}$  be the vector containing the expressed opinions of all players, so that  $z_i$  is the expressed opinion of player  $i$ . We define a cost-minimization  $n$ -player game in which the cost incurred by player  $i$  in the profile defined by  $\mathbf{z}$  is given by

$$c_i(\mathbf{z}) = \frac{\sum_{j \neq i} w_{ij}(\mathbf{z}) \cdot (z_i - z_j)^2}{\sum_{j \neq i} w_{ij}(\mathbf{z})} + \rho \cdot (s_i - z_i)^2,$$

where  $w_{ij}(\mathbf{z})$  is the social influence that  $j$  exercises over  $i$  which, being a function of  $\mathbf{z}$ , changes dynamically as the game evolves. More particularly, for a fixed

$k > 0$ , we set  $w_{ij}(\mathbf{z}) = (1 - |s_i - z_j|)^k$ . As it can be easily seen, the more  $z_j$  is close to  $s_i$ , the more  $j$  influences  $i$ 's opinion. The first term of  $c_i(\mathbf{z})$  is the cost that  $i$  incurs for disagreeing with the society and is defined as the average of the quadratic distances of  $i$ 's expressed opinion from the expressed opinion of the others weighted by their social influences. The second term of  $c_i(\mathbf{z})$ , instead, is the quadratic distance of  $i$ 's expressed opinion from her innate one, scaled by the player's stubbornness (we assume that all players have the same stubbornness). The higher  $\rho$ , the less a player is willing to deviate from her innate opinion because of the social pressure.

In this work, we focus on the case in which, for each player  $i$ , the innate opinion  $s_i \in [0, 1]$ , while the expressed opinion  $z_i \in \{0, 1\}$ . Despite their apparent simplicity, these games are able to capture several interesting scenarios. For instance, think of the situation in which one has to decide whether or not to buy a certain product given that she is not yet completely in favor of one of the two alternatives, or of the situation in which one has to choose between two candidates that might not both exactly reflect her own political ideas.

We show that any game in this class always admits an ordinal potential which implies the existence of pure Nash equilibria and convergence of better-response dynamics starting from any arbitrary strategy profile. Moreover, we prove that any pure Nash equilibrium and any social optimum (with respect to the problem of minimizing the sum of the players' costs) share the same structural property: if one numbers the players in non-decreasing order according to their innate opinions, the sequence of expressed opinions is also non-decreasing, i.e., it can be split into two (possibly empty) subsequences such that the first is made up of only zeroes and the second is made up of only ones. As a consequence, one obtains a simple and efficient algorithm for computing the set of pure Nash equilibria and social optima of a given game (since one has to discriminate among  $n + 1$  candidate strategy profiles only).

We also focus on the efficiency losses due to selfish behavior and give upper and lower bounds on the price of anarchy and lower bounds on the price of stability that only depend on the players' stubbornness, i.e., they neither involve the variable  $k$  nor the number of players  $n$ . In particular, we show that the price of anarchy is unbounded for  $\rho \in (0, 1]$ , while it is between  $\left(\frac{\rho+1}{\rho-1}\right)^2$  and  $2\left(\frac{\rho+1}{\rho-1}\right)^2$  for  $\rho > 1$ . For any value of  $\rho$ , the lower bound is attained in the situation in which both consensuses (i.e., all players expressing opinion 0, or all players expressing opinion 1) are pure Nash equilibria, but the players reach the wrong one, that is, the one yielding the highest social cost. We conjecture that our lower bound is tight, but proving a matching upper bound seems to be quite a challenging task, perhaps requiring tedious machineries. For such a reason, even if we are able to derive a better result than the above mentioned factor-2 upper bound, we decided to present a simpler (but still intricate) proof in this conference version. For the price of stability, instead, we only have some preliminary results, as we can just show a lower bound of  $\frac{\rho^2+6\rho+1}{(\rho+1)^2}$  for the case of  $\rho > 1$  (holding even when  $n = 2$ ), and that there is a 5-player game for which the price of stability is greater than one whenever  $\rho \in \left(\frac{217}{566}, 1\right]$ .

## 1.2 Related Work

To the best of our knowledge, Bindel et al. [2] have been the first to revisit opinion formation games under an Algorithmic Game Theory point of view. They consider the case in which both the innate and the expressed opinions are real values and the social influences are defined by an edge-weighted graph  $G$ . More formally, they analyze cost-minimization games in which the cost that player  $i$  incurs in the strategy profile  $\mathbf{z}$  is defined as

$$c_i(\mathbf{z}) = \sum_{j \neq i} w_{ij} \cdot (z_i - z_j)^2 + w_i \cdot (s_i - z_i)^2,$$

where  $w_{ij}$  is the weight of edge  $(i, j)$  in  $G$  and  $w_i$  is player  $i$ 's stubbornness. In this type of games, players converge to a unique pure Nash equilibrium, so that the prices of anarchy and stability coincide. For symmetric social influences, i.e., the case in which  $G$  is undirected, they show a tight price of anarchy of  $9/8$ . For the asymmetric case, instead, the price of anarchy can grow up to  $\Omega(n)$  and better, i.e. constant, upper bounds are shown when  $G$  belongs to two subclasses of Eulerian graphs.

Ferraioli et al. [6] investigate the above opinion formation games under the assumptions that  $z_i \in \{0, 1\}$  (as in our model). They show that these games are exact potential games [11], thus isomorphic to congestion games [12]. As to the efficiency of equilibria, the price of anarchy is shown to be unbounded, while, for the price of stability, exact bounds of 2 and 1 are proven for the cases in which the edge weights of  $G$  are integer and rational numbers, respectively. Then, they presents several results bounding the rate of convergence of decentralized best-response dynamics and logit dynamics.

Bhawalkar et al. [1] extend the model of Bindel et al. [2] by considering co-evolutionary opinion formation games in which the cost of player  $i$  in the strategy profile  $\mathbf{z}$  is defined as

$$c_i(\mathbf{z}) = \sum_{j \neq i} f_{ij}(z_i - z_j) + w_i \cdot g_i(s_i - z_i),$$

where  $f_{ij}$  and  $g_i$  are **fixed** real valued functions with the assumption that  $f_{ij} = f_{ji}$  (observe that the model of Bindel et al. [2] can be reobtained by setting  $f_{ij}(x) = w_{ij} \cdot x^2$  and  $g_i(x) = x^2$ ). They show that, when the  $f$  and  $g$  functions are either not convex or not differentiable, the price of anarchy is unbounded. When these functions are either convex and differentiable, instead, they show a tight bound of 2 on the price of anarchy. Moreover, for the particular case of  $f(x) = g(x) = |x|^\alpha$  a closed formula expressing the exact price of anarchy is derived.

Bhawalkar et al. [1] also consider the setting in which, for a given integer  $K > 0$ , each player  $i$  gets influenced by the  $K$  other players whose expressed opinions are closest to  $s_i$ . Denoted this set of acquaintances in the strategy profile  $\mathbf{z}$  as  $S(\mathbf{z}, i)$ , the cost of player  $i$  in  $\mathbf{z}$  is defined as

$$c_i(\mathbf{z}) = \sum_{j \in S(\mathbf{z}, i)} (z_i - z_j)^2 + \rho \cdot K \cdot (s_i - z_i)^2.$$

They show that, for  $\rho = 1 + \epsilon > 1$ , the robust price of anarchy of this game is at most  $\frac{(7+\epsilon)(2+\epsilon)}{\epsilon(1+\epsilon)}$  which is independent from  $K$ . Similarly to our results, the price of anarchy tends to 1 as  $\rho$  increases and it gets unbounded as  $\rho$  goes below 1, since they show a lower bound of at least  $1/\rho^2$  for  $\rho < 1$ .

Furthermore, Bhawalkar et al. [1] propose a model accounting for dynamic social influences by defining

$$c_i(\mathbf{z}) = \sum_{j \neq i} w_{ij}(\mathbf{z}) \cdot (z_i - z_j)^2 + \rho_i \cdot (s_i - z_i)^2,$$

where  $w_{ij}(\mathbf{z})$  is a continuous function depending on two variables, namely, the distance between  $s_i$  and  $z_j$ , and the total distance between  $s_i$  and the expressed opinions of all the players other than  $i$  and  $j$ . For these games, they prove existence of pure Nash equilibria.

Finally, Chierichetti et al. [3] studied the price of stability a similar model with an unweighted social graph  $G$ : each player tries to minimize the distance from her internal opinion and the sum of distances from the expressed opinions of the players adjacent to her in  $G$ .

### 1.3 Paper Organization

The paper is organized as follows. Next section contains the game definition and preliminary material. In Sect. 3, we derive the ordinal potential function for our games, while in Sect. 4, we show the structural properties of pure Nash equilibria and social optima. Section 5 describes our upper and lower bounds on the price of anarchy, and Sect. 6 contains lower bounds on the price of stability. Finally, in Sect. 7, we discuss open problems.

## 2 Model

The opinion formation games we consider in this paper are defined as follows. We are given a set of players  $N = \{1, 2, \dots, n\}$ . Every player  $i \in N$  has an *internal opinion*  $s_i \in [0, 1]$ . We will always assume  $s_i \leq s_j$  for every  $0 \leq i \leq j \leq n$ . The strategy of player  $i$  is a real number in  $[0, 1]$ , which is referred to as the *expressed opinion of player  $i$* . A state of the game is denoted by a strategy vector, that is a vector in  $[0, 1]^n$  whose  $i$ -th coordinate denotes the expressed opinion of player  $i$ . We denote by  $Z = [0, 1]^n$  the set of states of the game.

For every state  $\mathbf{z} = (z_j)_{j \in N} \in Z$ , we denote by  $\mathbf{z}_{-i}$  the strategy vector obtained from  $\mathbf{z}$  by removing the  $i$ -th coordinate, i.e.,  $\mathbf{z}_{-i} = (z_j)_{j \in N \setminus \{i\}}$ , and by  $(\mathbf{z}_{-i}, z'_i) \in Z$  the new state obtained from  $\mathbf{z}$  by replacing the  $i$ -th coordinate  $z_i$  with  $z'_i$ .

For every pair of players  $i, j \in N$  we have a *state-dependent weight function*  $w_{ij} : Z \mapsto \mathbb{R}_{\geq 0}$ , which denotes the influence of player  $j$  on  $i$  in a certain state. We assume  $w_{ii}(\mathbf{z}) = 0$ , for every  $\mathbf{z} \in Z$  and every  $i \in N$ . We define  $d_i(\mathbf{z}) = \sum_{j \in N} w_{ij}(\mathbf{z})$ , for every  $i \in N$ . We assume that, for any  $i \in N$ ,  $d_i(\mathbf{z}) > 0$  (i.e., for every players  $i$  there exists at least another player  $j$  such that  $w_{ij} > 0$ ).

The cost of player  $i$  in a state  $\mathbf{z} = (z_j)_{j \in N} \in Z$  is

$$c_i(\mathbf{z}) = \frac{1}{d_i(\mathbf{z})} \sum_{j \in N} w_{ij}(\mathbf{z}) \cdot (z_j - z_i)^2 + \rho \cdot (s_i - z_i)^2,$$

where  $\rho > 0$  is the *stubbornness* factor of each player.

A state  $\mathbf{e} \in Z$  is a *pure Nash equilibrium* if  $c_i(\mathbf{e}) \leq c_i(\mathbf{e}_{-i}, e'_i)$ , for every  $i \in N$  and every  $e'_i \in [0, 1]$ . We denote by  $E \subseteq Z$  the set of pure Nash equilibria of the game.

The *social cost* of a state  $\mathbf{z}$  is a measure of the social welfare. In this work we define the social cost of  $\mathbf{z}$  as  $C(\mathbf{z}) = \sum_{i \in N} c_i(\mathbf{z})$ . A state  $\mathbf{o} \in Z$  is a *social optimum* if it is one of the states minimizing the social cost, i.e.,  $C(\mathbf{o}) = \min_{\mathbf{z} \in Z} C(\mathbf{z})$ . We denote by  $O \subseteq Z$  the set of social optima, i.e.,  $O = \arg \min_{\mathbf{z} \in Z} C(\mathbf{z})$ , and by  $\text{OPT}$  the cost of any state in  $O$ .

The price of anarchy of the game is defined as  $\text{POA} = \max_{\mathbf{e} \in E} \frac{C(\mathbf{e})}{\text{OPT}}$ , if  $\text{OPT} > 0$ . If  $\text{OPT} = 0$  then  $\text{POA} = 1$  if  $E = O$ , and  $\text{POA} = +\infty$  otherwise. The price of stability of the game is defined as  $\text{POS} = \min_{\mathbf{e} \in E} \frac{C(\mathbf{e})}{\text{OPT}}$ , if  $\text{OPT} > 0$ . If  $\text{OPT} = 0$  then  $\text{POS} = 1$ .

## 2.1 Discrete Strategies and Polynomially Distance Decreasing Weights

In this paper we restrict to the *discrete setting* in which each player can choose only an element in  $\{0, 1\}$ ; hence, in this case,  $Z$  corresponds to the set of vectors  $\{0, 1\}^n$ . Moreover, if  $i \neq j$ , we define  $w_{ij}(\mathbf{z}) = (1 - |s_i - z_j|)^k$ , with  $k > 0$ , and we keep the assumption that  $w_{ii} = 0$ , for every  $i \in N$ . Notice that in general the weights, as defined herein, are asymmetric, i.e., it may happen that  $w_{ij} \neq w_{ji}$ . Given the two previous restrictions, it is easy to see that the cost function of each player  $i$  becomes

$$c_i(\mathbf{z}) = \frac{\sum_{j \in N: z_j \neq z_i} (1 - |s_i - z_j|)^k}{\sum_{j \in N \setminus \{i\}} (1 - |s_i - z_j|)^k} + \rho \cdot (s_i - z_i)^2. \tag{1}$$

For every state  $\mathbf{z} \in Z$ , we define  $\mathbf{1}(\mathbf{z}) = |\{j \in N : z_j = 1\}|$  and  $\mathbf{0}(\mathbf{z}) = |\{j \in N : z_j = 0\}|$ . It is useful to notice that  $\mathbf{1}(\mathbf{z}) + \mathbf{0}(\mathbf{z}) = n$ , and, for every player  $i$ ,  $\mathbf{1}(\mathbf{z}_{-i}) + \mathbf{0}(\mathbf{z}_{-i}) = n - 1$ . Finally  $\mathbf{1}(\mathbf{z}) = \mathbf{1}(\mathbf{z}_{-i})$  when  $z_i = 0$ , and  $\mathbf{0}(\mathbf{z}) = \mathbf{0}(\mathbf{z}_{-i})$  when  $z_i = 1$ . We will extensively use the previous relations in our calculations.

For instance, we trivially obtain that

$$\sum_{j \in N: z_j \neq z_i} (1 - |s_i - z_j|)^k = \begin{cases} \mathbf{1}(\mathbf{z})s_i^k = \mathbf{1}(\mathbf{z}_{-i})s_i^k = (n - 1)s_i^k - \mathbf{0}(\mathbf{z}_{-i})s_i^k & \text{if } z_i = 0 \\ \mathbf{0}(\mathbf{z})(1 - s_i)^k = \mathbf{0}(\mathbf{z}_{-i})(1 - s_i)^k & \text{if } z_i = 1. \end{cases}$$



Moreover,

$$\begin{aligned} \sum_{j \in N \setminus \{i\}} (1 - |s_i - z_j|)^k &= \sum_{\substack{j \in N \setminus \{i\}: \\ z_j=0}} (1 - |s_i - z_j|)^k + \sum_{\substack{j \in N \setminus \{i\}: \\ z_j=1}} (1 - |s_i - z_j|)^k \\ &= \mathbb{O}(\mathbf{z}_{-i})(1 - s_i)^k + \mathbb{1}(\mathbf{z}_{-i})s_i^k \\ &= \mathbb{O}(\mathbf{z}_{-i})\left[(1 - s_i)^k - s_i^k\right] + (n - 1)s_i^k. \end{aligned}$$

By combining the previous two equalities, according to (1), we get that the cost of player  $i$  can be written as follows

$$c_i(\mathbf{z}) = \begin{cases} \frac{(n-1)s_i^k - \mathbb{O}(\mathbf{z}_{-i})s_i^k}{\mathbb{O}(\mathbf{z}_{-i})\left[(1-s_i)^k - s_i^k\right] + (n-1)s_i^k} + \rho \cdot s_i^2 & \text{if } z_i = 0 \\ \frac{\mathbb{O}(\mathbf{z}_{-i})(1-s_i)^k}{\mathbb{O}(\mathbf{z}_{-i})\left[(1-s_i)^k - s_i^k\right] + (n-1)s_i^k} + \rho \cdot (1 - s_i)^2 & \text{if } z_i = 1. \end{cases} \tag{2}$$

In order to characterize a player’s best-response in a given state, we make use of the function  $f_{\rho,k} : [0, 1] \mapsto \mathbb{R}_{\geq 0}$  defined as follows

$$f_{\rho,k}(x) = \frac{x^k \left[1 + \rho(2x - 1)\right]}{(1 - x)^k \left[1 - \rho(2x - 1)\right] + x^k \left[1 + \rho(2x - 1)\right]}. \tag{3}$$

**Lemma 1.** *For every player  $i \in N$  and every state  $\mathbf{z} = (z_j)_{j \in N} \in Z$ , we have that  $c_i((\mathbf{z}_{-i}, 0)) \leq c_i((\mathbf{z}_{-i}, 1))$  if and only if  $f_{\rho,k}(s_i) \leq \frac{\mathbb{O}(\mathbf{z}_{-i})}{n-1}$ .*

*Proof.* By making use of Eqs. (2) and (3), elementary calculations show that  $c_i((\mathbf{z}_{-i}, 0)) \leq c_i((\mathbf{z}_{-i}, 1))$  if and only if  $f_{\rho,k}(s_i) \leq \frac{\mathbb{O}(\mathbf{z}_{-i})}{n-1}$ .

**Corollary 1.** *For every pure Nash equilibrium  $\mathbf{e} = (e_j)_{j \in N} \in E$  and for every player  $i \in N$ , if  $f_{\rho,k}(s_i) \leq 0$  then  $e_i = 0$ , while if  $f_{\rho,k}(s_i) \geq 1$  then  $e_i = 1$ .*

*Proof.* The claim easily follows from Lemma 1 and the fact that  $\frac{\mathbb{O}(\mathbf{e}_{-i})}{n-1}$  always gets values in the interval  $[0, 1]$ , for every player  $i \in N$ . □

### 3 Potential Function

In this section, we show that opinion formation games admit an ordinal potential function.

**Theorem 1.** *For every state  $\mathbf{z} = (z_j)_{j \in N} \in Z$ , let*

$$\Phi(\mathbf{z}) = \frac{1}{2} \cdot \mathbb{O}(\mathbf{z})\mathbb{1}(\mathbf{z}) - \frac{1}{4}(n - 1) \sum_{j \in N} (2z_j - 1)(2f_{\rho,k}(s_j) - 1).$$

*$\Phi$  is an ordinal potential function for the opinion formation game.*

*Proof.* We need to prove that, for every player  $i$ , it holds that

$$\Phi((\mathbf{z}_{-i}, 1)) - \Phi((\mathbf{z}_{-i}, 0)) > 0 \iff c((\mathbf{z}_{-i}, 1)) - c((\mathbf{z}_{-i}, 0)) > 0. \quad (4)$$

We trivially obtain that

$$\begin{aligned} \Phi((\mathbf{z}_{-i}, 1)) &= \frac{1}{2} \cdot \mathbb{O}(\mathbf{z}_{-i}) \left( \mathbb{1}(\mathbf{z}_{-i}) + 1 \right) - \frac{1}{4}(n-1)(2f(s_i) - 1) \\ &\quad - \frac{1}{4}(n-1) \sum_{\substack{j \in N: \\ j \neq i}} (2z_j - 1)(2f(s_j) - 1) \\ &= \frac{1}{2} \cdot \mathbb{O}(\mathbf{z}_{-i}) \mathbb{1}(\mathbf{z}_{-i}) + \frac{1}{2} \cdot \mathbb{O}(\mathbf{z}_{-i}) - \frac{1}{4}(n-1)(2f(s_i) - 1) \\ &\quad - \frac{1}{4}(n-1) \sum_{\substack{j \in N: \\ j \neq i}} (2z_j - 1)(2f(s_j) - 1), \end{aligned}$$

and

$$\begin{aligned} \Phi((\mathbf{z}_{-i}, 0)) &= \frac{1}{2} \left( \mathbb{O}(\mathbf{z}_{-i}) + 1 \right) \mathbb{1}(\mathbf{z}_{-i}) + \frac{1}{4}(n-1)(2f(s_i) - 1) \\ &\quad - \frac{1}{4}(n-1) \sum_{\substack{j \in N: \\ j \neq i}} (2z_j - 1)(2f(s_j) - 1) \\ &= \frac{1}{2} \cdot \mathbb{O}(\mathbf{z}_{-i}) \mathbb{1}(\mathbf{z}_{-i}) + \frac{1}{2} \cdot \mathbb{1}(\mathbf{z}_{-i}) + \frac{1}{4}(n-1)(2f(s_i) - 1) \\ &\quad - \frac{1}{4}(n-1) \sum_{\substack{j \in N: \\ j \neq i}} (2z_j - 1)(2f(s_j) - 1). \end{aligned}$$

From the previous two equalities, we get that

$$\begin{aligned} \Phi((\mathbf{z}_{-i}, 1)) - \Phi((\mathbf{z}_{-i}, 0)) &= \frac{1}{2} \left( \mathbb{O}(\mathbf{z}_{-i}) - \mathbb{1}(\mathbf{z}_{-i}) \right) - \frac{1}{2}(n-1)(2f(s_i) - 1) \\ &= \mathbb{O}(\mathbf{z}_{-i}) - \frac{1}{2}(n-1) - \frac{1}{2}(n-1)(2f(s_i) - 1) \\ &= \mathbb{O}(\mathbf{z}_{-i}) - (n-1)f(s_i), \end{aligned}$$

which implies that

$$\Phi((\mathbf{z}_{-i}, 1)) - \Phi((\mathbf{z}_{-i}, 0)) > 0 \iff \mathbb{O}(\mathbf{z}_{-i}) - (n-1)f_{\rho,k}(s_i) > 0. \quad (5)$$

On the other hand, from Lemma 1, we know that

$$c((\mathbf{z}_{-i}, 1)) - c((\mathbf{z}_{-i}, 0)) > 0 \iff \mathbb{O}(\mathbf{z}_{-i}) - (n-1)f_{\rho,k}(s_i) > 0. \quad (6)$$

The claim follows by combining (5) and (6).  $\square$

## 4 Social Optima and Pure Nash Equilibria

In this section, we characterize the social optima of opinion formation games, and provide a polynomial time algorithm for computing them.

**Proposition 1.** *For every  $\rho > 0$  and  $k > 0$ ,  $f_{\rho,k}$  satisfies the following properties:*

- (a)  $f_{\rho,k}(0) = 0$ ,  $f_{\rho,k}(\frac{1}{2}) = \frac{1}{2}$ , and  $f_{\rho,k}(1) = 1$ ;
- (b)  $1 - f_{\rho,k}(\frac{1}{2} + y) = f_{\rho,k}(\frac{1}{2} - y)$ , for every  $y > 0$ ;
- (c) if  $\rho > 1$ ,
  - (1)  $f_{\rho,k}(x) \leq 0$ , for every  $x \in [0, \frac{1}{2} - \frac{1}{2\rho}]$ ;
  - (2)  $f_{\rho,k}(x) \in [0, 1]$ , for every  $x \in [\frac{1}{2} - \frac{1}{2\rho}, \frac{1}{2} + \frac{1}{2\rho}]$ ;
  - (3)  $f_{\rho,k}(x) \geq 1$ , for every  $x \in [\frac{1}{2} + \frac{1}{2\rho}, 1]$ ;
- (d) if  $\rho \in (0, 1]$ ,  $f_{\rho,k}(x) \in [0, 1]$ , for every  $x \in [0, 1]$ ;
- (e) if  $\rho > 1$ ,  $f_{\rho,k}$  is increasing in the interval  $[\frac{1}{2} - \frac{1}{2\rho}, \frac{1}{2} + \frac{1}{2\rho}]$ ;
- (f) if  $\rho \in (0, 1]$ ,  $f_{\rho,k}$  is increasing in the interval  $[0, 1]$ .

**Proposition 2.** *For every pair of players  $u, v \in N$  with  $s_u < s_v$ , and every pair of states  $\mathbf{z} = (z_j)_{j \in N}, \mathbf{z}' = (z'_j)_{j \in N} \in Z$  we have that*

- (a) if  $z_u = z'_v = 1$  and  $\mathbb{O}(\mathbf{z}_{-u}) = \mathbb{O}(\mathbf{z}'_{-v})$  then  $c_u(\mathbf{z}) > c_v(\mathbf{z}')$ ;
- (b) if  $z'_u = z_v = 0$  and  $\mathbb{O}(\mathbf{z}'_{-u}) = \mathbb{O}(\mathbf{z}_{-v})$  then  $c_u(\mathbf{z}') < c_v(\mathbf{z})$ .

**Lemma 2.** *Every pure Nash equilibrium  $\mathbf{e} = (e_j)_{j \in N} \in E$  is such that there is an index  $t \in \{0, 1, \dots, n\}$  such that  $e_i = 0$  if and only if  $i \leq t$ .*

*Proof.* Let us assume  $\rho > 1$ . By combining Proposition 1.c with Corollary 1, we get that for every player  $i$  with  $s_i \in [0, \frac{1}{2} - \frac{1}{2\rho}]$  we have  $e_i = 0$ , and for every player  $i$  with  $s_i \in [\frac{1}{2} + \frac{1}{2\rho}, 1]$  we have  $e_i = 1$ . By the way of contradiction, let us assume that there is a player  $h$  with  $s_h \in [\frac{1}{2} - \frac{1}{2\rho}, \frac{1}{2} + \frac{1}{2\rho}]$ , such that  $e_h = 1$  and  $e_{h+1} = 0$ . We want to prove that if player  $h + 1$  is in equilibrium then player  $h$  is not in equilibrium. From Lemma 1 we know that if player  $h + 1$  is in equilibrium then  $f_{\rho,k}(s_{h+1}) \leq \frac{\mathbb{O}(\mathbf{e}_{-(h+1)})}{n-1}$ . Since from Proposition 1.e  $f_{\rho,k}$  is increasing in the interval  $[\frac{1}{2} - \frac{1}{2\rho}, \frac{1}{2} + \frac{1}{2\rho}]$ , this implies that  $f_{\rho,k}(s_h) \leq \frac{\mathbb{O}(\mathbf{e}_{-(h+1)})}{n-1}$ . Given that  $\mathbb{O}(\mathbf{e}_{-(h+1)}) = \mathbb{O}(\mathbf{e}_{-h}) - 1$ , we obtain that  $f_{\rho,k}(s_h) \leq \frac{\mathbb{O}(\mathbf{e}_{-h})}{n-1}$ , which contradicts Lemma 1.

By using claim (d) and (f) of Proposition (1), instead of (c) and (e), the same argument applies for  $\rho \in (0, 1]$ .  $\square$

**Lemma 3.** *Every social optimum  $\mathbf{o} = (o_j)_{j \in N} \in O$  is such that there is an index  $t \in \{0, 1, \dots, n\}$  such that  $o_i = 0$  if and only if  $i \leq t$ .*

*Proof.* By the way of contradiction, let us assume that there is a player  $i < n$  such that  $o_i = 1$  and  $o_{i+1} = 0$ . Let  $\mathbf{o}' = (o'_j)_{j \in N}$  be the state obtained from  $\mathbf{o}$  by changing the strategy of  $i$  to 0 and the strategy of  $i + 1$  to 1. We want to prove that  $\mathbf{o}'$  has a lower cost than  $\mathbf{o}$ . From the definition of cost in Eq. (2) we

notice that for every player  $h \in N \setminus \{i, i + 1\}$ , since  $\mathbb{O}(\mathbf{o}_{-h}) = \mathbb{O}(\mathbf{o}'_{-h})$ , the cost in  $\mathbf{o}$  is the same as the cost in  $\mathbf{o}'$ , i.e.,  $c_h(\mathbf{o}) = c_h(\mathbf{o}')$ . Hence,  $C(\mathbf{o}) - C(\mathbf{o}') = c_i(\mathbf{o}) - c_i(\mathbf{o}') + c_{i+1}(\mathbf{o}) - c_{i+1}(\mathbf{o}')$ . We can apply Proposition 2 by setting  $u = i$ ,  $v = i + 1$ ,  $\mathbf{z} = \mathbf{o}$  and  $\mathbf{z}' = \mathbf{o}'$ . In particular, from Proposition 2.a we obtain that  $c_i(\mathbf{o}) > c_i(\mathbf{o}')$ , and from Proposition 2.b we obtain that  $c_{i+1}(\mathbf{o}') < c_{i+1}(\mathbf{o})$ . We can conclude that  $C(\mathbf{o}) - C(\mathbf{o}') > 0$ , hence a contradiction.  $\square$

Since there are at most  $n + 1$  possible states with a suffix of  $x$  ones and a postfix of  $n - x$  zeros, from Lemmas 2 and 3 the following theorem easily follows.

**Theorem 2.** *All pure Nash equilibria and social optima of the game can be enumerated in polynomial time.*

### 5 Price of Anarchy

In this section, we provide asymptotically matching upper and lower bounds on the price of anarchy of opinion formation games. We first provide lower bounds holding for different values of  $\rho$  and then we focus on the corresponding upper bounds.

**Theorem 3.** *If  $\rho > 1$ , there exists an instance of opinion formation game such that  $\text{PoA} \geq \left(\frac{\rho+1}{\rho-1}\right)^2$ ; if  $\rho \leq 1$ , there exists an instance of opinion formation game such that  $\text{PoA} = \infty$ .*

*Proof.* Given  $s \in [0, 1]$ , consider an instance such that  $s_i = s$  for every  $1 \leq i \leq n$  and  $k = 1$ . Let  $\mathbf{z}^0 = (0, 0, \dots, 0)$  and  $\mathbf{z}^1 = (1, 1, \dots, 1)$  the states corresponding to the two possible consensuses. It trivially holds that  $C(\mathbf{z}^0) = n\rho s^2$  and  $C(\mathbf{z}^1) = n\rho(1 - s)^2$ . Clearly,  $\text{OPT} \leq C(\mathbf{z}^0)$ . By Lemma 1,  $\mathbf{z}^1 \in E$  if  $f_{\rho,1}(s) \geq 0$ . We consider two cases, depending on the value of  $\rho$ .

If  $\rho > 1$ , by Proposition 1.c,  $f_{\rho,1}(s) \geq 0$  if  $s \geq \frac{\rho-1}{2\rho}$ . Therefore, we obtain

$$\text{PoA} \geq \frac{C(\mathbf{z}^1)}{C(\mathbf{z}^0)} = \left(\frac{1-s}{s}\right)^2 = \left(\frac{\rho+1}{\rho-1}\right)^2,$$

where the last equality holds by choosing  $s = \frac{\rho-1}{2\rho}$ .

If  $\rho \leq 1$ , by Proposition 1.d,  $f_{\rho,1}(s) \geq 0$  for any value of  $s \geq 0$ . Therefore, we obtain that  $\text{OPT} = C(\mathbf{z}^0) = 0$  and  $C(\mathbf{z}^1) > 0$ , for an unbounded price of anarchy.  $\square$

**Theorem 4.** *If  $\rho > 1$  then  $\text{PoA} \leq 2 \left(\frac{\rho+1}{\rho-1}\right)^2$ .*

*Proof (Sketch).* Given an instance of opinion formation game  $G$  with  $\rho > 1$ , fix a pure Nash equilibrium  $\mathbf{e} \in E$  and a social optimum  $\mathbf{o} \in O$ . By Lemmas 3 and 2, both  $\mathbf{e}$  and  $\mathbf{o}$  can be determined by  $\mathbb{O}(\mathbf{e})$  and  $\mathbb{O}(\mathbf{o})$ , i.e., the number of players choosing strategy 0 (in fact, we know that in  $\mathbf{e}$  and  $\mathbf{o}$  all players choosing strategy 0 precede the ones choosing strategy 1). Moreover, without loss of generality, we

can assume that  $\mathbb{O}(\mathbf{o}) \geq \mathbb{O}(\mathbf{e})$ . In fact, if it holds that  $\mathbb{O}(\mathbf{o}) < \mathbb{O}(\mathbf{e})$ , we can build a “symmetric” game  $G'$  in which, if  $\alpha_i$  is internal opinion of player  $i$  in  $G$ , the internal opinion of player  $i$  in  $G'$  is  $1 - \alpha_i$ . It is easy to check that  $G'$  is such that any state  $\mathbf{z} = (z_j)_{j \in N}$  of  $G$  has the same social cost of the symmetric state  $\mathbf{z}' = (z'_j)_{j \in N}$  of  $G'$  ( $\mathbf{z}'$  is such that  $z'_i = 1 - z_i$  for every  $i \in N$ ), and moreover  $\mathbf{z}$  is a Nash equilibrium for  $G$  if and only if  $\mathbf{z}'$  is a Nash equilibrium for  $G'$ . Therefore, if  $\mathbb{O}(\mathbf{o}) < \mathbb{O}(\mathbf{e})$  we can study the symmetric instance  $G'$  in which  $\mathbb{O}(\mathbf{o}') \geq \mathbb{O}(\mathbf{e}')$ , where  $\mathbf{o}'$  and  $\mathbf{e}'$  are the symmetric states of  $\mathbf{o}$  and  $\mathbf{e}$ , respectively.

We partition  $N$  into three sets of consecutive players as follows:  $N_1$  contains the first  $\mathbb{O}(\mathbf{e})$  players, i.e., all players in  $N_1$  choose strategy 0 both in  $\mathbf{e}$  and in  $\mathbf{o}$ ;  $N_3$  contains the last  $n - \mathbb{O}(\mathbf{o})$  players, i.e., all players in  $N_3$  choose strategy 1 both in  $\mathbf{e}$  and in  $\mathbf{o}$ ;  $N_2$  contains the remaining  $\mathbb{O}(\mathbf{o}) - \mathbb{O}(\mathbf{e})$  players, i.e., all players in  $N_2$  choose strategy 0 in  $\mathbf{o}$  and strategy 1 in  $\mathbf{e}$ .

Since the case in which  $\mathbb{O}(\mathbf{e}) = 0$  is quite straightforward, we omit its discussion in this sketch.

We now analyze the case in which  $\mathbb{O}(\mathbf{e}) = 1$ . In this case,  $N_1 = \{1\}$  and, by Eq. (2),  $c_1(\mathbf{o}) \geq \rho s_1^2$  and  $c_1(\mathbf{e}) \leq 1 + \rho s_1^2$ . Moreover,  $c_i(\mathbf{e}) \leq c_i(\mathbf{o})$  for any player  $i \in N_3$ . In fact, for any of these players, by Eq. (2), the cost of disagreement with her innate opinion is the same in both  $\mathbf{e}$  and  $\mathbf{o}$ , while then cost of disagreement with the society is higher in  $\mathbf{o}$  since  $\mathbf{1}(\mathbf{e}) > \mathbf{1}(\mathbf{o})$ .

Given a state  $\mathbf{z} \in Z$  and a player  $i \in N$ , again by Eq. (2), we have  $c_i(\mathbf{z}_{-i}, 0) + c_i(\mathbf{z}_{-i}, 1) = 1 + 2\rho s_i^2 + \rho(1 - 2s_i)$ . Thus, for any player  $i \in N_2$ ,  $c_i(\mathbf{e}) \leq \frac{1+2\rho s_i^2+\rho(1-2s_i)}{2}$ , otherwise  $\mathbf{e}$  would not be a Nash equilibrium, and  $c_i(\mathbf{o}) \geq \rho s_i^2$ .

Therefore, we obtain that

$$\begin{aligned} \text{POA} &\leq \frac{1 + \rho s_1^2 + \frac{1}{2} \sum_{i \in N_2} (1 + 2\rho s_i^2 + \rho(1 - 2s_i))}{\rho s_1^2 + \sum_{i \in N_2} \rho s_i^2} \\ &\leq \frac{1 + \frac{1}{2} \sum_{i \in N_2} (1 + 2\rho s_i^2 + \rho(1 - 2s_i))}{\sum_{i \in N_2} \rho s_i^2} \\ &\leq \frac{|N_2| + \frac{1}{2} \sum_{i \in N_2} (1 + 2\rho s_i^2 + \rho(1 - 2s_i))}{\sum_{i \in N_2} \rho s_i^2} \\ &= \frac{\sum_{i \in N_2} (1 + \frac{1}{2} (1 + 2\rho s_i^2 + \rho(1 - 2s_i)))}{\sum_{i \in N_2} \rho s_i^2} \\ &\leq \max_{i \in N_2} \frac{1 + \frac{1}{2} (1 + 2\rho s_i^2 + \rho(1 - 2s_i))}{\rho s_i^2}. \end{aligned}$$

Since  $\frac{1+\frac{1}{2}(1+2\rho s^2+\rho(1-2s))}{\rho s^2}$  is a decreasing function in  $s$  (for  $0 \leq s \leq 1$  its partial first derivative with respect to variable  $s_i$  is negative), it is upper bounded by  $\frac{\rho^2+6\rho+1}{(\rho-1)^2}$  because, by Proposition 1.c, we know that for every player  $i \in N_2$  it holds that  $s_i \geq \frac{\rho-1}{2\rho}$ . Therefore, we obtain  $\text{POA} \leq \frac{\rho^2+6\rho+1}{(\rho-1)^2} \leq 2 \left( \frac{\rho+1}{\rho-1} \right)^2$ .

It remains to prove the claim for the case in which  $\mathbb{O}(\mathbf{e}) > 1$ .

By the same arguments used in the previous case, for any player  $i \in N_3$ ,  $c_i(\mathbf{e}) \leq c_i(\mathbf{o})$ .

We further partition  $N_1$  into two sets of consecutive players  $N_1^<$  and  $N_1^{\geq}$  such that every player  $i \in N_1^<$  is such that  $s_i < \frac{\rho-1}{2\rho}$  and every player  $i \in N_1^{\geq}$  is such that  $s_i \geq \frac{\rho-1}{2\rho}$ .

As in the previous case, for every player  $i \in N_1^{\geq} \cup N_2$ , we have that  $c_i(\mathbf{e}) \leq \frac{1+2\rho s_i^2 + \rho(1-2s_i)}{2}$  and  $c_i(\mathbf{o}) \geq \rho s_i^2$ . Since  $\frac{1+2\rho s_i^2 + \rho(1-2s_i)}{2\rho s_i^2}$  is a decreasing function in  $s$  (for  $0 \leq s \leq 1$  its partial first derivative with respect to variable  $s_i$  is negative), it holds that  $\max_{i \in N_1^{\geq} \cup N_2} \frac{1+2\rho s_i^2 + \rho(1-2s_i)}{2\rho s_i^2} \leq \left(\frac{\rho+1}{\rho-1}\right)^2 < 2 \left(\frac{\rho+1}{\rho-1}\right)^2$  because, by Proposition 1.c and by the definition of  $N_1^{\geq}$ , we know that for every player in  $i \in N_1^{\geq} \cup N_2$  it holds that  $s_i \geq \frac{\rho-1}{2\rho}$ .

It follows that, in order to prove the claim, it is sufficient to show that  $\frac{\sum_{i \in N_1^<} c_i(\mathbf{e})}{\sum_{i \in N_1^<} c_i(\mathbf{o})} \leq 2 \left(\frac{\rho+1}{\rho-1}\right)^2$ .

Given a player  $i \in N$  and a state  $\mathbf{z} \in Z$ , define  $\bar{c}_i(\mathbf{z}) = c_i(\mathbf{z}) - \rho(s_i - z_i)^2$ , i.e.  $\bar{c}_i(\mathbf{z})$  is the cost of player  $i$  for disagreeing with the society. Denote  $g(n, x, s, k) = \frac{(n-x)s^k}{(n-x)s^k + (x-1)(1-s)^k}$ . It is easy to check that, for every player  $i \in N_1^<$ ,  $\bar{c}_i(\mathbf{e}) = g(n, \mathbf{O}(\mathbf{e}), s_i, k)$  and  $\bar{c}_i(\mathbf{o}) = g(n, \mathbf{O}(\mathbf{o}), s_i, k)$ .

We obtain

$$\frac{\sum_{i \in N_1^<} c_i(\mathbf{e})}{\sum_{i \in N_1^<} c_i(\mathbf{o})} = \frac{\sum_{i \in N_1^<} (\rho s_i^2 + \bar{c}_i(\mathbf{e}))}{\sum_{i \in N_1^<} (\rho s_i^2 + \bar{c}_i(\mathbf{o}))} \leq \frac{\sum_{i \in N_1^<} \bar{c}_i(\mathbf{e})}{\sum_{i \in N_1^<} \bar{c}_i(\mathbf{o})} \leq \max_{i \in N_1^<} \frac{\bar{c}_i(\mathbf{e})}{\bar{c}_i(\mathbf{o})}.$$

Since  $\frac{\bar{c}_i(\mathbf{e})}{\bar{c}_i(\mathbf{o})}$  is a decreasing function in  $s_i$  (for  $0 \leq s_i \leq 1$  its partial first derivative with respect to variable  $s_i$  is negative), it is maximized when  $s_i$  tends to 0 (notice that if  $s_i = 0$ , player  $i$  can be discarded because she contributes 0 both to  $c(\mathbf{e})$  and  $c(\mathbf{o})$ ). By standard calculation,

$$\lim_{s_i \rightarrow 0^+} \frac{\bar{c}_i(\mathbf{e})}{\bar{c}_i(\mathbf{o})} = \frac{(\mathbf{O}(\mathbf{o}) - 1)(n - \mathbf{O}(\mathbf{e}))}{(\mathbf{O}(\mathbf{e}) - 1)(n - \mathbf{O}(\mathbf{o}))}.$$

Note that  $|N_2| = \mathbf{O}(\mathbf{o}) - \mathbf{O}(\mathbf{e}) > 0$ . Then,  $\lim_{s_i \rightarrow 0^+} \frac{\bar{c}_i(\mathbf{e})}{\bar{c}_i(\mathbf{o})} = \frac{(\mathbf{O}(\mathbf{e}) + |N_2| - 1)(n - \mathbf{O}(\mathbf{e}))}{(\mathbf{O}(\mathbf{e}) - 1)(n - \mathbf{O}(\mathbf{e}) - |N_2|)}$ . Define

$$\delta = \frac{(\rho^2 + 6\rho + 1)(\mathbf{O}(\mathbf{e}) - 1)(n - \mathbf{O}(\mathbf{e}))}{\mathbf{O}(\mathbf{e})(\rho^2 + 6\rho + 1) + n(\rho - 1)^2 - 2(\rho + 1)^2};$$

it can be checked that, if  $|N_2| \leq \delta$ ,  $\frac{(\mathbf{O}(\mathbf{e}) + |N_2| - 1)(n - \mathbf{O}(\mathbf{e}))}{(\mathbf{O}(\mathbf{e}) - 1)(n - \mathbf{O}(\mathbf{e}) - |N_2|)} \leq 2 \left(\frac{\rho+1}{\rho-1}\right)^2$ , thus proving the claim.

Therefore, in order to complete the proof, we can assume that  $|N_2| > \delta$ . The idea now is to spread the costs  $\bar{c}_1(\mathbf{e}), \dots, \bar{c}_{|N_1^<}(\mathbf{e})$  of the players in  $N_1^<$  on the players in  $N_2$ . More formally,

$$\begin{aligned} \text{POA} &\leq \frac{\sum_{i \in N_1^<} \bar{c}_i(\mathbf{e}) + \sum_{i \in N_2} \frac{1}{2} (1 + 2\rho s_i^2 + \rho(1 - 2s_i))}{\sum_{i \in N_1^< \cup N_2} \rho s_i^2} \\ &\leq \frac{\sum_{i \in N_1^<} g\left(n, \mathbf{O}(\mathbf{e}), \frac{\rho-1}{2\rho}, k\right) + \sum_{i \in N_2} \frac{1}{2} (1 + 2\rho s_i^2 + \rho(1 - 2s_i))}{\sum_{i \in N_2} \rho s_i^2} \end{aligned} \tag{7}$$

$$\begin{aligned} &= \frac{\sum_{i \in N_2} \left( \frac{|N_1^<|}{|N_2|} g\left(n, \mathbf{O}(\mathbf{e}), \frac{\rho-1}{2\rho}, k\right) + \frac{1+2\rho s_i^2 + \rho(1-2s_i)}{2} \right)}{\sum_{i \in N_2} \rho s_i^2} \\ &\leq \frac{\sum_{i \in N_2} \left( \frac{\mathbf{O}(\mathbf{e})}{\delta} g\left(n, \mathbf{O}(\mathbf{e}), \frac{\rho-1}{2\rho}, k\right) + \frac{1+2\rho s_i^2 + \rho(1-2s_i)}{2} \right)}{\sum_{i \in N_2} \rho s_i^2} \tag{8} \\ &\leq \max_{i \in N_2} \frac{\frac{\mathbf{O}(\mathbf{e})}{\delta} g\left(n, \mathbf{O}(\mathbf{e}), \frac{\rho-1}{2\rho}, k\right) + \frac{1+2\rho s_i^2 + \rho(1-2s_i)}{2}}{\rho s_i^2}, \end{aligned}$$

where inequality 7 holds because  $\bar{c}_i(\mathbf{e})$  is increasing in  $s_i$  and  $s_i < \frac{\rho-1}{2\rho}$  for every player in  $N_1^<$ , and inequality 8 holds because  $|N_1^<| \leq \mathbf{O}(\mathbf{e})$  and  $|N_2| > \delta$ .

By performing a smart and accurate calculation, in particular optimizing with respect to parameters  $n$  and  $s_i$ , it can be checked that

$$\frac{\frac{\mathbf{O}(\mathbf{e})}{\delta} g\left(n, \mathbf{O}(\mathbf{e}), \frac{\rho-1}{2\rho}, k\right) + \frac{1+2\rho s_i^2 + \rho(1-2s_i)}{2}}{\rho s_i^2} \leq 2 \left( \frac{\rho + 1}{\rho - 1} \right)^2.$$

□

## 6 Price of Stability

In this section, we provide (some) lower bounds on the price of stability of opinion formation games.

**Theorem 5.** *If  $\rho > 1$ , there exists an instance of opinion formation game such that  $\text{POS} \geq \frac{\rho^2+6\rho+1}{(\rho+1)^2}$ . Moreover, if  $\rho \in \left(\frac{217}{566}, 1\right]$ , there exists an instance of opinion formation game such that  $\text{POS} > 1$ .*

*Proof.* Consider an instance  $G$  of opinion formation game, with  $N = \{1, 2\}$  and  $\mathbf{s} = \left(\frac{\rho-1}{2\rho} - \epsilon, 1\right)$ . By Lemma 1 and Proposition 1.c, it follows that the unique pure Nash equilibrium for  $G$  is  $\mathbf{e} = (0, 1)$ ; it can be checked that, when  $\epsilon$  tends to 0,  $C(\mathbf{e}) = \frac{\rho^2+6\rho+1}{4\rho}$ , because  $c_1(\mathbf{e}) = 1 + \rho \left(\frac{\rho-1}{2\rho} - \epsilon\right)^2$  and  $c_2(\mathbf{e}) = 1$ . Consider state  $\mathbf{z} = (1, 1)$ ; it can be checked that, when  $\epsilon$  tends to 0,  $C(\mathbf{e}) = \frac{(\rho+1)^2}{4\rho}$ , because  $c_1(\mathbf{z}) = \rho \left(1 - \frac{\rho-1}{2\rho} - \epsilon\right)^2$  and  $c_2(\mathbf{z}) = 0$ . It follows that  $\text{POS} \geq \frac{C(\mathbf{e})}{C(\mathbf{z})} = \frac{\rho^2+6\rho+1}{(\rho+1)^2}$ .

For the case of  $\rho \in \left(\frac{217}{566}, 1\right]$  let us consider an instance  $G$  such that  $n = 5$ ,  $k = 1$  and  $\mathbf{s} = \left(\frac{1}{1000}, \frac{3}{20}, \frac{1}{2} + \epsilon, 1, 1\right)$ , where  $\epsilon > 0$  is an arbitrarily small number. As we have shown in Sect. 4, we can compute the price of stability of  $G$  by analyzing  $n + 1 = 6$  different strategy profiles. It is easy to see that, of all of these candidate profiles, only  $\mathbf{z}_1 = (0, 0, 1, 1, 1)$  and  $\mathbf{z}_2 = (1, 1, 1, 1, 1)$  are pure Nash equilibria for  $G$ . As to their social cost, as  $\epsilon$  goes to 0, we have  $C(\mathbf{z}_1) \approx \frac{3687}{4342} + 0.272501\rho$  and  $C(\mathbf{z}_2) \approx 1.970501\rho$ . We lower bound the price of stability of  $G$  by comparing the minimum of these two costs with the value  $C(\mathbf{z}^*) \approx 0.651 + 0.272501\rho$ , where  $\mathbf{z}^* = (0, 0, 0, 1, 1)$ . When  $C(\mathbf{z}_1) \leq C(\mathbf{z}_2)$ , we get  $\text{POS} \geq \frac{C(\mathbf{z}_1)}{C(\mathbf{z}^*)} \approx \frac{\frac{3687}{4342} + 0.272501\rho}{0.651 + 0.272501\rho} > 1$  for any value of  $\rho$ . When  $C(\mathbf{z}_1) > C(\mathbf{z}_2)$ , we get  $\text{POS} \geq \frac{C(\mathbf{z}_2)}{C(\mathbf{z}^*)} \approx \frac{1.970501\rho}{0.651 + 0.272501\rho} > 1$  for any value of  $\rho > \frac{217}{566}$ .  $\square$

## 7 Conclusions and Open Problems

Our work can be seen as an opening step in the investigation of opinion formation games with dynamic social influences, thus leaving a host of open questions to be solved. The most challenging ones are those regarding the characterization of the inefficiency of pure Nash equilibria. Bridging the gap between upper and lower bounds on the price of anarchy for the case of  $\rho > 1$  seems to require clever arguments. To this aim we conjecture that our lower bound is tight, that is, the price of anarchy equals the worst-case ratio between the social values of the strategy profiles realizing the two possible consensuses. Even more complicated appears to be the situation for the price of stability, for which, at the moment, no upper bounds are known (except for the one holding for the price of anarchy when  $\rho > 1$ ) and lower bounds are also missing for some values of  $\rho$ . Moreover, we only focused on the case in which the expressed opinion of each player is binary and the social influences are defined by the function  $w_{ij}(\mathbf{z}) = (1 - |s_i - z_j|)^k$  for any value of  $k > 0$ . Clearly, more general models can be considered, also with respect to the definition of the players' cost functions.

## References

1. Bhawalkar, K., Gollapudi, S., Munagala, V.: Coevolutionary opinion formation games. In: Proceedings of the 45th ACM Symposium on Theory of Computing (STOC), pp. 41–50. ACM Press (2013)
2. Bindel, D., Kleinberg, J.M., Oren, S.: How bad is forming your own opinion? In: Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Sciences (FOCS), pp. 57–66. IEEE Computer Society (2011)
3. Chierichetti, F., Kleinberg, J.M., Oren, S.: On discrete preferences and coordination In: Proceedings of the 14th ACM Conference on Electronic Commerce (EC), pp. 233–250. ACM Press (2013)
4. DeGroot, M.H.: Reaching a consensus. *J. Am. Stat. Assoc.* **69**, 118–121 (1974)
5. Durrett, R., Gleeson, J.P., Lloyd, A.L., Mucha, P.J., Shi, F., Sivakoff, D., Socolar, J.E.S., Varghese, C.: Graph fission in an evolving voter model. *Proc. Natl. Acad. Sci.* **109**(10), 3682–3687 (2012)



6. Ferraioli, D., Goldberg, P.W., Ventre, C.: Decentralized Dynamics for Finite Opinion Games. *Theoret. Comput. Sci.* **648**, 96–115 (2016)
7. Friedkin, N.E., Johnsen, E.C.: Social influence and opinions. *J. Math. Sociol.* **15**(3–4), 193–205 (1990)
8. Hegselmann, R., Krause, U.: Opinion dynamics and bounded confidence models, analysis, and simulation. *Artif. Soc. Soc. Simul.* **5**(3) (2002)
9. Henry, A.D., Pralat, P., Zhang, C.Q.: Emergence of segregation in evolving social networks. *Proc. Nat. Acad. Sci.* **108**, 8605–8610 (2011)
10. Holme, P., Newman, M.E.J.: Nonequilibrium phase transitions in the coevolution of networks and opinions. *Phys. Rev. E* **74**, 056108 (2006)
11. Monderer, D., Shapley, L.S.: Potential games. *Games Econ. Behav.* **14**(1), 124–143 (1996)
12. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **2**(1), 65–67 (1973)
13. Weisbuch, G., Deffuant, G., Amblard, F., Nadal, J.-P.: Meet, discuss and segregate!. *Complexity* **7**(3), 55–63 (2002)

# Complex Contagions on Configuration Model Graphs with a Power-Law Degree Distribution

Grant Schoenebeck and Fang-Yi Yu<sup>(✉)</sup>

University of Michigan, Ann Arbor, USA  
schoeneb@eecs.umich.edu, fayu@umich.edu

**Abstract.** In this paper we analyze  $k$ -complex contagions (sometimes called bootstrap percolation) on configuration model graphs with a power-law distribution. Our main result is that if the power-law exponent  $\alpha \in (2, 3)$ , then with high probability, the single seed of the highest degree node will infect a constant fraction of the graph within time  $O\left(\log^{\frac{\alpha-2}{3-\alpha}}(n)\right)$ . This complements the prior work which shows that for  $\alpha > 3$  bootstrap percolation does not spread to a constant fraction of the graph unless a constant fraction of nodes are initially infected. This also establishes a threshold at  $\alpha = 3$ .

The case where  $\alpha \in (2, 3)$  is especially interesting because it captures the exponent parameters often observed in social networks (with approximate power-law degree distribution). Thus, such networks will spread complex contagions even lacking any other structures.

We additionally show that our theorem implies that  $\omega\left(n^{\frac{\alpha-2}{\alpha-1}}\right)$  random seeds will infect a constant fraction of the graph within time  $O\left(\log^{\frac{\alpha-2}{3-\alpha}}(n)\right)$  with high probability. This complements prior work which shows that  $o\left(n^{\frac{\alpha-2}{\alpha-1}}\right)$  random seeds will have no effect with high probability, and this also establishes a threshold at  $n^{\frac{\alpha-2}{\alpha-1}}$ .

## 1 Introduction

Social behavior is one of the defining characteristics of us as a species. Social acts are influenced by the behavior of others while influencing them at the same time. These interactions have been observed in a wide array of activities including financial practices [8, 14], healthy/unhealthy habits [23], and voting practices [1]. Some of these are beneficial (e.g., adopting a healthy lifestyle) or profitable (e.g., viral marketing), while others are destructive and undesirable (e.g., teenager smoking, drug abuse).

To effectively promote desirable contagions and discourage undesirable ones, the first step is to understand how these contagions spread in networks and what are the important parameters that lead to fast spreading.

---

G. Schoenebeck—Supported by National Science Foundation Career Award #1452915

F.-Y. Yu—Supported by National Science Foundation Algorithms in the Field Award #1535912.

The high level objective is to understand how these behaviors spread in a social network. Two key factors determine the scope and rate of such diffusion: the model of contagions, i.e., how a node is influenced by its neighbors; and the network topology.

The copying of behaviors leading to a social cascade of behavioral changes is attributed to two effects: the informational benefit (inferring hidden, private information others may know) and direct benefit effects (resulting from coordinated actions or social pressure). In the *threshold model* [19], introduced by Granovetter, each agent has a *threshold* and when an agent's number of infected neighbors reaches her threshold, then she adopts the cascade.

We deal with a simplified version of this model where all agents have the same threshold. This is called *k-complex contagion* [18] or *bootstrap percolation* – the latter is generally used in the physics community where it was originally studied in the context of magnetic disordered systems [2, 13], but has since been applied to describe several complex phenomena including neuronal activity and the dynamics of the Ising model at zero temperature. In the context of social networks, bootstrap percolation provides a model of complex contagions [12] which model for the spread of ideas, beliefs, and behaviors.

A *k-complex contagion* is a deterministic process on a graph  $G$  that evolves in rounds. In each round every node has two possible states: it is either infected or uninfected. The network begins with a seed set  $I$  of infected nodes. In each subsequent round every uninfected node become infected if it has at least  $k$  edges incident on infected neighbors, otherwise it remains uninfected. Once a node has become infected, it remains infected forever.

A key trait of *k-complex contagions* is that they are not “submodular”. This implies that the marginal influence of an additional neighbor is not decreasing. While many cascade models, such as the Independent Cascade model and the Linear Threshold model have the submodularity property [21], many real-world cascades seem not to. Non-submodular contagions are observed by sociologists in the case of the adoption of pricey technology innovations, the change of social behaviors, and the decision to participate in a migration, etc [15, 22], and by data scientists on LiveJournal [6], DBLP [6], Twitter [26], and Facebook [28]. An additional confirmation is crucial, suggesting the model of complex contagion.

Janson et al. [20] show that *k-complex contagions* do not spread on sparse  $G(n, p)$  random graphs. Such cascades require  $\Omega(n)$  seeds to infect a constant fraction of vertices. [7] extended these results to configuration model graphs with regular degree distributions.

However, many networks do not have regular degree distributions. In a graph with power law degree distribution, the number of nodes having degree  $d$  is proportional to  $1/d^\alpha$ , for a positive constant  $\alpha$ . In 1965, Price [25] showed that the number of citations to papers follows a power law distribution. Later, studies of the World Wide Web reported that the network of webpages also has a power law degree distribution [9, 11]. Observations of many different types of social networks also found power law degree distributions, as well as biological, economic and semantic networks [3, 24, 27].

Additional work by [4] studies the configuration model with power-law degree distribution for  $\alpha > 3$  and showed and shows theorem which implies (see Sect. 7) that, with high probability, infecting a constant fraction of the nodes requires an initial seed that comprises a constant fraction of the graph.

Intuitively complex contagions spread well in the presence of additional community structure, and several networks with such structure have been analyzed including the Watts-Strogatz model [18], the Kleinberg Small World graph [17], and the preferential attachment graph [16].

Amini and Fountoulakis [5] also have examined the Chung-Lu model with power-law exponent  $2 < \alpha < 3$ . They show that there exists a function  $a(n) = o(n)$  such that if the number of initial seeds is  $\ll a(n)$ , the process does not evolve w.h.p.; and if the number of initial seeds is  $\gg a(n)$ , then a constant fraction of the graph is infected with high probability. However, this function is still super-constant— $n^{\Omega(1)}$ .

The question remained open, can non-submodular cascades spread and spread quickly from a constant-sized seed set on sparse graphs with no other structure imposed besides a skewed degree distribution.

## 1.1 Our Contributions

Our main result is that for a configuration model graph with power-law exponent  $\alpha \in (2, 3)$ , with high probability, the single seed of the highest degree node will infect a constant fraction of the graph within time  $O(\log^{\frac{\alpha-2}{3-\alpha}}(n))$ . This complements the prior work which showed that for  $\alpha > 3$  boot strap percolation does not spread to a constant fraction of the graph unless a constant fraction of nodes are initially infected. This also establishes a threshold at  $\alpha = 3$ .

The case where  $\alpha \in (2, 3)$  is especially interesting because it captures the exponent parameters often observed in social networks (with approximate power-law degree distribution). Thus, such networks will spread complex contagions even lacking any other structure.

We additionally show that our main theorem implies that  $\omega(n^{\frac{\alpha-2}{\alpha-1}})$  random seeds will infect a constant fraction of the graph within time  $O(\log^{\frac{\alpha-2}{3-\alpha}}(n))$ . This complements the prior work which shows that  $o(n^{\frac{\alpha-2}{\alpha-1}})$  random seeds will have no effect with high probability. This also establishes a threshold at  $n^{\frac{\alpha-2}{\alpha-1}}$ .

To prove these results, we provide new analysis that circumvents previous difficulties. While our results are similar to those of [16] (they study the preferential attachment model, while we study the configuration model), the techniques required are completely different. For example, it is an easy observation that  $k$ -complex contagions spread on the configuration model (if  $k$  is greater than the minimum degree), but much more difficult to show it spreads quickly.

The previous analyses on the configuration model required that the graph was locally tree-like, an assumption that fails in our case, and then were able to approximate the process using differential equations and obtain rigorous results by applying Wormald's Theorem [30]. However, their analysis fails when the degree distribution is power-law with exponent between 2 and 3.

## 2 Preliminaries

Let  $[m] := \{1, 2, \dots, m\}$ . We say an event would be true *with high probability* if its probability of being true is  $1 - o(1)$ . When we use  $\Theta(1)$ , the constant may depend on various constant parameters, but should not depend on  $n$ .

**Definition 1.** A  $k$ -complex contagion  $CC(G, k, I)$  is a contagion that initially infects vertices of  $I \subseteq V(G)$  and spreads over the graph  $G$ . The contagion proceeds in rounds. At each round, each vertex with at least  $k$  edges incident on infected neighbors becomes infected. The vertices of  $I$  are called the initial seeds. Let  $|CC(G, k, I)|$  denote the random variable of the final size of such a cascade.

We use the *configuration model* introduced by [10] to define a distributions over multigraphs.

**Definition 2.** Let  $\mathbf{d} = (d_1, \dots, d_n)$  be a decreasing degree sequence where the sum of the terms is even. Define  $V = [n]$  (Here we use integers  $\{1, 2, \dots, n\}$  to denote the vertices, and call nodes with lower indexes “earlier”. Because the degrees decrease, earlier nodes have higher degrees). Let  $m$  be such that  $2m = \sum_i d_i$ . To create the  $m$  (multi-)edges, we first assign each node  $i$   $d_i$  stubs. Next we choose a perfect matching of the stubs uniformly at random and for each pair of matched stubs construct an edge connecting the corresponding nodes.

We use  $CM(\mathbf{d})$  to denote the Configuration Model with respect to the degree distribution  $\mathbf{d}$ .

### 2.1 Power-Law Degree Distributions

For any decreasing degree sequence  $\mathbf{d} = (d_1, \dots, d_n)$  where the sum of the terms is even, we define

- the empirical distribution function of the degree distribution  $F_{\mathbf{d}}$

$$F_{\mathbf{d}}(x) = \frac{1}{n} \sum_{i=1}^n \mathbb{I}[d_i \leq x] \quad \forall x \in [1, \infty)$$

- the fraction of nodes that have degree less than  $x$ .
- Let  $N_{\mathbf{d}}(x) = n(1 - F_{\mathbf{d}}(x))$  be the number of nodes with degree at least  $x$ .
- Let  $S_{\mathbf{d}}(x)$  be the number of stubs from nodes with degree at least  $x$ .
- Let  $s_{\mathbf{d}}(x)$  be the number of stubs from nodes with index less than  $x$ .

We will omit the index  $\mathbf{d}$  when there is no ambiguity.

**Definition 3 (Power-Law Degree Distributions).** Adopting the notation of [29], we say a series  $\mathbf{d}$  has power-law distribution with exponent  $\alpha$  if there exists  $0 < C_1 < C_2$  and  $0 < x_0$  such that (1)  $F_{\mathbf{d}}(x) = 0$  for  $x < x_0$ ; (2)  $F_{\mathbf{d}} = 1$  for  $x > d_1 = n^{2/(\alpha+1)}$ , and (3) for all  $x_0 \leq x \leq d_1$  then

$$C_1 x^{-\alpha+1} \leq 1 - F_{\mathbf{d}}(x) \leq C_2 x^{-\alpha+1}$$

Let  $\mathbf{d}$  have power-law distribution of *power law with exponent*  $\alpha$  then it is easy to check that:

1.  $N(x) = \Theta(nx^{-\alpha+1})$
2.  $S(x) = \Theta(nx^{-\alpha+2})$ .
3.  $d(i) = \Theta\left(\left(\frac{n}{i}\right)^{1/(\alpha-1)}\right)$
4.  $s(i) = \Theta(n^{1/(\alpha-1)}i^{\frac{\alpha-2}{\alpha-1}})$

### 3 Main Theorem

In this section, we state and prove our main theorem: in a configuration model graph with the power-law exponent  $\alpha \in (2, 3)$ , with high probability, the single seed of the highest degree node will infect a constant fraction of the graph within time  $O(\log^{\frac{\alpha-2}{3-\alpha}}(n))$ .

**Theorem 1.** *Given a power law distribution  $\mathbf{d} = (d_1, \dots, d_n)$  with exponent  $\alpha \in (2, 3)$  and  $d_1 > n^{\frac{3-\alpha}{\alpha+1}}$ , with probability  $1 - O\left(\frac{\log^{\frac{\alpha-1}{3-\alpha}} n}{n}\right)$ , the  $k$ -complex contagion on configuration model  $CM(\mathbf{d})$  with constant  $k$  and initial infection being the highest degree node  $I = \{1\}$ ,  $CC(CM(\mathbf{d}), k, I)$ , infects  $\Omega(n)$  vertices within time  $O(\log^{\frac{\alpha-2}{3-\alpha}} n)$ .*

#### 3.1 Proof Setup

We consider a restricted form of contagion where nodes can only be infected by those preceding them in the ordering. Formally, recall the nodes  $\{d_i\}$  are ordered in terms of their degree. Node  $d_i$  will only be infected if  $|\{j : j < i \text{ and } d_j \text{ is infected}\}| \geq k$  neighbors are infected. Hence, the total number of infected nodes in this process will be fewer than the number of infected nodes in original complex contagions, and it is sufficient to prove that a constant fraction of nodes become infected in this restricted contagion with high probability.

*Buckets* We first partition the nodes  $V = [n]$  into buckets. We design the buckets to have at least (and about the same number of) stubs  $b = \Theta\left(\frac{n^{\frac{\alpha-2}{3-\alpha}}}{\log^{\frac{\alpha-2}{3-\alpha}} n}\right)$ .

We can define  $N_\ell$  as follows

$$N_1 = \frac{n}{\log^{\frac{\alpha-1}{3-\alpha}} n}, \text{ and } N_{\ell+1} = \arg \min_{i > N_\ell} \{s(i) - s(N_\ell) \geq b\}$$

Since  $d(N_1) = \Theta(\log^{1/(3-\alpha)} n) = o(b)$  and  $\forall i > N_1, d(i) \leq d(N_1)$ ,

$$b < s(N_{\ell+1}) - s(N_\ell) \leq b + o(b) < 2b.$$

Therefore, we have  $\ell b \leq s(N_\ell) \leq 2\ell b$  and  $N_\ell = \Theta\left(\frac{n}{\log^{\frac{\alpha-1}{3-\alpha}} n} \ell^{\frac{\alpha-1}{\alpha-2}}\right)$  by (4), and so the total number of buckets is  $L \leq \frac{s(n)}{b} = O(\log^{\frac{\alpha-2}{3-\alpha}} n)$ .

We define our buckets to be  $B_1 = \{1, \dots, N_1\}, B_2 = \{N_1 + 1, \dots, N_2\}, \dots, B_{\ell+1} = \{N_\ell + 1, \dots, N_{\ell+1}\}, \dots, B_L = \{N_{L-1} + 1, \dots, N_L\}$ .

*Filtration.* We now state our filtration.

$\mathcal{F}_0$ : The node  $i$  starts with  $d_i$  stubs of edges without revealing any edges.

$\mathcal{F}_1$ : In the first stage we reveal all edges within the first bucket  $B_1$ ,

$\mathcal{F}_\ell, 1 \leq \ell \leq L$ : In the stage  $\ell > 1$ , we reveal/match all the edges from  $B_\ell$  to early nodes in  $B_{<\ell}$ .

### 3.2 Proof Summary

There are two parts of the proof.

1. All of the nodes in the first bucket would be infected with high probability.
2. For some constant  $\rho > 0$ , in the first  $L' = \rho L$  buckets  $B_1, \dots, B_{L'}$  a constant fraction  $\epsilon$  of nodes will be infected. Because  $N_{L'} = \Omega(n)$  nodes, the total number of infection also constant fraction.

In the first part of the proof is capture by the following lemma:

**Lemma 1 (Base).** *Given at  $\mathcal{F}_0$   $d_1 > n^{\frac{3-\alpha}{\alpha+1}}$ , at  $\mathcal{F}_1$  all the nodes in  $B_1$  will be infected within  $O(\log \log(n))$  steps with probability greater than  $1 - O(\frac{1}{n})$ .*

To prove this lemma we further decompose the first bucket into  $O(\log \log(n))$  finer intervals, which we call bins. We first argue that every node in the first bin will have at least  $k$  multi-edges to the first node, and we inductively show the nodes in following bin will have at least  $k$  edges to the previous bins. The analysis is by straight-forward probabilistic techniques.

The time for the first bucket's infection is at most the number of the bins because inclusion of each bin only costs 1 step.

We need some additional notation to state the lemma which will imply the second part. Let  $X_\ell$  be the number of stubs from buckets  $B_{<\ell}$  to  $B_{\geq\ell}$ . Let  $Y_\ell$  be the number of uninfected stubs from  $B_{<\ell}$  to  $B_{\geq\ell}$  before stage  $\ell$ , of which  $Y_\ell^{(1)}$  issue from  $B_{<\ell-1}$  and the remaining  $Y_\ell^{(2)}$  issue from  $B_{\ell-1}$ . We use  $\mathbb{I}_i$  as the indicator variable that node  $i \in B_\ell$  is not infected after stage  $\ell$ . Let  $\epsilon > 0$  be some constant we define later. Let  $\delta_n = \Theta(\frac{1}{\log^{\frac{\alpha-2}{3-\alpha}} n})$ .

Now we can formally define  $\mathcal{A}_\ell$  as the intersection of the following three events:

1. connection:  $(1 - \delta_n)\mathbb{E}[X_\ell] \leq X_\ell \leq (1 + \delta_n)\mathbb{E}[X_\ell]$ ;
2. number of uninfected nodes:  $\sum_{i \in B_{\ell-1}} \mathbb{I}_i \leq 2\mu_H$  where  $\mu_H = K \frac{|B_\ell|^\frac{3-\alpha}{\alpha-2}}{\log n}$  for some constant  $K$  independent of  $\ell$  and  $n$ ;
3. number of uninfected stubs:  $Y_\ell \leq \epsilon X_\ell$ .

**Lemma 2 (Induction).** *Fix sufficiently small  $\epsilon > 0, \rho > 0$ . Let  $\ell < \rho L$ , and suppose  $\Pr[\mathcal{A}_\ell] > 0.5$ , then we have*

$$\Pr[\mathcal{A}_{\ell+1} | \mathcal{A}_\ell] = 1 - O(1) \frac{(\log n)^{\frac{\alpha-1}{3-\alpha}}}{n\ell^{1/(\alpha-2)}}$$

This lemma will be proved by showing that each of three events happens with high probability conditioned on  $\mathcal{A}_\ell$ . The most technically challenging of these is the second event, where we need to apply Chebychev’s Inequality twice. One challenging is that the edges from  $B_{<\ell}$  to  $B_\ell$  are not independent. Another challenge is that if the buckets are too small, we fail to have concentration properties, but if they are too large, then the fraction of infected nodes at each stage will drop too quickly.

### 3.3 Proof of Theorem 1

*Proof.* If  $\bigcap_{\ell=1}^{L'} \mathcal{A}_\ell$  happens, then the total fraction of infected nodes is  $\Omega(n)$ .

Using Lemma 1 as the base case and Lemma 2 as the induction steps we see that

$$\Pr \left[ \bigcap_{\ell=1}^{L'} \mathcal{A}_\ell \right] \geq 1 - \sum_{\ell=1}^{L'} O(1) \frac{(\log n)^{\frac{\alpha-1}{3-\alpha}}}{n\ell^{1/(\alpha-2)}} - O\left(\frac{1}{n}\right) = 1 - O\left(\frac{\log^{\frac{\alpha-1}{3-\alpha}} n}{n}\right)$$

which is arbitrarily close to 1.

Moreover, the total time spent is the time in first bucket plus the number of buckets (because the infection spreads from bucket to bucket in only 1 step). Therefore the total time spent is

$$O(\log \log n) + O(\log^{\frac{\alpha-2}{3-\alpha}} n) = O(\log^{\frac{\alpha-2}{3-\alpha}} n)$$

which completes our proof.

## 4 Proof of Lemma 1: Contagion in the First Bucket

In this section, we will show that with high probability, the contagion process infects all nodes within the first bucket. Recall that  $N_1 = \frac{n}{\log^{\frac{\alpha-1}{3-\alpha}} n}$  and the number of stubs within the first bucket is  $S(N_1) = b$ .

We partition the first bucket into finer bins such that  $B_1 = \bigcup_{t=1}^T V_t$  and  $V_t = \{v_{t-1} + 1, \dots, v_t\}, t = 1, \dots, T$  with ascending order and  $v_0 = 1$ . The  $v_t$  will be specified in Lemma 4. We define the event that every nodes in bin  $V_t$  is infected as  $E_t$ , then the event that all the nodes in  $B_1$  are infected is equal to  $\bigcap_{t=1}^T E_t$ .

We recall Lemma 1:

**Lemma 1 (Base).** *Given at  $\mathcal{F}_0$   $d_1 > n^{\frac{3-\alpha}{\alpha+1}}$ , at  $\mathcal{F}_1$  all the nodes in  $B_1$  will be infected within  $O(\log \log(n))$  steps with probability greater than  $1 - O(\frac{1}{n})$ .*

We will use two Lemmas in the proof of Lemma 1, which will be a proof by induction. The first lemma will form the base case of the induction. It states the high degree nodes will all be infected by the first node by showing any high degree node forms  $k$  multi-edges to the first node.



**Lemma 3.** *Given  $d_1 > n^{\frac{3-\alpha}{\alpha+1}}$  we define node  $v_1 = \max\{v : d(v) \geq n^{\frac{3-\alpha}{\alpha+1}}\}$ . (Recall nodes are ordered by degree.) Then all the nodes in  $V_1 = \{1, \dots, v_1\}$  will be infected in one step with probability*

$$\Pr[E_1] = 1 - n^{\frac{3-\alpha}{\alpha+1}} \exp(-\Theta(1)n^{\frac{3-\alpha}{\alpha+1}}).$$

The second Lemma will form the inductive step in the proof of Lemma 1. It can be proved by induction itself.

**Lemma 4.** *Let  $v \in V_t = \{v_{t-1} + 1, \dots, v_t\}$  and  $v_t = \max\{v : d(v) \geq \frac{n^{\frac{\alpha-1}{3-\alpha}}}{\log^{\frac{\alpha-1}{3-\alpha}} n^{(\alpha-2)^t}}\}$ , then*

$$\Pr \left[ u \text{ is not infected} \mid \bigcap_{s=1}^{t-1} E_s \right] \leq \frac{1}{n^2}$$

Moreover,  $T = O(\log \log n)$ .

The proofs for Lemmas 3 and 4 are the simple application of a Chernoff bound and a union bound which is in the full version.

*Proof (Lemma 1).* The proof is by induction. For the base case, Lemma 3 ensures every node in the first bin will be infected. Suppose all nodes before  $v_{t-1}$  are infected. We can use a union bound to show every node in  $V_t$  will be also infected. Moreover, in each bin the contagion only takes one time step which implies that the infection time for the first bucket is at most  $O(\log \log n)$ .

For the probability that all these events hold, we apply a union bound.

$$\begin{aligned} & \Pr[\text{all the nodes in } B_1 \text{ are infected}] \\ &= \Pr \left[ \bigcap_{t=1}^T E_t \right] \\ &\geq 1 - \Pr[\neg E_1] - \sum_{t=2}^T \Pr \left[ \neg E_t \mid \bigcap_{s=1}^{t-1} E_s \right] \quad (\text{union bound}) \\ &\geq 1 - n^{\frac{3-\alpha}{\alpha+1}} \exp(-\Theta(1)n^{\frac{3-\alpha}{\alpha+1}}) - \frac{1}{n^2} |B_1| \quad \text{by Lemmas 3 and 4} \end{aligned}$$

## 5 Proof of Lemma 2: Contagion from Buckets to Bucket

In this section we prove Lemma 2.

**Lemma 2 (Induction).** *Fix sufficiently small  $\epsilon > 0$ ,  $\rho > 0$ . Let  $\ell < \rho L$ , and suppose  $\Pr[\mathcal{A}_\ell] > 0.5$ , then we have*

$$\Pr[\mathcal{A}_{\ell+1} | \mathcal{A}_\ell] = 1 - O(1) \frac{(\log n)^{\frac{\alpha-1}{3-\alpha}}}{n \ell^{1/(\alpha-2)}}$$

Recall that  $\mathcal{A}_\ell$  is the intersection of the three events, we will show that at stage  $\ell$  if these three events happen, then the requirements in Lemma 2 will be met, and those events would be proven in Lemmas 5, 6 and 8 respectively.

**5.1 First Event: Connection**

We first note that the first event holds with high probability. This follows almost immediately from a standard Chernoff bound application, and the proof is in the full version.

**Lemma 5.** *Let  $\delta_n = \Theta\left(\frac{1}{\log^{\frac{\alpha-2}{3-\alpha}} n}\right)$ , if  $\Pr[\mathcal{A}_\ell] \geq 0.5$*

$$\Pr\left[|X_{\ell+1} - \mathbb{E}[X_{\ell+1}]| \leq \delta_n \mathbb{E}[X_{\ell+1}] \mid \mathcal{A}_\ell\right] \geq 1 - 4 \exp\left(-\Theta\left(\frac{n}{\log^{6 \cdot \frac{\alpha-2}{3-\alpha}} n}\right)\right).$$

Here the constant only depends on the product of  $\delta_n$  and  $L$ .

**5.2 Second Event: Number of Infected Nodes**

Now we will prove the second events holds with high probability.

**Lemma 6 (Number of Uninfected Nodes in a Single Bucket).** *For sufficiently small  $\epsilon > 0$ , conditioned on  $\mathcal{A}_\ell$*

$$\Pr\left[\sum_{i \in B_\ell} \mathbb{I}_i \geq 2\mu_H \mid \mathcal{A}_\ell\right] \leq O(1) \frac{(\log n)^{\frac{\alpha-1}{3-\alpha}}}{n\ell^{1/(\alpha-2)}}$$

where  $\mu_H = K \frac{|B_\ell|\ell^{\frac{3-\alpha}{\alpha-2}}}{\log n}$  and  $K$  is independent of  $\ell$  and  $n$ .

The proof relies on an application of Chebyshev’s inequality and the following Lemma, which is in turn proved using Chebyshev’s inequality. The full proof is in the full version.

**Lemma 7 (Infection of a single node).** *If  $\mathcal{F}_\ell \subseteq \mathcal{A}_\ell$  for some constant  $0 < \epsilon < 1/2$  and  $\delta_n = \Theta\left(\frac{1}{\log^{\frac{\alpha-2}{3-\alpha}} n}\right) < 1/2$ , then the probability any node  $i \in B_\ell$  is not infected is*

$$\Pr[\mathbb{I}_i \mid \mathcal{A}_\ell] \leq O(1) \frac{\ell^{\frac{3-\alpha}{\alpha-2}}}{\log n}$$

where the constant  $O(1)$  only depends on  $\alpha, k, \rho$  if  $\delta_n, \epsilon$  is small enough, and  $\rho \leq 0.3 \frac{\alpha-1}{\alpha-2} k^{\frac{\alpha-2}{3-\alpha}}$ .

The main proof idea of Lemma 7 is that because the events that a infected stub from  $B_{<\ell}$  to a node  $i$  in  $B_\ell$  are negative dependent, the variance of the number of infected stubs from  $B_{<\ell}$  to node  $i$  is small, and we can use Chebyshev’s inequality to show each node has a high chance of being infected when fraction of uninfected stubs from  $B_{<\ell}$ ,  $\epsilon$ , is small. The full proof is in the full version.

**5.3 Third Event: Number of Uninfected Stubs**

**Lemma 8.** *Suppose  $A_\ell$ , the first event,  $(1 - \delta_n)\mathbb{E}[X_{\ell+1}] \leq X_{\ell+1} \leq (1 + \delta_n)\mathbb{E}[X_{\ell+1}]$  and the second event,  $\sum_{i \in B_\ell} \mathbb{I}_i \leq 2\mu_H$  is true (this is the conclusion of Lemma 6), then*

$$\Pr \left[ Y_{\ell+1} \leq \epsilon X_{\ell+1} \mid |X_{\ell+1} - \mathbb{E}[X_{\ell+1}]| \leq \delta_n \mathbb{E}[X_{\ell+1}] \wedge \sum_{i \in B_\ell} \mathbb{I}_i \leq 2\mu_H \wedge \mathcal{A}_\ell \right]$$

is greater than  $1 - \exp\left(-\Theta\left(\frac{n}{\log^{5. \frac{\alpha-2}{3-\alpha}}}\right)\right)$  when  $\rho > 0$  is small enough and  $\delta_n > 0$  is smaller than some constant.

For the third event, in Lemma 8 we want to argue the fraction of uninfected stubs is smaller than  $\epsilon$  after stage  $\ell$ . That requires both that  $X_{\ell+1}$  is large and that  $Y_{\ell+1}$ —which is the summation of  $Y_{\ell+1}^{(1)}$  and  $Y_{\ell+1}^{(2)}$ —is small. Upper bounds on  $Y_{\ell+1}^{(1)}$  and  $Y_{\ell+1}^{(2)}$  will be proven by Lemmas 9 and 10 respectively. The full proof for Lemma 8 is in the full version.

**Lemma 9.** *Let  $Y_\ell^{(1)}$  be the number of free uninfected stubs from  $B_{<\ell}$  to  $B_{>\ell}$  over the probability space  $\mathcal{F}_{\ell+1} | \mathcal{F}_\ell$ , then*

$$\Pr \left[ Y_{\ell+1}^{(1)} \geq (1 + \delta_n)\epsilon X_\ell | \mathcal{A}_\ell \right] \leq \exp\left(-\Theta\left(\frac{n}{\log^{5. \frac{\alpha-2}{3-\alpha}}}\right)\right)$$

Here the constant only depends on  $\delta_n \cdot L$ ,  $\epsilon$  and  $\rho$ .

**Lemma 10.** *Suppose  $A_\ell$  and the  $\sum_{i \in B_\ell} \mathbb{I}_i \leq 2\mu_H$  is true (this is the conclusion of Lemma 6), then  $Y_{\ell+1}^{(2)}$ , the total number of uninfected stubs from  $B_\ell$  to  $B_{>\ell}$  is*

$$Y_{\ell+1}^{(2)} = O(1) \frac{(\log n)^{\frac{3-\alpha}{\alpha-2}}}{n \ell^{2/(\alpha-2)}}$$

The full proofs for Lemmas 9 and 10 are in the full version.

**5.4 Proof of Lemma 2**

*Proof.* Recall the the event  $\mathcal{A}_{\ell+1}$  is the intersection of the three events, so

$$\Pr[\mathcal{A}_{\ell+1} | \mathcal{A}_\ell] \geq 1 - \Pr[\neg(|X_\ell - \mathbb{E}[X_\ell]| \leq \delta_n \mathbb{E}[X_\ell]) | \mathcal{A}_\ell] \tag{1}$$

$$- \Pr \left[ \sum_{i \in B_{\ell-1}} \mathbb{I}_i \geq 2\mu_H | \mathcal{A}_\ell \right] \tag{2}$$

$$- \Pr \left[ Y_\ell \leq \epsilon X_\ell (|X_\ell - \mathbb{E}[X_\ell]| \leq \delta_n \mathbb{E}[X_\ell]) \wedge \sum_{i \in B_{\ell-1}} \mathbb{I}_i \leq 2\mu_H \wedge \mathcal{A}_\ell \right] \tag{3}$$

Applying Lemma 5 to Eq. 1, Lemma 6 to Eq. 2, and Lemma 8 to Eq. 3, and we have

$$\begin{aligned} \Pr[\mathcal{A}_{\ell+1}|\mathcal{A}_\ell] &\geq 1 - 4 \exp\left(-\Theta\left(\frac{n}{\log^{6 \cdot \frac{\alpha-2}{3-\alpha}} n}\right)\right) \\ &\quad - O(1) \frac{(\log n)^{\frac{\alpha-1}{3-\alpha}}}{n\ell^{1/(\alpha-2)}} \\ &\quad - \exp\left(-\Theta\left(\frac{n}{\log^{5 \cdot \frac{\alpha-2}{3-\alpha}}}\right)\right) \end{aligned}$$

Therefore

$$\Pr[\mathcal{A}_{\ell+1}|\mathcal{A}_\ell] \geq 1 - O(1) \frac{(\log n)^{\frac{\alpha-1}{3-\alpha}}}{n\ell^{1/(\alpha-2)}}$$

### 6 Infection with Random Seeds

Theorem 1 together with prior results in Ebrahimi et al. [16] immediately implies the following corollary:

**Corollary 1.** *For a configuration model graph with power-law exponent  $\alpha$ , if  $\Omega(n^{\frac{\alpha-2}{\alpha-1}})$  initially random seeds are chosen, then with probability  $1 - o(1)$   $k$ -complex contagion infects a constant fraction of nodes.*

We first restate two results from [16].

**Proposition 2 [16].** *For any graph, let  $u$  be a node with degree  $d$ . If  $\Omega(d/n)$  initial random seeds are chosen, then with probability  $1 - o(1)$   $u$  is infected after one round.*

*Proof.* The initial node has  $\Theta(n^{\frac{1}{\alpha-1}})$  neighbors. If there are  $\Omega(n^{\frac{\alpha-2}{\alpha-1}})$  initial seeds then by Proposition 2 the first seed is infected with probability  $1 - o(1)$ . However, then by Theorem 1 a constant fraction of the remaining nodes are infected in  $\log^{O(\alpha)}(1)$  rounds.

This is tight as in Ebrahimi et al. [16] the following was proven:

**Proposition 3 [16].** *For any graph, with power law distribution  $\alpha$ , if  $o(n^{\frac{\alpha-2}{\alpha-1}})$  initially random seeds are chosen, then with probability  $1 - o(1)$ , no additional nodes are infected.*

### 7 $\alpha > 3$

For the case of power-law degree distribution with  $\alpha > 3$ , Amini [4] shows how to analyze  $k$ -complex contagions using a differential equation method [30]. This approach heavily depends on the variance of the degree distribution and fails when  $\alpha < 3$ . For the case where the seed set contains all nodes with degree greater than  $\rho > 0$  we can state their theorem as follows:

**Theorem 4** [4]. *Given a power law distribution  $\mathbf{d}$  with exponent  $\alpha > 3$  and  $d_1 < n^{1/\alpha-1}$ , the  $k$ -complex contagion on configuration  $CM(\mathbf{d})$  with constant  $k$  and seed set  $I_\rho = \{i | d_i \geq \rho\}$  where  $0 \leq \rho \leq n$ . Then with high probability*

$$|CC(CM(\mathbf{d}), k, I_\rho)| = n \left( 1 - \sum_{\substack{1 \leq d < \rho, \\ 0 \leq j < k}} p_{\mathbf{d}}(d) \binom{d}{j} (y^*)^{d-j} (1 - y^*)^j + o(1) \right) \quad (4)$$

where  $p_{\mathbf{d}}(d) = (F_{\mathbf{d}}(d + 1) - F_{\mathbf{d}}(d))$  and  $0 < y^* \leq 1$  is the largest root such that  $f(y) = 0$  and

$$f(y) = y^2 \left( \sum_{1 \leq d} d p_{\mathbf{d}}(d) \right) - y \left( \sum_{\substack{1 \leq d < \rho, \\ 0 \leq j < k}} d p_{\mathbf{d}}(d) \binom{d-1}{j} y^{d-1-j} (1 - y)^j \right) \quad (5)$$

Before stating our corollary, we wish to give a brief idea of the proof of Theorem 4. They consider a Markov chain which results in the same number infected nodes as a  $k$ -complex contagion, but proceeds using the randomness of the configuration model. The Markov chain starts with the initially infected nodes and at each step the process reveals one of the unmatched edges from the set of infected nodes. This process needs only track: the number of unmatched edges, and the number of  $d$ -degree uninfected nodes with  $j$  infected neighbors, for each  $j < k$ . The Markov chain stops when all the agent are infected, or there are no unmatched edges from already infected nodes. It turns out, that if  $\alpha > 3$ , the process is smooth and we can use the corresponding differential equations to approximate this Markov chain and derive the fraction of infections.

With their results we can prove that to infect a constant fraction of nodes requires the initial seed need to also be a constant fraction of nodes. Note that if our initial seed set infects the highest degree nodes, but does not infect a constant fraction of the nodes, then the greatest degree node not in the initially infected set has degree  $\omega(1)$ .

**Corollary 2.** *Given a power law distribution  $\mathbf{d}$  with exponent  $\alpha > 3$  and  $d_1 < n^{1/\alpha-1}$ , the  $k$ -complex contagion on configuration  $CM(\mathbf{d})$  with constant  $k$  and seed set  $I_\rho = \{i | d_i \geq \rho\}$  where  $\rho = \omega(1)$ , the  $|CC(CM(\mathbf{d}), k, I_\rho)| = o(n)$  with high probability.*

The proof of the corollary requires some delicate calculations and is in the full version.

## References

1. Adamic, L.A., Glance, N.: The political blogosphere, the: divided they blog. In: Proceedings of the 3rd International Workshop on Link Discovery, pp. 36–43. ACM (2005)

2. Adler, J.: Bootstrap percolation. *Phys. A: Stat. Theor. Phys.* **171**(3), 453–470 (1991)
3. Albert, R., Barabási, A.-L.: Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74**, 47–97 (2002). doi:[10.1103/RevModPhys.74.47](https://doi.org/10.1103/RevModPhys.74.47). <http://link.aps.org/doi/10.1103/RevModPhys.74.47>
4. Amini, H.: Bootstrap percolation and diffusion in random graphs with given vertex degrees. *Electron. J. Comb.* **17**(1), 1–20 (2010)
5. Amini, H., Fountoulakis, N.: What I tell you three times is true: bootstrap percolation in small worlds. In: Goldberg, P.W. (ed.) *WINE 2012*. LNCS, vol. 7695, pp. 462–474. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-35311-6\\_34](https://doi.org/10.1007/978-3-642-35311-6_34)
6. Backstrom, L., Huttenlocher, D., Kleinberg, J., Lan, X.: Group formation in large social networks: membership, growth, and evolution. In: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 44–54 (2006)
7. Balogh, J., Pittel, B.: Bootstrap percolation on the random regular graph. *Random Struct. Algorithms* **30**, 257–286 (2007)
8. Banerjee, A., Chandrasekhar, A.G., Duflo, E., Jackson, M.O.: The diffusion of microfinance. *Science* **341**(6144), 1236498 (2013)
9. Barabási, A., Albert, R.: Emergence of scaling in random networks. *Science* **286**, 509–512 (1999)
10. Bollobás, B., McKay, B.D.: The number of matchings in random regular graphs and bipartite graphs. *J. Comb. Theory, Series B* **41**(1), 80–91 (1986)
11. Broder, A., Kumar, R., Maghoul, F., Raghavan, P., Rajagopalan, S., Stata, R., Tomkins, A., Wiener, J.: Graph structure in the web. In: *Proceedings of the 9th International World Wide Web Conference on Computer Networks*, pp. 309–320 (2000)
12. Centola, D., Macy, M.: Complex contagions and the weakness of long ties<sup>1</sup>. *Am. J. Sociol.* **113**(3), 702–734 (2007)
13. Chalupa, J., Leath, P.L., Reich, G.R.: Bootstrap percolation on a bethe lattice. *J. Phys. C: Solid State Phys.* **12**(1), L31 (1979)
14. Coleman, J., Katz, E., Menzel, H.: The diffusion of an innovation among physicians. *Sociometry* **20**(4), 253–270 (1957)
15. Coleman, J.S., Katz, E., Menzel, H.: *Medical Innovation: A Diffusion Study*. Bobbs-Merrill Co., Indianapolis (1966)
16. Ebrahimi, R., Gao, J., Ghasemiefteh, G., Schoenebeck, G.: How complex contagions spread quickly in the preferential attachment model, other time-evolving networks. arXiv preprint [arXiv:1404.2668](https://arxiv.org/abs/1404.2668) (2014)
17. Ebrahimi, R., Gao, J., Ghasemiefteh, G., Schoenebeck, G.: Complex contagions in Kleinberg’s small world model. In: *Proceedings of the Conference on Innovations in Theoretical Computer Science*, pp. 63–72. ACM (2015)
18. Ghasemiefteh, G., Ebrahimi, R., Gao, J.: Complex contagion, the weakness of long ties in social networks: revisited. In: *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, pp. 507–524, June 2013
19. Granovetter, M.: Threshold models of collective behavior. *Am. J. Sociol.* **83**(6), 1420–1443 (1978)
20. Janson, S., Łuczak, T., Turova, T., Vallier, T.: Bootstrap percolation on the random graph  $G_{n,p}$ . *Ann. Appl. Prob.* **22**(5), 1989–2047 (2012)
21. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 137–146 (2003)

22. Macdonald, J.S., Macdonald, L.D.: Chain migration, ethnic neighborhood formation and social networks. *Milbank Meml. Fund Q.* **42**(1), 82–97 (1964)
23. Mermelstein, R., Cohen, S., Lichtenstein, E., Baer, J.S., Kamarck, T.: Social support and smoking cessation and maintenance. *J. Consult. Clin. Psychol.* **54**(4), 447 (1986)
24. Newman, M.E.J.: The structure and function of complex networks. *SIAM Rev.* **45**, 167–256 (2003)
25. De Solla Price, D.: Networks of scientific papers. *Science* **149**(3683), 510–515 (1965). doi:[10.1126/science.149.3683.510](https://doi.org/10.1126/science.149.3683.510)
26. Romero, D.M., Meeder, B., Kleinberg, J.: Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In: *Proceedings of the 20th International Conference on World Wide Web*, pp. 695–704 (2011)
27. Steyvers, M., Tenenbaum, J.B.: The large-scale structure of semantic networks: statistical analyses and a model of semantic growth. *Cogn. Sci.* **29**, 41–78 (2005)
28. Ugander, J., Backstrom, L., Marlow, C., Kleinberg, J.: Structural diversity in social contagion. *Proc. Natl. Acad. Sci.* **109**(16), 5962–5966 (2012)
29. Van Der Hofstad, R.: Random graphs and complex networks, p. 11 (2009). <http://www.win.tue.nl/rhofstad/NotesRGCN.pdf>
30. Wormald, N.C.: Differential equations for random processes and random graphs. *Ann. Appl. Prob.* **5**(4), 1217–1235 (1995)

# **Abstracts**



# The Magician's Shuffle: Reusing Lottery Numbers for School Seat Redistribution

Itai Feigenbaum<sup>1,2(✉)</sup>, Yash Kanoria<sup>1</sup>, Irene Lo<sup>1</sup>, and Jay Sethuraman<sup>1</sup>

<sup>1</sup> Columbia University, New York, USA

<sup>2</sup> Lehman College and the Graduate Center, City University of New York,  
New York, USA

[itai.feigenbaum@lehman.cuny.edu](mailto:itai.feigenbaum@lehman.cuny.edu)

**Abstract.** In many centralized school admission systems, a significant fraction of allocated seats are later vacated, often due to students obtaining better outside options. We consider the problem of reassigning these seats in a fair and efficient manner while also minimizing the movement of students between schools. Centralized admissions are typically conducted using the deferred acceptance (DA) algorithm, with a lottery used to break ties caused by indifference in school priorities. For reassignment, we propose a class of mechanisms called Permuted Lottery Deferred Acceptance (PLDA). After the initial (first-round) assignment is computed via DA, students' preferences change (get truncated) due to the revelation of their outside options. A PLDA mechanism then computes a reassignment of the students by re-running DA; however, students are guaranteed to get at least their first-round assignment (if they still want it) or a school they prefer, and ties are broken according to a permutation of the first-round lottery order. We show that a PLDA based on a *reversal* of the first-round lottery order performs well.

Our theoretical analysis takes place in a continuum model with no school priorities. We characterize PLDA mechanisms as the class of mechanisms that satisfy a few natural properties, which include not removing students from their first-round assignments against their will, a strong form of strategyproofness (against manipulations involving misreporting both the original and changed preferences), and certain efficiency and fairness axioms. We then identify a technical condition, called the *order condition*, essentially requiring that the change in preferences does not modify the relative overdemand for schools. When the order condition is satisfied, all PLDA mechanisms yield identical allocative efficiency, and among all of them, the lottery-reversal based PLDA reassigns the minimal amount of students (from their first-round assignments). Finally, we conduct computational experiments and obtain results that support our theoretical findings. Specifically, we use data from NYC's school choice program to simulate the performance of different PLDA mechanisms in the presence of school priorities, and find that all simulated PLDAs have similar allocative efficiency, while the lottery-reversal based PLDA minimizes the number of reassigned students.

A full version is available at: <http://www.columbia.edu/~yk2577/reallocation.pdf>

# Near-Efficient Allocation Using Artificial Currency in Repeated Settings (Extended Abstract)

Artur Gorokh<sup>1</sup>, Siddhartha Banerjee<sup>2(✉)</sup>, and Krishnamurthy Iyer<sup>2</sup>

<sup>1</sup> Center for Applied Mathematics, Ithaca, USA  
ag2282@cornell.edu

<sup>2</sup> School of Operations Research and Information Engineering, Cornell University,  
Ithaca, NY 14853, USA  
{sbanerjee, kriyer}@cornell.edu

**Abstract.** We study the design of mechanisms without money for repeated allocation of resources among competing agents. Such mechanisms are gaining widespread use in allocating computing resources in universities and companies, and also distributing of public goods like vaccines among hospitals and food donations among food banks. We consider repeated allocation mechanisms based on *artificial currencies*, wherein we first allot each agent a chosen endowment of credits, which they can then use over time to bid for the item in a chosen auction format. Our main contribution is in showing that a simple mechanism, based on a repeated all-pay auction with personalized endowments and static pricing rules, simultaneously guarantees *vanishing gains from non-truthful bidding* as well as *vanishing loss in efficiency*. Our work lies at the intersection of dynamic mechanism design and mechanisms without money, and the techniques we develop here may prove of independent interest in these settings.

Our work studies the question of whether the *incentive properties and allocative efficiency* of mechanisms with money can be approximated via mechanisms based on an *artificial currency* – one which has no independent valuation outside the setting of the mechanism. This has attracted a lot of attention in recent times due to the establishment of platforms that use artificial-currency systems to solve real-world problems such as university course allocation and food banks.

We consider a problem of allocating a single item between 2 agents  $\{a, b\}$  in  $T$  consecutive periods  $t = 1, 2, \dots, T$ . At time  $t$ , agent  $s \in \{a, b\}$  has i.i.d valuation  $V_{s,t} = v$  with probability  $\{\mathbf{q}_s(v)\}$ ; valuations are independent across agents, and distributions are known publicly. Given any mechanism  $\mathcal{M}$  not involving money, agent  $s$ 's utility is  $U_s^{\mathcal{M}} \triangleq \sum_{t=1}^T V_{s,t} x_{s,t}^{\mathcal{M}}$ , where  $x_{s,t}^{\mathcal{M}}$  is the allocation to  $s$  at time  $t$ . In this setting, it is easy to see that for  $T = 1$ , no mechanism can be both incentive compatible and efficient; our aim is to use the repeated nature of the process to ensure approximate efficiency and incentive compatibility.

Formally, we define a mechanism to be an  $(\alpha, \beta)$ -approximate mechanism if it simultaneously guarantees that (i) *truthful play is an  $\alpha$ -equilibrium*, i.e.,

for any agent  $s$ , assuming all other agents play truthfully, the utility gain from deviating from truthful play is at most  $\alpha T$ , and (ii) *the mechanism is  $\beta$ -efficient*, i.e., assuming all agents play truthfully, the loss in welfare from the optimal is at most  $\beta T$ . For example, a uniform lottery achieves  $(\alpha, \beta) = (0, \Omega(1))$ ; on the other hand, we show that a second-price auction with artificial currency has  $(\alpha, \beta) = (\Omega(1), 0)$ . This raises the question as to whether there are mechanisms where both  $\alpha$  and  $\beta$  are  $o(1)$ . To this end, we propose the *Repeated Endowed All-Pay* (or REAP) mechanism, wherein we first give each agent an endowment of credits, and then in each period, agents are charged credits to report a valuation according to a personalized price function; the item is then allocated to the highest reported valuation. Our main result is the following:

**Theorem 1.** *REAP is an  $(\alpha, \beta)$ -approximate mechanism with  $\alpha = O\left(\sqrt{\frac{\log T}{T}}\right)$ ,  $\beta = O\left(\frac{1}{T}\right)$ .*

Our result is based on setting prices via a novel LP-based analysis of an auxiliary game, and then showing the sample paths of the mechanism concentrate close to this auxiliary game. In addition, our work suggests several future directions for research on the scope and practicality of mechanisms without money. For details, refer to our full version: [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2852895](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2852895).

**Acknowledgment.** K. Iyer gratefully acknowledges support from the National Science Foundation under grant CMMI-1462592.

# Multi-unit Facility Location Games

Omer Ben-Porat<sup>(✉)</sup> and Moshe Tennenholtz

Faculty of Industrial Engineering and Management,  
Technion - Israel Institute of Technology, Haifa, Israel  
omerbp@campus.technion.ac.il, moshet@ie.technion.ac.il

Motivated by applications in clustering and information retrieval, we extend the classical Hotelling setting (see [1]) to address the scenario where players may control more than one facility. In his seminal work, Hotelling considers a duel between two parties who compete over consumers distributed uniformly over the interval  $[0, 1]$ ; each party locates its facility on that interval, and grabs the proportion of the population closer to it. As it turns out, the only equilibrium in that setting is for both parties to locate their facility at  $\frac{1}{2}$ . Interestingly, while overwhelming many extensions of that basic setting exist, the economic studies refer to competition between single-facility owners only, which make that work non-applicable to many applications, e.g. in clustering we are typically after selecting several centroids/clusters.

Consider for example the strategic behavior of publishers in the web. Assume a “strong” publisher who controls several outlets of its site which it can maintain, e.g. two different Internet versions of its newspaper. This publisher can be viewed as being able to locate two “facilities” in the space of published data rather than only one; however, a “weak” publisher who can not maintain two such versions will need and be able to locate only one “facility”. How would these different powers effect the behavior of the publishers? What would be optimal strategies for the different publishers? This is a novel challenge and question, which illustrates how valuable and deep the understanding of these games may be for theory and practice.

We extend the Hotelling setting to *multi-unit facility location games*, where there are  $n$  players, where player  $i$  may control **several** facilities. We first analyze competition among the owner of  $k$  facilities to the owner of  $l$  facilities, for arbitrary  $(l, k)$ , where  $l \leq k$ . Our message for this extended Hotelling duel is quite striking: in **no** equilibrium of any such  $(l, k)$  facility location duel a facility will materialize in a location which is not part of the social welfare maximizing locations of the player who has  $k$  facilities, if she were to locate her facilities under no competition. This is obtained despite the lack of pure strategy equilibrium in any  $(l, k)$  duel whenever  $l \neq k$ .

Moreover, for the  $n$ -player setting, we provide sufficient and necessary conditions for a pure strategy profile to be an equilibrium in such game. In particular, we show that a pure-strategy equilibrium exist if and only if there is no

---

Full paper available at <http://arxiv.org/abs/1602.03655>.

dominant player who controls more than half of the facilities; in the latter case, under some conditions, a mixed strategy equilibrium of the form obtained in the  $(l, k)$  duel does exist.

## Reference

1. Hotelling, H.: Stability in competition. *Econ. J.* **39**(153), 4157 (1929)

## Author Index

- Albers, Susanne 309  
Amanatidis, Georgios 414  
Anagnostopoulos, Aris 102  
Anshelevich, Elliot 145, 265  
Arias, Emilio Jesús Gallego 279
- Banerjee, Siddhartha 476  
Barthe, Gilles 279  
Ben-Porat, Omer 478  
Bilò, Vittorio 444  
Birmpas, Georgios 414  
Blumrosen, Liad 400  
Brunsch, Tobias 339
- Caragiannis, Ioannis 236  
Cavallo, Ruggiero 102  
Chan, Hau 44  
Colini-Baldeschi, Riccardo 207  
Conitzer, Vincent 1  
Czumaj, Artur 15
- D., Thirumulanathan 174  
Deligkas, Argyrios 15, 29  
Deng, Xiaotie 59
- Etscheid, Michael 339
- Fain, Brandon 384  
Fanelli, Angelo 444  
Fasoulakis, Michail 15  
Fearnley, John 15, 29  
Feigenbaum, Itai 475  
Feldman, Michal 131  
Feng, Zhe 59  
Filos-Ratsikas, Aris 236  
Fotakis, Dimitris 221  
Frederiksen, Søren Kristoffer Stiil 236  
Freeman, Rupert 354
- Gaboardi, Marco 279  
Goel, Ashish 384  
Goldner, Kira 160  
Gorokh, Artur 476
- Gourvès, Laurent 221  
Goyal, Sanjeev 429
- Haney, Samuel 354  
Hansen, Kristoffer Arnsfelt 236  
Hoefler, Martin 294  
Hoy, Darrell 73  
Hsu, Justin 279
- Immorlica, Nicole 73  
Iyer, Krishnamurthy 476
- Jabbari, Shahin 429  
Jiang, Albert Xin 44  
Jurdziński, Marcin 15
- Kanoria, Yash 475  
Kar, Koushik 145  
Karlin, Anna R. 160  
Kearns, Michael 429  
Kesselheim, Thomas 294  
Khanna, Sanjeev 429  
Kodric, Bojana 294  
Kong, Yuqing 251  
Kraft, Dennis 309
- Leonardi, Stefano 102, 207  
Leyton-Brown, Kevin 44  
Ligett, Katrina 251  
Lo, Irene 475  
Lucier, Brendan 73, 131
- Markakis, Evangelos 414  
Mehta, Ruta 44  
Mizrahi, Yehonatan 400  
Monnot, Jérôme 221  
Morgenstern, Jamie 429  
Moscaredelli, Luca 444  
Munagala, Kamesh 384
- Narahari, Y. 174  
Nath, Swaprava 369

Niazadeh, Rad 116

Nisan, Noam 131

Panigrahi, Debmalya 354

Papadimitriou, Christos H. 59

Röglin, Heiko 339

Roth, Aaron 279

Sandholm, Tuomas 369

Savani, Rahul 15, 29

Schoenebeck, Grant 251, 459

Sekar, Shreyas 145, 265

Sethuraman, Jay 475

Skowron, Piotr 324

Strub, Pierre-Yves 279

Sundaresan, Rajesh 174

Sviridenko, Maxim 102

Tan, Zihan 236

Tennenholtz, Moshe 478

Tzamos, Christos 188

Wang, Weina 87

Wilkens, Christopher A. 116, 188

Ying, Lei 87

Yu, Fang-Yi 459

Zhang, Junshan 87

Zhang, Qiang 207