
DEPARTMENT OF COMPUTER ENGINEERING AND INFORMATICS

UNIVERSITY OF PATRAS



Design and analysis of algorithms for non-cooperative environments

DOCTORAL THESIS

Alexandros A. Voudouris

Advisor: Ioannis Caragiannis, Associate Professor

Patras, September 2018

Τμήμα Μηχανικών Ηλεκτρονικών Υπολογιστών και Πληροφορικής

Πανεπιστήμιο Πατρών

Αλέξανδρος Α. Βουδούρης

© 2018 – Με την επιφύλαξη παντός δικαιώματος

Department of Computer Engineering and Informatics

University of Patras

Design and analysis of algorithms for non-cooperative environments

Doctoral Thesis

Alexandros A. Voudouris

Thesis committee:

Ioannis Caragiannis (Advisor)
Associate Professor, University of Patras

Christos Kaklamanis (Member of the advisory committee)
Professor, University of Patras

Sotiris Nikolettseas (Member of the advisory committee)
Professor, University of Patras

Stavros Cosmadakis
Professor, University of Patras

Evangelos Markakis
Assistant Professor, Athens University of Economics and Business

Paul Spirakis
Professor, University of Patras

Christos Zaroliagis
Professor, University of Patras

Design and analysis of algorithms for non-cooperative environments

Alexandros A. Voudouris

Abstract

This thesis studies issues related to problems that arise in large-scale distributed environments with non-cooperative users, who act strategically and compete with each other to maximize their personal payoffs.

For instance, imagine a scenario where a set of users compete over a resource, such as the bandwidth of a communication link or advertisement slots when keywords are queried in search engines on the Internet. A mechanism takes input from all participating users (which represents their preferences) and outputs an allocation of the resource to them (it distributes the bandwidth or assigns slots). Each user aims to select her input to the mechanism in order to satisfy her personal objectives (possibly by misreporting her true preferences), without caring about the social welfare which we would like to maximize as the designers of the mechanism. Therefore, this behavior induces a strategic game among the users who act as players and sequentially change their strategies until they reach an equilibrium state (if one exists) from which no one has any incentive to deviate. Due to the strategic behavior of the users, the equilibrium that is reached may be of low quality in terms of some objective function like the social welfare, compared to what could happen if a central authority dictated the strategies of the users. The price of anarchy and stability are two quantification measures of this kind of inefficiency at equilibrium.

Our main goal in this thesis is to understand the advantages and constraints of the strategic games that arise in non-cooperative environments as means of computation. What can they compute and how well can they compute it? Is it possible to alter the rules of the game and incentivize the players to truthfully report their preferences? We answer to such questions related to equilibrium computation, price of anarchy and stability estimation, and truthful mechanism design for many interesting and important classes of problems. In particular, we study resource allocation with budget constraints, opinion formation, ownership transfer with expert advice, and revenue maximization in randomized combinatorial sales.

Σχεδιασμός και ανάλυση αλγορίθμων για μη συνεργατικά περιβάλλοντα

Αλέξανδρος Ανδρέας Βουδούρης

Περίληψη

Η παρούσα Διατριβή μελετά προβλήματα που προκύπτουν σε περιβάλλοντα μεγάλης κλίμακας με εγωκεντρικούς χρήστες, οι οποίοι συμπεριφέρονται στρατηγικά και ανταγωνίζονται μεταξύ τους με σκοπό να μεγιστοποιήσουν το ατομικό τους κέρδος.

Για παράδειγμα, φανταστείτε ένα σύνολο από χρήστες που ανταγωνίζονται για έναν πόρο, όπως το εύρος ζώνης ενός τηλεπικοινωνιακού καναλιού ή θέσεις διαφήμισης σε αποτελέσματα αναζήτησης στο Διαδίκτυο. Ένας μηχανισμός δέχεται εισοδο από όλους τους χρήστες και παράγει ως έξοδο μια κατανομή του πόρου σε αυτούς. Κάθε χρήστης προσπαθεί να επιλέξει την εισοδο του έτσι ώστε να εξυπηρετήσει τα προσωπικά του συμφέροντα (ενδεχομένως αποκρύπτοντας τις αληθινές του προτιμήσεις), χωρίς να νοιάζεται για το κοινωνικό όφελος το οποίο εμείς επιθυμούμε να μεγιστοποιήσουμε ως σχεδιαστές τους μηχανισμούς. Η συμπεριφορά αυτή ορίζει ένα στρατηγικό παιχνίδι μεταξύ των χρηστών οι οποίοι αλλάζουν στρατηγικές μέχρι το παιχνίδι να φτάσει σε κατάσταση ισορροπίας από την οποία κανείς δεν έχει κίνητρο να αποκλίνει. Η ισορροπία ενδέχεται να έχει μικρή απόδοση (σύμφωνα με κάποια αντικειμενική συνάρτηση όπως το κοινωνικό όφελος) σε σχέση με το τι θα μπορούσε να συμβεί αν κάποια κεντρική αρχή διέταζε τους χρήστες για το πώς να συμπεριφερθούν. Το κόστος της αναρχίας και της ευστάθειας είναι δύο μετρικές που χρησιμοποιούνται για την ποσοτικοποίηση αυτής της μη-αποδοτικότητας.

Κύριος στόχος μας είναι η κατανόηση των δυνατοτήτων καθώς και των περιορισμών των στρατηγικών παιχνιδιών ως μέσα υπολογισμού. Τι μπορούν να υπολογίσουν και πόσα καλά μπορούν να το υπολογίσουν; Είναι δυνατόν να μεταβάλλουμε τους κανόνες του παιχνιδιού έτσι ώστε οι παίκτες να έχουν κίνητρο να λένε πάντα την αλήθεια; Απαντάμε σε τέτοιου είδους ερωτήσεις μελετώντας ζητήματα υπολογισμού ισορροπιών, εκτίμησης κόστους αναρχίας και ευστάθειας, καθώς και σχεδίασης φιλαληθών μηχανισμών για ενδιαφέρουσες και σημαντικές κλάσεις προβλημάτων. Πιο συγκεκριμένα, ασχολούμαστε με προβλήματα ανάθεσης πόρων υπό περιορισμούς προϋπολογισμού, διαμόρφωσης απόψεων σε κοινωνικά δίκτυα, μεταφοράς ιδιοκτησίας, και μεγιστοποίησης εσόδων σε συνδυαστικές αγορές.

Acknowledgements

To begin, I would like to thank my advisor, Iannis Caragiannis. Almost five years ago, I was looking to find a constructive way to spend the last semester of my undergraduate studies and was unsure of whether I should continue towards a Masters degree (and a PhD) or go after a software developer position at some company. This dilemma was resolved soon after the first few meetings that me and Iannis had about a possible side-project collaboration. His passion about research and his way of thinking about problem solving quickly convinced me that research was something that I really wanted to do for the rest of my life. I am forever grateful to Iannis for his unstoppable support, his valuable guidance throughout my postgraduate studies, his general belief in me, and the countless hours that we spent together trying to solve the problems considered in this thesis (and many more).

I am also thankful to this dissertation's committee members: Stavros Cosmadakis, Christos Kaklamanis, Vangelis Markakis, Sotiris Nikolettseas, Paul Spirakis, and Christos Zaroliagis for their useful comments on the presentation and structure of this thesis.

Special thanks go to all my co-authors: Fidaa Abed, Xenophon Chatzigeorgiou, Aris Filos-Ratsikas, Panagiotis Kanellopoulos, George Krimpas, Adelina Madhja, Swaprava Nath, Sotiris Nikolettseas, Marianna Panteli, Nikos Protopapas, and Dimitris Tsolovos. Their contribution in our joint projects as well as the many interesting and fruitful discussions that we had on research and other topics have deeply helped me define myself as a researcher and a team player.

I would like to thank my entire family for always encouraging me to pursue my dreams in research and for always being around to support me when things didn't go as planned. Finally, I am grateful to my wife Adelina for choosing to follow me along this journey, and for her love and support that helped me overcome the many obstacles that I came across.

This PhD thesis was partially supported by a scholarship from the Caratheodory research grant E.114 "Strategic games as computational machines – Design and analysis of algorithms for non-cooperative environments" from the University of Patras (October 2014 – September 2016), and by a PhD scholarship from the Alexander S. Onassis Foundation (October 2016 – September 2018).

Alexandros Voudouris
Patras, September 2018

Contents

1	Introduction	1
1.1	Resource allocation with budget constraints	4
1.2	Opinion formation and compromise	6
1.3	Ownership transfer	8
1.4	Asymmetry of information in revenue maximization	10
2	The efficiency of resource allocation mechanisms for budget-constrained users	13
2.1	Overview of contribution and techniques	13
2.1.1	Chapter roadmap	14
2.2	Related work	15
2.3	Definitions and notation	15
2.3.1	Examples of resource allocation mechanisms	18
2.4	A lower bound for all mechanisms	19
2.5	The structure of worst-case games and equilibria	20
2.6	Pay-your-signal mechanisms	26
2.7	Two-player mechanisms	29
2.7.1	The E2-PYS mechanism	29
2.7.2	The E2-SR mechanism	31
2.8	Some extensions	33
2.8.1	Budget-aware mechanisms	34
2.8.2	Higher expressiveness	35
2.9	Conclusion	35

3	Bounding the inefficiency of compromise in opinion formation	37
3.1	Overview of contribution and techniques	37
3.1.1	Chapter roadmap	38
3.2	Related work	38
3.3	Definitions and notation	40
3.4	Some properties about equilibria	43
3.5	Existence and quality of equilibria	47
3.5.1	Existence of equilibria	47
3.5.2	Price of stability	48
3.6	Complexity of equilibria	51
3.7	Upper bounds on the price of anarchy	53
3.8	An improved bound on the price of anarchy for 1-COF games	65
3.9	Lower bounds on the price of anarchy	70
3.9.1	The case of 1-COF games	70
3.9.2	The general case of k -COF games with $k \geq 2$	73
3.10	Conclusion	76
4	Truthful mechanisms for ownership transfer with expert advice	77
4.1	Overview of contribution and techniques	77
4.1.1	Chapter roadmap	79
4.2	Related work	79
4.3	Definitions and notation	81
4.3.1	An alternative view of profiles	83
4.3.2	Reasoning about truthfulness	84
4.4	Ordinal mechanisms	88
4.5	Bid-independent mechanisms	90
4.6	Expert-independent mechanisms	95
4.7	Beyond expert-independent mechanisms	98

4.8	Unconditional lower bounds	102
4.9	Conclusion	103
5	Near-optimal asymmetric binary matrix partitions	105
5.1	Problem definition and overview of contribution	105
5.1.1	Chapter roadmap	106
5.2	Related work	107
5.3	Definitions, examples and structural observations	108
5.4	The uniform case	112
5.5	Asymmetric binary matrix partition as welfare maximization	121
5.6	Conclusion	128
6	Conclusions and open problems	131
6.1	Resource allocation and auctions for budget-constrained users	131
6.2	Compromising opinion formation	133
6.3	Ownership transfer	134
6.4	Asymmetric matrix partition	136
	References	136
A	Extended abstract in Greek	149

Chapter 1

Introduction

Over the last two decades, the rapid and continuously increasing development of large-scale distributed systems and social networks, has led to the implementation of non-cooperative environments, where multiple self-interested agents may compete with each other in many different contexts. For instance, such agents could be the users of a communication link that compete over the limited available bandwidth, advertisers that compete over advertising space when keywords are queried in search engines, potential buyers that compete over acquiring government assets, or even simple people that debate with their social acquaintances over an issue by expressing opinions.

In such scenarios, each agent aims to select the best possible strategy in order to optimize various personal objectives (for example, she might want to maximize some utility function or minimize some cost function, depending on context), which are not only affected by the underlying structure of the environment, but also by the other agents and the strategies that they choose. Consequently, the agents engage as players into a *non-cooperative strategic game* [Nash, 1951, Nisan et al., 2007], which is defined by the ground rules of the environment (the underlying mechanism that is used), as well as by the different possible strategies and the personal objectives of the participating players.¹ When all players have chosen strategies such that they simultaneously maximize their utility (in the sense that none of them has any incentive to deviate to a different strategy in order to even slightly increase her personal utility), then we say that the corresponding strategic game has reached a stable state, an *equilibrium*. There are many important computational and mathematical questions regarding stability, computational complexity and efficiency in strategic games.

¹It is worth remarking that non-cooperative games significantly differ from cooperative games [Chalkiadakis et al., 2011], where the players are allowed to cooperate with each other and form coalitions in order to collectively achieve to optimize their personal objectives.

Existence of equilibria and complexity

The first apparent question is the following one: Do equilibria always exist? In his famous 1951 paper, John F. Nash proved that any finite strategic non-cooperative game has at least one equilibrium if the players are allowed to choose probability distributions over their strategies; these distributions are called *mixed* strategies. However, this is not always the case when the players choose their strategies deterministically; these are called *pure* strategies. If equilibria do exist, then what is the complexity of computing them? There has been much research devoted to this question, especially for the case of mixed equilibria for which the existential aspect has already been answered positively. In general, Daskalakis et al. [2009] and Chen et al. [2009] proved that the problem of computing a mixed equilibrium in reasonable time is hard (PPAD-complete) even for two players only.

Efficiency at equilibrium

Of course, one of the most important issues in computing systems are related to efficiency. To measure efficiency in a strategic game, we can define a social objective function over all possible states of the game; the value of this social function for a particular state of the game can be thought of as representing the total happiness (or unhappiness, depending on context) of the players for this state. As system designers, we would like the game to end up in a state that globally maximizes the social function in order to maximize the total happiness of the participating players. However, this goal is not always totally aligned to the selfishness of the players, which may only lead to local maxima of the social function instead.

To quantify the worst-case degradation of quality in equilibria, in their celebrated 1999 paper, Koutsoupias and Papadimitriou introduced the notion of the *price of anarchy*, which is defined as the ratio between the maximum value of the social function (attained at any state of the game) and the minimum value of the social function attained at any equilibrium; essentially, the price of anarchy is an analogue to the approximation ratio in combinatorial optimization [Vazirani, 2001, Williamson and Shmoys, 2011]. The similar notion of the *price of stability* was later introduced by Anshelevich et al. [2008] in order to quantify the best-case efficiency loss at equilibria (using the maximum function value equilibrium instead of the minimum one in the definition of the ratio).

Apart from the aforementioned papers that introduced the price of anarchy and the price of stability, these notions have been used extensively in order to bound the inefficiency of

equilibria in many important classes of strategic games that naturally arise in distributed systems. Indicatively, they have been applied in the context of congestion games (for instance, see the papers by Roughgarden and Tardos [2002] and Christodoulou and Koutsoupias [2005], which are among the very first ones on this topic) and auctions (for example, see the seminal work of Christodoulou et al. [2016a], which initiated the analysis of price of anarchy in Bayesian auction settings, as well as the recent survey by Roughgarden et al. [2017], which goes through almost all recent developments on issues related to efficiency in auctions).

Mechanism design

Another rich line of research that focuses on mechanism design deals with the question of whether we can guide the strategic behavior of the players by altering the rules of the game so that they have the motive to act truthfully and the game is able to reach more efficient equilibria. A prime example of such a mechanism is the single-item second price auction of Vickrey [1961] which allocates the item to the highest bidder and requires a payment from her that is equal to the second highest bid. This auction format achieves to allocate the item to the player that values it the most and is truthful in the sense that all participants have incentive to simply bid the value they have for the item.

For multiple items, the ideas of the second price auction are adapted by the well-known Vickrey-Clarke-Groves (VCG) mechanism, which allocates the items in order to maximize the total value of the players, while it requires from each of them a payment that is equal to the value of the allocation that would be computed if they did not participate in the auction. Even though the VCG mechanism looks like the ideal solution, it cannot always be applied. In many scenarios, it may require to search over an exponentially large space in order to identify the optimal allocation, while the players need to report their whole valuation functions, which may lead to exponential communication complexity.

Given these limitations of the VCG mechanism in many interesting scenarios as well as the fact that it requires the use of monetary transfers in order to operate,² a plethora of researchers have instead focused on the design of simple truthful mechanisms that are approximately optimal and may use money [Nisan and Ronen, 2001] or not [Procaccia and Tennenholtz, 2013]. Subsequently, approximate mechanism design has been applied in many different settings like in combinatorial auctions [Dobzinski et al., 2012, Mualem and Nisan, 2008], keyword search

²Actually, the use of money is essential in order to avoid well-known impossibilities from social choice theory, which state that optimal and truthful mechanisms are necessarily dictatorial [Gibbard, 1973, Satterthwaite, 1975].

auctions [Aggarwal et al., 2006], fair division [Cole et al., 2013], social welfare maximization problems [Briest et al., 2011, Filos-Ratsikas et al., 2014, Filos-Ratsikas and Miltersen, 2014], scheduling problems [Archer and Tardos, 2001], and even kidney exchange [Ashlagi et al., 2015, Caragiannis et al., 2015].

Problems considered in this thesis

In this thesis, we consider issues related to stability, computational complexity, efficiency, and mechanism design for four different problems that emerge due to the strategic behavior of the participating agents. In particular, we first focus on the efficiency of mechanisms for the allocation of a single divisible resource among users that have budget constraints. Second, we turn our attention to a particular class of compromising opinion formation games. Third, we consider a novel mechanism design problem related to ownership transfer. Finally, we also design approximation algorithms for revenue maximization in combinatorial sales. In the rest of this chapter, we will give a comprehensive introduction and motivation for each of these problems, and shortly discuss our contribution and techniques.

1.1 Resource allocation with budget constraints

Resource allocation is an ubiquitous task in computing systems and usually sets non-trivial algorithmic challenges to their design. As such, resource allocation problems have received much attention by the algorithmic community for decades. The recent emergence of large-scale distributed systems with non-cooperative users that compete for access to scarce resources has led to game-theoretic treatments of resource allocation.

In this thesis, we study a particular simple class of *resource allocation mechanisms* that aim to distribute a divisible resource (such as bandwidth of a communication link, CPU time, storage space, etc.) by *auctioning it off* to different users as follows. Each user is asked to submit a scalar *signal*. Given the submitted signals, the mechanism decides the fraction of the resource that will be allocated to each user, as well as the payment that will be received from each of them. A typical example is a mechanism that has been proposed by Kelly [1997] (henceforth called the Kelly mechanism; see also Kelly et al. [1998]), according to which the fraction of the resource allocated to each user is proportional to the user's signal, and the signal itself is her payment.

Following the standard modeling assumptions in the related literature, the value of each user for a resource fraction is given by a private *valuation function*. The above definition of resource allocation mechanisms allows the users to act strategically in the sense that the

signal they select to submit is such that their utility (value for the fraction of the resource they receive minus payment) is maximized. Naturally, this behavior defines a strategic game among the users, who act as players. Soon after the definition of the Kelly mechanism, a series of papers studied the existence and uniqueness of *pure Nash equilibria* (snapshots of player strategies, in which the signal of each player maximizes her own utility) of the induced games [Hajek and Gopalakrishnan, 2002, La and Anantharam, 2000, Maheswaran and Basar, 2003] and quantified their inefficiency [Johari and Tsitsiklis, 2004] using the notion of the *price of anarchy* [Koutsoupias and Papadimitriou, 1999].

In particular, Johari and Tsitsiklis [2004] used the *social welfare* — the total value of the players for their received fraction of the resource — as an efficiency benchmark and proved that the social welfare at any equilibrium is at least $3/4$ times the optimal social welfare. This translates into a price of anarchy bound of $4/3$, which is tight. The paper of Johari and Tsitsiklis [2004] sparked subsequent research on other resource allocation mechanisms, that use different allocation rules or payments.

A first apparent question was whether improved price of anarchy bounds are possible by changing the proportional allocation function, but keeping the simple pay-your-signal (PYS, for short) payment rule. Sanghavi and Hajek [2004] showed that no PYS mechanism has price of anarchy better than $8/7$, designed an allocation function that achieves this bound for two players, and provided strong experimental evidence that a slightly inferior bound holds for arbitrarily many players. Surprisingly, full efficiency at equilibria (i.e., a price of anarchy equal to 1) is possible via different allocation/payment functions. This discovery was made in three independent papers by Maheswaran and Basar [2006], Yang and Hajek [2007], and Johari and Tsitsiklis [2009]. The mechanism of Maheswaran and Basar [2006] uses proportional allocation but different payments (see Section 2.3 for its description), while the mechanisms of Johari and Tsitsiklis [2009] and Yang and Hajek [2007] are adaptations of the well-known VCG paradigm (see also the survey by Johari [2007] on these results).

Our focus is on the — arguably, more realistic — setting, in which each player has a private *budget* that restricts the payments that she can afford and, consequently, narrows her strategy space. As resource allocation mechanisms do not have direct access to budgets, the set of equilibria can drastically change and their social welfare can be extremely low compared to the optimal social welfare, which in turn is not related to player strategies, payments, or budgets. An efficiency benchmark that is suitable for budget-constrained players is known as *liquid*

welfare (introduced by Dobzinski and Paes Leme [2014] and, independently, by Syrgkanis and Tardos [2013] who call it *effective welfare*) and is obtained by slightly changing the definition of the social welfare, taking budgets into account. Informally, the liquid welfare is the total value of the players for the resource fraction they receive, with the value of each player capped by her budget. Following the recent paper of Azar et al. [2017], we use the term *liquid price of anarchy* (and abbreviate it as LPOA) to refer to the price of anarchy with respect to the liquid welfare, i.e., the ratio between the optimal liquid welfare of a game induced by a resource allocation mechanism and the worst liquid welfare over all equilibria of the game.

Our contribution

In chapter 2, we show a tight bound of 2 on the liquid price of anarchy of the Kelly mechanism and an unconditional lower bound of $2 - 1/n$ for any n -player mechanism, essentially proving that Kelly is best possible among all multi-user resource allocation mechanisms. In our proofs, we exploit the particular structure of worst-case games and equilibria, which also allows us to design (nearly) optimal two-player mechanisms by solving simple differential equations. These results have been published in [Caragiannis and Voudouris, 2018].

1.2 Opinion formation and compromise

Opinion formation has been the subject of much research in sociology, economics, physics, and epidemiology for decades. The widespread adoption of the Internet has allowed the recent blossoming of social networks, which have facilitated information dissemination in ways that have been beneficial for their users, but they are often used strategically in order to spread information that only serves the objectives of particular parties. These properties have recently attracted the interest of researchers in artificial intelligence [Auletta et al., 2016, Schwind et al., 2015, Tsang and Larson, 2014] as well as in computer science at large [Bindel et al., 2015, Mossel and Tamuz, 2014, Olshevsky and Tsitsiklis, 2009], and has led to revisions of classical opinion formation models from sociology using game-theoretic notions and tools.

An influential model that captures the adoption of opinions in a social context has been proposed by Friedkin and Johnsen [1990]. According to this, each individual has an internal belief on an issue and publicly expresses a (possibly different) opinion; internal beliefs and public opinions are modeled as real numbers. In particular, the opinion that an individual expresses follows by averaging between her internal belief and the opinions expressed by her social acquaintances. Recently, Bindel et al. [2015] showed that this behavior can be interpreted

through a game-theoretic lens: averaging between the internal belief of an individual and the opinions in her social circle is simply a *strategy* that minimizes an implicit cost for the individual. This cost is defined using a quadratic function which is equal to the total squared distance of the opinion that the individual expresses from her belief and the opinions expressed in her social circle. In a sense, the strategic behavior of the individual leads to opinions that follow the majority of her social acquaintances.

Bindel et al. [2015] considered a static snapshot of the social network and assumed that the opinion of each individual is affected by *all* of her social acquaintances. However, in reality, as opinions evolve, people usually tend to disregard opinions that are far away from their own personal beliefs, even if these are expressed by their best friends. Following such a reasoning, Bhawalkar et al. [2013] implicitly assumed that the opinion of an individual depends only on a *small* number of people in her social circle, her *neighbors*. So, in their model, opinion formation *co-evolves* with the neighborhood for each individual, which consists of those people who have opinions that are similar to her belief. Then, the opinion expressed is assumed to minimize the same quadratic cost function that was previously used by Bindel et al. [2015], taking now into account the neighborhood instead of the whole social circle.

Both Bindel et al. [2015] and Bhawalkar et al. [2013] were able to prove small constant ($9/8$ and approximately 14, respectively) bounds on the price of anarchy of the strategic games that may be induced by the assumptions of their models. These bounds essentially indicate that an abnormally high fraction of the population of people expresses opinions that are close to their personal beliefs. Unfortunately, this is hard to rationalize given the so many different and, in some cases, extreme opinions that are expressed, for example, in discussions regarding politics or religion.

We follow the co-evolutionary model introduced by Bhawalkar et al. [2013], and assume that the neighborhood of each individual consists of the k other individuals whose opinions are the closest ones to her belief. However, we deviate from the quadratic cost definition and, instead, consider individuals that seek to *compromise* more with their neighbors, by assuming that each individual aims to minimize the *maximum* distance of the opinion she expresses from her internal belief and each of the opinions expressed in her neighborhood. Naturally, these modeling decisions lead to the definition of strategic games, which we call *k-compromising opinion formation (k-COF) games*, where each individual is a cost-minimizing player with the opinion expressed as her strategy.

Our contribution

In chapter 3 of this thesis, we quantify the inefficiency of equilibria arising in k -COF games and show that compromise comes at a cost that strongly depends on the neighborhood size. Specifically, we prove (both upper and lower) bounds on the price of anarchy and stability [Anshelevich et al., 2008], which depend linearly on the neighborhood size. For the special case of $k = 1$ we also design a simple algorithm that is based on path computations on particularly defined directed acyclic graphs, which can verify whether there exists a pure equilibrium or not, and in case it does, it can compute both the best and the worst equilibrium (in terms of social cost). These results have been published in [Caragiannis et al., 2017a].

1.3 Ownership transfer

Most well-studied problems in *computational social choice* [Brandt et al., 2016] deal with the task of merging individual preferences over alternatives – often expressed as *rankings* – into a collective choice [Caragiannis et al., 2017,b, Procaccia et al., 2012, Skowron et al., 2016]. More often than not, the mechanisms employed for this aggregation task are *ordinal* and do not utilize the intensities of the preferences of the individuals. Further, due to several well-known impossibility theorems [Gibbard, 1973, 1977, Satterthwaite, 1975], these mechanisms are also *non-truthful*, meaning that some of the participating individuals may have strong incentives to misreport their preferences to manipulate the mechanism to output an alternative that they prefer more.

In contrast, the class of truthful *cardinal* mechanisms has been shown to be much richer [Barbera et al., 1998, Feige and Tennenholtz, 2010, Freixas, 1984] and exploiting the additional information provided by the numerical values (expressing individual preferences) can notably increase the overall well-being of the society [Cheng, 2016, Filos-Ratsikas and Miltersen, 2014, Guo and Conitzer, 2010]. At the same time, truthful mechanisms *with money* are pretty well-understood by now and the welfare-maximizing mechanisms for a wide class of problems are known [Nisan et al., 2007]. A celebrated such example is the family of VCG mechanisms [Clarke, 1971, Groves, 1973, Vickrey, 1961].

However, in a rich set of *hybrid social choice* problems, monetary transfers are possible only for a fraction of the participating individuals. This naturally renders solutions like the VCG mechanism insufficient. Therefore, designing *truthful, cardinal* mechanisms is a much more challenging task and one needs to combine elements of *mechanism design with money* and *social*

choice.

We provide a few examples of such hybrid social choice scenarios. Government agencies routinely sell public assets such as spectrum, land, or government securities, by transferring their ownership (or usage rights). As such transfers may have huge impact to citizens, the decision about the new ownership is not simply the outcome of some competitive process among the potential buyers (for instance, through an auction), but it usually also involves experts from the citizen community who provide advice regarding the societal impact of each potential ownership transfer [Janssen, 2004]. In contrast to each potential buyer who faces a value-for-money trade-off, the experts care only about societal value; their compensation is unrelated to the ownership decision and instead depends on their reputation and experience only. The government needs both parties for a successful transfer of the public assets and a reasonable goal would be to maximize the *social welfare*, which aggregates the values of buyers and experts for the ownership transfer.

A very similar situation occurs for private ownership transfers. Mergers and acquisitions play a central role in the competition among private players in a market, and the rules or the policies that dictate the mergers are often up for debate.³ There is ample evidence to support the fact that the transfer of ownership of an organization has a significant impact on the economy of the employees and the customers [Auerbach, 2008, Hitt et al., 2001]. The current owner or the administration can employ industry experts for their opinion on the transfer and ask the potential buyers to quote their values. Similarly to the previous example, the administration takes into account the input of both parties and social welfare maximization among them is a reasonable goal. Furthermore, in the organization of sporting events, the bids of the potential hosts are taken into consideration along with the recommendations of a respective sports administrative body (for example, IOC for the Olympic Games, FIFA for the World Cup, and FIA for Formula One).

Motivated by examples like the ones described above, we consider a setting where the bidders offer monetary compensations (to buy into a new company or a government asset), but the experts (the citizen representatives or the administrative body) do not. The objective is to achieve the decision that maximizes the *social welfare*, which includes the *cardinal* values of both the expert and the bidders. This is a hybrid social choice setting that blends together classical social choice and classical mechanism design with money, but is distinct from both of

³EU data show that more than 6500 mergers have taken place in the EU since 1990, and strict rules are in effect for mergers [European Commission, 2018].

them, thereby rendering celebrated solutions like the VCG mechanism insufficient.

Our contribution

In chapter 4, we study the fundamental version of the aforementioned ownership transfer problem with one expert and two potential buyers, and provide *tight approximation guarantees* of the optimal social welfare for many classes of truthful mechanisms. We distinguish between mechanisms that use ordinal and cardinal information, as well as between mechanisms that base their decisions on one of the two sides (either the buyers or the expert) or both. Our analysis shows that the cardinal setting is quite rich and admits several non-trivial randomized truthful mechanisms, and also allows for closer-to-optimal-welfare guarantees. These results can be found in [Caragiannis et al., 2018].

1.4 Asymmetry of information in revenue maximization

Exploiting information asymmetries to maximize revenue dates all the way back to the seminal work of Akerlof [1970] who considered such issues in the so-called *market for lemons*. Suppose a market for cars including high-quality ones (which are known as *peaches* in American slang) as well as low-quality ones that can be found to be defective only after they have been bought; these are known as *lemons* (due to the sourness they cause to their buyers). In such a market, the seller has much more accurate information about the quality of the cars, while the potential buyers do not and cannot distinguish between peaches and lemons. This boils down to an interesting strategic decision making problem from the seller's side to find the best possible way to exploit the situation and sell the items at a higher price than the one that could be set if the buyers knew the whole truth. As expected, the work of Akerlof [1970] on information asymmetry sparked subsequent research in economics [Crawford and Sobel, 1982, Levin and Milgrom, 2010, Milgrom, 2010, Milgrom and Weber, 1982] and, recently, in computer science as well [Dughmi, 2014, Emek et al., 2012, Ghosh et al., 2007, Guo and Deligkas, 2013, Miltersen and Sheffet, 2012].

Following the work of Alon et al. [2013], we focus on randomized *take-it-or-leave-it sales*. There are m items and n potential buyers. Each buyer has a value for each item, and she is generally unaware of the existence of the other buyers and their values. In contrast, the seller is assumed to know the values of the buyers for the items. According to some probability distribution, nature selects a single item for sale at random, and this random choice is revealed to the seller, but not to the buyers. Then, the seller approaches the highest value buyer and

offers the item to her at a price that is equal to her value for it. A specific instantiation of this setting could be the following: the items correspond to keywords and the potential buyers correspond to advertisers. Every advertiser has a value for each keyword which represents the maximum amount of money she is willing to pay in order to occupy the advertising space that is allocated when the particular keyword is queried. The role of nature is played by users who submit queries, and the role of the seller is played by the search engine, which allocates the advertising space according to the keyword queried each time, and in such a way that its revenue is maximized.

Can the seller exploit the fact that she has much more accurate information about the items for sale compared to the potential buyers? In particular, information asymmetry arises since the seller knows the realization of the randomly selected item whereas the buyers do not. A possible approach is to let the seller define a buyer-specific signalling scheme. That is, for each buyer, the seller can partition the set of items into disjoint subsets (bundles) and report this partition to the buyer. For example, the search engine could bundle together keywords that are closely related to each other. After nature's random choice, the seller can reveal to each buyer the bundle that contains the realization, thus enabling her to re-evaluate her beliefs for the particular bundle (i.e., compute her expected value for the whole bundle and each item therein).

Alon et al. [2013] introduced the *asymmetric matrix partition* problem as an abstraction of revenue maximization in take-it-or-leave-it sales. Instances of the problem consist of an $n \times m$ matrix A containing non-negative real values and a probability distribution over its columns. A *partition scheme* $B = (B_1, \dots, B_n)$ consists of a partition B_i for each row i of A . The partition B_i acts as a smoothing operator on row i that distributes the expected value of each partition subset proportionally to all its entries. Given a scheme B that induces a smooth matrix A^B , the partition value is the expected maximum column entry of A^B . The objective is to compute a partition scheme such that the resulting partition value is maximized. The relation to take-it-or-leave-it sales should be apparent: the columns of the input matrix correspond to items, the rows correspond to potential buyers, and the value of the entry (i, j) corresponds to the value that buyer i has for item j . After the bundling of the items for a specific buyer, the smooth value of a bundle corresponds to the expected value the buyer has for each item included in the bundle. Finally, the partition value corresponds to the expected revenue of the seller. Among other results, Alon et al. [2013] proved that the problem is APX-hard even for the simplest case

of binary matrices, and designed 0.563 - and $1/13$ -approximation algorithms for the cases of uniform and non-uniform probability distributions, respectively.

Our contribution

In chapter 5, we significantly improve both results of Alon et al. [2013]. We present a $9/10$ -approximation algorithm for the case where the probability distribution is uniform and a $(1 - 1/e)$ -approximation algorithm for non-uniform distributions. Although our first algorithm is combinatorial (and very simple), the analysis is based on linear programming and duality arguments. In our second result we exploit a nice relation of the problem to submodular welfare maximization. These results have been published in [Abed et al., 2018].

Chapter 2

The efficiency of resource allocation mechanisms for budget-constrained users

In this chapter we present our results on the efficiency of resource allocation mechanisms for users with budget constraints, as they were published in [Caragiannis and Voudouris, 2018]; see the discussion in Section 1.1 for a comprehensive introduction to the problem.

2.1 Overview of contribution and techniques

We aim to explore *all* resource allocation mechanisms to find the one with the best possible LPoA. Our results suggest a drastically different picture compared to the no-budget setting. First, the analogue of full efficiency is not achievable; we show a lower bound of $2 - 1/n$ on the LPoA of any n -player resource allocation mechanism (under standard technical assumptions for player valuations and mechanism characteristics). We prove that the Kelly mechanism has an almost best possible LPoA of exactly 2, while the Sanghavi and Hajek (SH) mechanism has an LPoA of 3. Improved bounds are possible for two players. We design the two-player pay-your-signal (PYS) resource allocation mechanism E2-PYS that has an LPoA of 1.792; this bound is optimal among a very broad class of mechanisms. We also design the two-player mechanism E2-SR that achieves an almost optimal LPoA bound of at most 1.529; this mechanism uses different payments. See Table 2.1 for a summary.

Our results exploit a particular structure of worst-case (in terms of LPoA) games and their equilibria. We prove that for every resource allocation mechanism, the worst-case LPoA is obtained at instances in which players have *affine* valuation functions. In addition, all players besides one have finite budgets and play strategies that imply payments that are either zero or

Mechanism	LPoA	Comment
all	$\geq 2 - 1/n$	No mechanism can achieve full efficiency (Theorem 2.1)
Kelly	2	Tight bound; almost optimal among all n -player mechanisms (Theorem 2.5)
SH	3	Tight bound (Theorems 2.6 and 2.7)
E2-PYS	1.792	Tight bound (Theorem 2.8); optimal among all 2-player PYS mechanisms with concave allocation functions (Theorem 2.9)
E2-SR	1.529	Almost optimal among all 2-player mechanisms (Theorem 2.10)

Table 2.1: Summary of our liquid price of anarchy bounds for resource allocation mechanisms for budget-constrained users; see [Caragiannis and Voudouris, 2018].

equal to their budget, while a single player has infinite budget and a signal that nullifies the derivative of her utility. Compared to an analogous characterization for the no-budget case (with linear valuation functions and player signals that all nullify their utility derivatives), first observed by Johari and Tsitsiklis [2004] for the Kelly mechanism and later extended to all resource allocation mechanisms, the structure in our characterization is much richer and the proof is considerably more complicated. The characterization contains so much information that the LPoA bounds follow rather easily; the extreme example is the proof of our best LPoA bound of 2 for the Kelly mechanism which is only a few lines long. It can also be used in the design of new mechanisms; for example, the design and analysis of our two-player mechanisms E2-PYS and E2-SR follow by simple first-order differential equations, which would never have been identified without our characterization. And, furthermore, under assumptions about the resource allocation mechanisms (e.g., concave allocations and convex payments), the LPoA bound is automatically proved to be tight without the need to provide any explicit lower bound instance.

2.1.1 Chapter roadmap

The rest of the chapter is structured as follows. We begin with a discussion of other related work in Section 2.2. Then, we continue with preliminary definitions, notation and examples in Section 2.3. Our unconditional lower bound on the liquid price of anarchy of any resource allocation mechanism appears in Section 2.4. Section 2.5 is devoted to proving the structural characterization of worst-case resource allocation games and equilibria. Then, in Section 2.6 we present tight bounds on the liquid price of anarchy for the Kelly and SH mechanisms. In Section 2.7, we present our two-player mechanisms E2-PYS and E2-SR. Finally, we present some interesting extensions of our work in Section 2.8 and conclude in Section 2.9.

2.2 Related work

As an efficiency benchmark, the liquid welfare has been studied recently in many different contexts such as in the design of truthful mechanisms (see [Dobzinski and Paes Leme, 2014, Lu and Xiao, 2015, 2017]) and in the analysis of combinatorial Walrasian equilibria with budgets [Dughmi et al., 2016]. In the context of the price of anarchy, it was considered in simultaneous first price auctions by Azar et al. [2017] and in position auctions by Voudouris [2018].

Caragiannis and Voudouris [2016] were the first to prove that the LPoA of Kelly is constant. In particular, they showed upper and lower bounds of 2.78 and 2, respectively. The lower bound is essentially proved again here (see Theorem 2.5) with a completely different and more interesting technique. Christodoulou et al. [2016b] improved the LPoA upper bound to 2.618 and extended the results to more general settings involving multiple resources. Prior to these two papers, Syrgkanis and Tardos [2013] proved that the social welfare at equilibria of the Kelly mechanism is at most a constant factor away from the optimal liquid welfare.

In contrast to the analysis techniques of this chapter, the analysis of the Kelly mechanism by Caragiannis and Voudouris [2016], Christodoulou et al. [2016b] and Syrgkanis and Tardos [2013] is closer in spirit to the smoothness template [Roughgarden, 2015, Roughgarden et al., 2017] and is based on bounding the utility of each player by the utility she would have when deviating to appropriate signals. Their results extend to more general equilibrium concepts such as coarse-correlated or Bayes-Nash equilibria. Our LPoA bounds here hold specifically for pure Nash equilibria, but are superior and tight.

2.3 Definitions and notation

We consider a single divisible resource of unit size that is distributed among n users by a *resource allocation mechanism* M . The mechanism M consists of

- an *allocation function* $g^M : \mathbb{R}_{\geq 0}^n \rightarrow \mathcal{Q} \cup \mathbf{0}$, where $\mathcal{Q} = \{\mathbf{d} \in [0, 1]^n : \sum_{i=1}^n d_i = 1\}$ is the unit n -simplex and $\mathbf{0} = (0, \dots, 0)$, and
- a *payment function* $p^M : \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}^n$,

and works as follows. Each user i submits a *signal* $s_i \in \mathbb{R}_{\geq 0}$, and the mechanism M allocates a fraction of $g_i^M(\mathbf{s})$ of the resource to each user i and asks her for a payment of $p_i^M(\mathbf{s})$, where $\mathbf{s} = (s_1, \dots, s_n)$ denotes the vector formed by all signals.

Some important properties of allocation and payment functions are as follows:

- They are *anonymous*: any permutation of the entries of the input signal vector results in the same permutation of the output. So, all users get equal resource shares and are asked for equal payments when they submit identical signals;
- The mechanism does not allocate any fraction and does not ask for any payment from a user that submits a zero signal;
- By convention, when some user is the only one with a non-zero signal, she gets the whole resource and is asked for a payment of zero.

Let (y, \mathbf{s}_{-i}) denote the signal vector in which user i has signal y and the remaining users have their signals as in \mathbf{s} . Viewed as univariate functions (of variable y), the functions $g_i^M(y, \mathbf{s}_{-i})$ and $p_i^M(y, \mathbf{s}_{-i})$ are *increasing* and *differentiable* in $\mathbb{R}_{\geq 0}$ (with the exception of $(y, \mathbf{s}_{-i}) = \mathbf{0}$).

Each user i has

- a monotone non-decreasing, concave, and differentiable¹ *valuation function* $v_i : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$; $v_i(x)$ represents the value that user i has for a resource fraction of x ;
- a *budget* $c_i \in \mathbb{R}_{\geq 0} \cup \{+\infty\}$, which restricts (upper-bounds) her payment to the mechanism.

Her *utility* from the mechanism is defined as the value she gets for the fraction she is given minus her payment, i.e.,

$$u_i^M(\mathbf{s}) = v_i(g_i^M(\mathbf{s})) - p_i^M(\mathbf{s}).$$

To capture the fact that budgets impose hard constraints to the users, we technically assume that $u_i^M(\mathbf{s}) = -\infty$ when $p_i^M(\mathbf{s}) > c_i$.

The users act strategically as utility maximizers and, therefore, engage as players into a *strategic resource allocation game* \mathcal{G}^M that is *induced* by mechanism M . A (pure Nash) *equilibrium* is a signal vector \mathbf{s} such that, when viewed as a univariate function of variable y , $u_i^M(y, \mathbf{s}_{-i})$ is maximized for $y = s_i$, i.e., no player can increase her utility by unilaterally deviating to submitting a different signal. We denote by $\text{eq}(\mathcal{G}^M)$ the set of all equilibria of game \mathcal{G}^M . By the definition and properties of the allocation and payment functions, the signal vector $\mathbf{0}$ cannot be an equilibrium as (by the conventions mentioned above) any player has the incentive to

¹We remark that our results hold for semi-differentiable valuation functions as well. However, the proof of our characterization (Lemma 2.2) is technically more involved. So, the differentiability assumption keeps the exposition simple.

unilaterally deviate and get the whole resource without paying anything. We use \mathbb{X}_n as an abbreviation of the set $\mathbb{R}_{\geq 0}^n \setminus \{\mathbf{0}\}$.

Due to the budget constraints, we have three different cases for the strategy of player i at an equilibrium $\mathbf{s} \in \text{eq}(\mathcal{G}^M)$ (assuming a non-trivial budget $c_i > 0$) and for the corresponding value of the derivative of her utility. In particular, the derivative $\left. \frac{\partial u_i^M(y, \mathbf{s}_{-i})}{\partial y} \right|_{y=s_i}$ is equal to zero in case s_i is such that $0 < p_i^M(\mathbf{s}) < c_i$, non-positive in case $s_i = 0$, and non-negative in case s_i is such that $p_i^M(\mathbf{s}) = c_i$. Note that nullification of the utility derivative does not necessarily imply maximization of utility.

We are interested in studying the effect of the strategic behavior to the efficiency of resource allocation mechanisms. An efficiency benchmark that has been used extensively in the related literature is the *social welfare*. For an allocation $\mathbf{d} \in \mathcal{Q} \cup \mathbf{0}$ of a resource allocation game \mathcal{G}^M , the social welfare is defined as

$$\text{SW}(\mathbf{d}, \mathcal{G}^M) = \sum_{i=1}^n v_i(d_i),$$

where n is the number of players in \mathcal{G}^M and v_i is the valuation function of player i . Then, the inefficiency of equilibria of game \mathcal{G}^M can be measured by its *price of anarchy* which is defined as

$$\text{PoA}(\mathcal{G}^M) = \sup_{\mathbf{s} \in \text{eq}(\mathcal{G}^M)} \frac{\text{SW}^*(\mathcal{G}^M)}{\text{SW}(\mathbf{s}, \mathcal{G}^M)},$$

where $\text{SW}^*(\mathcal{G}^M)$ denotes the maximum social welfare over all allocations of \mathcal{G}^M .

However, the definition of the social welfare does not take into account the possibly finite budgets that the players may have. Therefore, we instead use the *liquid welfare* as our efficiency benchmark. The liquid welfare of an allocation \mathbf{d} is defined as

$$\text{LW}(\mathbf{d}, \mathcal{G}^M) = \sum_{i=1}^n \min\{v_i(d_i), c_i\},$$

where c_i is the budget of player i . Clearly, when players have no budget constraints, the liquid welfare coincides with the social welfare. The *liquid price of anarchy* of a resource allocation game \mathcal{G}^M is then defined as

$$\text{LPoA}(\mathcal{G}^M) = \sup_{\mathbf{s} \in \text{eq}(\mathcal{G}^M)} \frac{\text{LW}^*(\mathcal{G}^M)}{\text{LW}(\mathbf{s}, \mathcal{G}^M)},$$

where $\text{LW}^*(\mathcal{G}^M)$ denotes the maximum liquid welfare over all allocations of game \mathcal{G}^M . We use the overloaded term $\text{LPoA}(M)$ to denote the liquid price of anarchy of the resource allocation mechanism M . This is defined as the maximum (or, more formally, the supremum) liquid price of anarchy over all games that are induced by mechanism M .

2.3.1 Examples of resource allocation mechanisms

Let us devote some space to the definition of some well-known mechanisms from the literature. An important class of resource allocation mechanisms is that of *pay-your-signal* mechanisms (PYS, for short). When at least two players submit non-zero signals, a PYS mechanism charges each player i a payment equal to the signal s_i that she submits. Otherwise, PYS mechanisms follow the general convention that we have defined at the beginning of Section 2.3, and do not charge any payment to any player.

The most popular PYS mechanism is the Kelly mechanism that was introduced in Kelly [1997]. This mechanism allocates the resource proportionally to the players' signals (this is why it is also known as the proportional allocation mechanism in the related literature), i.e.,

$$g_i^{\text{Kelly}}(\mathbf{s}) = \frac{s_i}{\sum_{j=1}^n s_j}.$$

The Kelly mechanism has played a central role in the related literature; for the no-budget setting, Johari and Tsitsiklis [2004] proved that its price of anarchy is $4/3$. In their attempt to design the PYS mechanism with the lowest possible price of anarchy, Sanghavi and Hajek [2004] defined the allocation function

$$g_i^{\text{SH}}(\mathbf{s}) = \frac{s_i}{\max_{\ell} \{s_{\ell}\}} \int_0^1 \prod_{j \neq i} \left(1 - \frac{s_j}{\max_{\ell} \{s_{\ell}\}} t\right) dt.$$

We will refer to the PYS mechanism that uses this allocation function as SH. For two players, the allocation function has a very simple definition as $g_1^{\text{SH}}(\mathbf{s}) = \frac{s_1}{2s_2}$ when $s_1 \leq s_2$, and $g_1^{\text{SH}}(\mathbf{s}) = 1 - \frac{s_2}{2s_1}$ otherwise. Sanghavi and Hajek [2004] proved that the two-player version of the SH mechanism has an optimal (among all PYS mechanisms) price of anarchy of $8/7$ and provided experimental evidence that the price of anarchy of the n -player version is only marginally higher. As we will see later in Section 2.6, the comparison between Kelly and SH yields a drastically different result when players have budgets and the liquid welfare is used as the efficiency benchmark.

Other interesting classes of mechanisms use proportional allocation, but different kinds of payments. Among them, a mechanism defined by Maheswaran and Basar [2006] uses the class of payment functions

$$p_i^M(\mathbf{s}) = \left(\sum_{j \neq i} s_j \right) \cdot \int_0^{s_i} \frac{h^M(t + \sum_{j \neq i} s_j)}{(t + \sum_{j \neq i} s_j)^2} dt,$$

where $h^M : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is an increasing function (such as $h^M(z) = z$; Maheswaran and Basar [2006] suggest several other choices for h^M). These mechanisms have the remarkable property of full efficiency at equilibria in the no-budget setting (i.e., they have price of anarchy equal to 1). Independently from Maheswaran and Basar [2006], Johari and Tsitsiklis [2009] as well as Yang and Hajek [2007] presented resource allocation mechanisms that achieve full efficiency in the no-budget setting. All these mechanisms can be thought of as adaptations of the well-known VCG paradigm.

2.4 A lower bound for all mechanisms

The fact that the mechanisms of Maheswaran and Basar [2006], Johari and Tsitsiklis [2009], and Yang and Hajek [2007] achieve full efficiency seems quite surprising, since resource allocation mechanisms do not have direct access to the valuation functions of the players. The definition of these mechanisms is such that the incentives of the players are fully aligned to the global goal of maximizing the social welfare. In a sense, these mechanisms manage to achieve access to the valuation functions indirectly. In contrast, when players have budget constraints, we show below that a liquid price of anarchy equal to 1 is not possible. This means that resource allocation mechanisms fail to “mine” any kind of information about the budget values of the players, while budgets affect the strategic behavior of the players crucially.

Theorem 2.1. *Every n -player resource allocation mechanism has liquid price of anarchy at least $2 - 1/n$.*

Proof. Let M be any n -player resource allocation mechanism that uses an allocation function g^M and a payment function p^M . Let $\mathbf{s} = (s_1, \dots, s_n)$ be an equilibrium of the game \mathcal{G}_1^M induced by M for players with valuations $v_i(x) = x$ and budgets $c_i = +\infty$, for every $i \in [n]$. Assume that the allocation returned by M at this equilibrium is $\mathbf{d} = (d_1, \dots, d_n)$. Since all players have the same valuation function and budget, the liquid (or social) welfare at equilibrium is optimal and, hence, $\text{LPoA}(\mathcal{G}_1^M) = 1$.

Recall that, for every signal vector $\mathbf{y} = (y_1, \dots, y_n)$, the utility of player i is defined as $u_i^M(\mathbf{y}) = v_i(g_i^M(\mathbf{y})) - p_i^M(\mathbf{y})$. Now, let $i^* = \arg \min_i d_i$ (hence, $d_{i^*} \leq 1/n$) and consider the game \mathcal{G}_2^M where each player $i \neq i^*$ has the modified valuation function $\tilde{v}_i(x) = d_i + x$ and budget $\tilde{c}_i = d_i$, while player i^* is as in \mathcal{G}_1^M (see Figure 2.1). Observe that the modified utility of player $i \neq i^*$ as a function of a signal vector \mathbf{y} is now $\tilde{u}_i^M(\mathbf{y}) = \tilde{v}_i(g_i^M(\mathbf{y})) - p_i^M(\mathbf{y}) = u_i^M(\mathbf{y}) + d_i$. Also, since the utility of player $i \neq i^*$ is non-negative at the equilibrium \mathbf{s} of game \mathcal{G}_1^M , we have that $p_i^M(\mathbf{s}) \leq d_i = \tilde{c}_i$, meaning that player i can also afford this payment in game \mathcal{G}_2^M . Hence, \mathbf{s}




Figure 2.1: A graphical representation of the games used in the proof of Theorem 2.1. The two figures depict the valuation functions of players i^* and $i \neq i^*$ in games \mathcal{G}_1^M and \mathcal{G}_2^M . The blue points (i.e., point (d_{i^*}, d_{i^*}) in the left figure, and points (d_i, d_i) and $(d_i, 2d_i)$ in the right figure) represent the equilibrium in both games, and the optimal allocation in game \mathcal{G}_1^M . The optimal allocation in \mathcal{G}_2^M is represented by the red points (i.e., point $(1, 1)$ in the left figure and point $(0, d_i)$ in the right one).

is an equilibrium in \mathcal{G}_2^M as well (and, again, M returns the same allocation \mathbf{d}).

Its liquid welfare is $\sum_i \min\{\tilde{v}_i(d_i), \tilde{c}_i\} = \sum_i d_i = 1$ while the optimal liquid welfare is at least $1 + \sum_{i \neq i^*} d_i$, achieved at the allocation according to which the whole resource is given to player i^* . Hence, we conclude that the liquid price of anarchy of M is $\text{LPoA}(M) \geq \text{LPoA}(\mathcal{G}_2^M) \geq 1 + \sum_{i \neq i^*} d_i = 2 - d_{i^*} \geq 2 - 1/n$, as desired. \square

2.5 The structure of worst-case games and equilibria

In this section, we prove our structural characterization. Given an n -player resource allocation mechanism M (with allocation and payment functions g^M and p^M , respectively), signal vector $\mathbf{s} \in \mathbb{X}_n$, and an integer $j \in [n]$, define the n -player game $\mathcal{G}^M(\mathbf{s}, j)$ as follows. Every player has the affine valuation function $\tilde{v}_i(z) = \lambda_i^M(\mathbf{s}) \cdot z + \kappa_i^M(\mathbf{s})$ and budget \tilde{c}_i , where

$$\lambda_i^M(\mathbf{s}) = \left(\frac{\partial g_i^M(y, \mathbf{s}_{-i})}{\partial y} \Big|_{y=s_i} \right)^{-1} \cdot \frac{\partial p_i^M(y, \mathbf{s}_{-i})}{\partial y} \Big|_{y=s_i}$$

and $\kappa_j^M(\mathbf{s}) = 0$, $\tilde{c}_j = +\infty$, and $\kappa_i^M(\mathbf{s}) = \tilde{c}_i = p_i^M(\mathbf{s})$ for every player $i \neq j$.

In the following, we show that the games defined in this way are in a sense extreme in terms of the liquid price of anarchy of mechanism M .

Lemma 2.2. *Let \mathcal{G}_1^M be an n -player resource allocation game that is induced by a mechanism M with $\text{LPoA}(\mathcal{G}_1^M) > 1$. Let $\mathbf{s} \in \mathbb{X}_n$ be an equilibrium of \mathcal{G}_1^M of minimum liquid welfare. Then, there exists*

integer $i^* \in [n]$ such that

$$\begin{aligned} \text{LPoA}(\mathcal{G}_1^M) &\leq \frac{\text{LW}(\tilde{\mathbf{x}}, \mathcal{G}^M(\mathbf{s}, i^*))}{\text{LW}(g^M(\mathbf{s}), \mathcal{G}^M(\mathbf{s}, i^*))} \\ &= \frac{\sum_{i \neq i^*} p_i^M(\mathbf{s}) + \lambda_{i^*}^M(\mathbf{s})}{\sum_{i \neq i^*} p_i^M(\mathbf{s}) + \lambda_{i^*}^M(\mathbf{s}) \cdot g_{i^*}^M(\mathbf{s})}, \end{aligned}$$

where $\tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_n)$ denotes the allocation with $\tilde{x}_{i^*} = 1$ and $\tilde{x}_i = 0$ for $i \neq i^*$.

Proof. Consider an n -player resource allocation game \mathcal{G}_1^M that is induced by mechanism M . Let v_i and c_i be the valuation function and budget of player i , respectively. Let $\mathbf{s} \in \mathbb{X}_n$ be the equilibrium of game \mathcal{G}_1^M of minimum liquid welfare. We denote by \mathbf{x} the optimal allocation in \mathcal{G}_1^M . Without loss of generality, we assume that, for every player i , $x_i = 0$ if $v_i(0) > c_i$ and $v_i(x_i) \leq c_i$ otherwise, and we relax the allocation definition to $\sum_{i=1}^n x_i \leq 1$; this does not constrain the optimal liquid welfare which is $\text{LW}(\mathbf{x}, \mathcal{G}_1^M) = \sum_i \min\{v_i(x_i), c_i\}$. We use $d_i = g_i^M(\mathbf{s})$ for the resource fraction allocated to player i in \mathbf{s} ; let $\mathbf{d} = (d_1, \dots, d_n)$.

We partition the players into the following three sets:

- Set A consists of players i with $v_i(d_i) < c_i$ and signal s_i such that the derivative of their utility is equal to 0.
- Set B consists of players i with signal $s_i = 0$ (hence, $d_i = 0$) and negative utility derivative such that $v_i(0) < c_i$.
- Set Γ consists of players i with signal s_i such that $v_i(d_i) \geq c_i$.

First, observe that sets A and B cannot be both empty, since it would then be $\text{LW}(\mathbf{d}, \mathcal{G}_1^M) = \sum_{i \in [n]} c_i \geq \text{LW}(\mathbf{x}, \mathcal{G}_1^M)$, and the liquid price of anarchy of \mathcal{G}_1^M would be exactly 1, contradicting the assumption of the lemma. So, in the following, we assume that at least one of A and B is non-empty.

Now consider the games $\mathcal{G}^M(\mathbf{s}, j)$ for $j \in [n]$ and let $i^* = \arg \max_{j \in A \cup B} \{\lambda_j^M(\mathbf{s})\}$. We will show that

$$\text{LW}(\mathbf{d}, \mathcal{G}_1^M) \geq \text{LW}(\mathbf{d}, \mathcal{G}^M(\mathbf{s}, i^*)) \quad (2.1)$$

and we will furthermore show that the allocation $\tilde{\mathbf{x}}$ satisfies

$$\text{LW}(\mathbf{x}, \mathcal{G}_1^M) - \text{LW}(\tilde{\mathbf{x}}, \mathcal{G}^M(\mathbf{s}, i^*)) \leq \text{LW}(\mathbf{d}, \mathcal{G}_1^M) - \text{LW}(\mathbf{d}, \mathcal{G}^M(\mathbf{s}, i^*)). \quad (2.2)$$

In this way (recall that \mathbf{s} is the equilibrium of minimum liquid welfare in game \mathcal{G}_1^M and \mathbf{d} is the resulting allocation), we will have

$$\begin{aligned}
\text{LPoA}(\mathcal{G}_1^M) &= \frac{\text{LW}(\mathbf{x}, \mathcal{G}_1^M)}{\text{LW}(\mathbf{d}, \mathcal{G}_1^M)} \\
&\leq \frac{\text{LW}(\mathbf{x}, \mathcal{G}_1^M) - (\text{LW}(\mathbf{x}, \mathcal{G}_1^M) - \text{LW}(\tilde{\mathbf{x}}, \mathcal{G}^M(\mathbf{s}, i^*)))}{\text{LW}(\mathbf{d}, \mathcal{G}_1^M) - (\text{LW}(\mathbf{d}, \mathcal{G}_1^M) - \text{LW}(\mathbf{d}, \mathcal{G}^M(\mathbf{s}, i^*)))} \\
&= \frac{\text{LW}(\tilde{\mathbf{x}}, \mathcal{G}^M(\mathbf{s}, i^*))}{\text{LW}(\mathbf{d}, \mathcal{G}^M(\mathbf{s}, i^*))} \\
&= \frac{\sum_{i \neq i^*} p_i^M(\mathbf{s}) + \lambda_{i^*}^M(\mathbf{s})}{\sum_{i \neq i^*} p_i^M(\mathbf{s}) + \lambda_{i^*}^M(\mathbf{s}) \cdot g_{i^*}^M(\mathbf{s})},
\end{aligned}$$

as desired. The inequality follows by (2.1) and (2.2). The last equality follows since all players in $\mathcal{G}^M(\mathbf{s}, i^*)$ besides i^* have always their value capped by their budget, which is equal to their payment.

Inequality (2.1) is due to the fact that the contribution of each player to the liquid welfare at \mathbf{s} can only decrease between the two games. Indeed, if player i^* belongs to B , she has zero value in game $\mathcal{G}^M(\mathbf{s}, i^*)$. If she belongs to A , then her utility derivative is nullified and, hence, $v'_{i^*}(d_{i^*}) = \lambda_{i^*}^M(\mathbf{s})$. Due to the concavity of v_i , we get $v_{i^*}(d_{i^*}) \geq d_{i^*} v'_{i^*}(d_{i^*}) = d_{i^*} \lambda_{i^*}^M(\mathbf{s}) = \tilde{v}_{i^*}(d_{i^*})$. Moreover, the contribution of player $i \neq i^*$ in $\text{LW}(\mathbf{d}, \mathcal{G}^M(\mathbf{s}, i^*))$ is $\tilde{c}_i = p_i^M(\mathbf{s})$ which is at most her contribution $\min\{v_i(d_i), c_i\}$ in $\text{LW}(\mathbf{d}, \mathcal{G}_1^M)$ since the payment of player i cannot exceed her budget in \mathcal{G}_1^M and her utility at equilibrium \mathbf{s} is non-negative. See Figure 2.2 for a graphical representation of valuation functions and budgets in games \mathcal{G}_1^M and $\mathcal{G}^M(\mathbf{s}, i^*)$.

Let

$$\delta(i) = \min\{v_i(x_i), c_i\} - \min\{\tilde{v}_i(\tilde{x}_i), \tilde{c}_i\} - \min\{v_i(d_i), c_i\} + \min\{\tilde{v}_i(d_i), \tilde{c}_i\}$$

denote the contribution of player i to the expression

$$\text{LW}(\mathbf{x}, \mathcal{G}_1^M) - \text{LW}(\tilde{\mathbf{x}}, \mathcal{G}^M(\mathbf{s}, i^*)) - \text{LW}(\mathbf{d}, \mathcal{G}_1^M) + \text{LW}(\mathbf{d}, \mathcal{G}^M(\mathbf{s}, i^*)).$$

Then, in order to prove inequality (2.2) it suffices to prove that $\sum_i \delta(i) \leq 0$.

- For player i^* , we have that $v_{i^*}(d_{i^*}) < c_{i^*}$. Using the inequality $v_{i^*}(x_{i^*}) \leq v_{i^*}(d_{i^*}) + v'_{i^*}(d_{i^*})(x_{i^*} - d_{i^*})$ due to the concavity of the valuation function v_{i^*} and the fact $\tilde{x}_{i^*} = 1$, we have that

$$\begin{aligned}
\delta(i^*) &= \min\{v_{i^*}(x_{i^*}), c_{i^*}\} - \lambda_{i^*}^M(\mathbf{s}) \tilde{x}_{i^*} - v_{i^*}(d_{i^*}) + \lambda_{i^*}^M(\mathbf{s}) d_{i^*} \\
&\leq v_{i^*}(x_{i^*}) - \lambda_{i^*}^M(\mathbf{s}) - v_{i^*}(d_{i^*}) + \lambda_{i^*}^M(\mathbf{s}) d_{i^*}
\end{aligned}$$

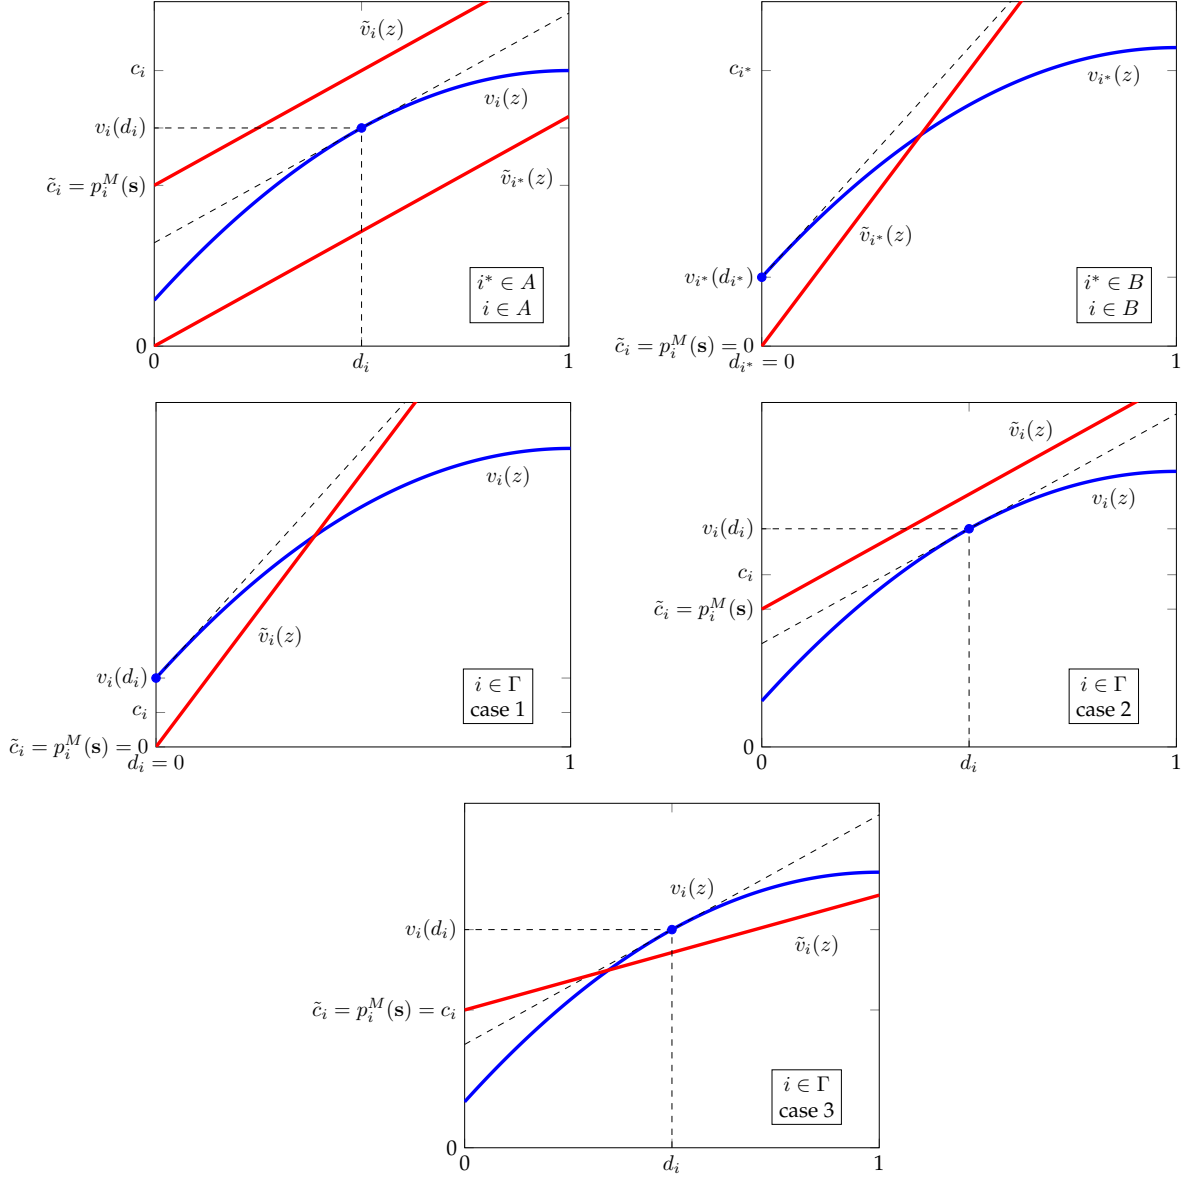


Figure 2.2: Relation between the two games \mathcal{G}_1^M and $\mathcal{G}^M(\mathbf{s}, i^*)$ that are used in the proof of Lemma 2.2. In the first two plots, player i is different than i^* and the budget \tilde{c}_{i^*} is infinite by definition. The dashed line is the tangent of v_i at d_i . The slope $\lambda_i^M(\mathbf{s})$ of the affine valuation function of player i in $\mathcal{G}^M(\mathbf{s}, i^*)$ is greater than (upper right and middle left plots), equal to (upper left and middle right plots), or smaller than (lower plot) $v'_i(d_i)$ depending on whether the utility derivative of the player is negative, zero, or positive, respectively (in particular, these are the three cases identified in the plots for $i \in \Gamma$). This follows by the definition of games \mathcal{G}_1^M and $\mathcal{G}^M(\mathbf{s}, i^*)$ and the fact that, as the utility of player i in game \mathcal{G}_1^M has derivative $v'_i(d_i) \frac{\partial g_i^M(y, \mathbf{s}_{-i})}{\partial y} \Big|_{y=s_i} - \frac{\partial p_i^M(y, \mathbf{s}_{-i})}{\partial y} \Big|_{y=s_i}$ at equilibrium, the sign of this derivative coincides with the sign of $v'_i(d_i) - \lambda_i^M(\mathbf{s})$.

$$\leq v'_{i^*}(d_{i^*})(x_{i^*} - d_{i^*}) - \lambda_{i^*}^M(\mathbf{s}) + \lambda_{i^*}^M(\mathbf{s})d_{i^*}.$$

Now, we observe that (for such observations, we follow the reasoning in the caption of Figure 2.2) if player i^* belongs to A , then $\lambda_{i^*}^M(\mathbf{s}) = v'_{i^*}(d_{i^*})$, while if she belongs to B , then $\lambda_{i^*}^M(\mathbf{s}) \geq v'_{i^*}(d_{i^*})$ and $d_{i^*} = 0$. In any case, we have that $v'_{i^*}(d_{i^*})(x_{i^*} - d_{i^*}) \leq \lambda_{i^*}^M(\mathbf{s})(x_{i^*} - d_{i^*})$, and we obtain

$$\delta(i^*) \leq \lambda_{i^*}^M(\mathbf{s})(x_{i^*} - 1). \quad (2.3)$$

- For all players $i \neq i^*$, observe that their value is always capped by their budget in $\mathcal{G}^M(\mathbf{s}, i^*)$. For player $i \neq i^*$ belonging to A or to B we have that either $\lambda_i^M(\mathbf{s}) = v'_i(d_i)$ (if $i \in A$), or $\lambda_i^M(\mathbf{s}) \geq v'_i(d_i)$ and $d_i = 0$ (if $i \in B$). Hence, using the concavity of v_i and the fact that $\tilde{x}_i = 0$, we obtain that

$$\begin{aligned} \delta(i) &\leq v_i(x_i) - \tilde{c}_i - v_i(d_i) + \tilde{c}_i \\ &\leq v_i(d_i) + \lambda_i^M(\mathbf{s})(x_i - d_i) - v_i(d_i) \\ &\leq \lambda_i^M(\mathbf{s})x_i, \end{aligned} \quad (2.4)$$

where the last inequality follows since $\lambda_i^M(\mathbf{s}) \leq \lambda_{i^*}^M(\mathbf{s})$, due to the definition of player i^* . Otherwise, if $i \in \Gamma$, we have

$$\delta(i) = \min\{v_i(x_i), c_i\} - \tilde{c}_i - c_i + \tilde{c}_i \leq 0. \quad (2.5)$$

Hence, summing over all players, and using inequalities (2.3), (2.4) and (2.5) as well as the fact that $\sum_i x_i \leq 1$, we obtain $\sum_i \delta(i) \leq 0$, and the proof is complete. \square

We are now ready to prove the main result of this section.

Lemma 2.3. *Let M be an n -player resource allocation mechanism with allocation and payment functions g^M and p^M , respectively. Then, its liquid price of anarchy is*

$$\text{LPoA}(M) \leq \sup_{\mathbf{s} \in \mathbb{X}_n} \left\{ \frac{\sum_{i \geq 2} p_i^M(\mathbf{s}) + \lambda_1^M(\mathbf{s})}{\sum_{i \geq 2} p_i^M(\mathbf{s}) + \lambda_1^M(\mathbf{s}) g_1^M(\mathbf{s})} \right\}, \quad (2.6)$$

where

$$\lambda_1^M(\mathbf{s}) = \left(\frac{\partial g_1^M(y, \mathbf{s}_{-1})}{\partial y} \Big|_{y=s_1} \right)^{-1} \cdot \frac{\partial p_1^M(y, \mathbf{s}_{-1})}{\partial y} \Big|_{y=s_1}.$$

If, in addition, $\mathbf{s} \in \mathbb{X}_n$ is always an equilibrium of game $\mathcal{G}^M(\mathbf{s}, 1)$, (2.6) holds with equality.

Proof. Let $\text{weq}(\mathcal{G}^M)$ be the set of equilibria of minimum liquid welfare in game \mathcal{G}^M . Using the definition of the liquid price of anarchy, Lemma 2.2, and the anonymity of resource allocation mechanisms, we have

$$\begin{aligned}
\text{LPoA}(M) &= \sup_{\mathcal{G}^M} \text{LPoA}(\mathcal{G}^M) \\
&= \sup_{\mathcal{G}^M} \sup_{\mathbf{s} \in \text{weq}(\mathcal{G}^M)} \frac{\text{LW}^*(\mathcal{G}^M)}{\text{LW}(g^M(\mathbf{s}), \mathcal{G}^M)} \\
&= \sup_{\mathbf{s} \in \mathbb{X}_n} \sup_{\mathcal{G}^M: \mathbf{s} \in \text{weq}(\mathcal{G}^M)} \frac{\text{LW}^*(\mathcal{G}^M)}{\text{LW}(g^M(\mathbf{s}), \mathcal{G}^M)} \\
&\leq \sup_{\mathbf{s} \in \mathbb{X}_n} \max_{i^* \in [n]} \frac{\sum_{i \neq i^*} p_i^M(\mathbf{s}) + \lambda_{i^*}^M(\mathbf{s})}{\sum_{i \neq i^*} p_i^M(\mathbf{s}) + \lambda_{i^*}^M(\mathbf{s}) g_{i^*}^M(\mathbf{s})} \\
&= \sup_{\mathbf{s} \in \mathbb{X}_n} \frac{\sum_{i \geq 2} p_i^M(\mathbf{s}) + \lambda_1^M(\mathbf{s})}{\sum_{i \geq 2} p_i^M(\mathbf{s}) + \lambda_1^M(\mathbf{s}) g_1^M(\mathbf{s})}.
\end{aligned}$$

Now, if $\mathbf{s} \in \text{eq}(\mathcal{G}^M(\mathbf{s}, 1))$ for every $\mathbf{s} \in \mathbb{X}_n$, by just considering the games $\mathcal{G}^M(\mathbf{s}, 1)$ induced by mechanism M , we have

$$\begin{aligned}
\text{LPoA}(M) &\geq \sup_{\mathbf{s} \in \mathbb{X}_n} \text{LPoA}(\mathcal{G}^M(\mathbf{s}, 1)) \\
&\geq \sup_{\mathbf{s} \in \mathbb{X}_n} \frac{\sum_{i \geq 2} p_i^M(\mathbf{s}) + \lambda_1^M(\mathbf{s})}{\sum_{i \geq 2} p_i^M(\mathbf{s}) + \lambda_1^M(\mathbf{s}) g_1^M(\mathbf{s})}
\end{aligned}$$

and (2.6) holds with equality. The last inequality follows by comparing the liquid welfare at \mathbf{s} to the liquid welfare of the allocation which gives the whole resource to player 1. Recall that all players besides player 1 have always their value capped by their budget in game $\mathcal{G}^M(\mathbf{s}, 1)$. \square

Lemma 2.3 is extremely powerful. It essentially says that no game-theoretic reasoning is needed anymore for proving upper bounds on the LPoA and, instead, all we have to do is to solve the corresponding mathematical program. Furthermore, it can be used to prove lower bounds on the LPoA without providing any explicit construction. In this case, we just need to show that the condition $\mathbf{s} \in \text{eq}(\mathcal{G}^M(\mathbf{s}, 1))$ holds; then the tight lower bound follows by solving the same mathematical program.

Before we continue with the rest of our results, we define the class \mathcal{C} of mechanisms M that use concave allocation functions g^M and convex payment functions p^M . Observe that both Kelly and SH (as well as the E2-PYS mechanism presented in Section 2.7) are members of this class. With our next lemma, we prove that the condition $\mathbf{s} \in \text{eq}(\mathcal{G}^M(\mathbf{s}, 1))$ is satisfied for any \mathcal{C} mechanism M . This will allow us to prove lower bounds in the upcoming sections.

Lemma 2.4. For any n -player resource allocation mechanism $M \in \mathcal{C}$ and $\mathbf{s} \in \mathbb{X}_n$, $\mathbf{s} \in \text{eq}(\mathcal{G}^M(\mathbf{s}, 1))$.

Proof. Consider any \mathcal{C} mechanism M that uses a concave allocation function g^M and a convex payment function p^M . By the definition of game $\mathcal{G}^M(\mathbf{s}, 1)$, the utility of any player i , as a function of her signal y , is $u_i^M(y, \mathbf{s}_{-i}) = \lambda_i^M(\mathbf{s}) \cdot g_i^M(y, \mathbf{s}_{-i}) + \kappa_i^M(\mathbf{s}) - p_i^M(y, \mathbf{s}_{-i})$ and its derivative is

$$\frac{\partial u_i^M(y, \mathbf{s}_{-i})}{\partial y} = \lambda_i^M(\mathbf{s}) \frac{\partial g_i^M(y, \mathbf{s}_{-i})}{\partial y} - \frac{\partial p_i^M(y, \mathbf{s}_{-i})}{\partial y}.$$

Observe that, by the definition of $\lambda_i^M(\mathbf{s})$, the signal s_i nullifies the utility derivative of player i , and since

$$\frac{\partial^2 u_i^M(y, \mathbf{s}_{-i})}{\partial y^2} = \lambda_i^M(\mathbf{s}) \frac{\partial^2 g_i^M(y, \mathbf{s}_{-i})}{\partial y^2} - \frac{\partial^2 p_i^M(y, \mathbf{s}_{-i})}{\partial y^2} \leq 0,$$

this signal actually maximizes the player's utility. \square

2.6 Pay-your-signal mechanisms

In this section, we will exploit Lemma 2.3 to prove tight bounds on the liquid price of anarchy of the Kelly and SH mechanisms. Our LPoA bounds are 2 for Kelly (Theorem 2.5) and 3 for SH (Theorems 2.6 and 2.7). Recall that both of these mechanisms belong to class \mathcal{C} and, by Lemma 2.4, the condition $\mathbf{s} \in \text{eq}(\mathcal{G}^M(\mathbf{s}, 1))$ is satisfied.

Theorem 2.5. The liquid price of anarchy of the Kelly mechanism is 2.

Proof. Consider any signal vector $\mathbf{s} \in \mathbb{X}_n$, and let $C = \sum_{i \geq 2} s_i$. Since Kelly is a PYS mechanism, we have that $\sum_{i \geq 2} p_i^{\text{Kelly}}(\mathbf{s}) = C$ and

$$\frac{\partial p_1^{\text{Kelly}}(y, \mathbf{s}_{-1})}{\partial y} = 1.$$

By the definition of the allocation function $g_1^{\text{Kelly}}(y, \mathbf{s}_{-1}) = \frac{y}{y+C}$, we have that

$$\frac{\partial g_1^{\text{Kelly}}(y, \mathbf{s}_{-1})}{\partial y} = \frac{C}{(y+C)^2}.$$

Also, since the mechanism belongs to class \mathcal{C} , by Lemma 2.4, we have that $\mathbf{s} \in \text{eq}(\mathcal{G}^{\text{Kelly}}(\mathbf{s}, 1))$.

Hence,

$$\lambda_1^{\text{Kelly}}(\mathbf{s}) = \frac{(s_1 + C)^2}{C}$$

and Lemma 2.3 yields

$$\text{LPoA}(\text{Kelly}) = \sup_{s_1, C \geq 0} \frac{C + (s_1 + C)^2/C}{C + (s_1 + C)s_1/C}$$

$$\begin{aligned}
&= \sup_{s_1, C \geq 0} \frac{2C^2 + 2s_1C + s_1^2}{C^2 + s_1C + s_1^2} \\
&= \sup_{s_1, C \geq 0} \left(2 - \frac{s_1^2}{C^2 + s_1C + s_1^2} \right) \\
&= 2,
\end{aligned}$$

as desired. \square

Notice that our proof of Theorem 2.5 is surprisingly short. The proof exploits Lemma 2.3 with (2.6) holding with equality and, as such, it simultaneously provides a tight (upper and lower) bound. In contrast, our analysis for the SH mechanism is slightly more involved. This is mainly due to the more complicated definition of the allocation function (see Section 2.3), which requires us to distinguish between two cases, depending on whether $s_1 < \max_{\ell} s_{\ell}$ or not. Both cases lead to inequalities that provide only an upper bound on the LPoA of SH in the proof of Theorem 2.6. In Theorem 2.7, we easily prove a matching lower bound by restricting our attention to the 2-player version of the mechanism. Actually, the proof can be thought of as providing a tight (i.e., not only lower, but also upper) bound on the LPoA of the 2-player version of the SH mechanism.

Theorem 2.6. *The liquid price of anarchy of the SH mechanism is at most 3.*

Proof. We will use Lemma 2.3 and upper-bound the ratio in the RHS of (2.6) by 3. Define $C = \sum_{i \geq 2} s_i$. First, let $\mathbf{s} \in \mathbb{X}_n$ with $s_1 < \max_{\ell} s_{\ell}$. Let $\arg \max_{\ell} s_{\ell} = i^* \neq 1$. Then, by the definition of SH and the definition of $\lambda_1^{\text{SH}}(\mathbf{s})$ in (2.6), we have

$$\lambda_1^{\text{SH}}(\mathbf{s}) = \frac{s_{i^*}}{\int_0^1 \prod_{i \geq 2} \left(1 - \frac{s_i}{s_{i^*}} t \right) dt} \quad (2.7)$$

and using the Bernoulli inequality stating that $1 - \gamma t \geq (1 - t)^{\gamma}$ for $t \leq 1$ and $\gamma \in [0, 1]$, (2.7) yields

$$\lambda_1^{\text{SH}}(\mathbf{s}) \leq \frac{s_{i^*}}{\int_0^1 \prod_{i \geq 2} (1 - t)^{\frac{s_i}{s_{i^*}}} dt} = \frac{s_{i^*}}{\int_0^1 (1 - t)^{\frac{C}{s_{i^*}}} dt} = s_{i^*} + C.$$

Since SH is PYS, $\sum_{i \geq 2} p_i^{\text{SH}}(\mathbf{s}) = C$. Using this observation together with the last inequality, we obtain

$$\frac{\sum_{i \geq 2} p_i^{\text{SH}}(\mathbf{s}) + \lambda_1^{\text{SH}}(\mathbf{s})}{\sum_{i \geq 2} p_i^{\text{SH}}(\mathbf{s}) + \lambda_1^{\text{SH}}(\mathbf{s}) g_1^{\text{SH}}(\mathbf{s})} \leq \frac{2C + s_{i^*}}{C} \leq 3. \quad (2.8)$$

The inequalities follow since $\lambda_1^{\text{SH}}(\mathbf{s}) g_1^{\text{SH}}(\mathbf{s}) \geq 0$, $s_1 \geq 0$, and $s_{i^*} \leq C$.

Now, let $\mathbf{s} \in \mathbb{X}_n$ with $s_1 = \max_{\ell} s_{\ell}$. In this case, $g_1^{\text{SH}}(\mathbf{s})$ is defined as

$$g_1^{\text{SH}}(\mathbf{s}) = \int_0^1 \prod_{i \geq 2} \left(1 - \frac{s_i}{s_1} t\right) dt$$

and

$$\begin{aligned} \left. \frac{\partial g_1^{\text{SH}}(y, \mathbf{s}_{-1})}{\partial y} \right|_{y=s_1} &= \int_0^1 \sum_{i \geq 2} \frac{s_i}{s_1^2} t \prod_{j \neq 1, i} \left(1 - \frac{s_j}{s_1} t\right) dt \\ &\geq \sum_{i \geq 2} \frac{s_i}{s_1^2} \int_0^1 t \prod_{j \neq 1, i} (1 - t)^{\frac{s_j}{s_1}} dt \\ &= \sum_{i \geq 2} \frac{s_i}{s_1^2} \int_0^1 t(1 - t)^{\frac{C-s_i}{s_1}} dt \\ &= \sum_{i \geq 2} \frac{s_i}{(C - s_i + s_1)(C - s_i + 2s_1)} \\ &\geq \frac{C}{(C + s_1)(C + 2s_1)}. \end{aligned}$$

Using the definition of $\lambda_1^{\text{SH}}(\mathbf{s})$ in (2.6), this last inequality implies that

$$\lambda_1^{\text{SH}}(\mathbf{s}) \leq \frac{(C + s_1)(C + 2s_1)}{C}. \quad (2.9)$$

Also, by applying the Bernoulli inequality to the RHS of the definition of $g_1^{\text{SH}}(\mathbf{s})$, we obtain

$$g_1^{\text{SH}}(\mathbf{s}) \geq \int_0^1 \prod_{i \geq 2} (1 - t)^{\frac{s_i}{s_1}} dt = \int_0^1 (1 - t)^{\frac{C}{s_1}} dt = \frac{s_1}{C + s_1}. \quad (2.10)$$

Now, we have

$$\begin{aligned} \frac{\sum_{i \geq 2} p_i^{\text{SH}}(\mathbf{s}) + \lambda_1^{\text{SH}}(\mathbf{s})}{\sum_{i \geq 2} p_i^{\text{SH}}(\mathbf{s}) + \lambda_1^{\text{SH}}(\mathbf{s}) g_1^{\text{SH}}(\mathbf{s})} &\leq \frac{C^2 + (C + s_1)(C + 2s_1)}{C^2 + (C + s_1)(C + 2s_1) g_1^{\text{SH}}(\mathbf{s})} \\ &\leq \frac{2C^2 + 3s_1 C + 2s_1^2}{C^2 + s_1 C + 2s_1^2} \leq 3. \end{aligned} \quad (2.11)$$

The two first inequalities follow by (2.9) and (2.10), respectively, and the last one is obvious since $s_1, C \geq 0$.

Now, the upper bound follows by Lemma 2.3 using (2.8) and (2.11). \square

Theorem 2.7. *The liquid price of anarchy of the SH mechanism is at least 3.*

Proof. It suffices to restrict our attention to the 2-player version of the mechanism. Let $\mathbf{s} \in \mathbb{X}_2$ with $s_1 \leq s_2$. In this case $g_1^{\text{SH}}(\mathbf{s}) = \frac{s_1}{2s_2}$ which implies that $\lambda_1^{\text{SH}}(\mathbf{s}) = 2s_2$. Since the SH mechanism belongs to class \mathcal{C} , by Lemma 2.4, we have that $\mathbf{s} \in \text{eq}(\mathcal{G}^{\text{SH}}(\mathbf{s}, 1))$. Using Lemma 2.3, we obtain

$$\text{LPoA}(\text{SH}) \geq \sup_{\mathbf{s} \in \mathbb{X}_2: s_1 \leq s_2} \frac{3s_2}{s_2 + s_1} = 3.$$

The proof is complete. \square

2.7 Two-player mechanisms

As we saw in Theorem 2.5, the Kelly mechanism has an LPoA of exactly 2 even in the case of two players. In contrast, our lower bound of $3/2$ for 2-player mechanisms in Theorem 2.1 seems to leave room for improvements. Such improvements are indeed possible as we show with the mechanisms that we present in this section. Interestingly, the E2-PYS mechanism that is defined in the following is also proved to have optimal LPoA among all 2-player PYS mechanisms with concave allocation functions.

2.7.1 The E2-PYS mechanism

Let $\beta \approx 1.792$ be the solution of the equation $\frac{1}{\beta} - \frac{1}{\beta} \exp\left(-\frac{\beta}{\beta-1}\right) = \frac{1}{2}$ and define mechanism E2-PYS to be the PYS 2-player mechanism that uses the allocation function

$$g_i^{\text{E2-PYS}}(\mathbf{s}) = \begin{cases} \frac{1}{\beta} - \frac{1}{\beta} \exp\left(-\frac{\beta}{\beta-1} \cdot \frac{s_i}{s_{3-i}}\right) & s_i \leq s_{3-i} \\ \frac{\beta-1}{\beta} + \frac{1}{\beta} \exp\left(-\frac{\beta}{\beta-1} \cdot \frac{s_{3-i}}{s_i}\right) & s_i > s_{3-i} \end{cases}$$

for player $i \in \{1, 2\}$ and (non-zero) signal vector $\mathbf{s} = (s_1, s_2)$. Due to the definition of β , E2-PYS is a well-defined resource allocation mechanism: it is anonymous, with an increasing and differentiable allocation function, which allocates the whole resource when some player has non-zero signal. Moreover, E2-PYS belongs to class \mathcal{C} : the allocation function can be seen to be concave (see also Figure 2.3) and the payment function is, of course, convex. The LPoA bound statement for E2-PYS follows.

Theorem 2.8. *The liquid price of anarchy of the E2-PYS mechanism is $\beta \approx 1.792$.*

Proof. We will prove the theorem using Lemma 2.3. Let $\mathbf{s} \in \mathbb{X}_2$. Due to Lemma 2.4, we have that $\mathbf{s} \in \text{eq}(\mathcal{G}^{\text{E2-PYS}}(\mathbf{s}, 1))$. Since E2-PYS is a PYS mechanism, we have that $p_1^{\text{E2-PYS}}(\mathbf{s}) = s_1$, which yields

$$\left. \frac{\partial p_1^{\text{E2-PYS}}(y, s_2)}{\partial y} \right|_{y=s_1} = 1.$$

Next, we distinguish between two cases. First, assume that $s_1 \leq s_2$; in this case, the allocation of player 1 is

$$g_1^{\text{E2-PYS}}(s_1, s_2) = \frac{1}{\beta} - \frac{1}{\beta} \exp\left(-\frac{\beta}{\beta-1} \cdot \frac{s_1}{s_2}\right)$$

and, thus, the derivative is equal to

$$\left. \frac{\partial g_1^{\text{E2-PYS}}(y, s_2)}{\partial y} \right|_{y=s_1} = \frac{1}{(\beta-1)s_2} \exp\left(-\frac{\beta}{\beta-1} \cdot \frac{s_1}{s_2}\right).$$

Therefore, $\lambda_1^{\text{E2-PYS}}(\mathbf{s})$ is defined as

$$\lambda_1^{\text{E2-PYS}}(\mathbf{s}) = (\beta - 1)s_2 \exp\left(\frac{\beta}{\beta - 1} \cdot \frac{s_1}{s_2}\right).$$

By substituting $p_2^{\text{E2-PYS}}(\mathbf{s})$, $\lambda_1^{\text{E2-PYS}}(\mathbf{s})$, and $g_1^{\text{E2-PYS}}(\mathbf{s})$ in (2.6), we obtain

$$\frac{p_2^{\text{E2-PYS}}(\mathbf{s}) + \lambda_1^{\text{E2-PYS}}(\mathbf{s})}{p_2^{\text{E2-PYS}}(\mathbf{s}) + \lambda_1^{\text{E2-PYS}}(\mathbf{s}) g_1^{\text{E2-PYS}}(\mathbf{s})} = \frac{s_2 + (\beta - 1)s_2 \exp\left(\frac{\beta}{\beta - 1} \cdot \frac{s_1}{s_2}\right)}{s_2 + \frac{\beta - 1}{\beta} s_2 \exp\left(\frac{\beta}{\beta - 1} \cdot \frac{s_1}{s_2}\right) \left(1 - \exp\left(-\frac{\beta}{\beta - 1} \cdot \frac{s_1}{s_2}\right)\right)} = \beta. \quad (2.12)$$

For the second case where $s_1 > s_2$, we have that

$$g_1^{\text{E2-PYS}}(s_1, s_2) = \frac{\beta - 1}{\beta} + \frac{1}{\beta} \exp\left(-\frac{\beta}{\beta - 1} \cdot \frac{s_2}{s_1}\right)$$

and

$$\left. \frac{\partial g_1^{\text{E2-PYS}}(y, s_2)}{\partial y} \right|_{y=s_1} = \frac{s_2}{(\beta - 1)s_1^2} \exp\left(-\frac{\beta}{\beta - 1} \cdot \frac{s_2}{s_1}\right).$$

Now, it is

$$\lambda_1^{\text{E2-PYS}}(\mathbf{s}) = \frac{(\beta - 1)s_1^2}{s_2} \exp\left(\frac{\beta}{\beta - 1} \cdot \frac{s_2}{s_1}\right).$$

By substituting $p_2^{\text{E2-PYS}}(\mathbf{s})$, $\lambda_1^{\text{E2-PYS}}(\mathbf{s})$, and $g_1^{\text{E2-PYS}}(\mathbf{s})$ in (2.6), we obtain

$$\frac{p_2^{\text{E2-PYS}}(\mathbf{s}) + \lambda_1^{\text{E2-PYS}}(\mathbf{s})}{p_2^{\text{E2-PYS}}(\mathbf{s}) + \lambda_1^{\text{E2-PYS}}(\mathbf{s}) g_1^{\text{E2-PYS}}(\mathbf{s})} = \beta \frac{1 + (\beta - 1) \left(\frac{s_1}{s_2}\right)^2 \exp\left(\frac{\beta}{\beta - 1} \cdot \frac{s_2}{s_1}\right)}{\beta + (\beta - 1)^2 \left(\frac{s_1}{s_2}\right)^2 \exp\left(\frac{\beta}{\beta - 1} \cdot \frac{s_2}{s_1}\right) + (\beta - 1) \left(\frac{s_1}{s_2}\right)^2} \leq \beta. \quad (2.13)$$

The inequality follows since the quantity at its left is decreasing in s_1/s_2 (its derivative with respect to s_1/s_2 can be shown by tedious calculations to be non-positive for $s_1/s_2 \geq 1$) and, hence, it is upper-bounded by its value for $s_1/s_2 = 1$; this is equal to β by its definition.

The theorem follows by Lemma 2.3 using (2.12) and (2.13). \square

We remark that a preliminary analysis similar to the first half of the proof of Theorem 2.8 inspired the design of the E2-PYS mechanism (as well as that of E2-SR mechanism that is defined later) at first place. By keeping the allocation function as the unknown and requiring that the RHS of (2.6) is equal to some value α for all signal vectors $\mathbf{s} \in \mathbb{X}_2$ with $s_1 \leq s_2$ (this is essentially what (2.12) captures), we obtained a first-order differential equation which, using the appropriate conditions so that the resulting mechanism is valid, led to E2-PYS (for $\alpha = \beta$). Luckily, for signal vectors $\mathbf{s} \in \mathbb{X}_2$ with $s_1 > s_2$, we were able to show that the RHS of (2.6) is at most α ; see inequality (2.13).

We now show that E2-PYS has optimal LPoA among 2-player PYS mechanisms in class \mathcal{C} . The proof makes use of Lemma 2.3 and a simple differential inequality that involves the allocation function.

Theorem 2.9. *Any 2-player PYS mechanism with concave allocation function has liquid price of anarchy at least $\beta \approx 1.792$.*

Proof. For the sake of contradiction, assume that there exists a PYS mechanism M that has liquid price of anarchy $\beta' < \beta$. Denote by $f : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ the function defined as $f(y) = g_1^M(y, 1)$. Then, by applying Lemma 2.3 with $\mathbf{s} = (y, 1) \in \mathbb{X}_2$ to M we have $\lambda_1^M(y, 1) = 1/f'(y)$ and $\text{LPoA}(M) \geq \frac{1+1/f'(y)}{1+f(y)/f'(y)}$ for every $y \in [0, 1]$. By our assumption $\text{LPoA}(M) \leq \beta'$, we get the differential inequality

$$(\beta' - 1)f'(y) + \beta' f(y) \geq 1$$

for every $y \in [0, 1]$. Using Grönwall's inequality, $f(y)$ is lower-bounded by the solution of the corresponding differential equation. Due to the condition $f(0) = 0$, this yields

$$f(y) \geq \frac{1}{\beta'} - \frac{1}{\beta'} \exp\left(-\frac{\beta'}{\beta' - 1}y\right)$$

and, hence,

$$\frac{1}{2} = f(1) \geq \frac{1}{\beta'} - \frac{1}{\beta'} \exp\left(-\frac{\beta'}{\beta' - 1}\right) > \frac{1}{\beta} - \frac{1}{\beta} \exp\left(-\frac{\beta}{\beta - 1}\right),$$

which contradicts the definition of β . The last inequality follows since the function $\frac{1}{z} - \frac{1}{z} \exp\left(-\frac{z}{z-1}\right)$ is decreasing in the interval $[1, 2]$. \square

2.7.2 The E2-SR mechanism

Let us now define a non-PYS mechanism that has considerably better LPoA than E2-PYS and almost matches the lower bound of $3/2$ from Theorem 2.1 for 2-player mechanisms. Let $\gamma \approx 1.529$ be the solution of the equation $\frac{1}{\gamma} - \frac{1}{\gamma} \exp\left(-\frac{\gamma}{2(\gamma-1)}\right) = \frac{1}{2}$ and define mechanism E2-SR to be the 2-player mechanism that uses the allocation function (see Figure 2.3 for a comparison of the allocation functions of Kelly, SH, E2-PYS, and E2-SR)

$$g_i^{\text{E2-SR}}(\mathbf{s}) = \begin{cases} \frac{1}{\gamma} - \frac{1}{\gamma} \exp\left(-\frac{\gamma}{2(\gamma-1)} \cdot \left(\frac{s_i}{s_{3-i}}\right)^2\right) & s_i \leq s_{3-i} \\ \frac{\gamma-1}{\gamma} + \frac{1}{\gamma} \exp\left(-\frac{\gamma}{2(\gamma-1)} \cdot \left(\frac{s_{3-i}}{s_i}\right)^2\right) & s_i > s_{3-i} \end{cases}$$

and the payment function $p_i^{\text{E2-SR}}(\mathbf{s}) = s_i/s_{3-i}$ for player $i \in \{1, 2\}$ and (non-zero) signal vector $\mathbf{s} = (s_1, s_2)$. By the general conventions of Section 2.3, the payments are 0 when some of the




Figure 2.3: A comparison of the allocation function g_i^M used by E2-PYS (in green), (the 2-player version of) Kelly (in blue), SH (dashed), and E2-SR (in red) as a function of s_i/s_{3-i} for $s_i \leq s_{3-i}$. Among these mechanisms, E2-SR is the only one with a non-concave allocation function.

signals is equal to zero. Due to the definition of γ , E2-SR is a well-defined resource allocation mechanism. However, observe that E2-SR does not belong to class \mathcal{C} (the allocation function is not concave; see Figure 2.3) and the condition $\mathbf{s} \in \text{eq}(\mathcal{G}^{\text{E2-SR}}(\mathbf{s}, 1))$ is not guaranteed to be satisfied. Next, we will prove an upper bound on the LPoA of E2-SR. The proof follows in a similar way to the proof of Theorem 2.8, but it does not provide a tight bound.

Theorem 2.10. *The liquid price of anarchy of the E2-SR mechanism is at most $\gamma \approx 1.529$.*

Proof. We will prove the theorem by mimicking the proof of Theorem 2.8. Let $\mathbf{s} \in \mathbb{X}_2$. Again, we distinguish between two cases. First, assume that $s_1 \leq s_2$. Then, since

$$g_1^{\text{E2-SR}}(s_1, s_2) = \frac{1}{\gamma} - \frac{1}{\gamma} \exp\left(-\frac{\gamma}{2(\gamma-1)} \cdot \left(\frac{s_1}{s_2}\right)^2\right),$$

the derivative of the allocation for player 1 is

$$\left. \frac{\partial g_1^{\text{E2-SR}}(y, s_2)}{\partial y} \right|_{y=s_1} = \frac{s_1}{(\gamma-1)s_2^2} \exp\left(-\frac{\gamma}{2(\gamma-1)} \cdot \left(\frac{s_1}{s_2}\right)^2\right).$$

Also, since the payment function used by E2-SR is equal to the signal ratio, its derivative for player 1 is

$$\left. \frac{\partial p_1^{\text{E2-SR}}(y, s_2)}{\partial y} \right|_{y=s_1} = \frac{1}{s_2}.$$

Therefore,

$$\lambda_1^{\text{E2-SR}}(\mathbf{s}) = (\gamma-1) \frac{s_2}{s_1} \exp\left(\frac{\gamma}{2(\gamma-1)} \cdot \left(\frac{s_1}{s_2}\right)^2\right).$$

By substituting $p_2^{\text{E2-SR}}(\mathbf{s})$, $\lambda_1^{\text{E2-SR}}(\mathbf{s})$, and $g_1^{\text{E2-SR}}(\mathbf{s})$ to (2.6), we can now easily verify that

$$\frac{p_2^{\text{E2-SR}}(\mathbf{s}) + \lambda_1^{\text{E2-SR}}(\mathbf{s})}{p_2^{\text{E2-SR}}(\mathbf{s}) + \lambda_1^{\text{E2-SR}}(\mathbf{s}) g_1^{\text{E2-SR}}(\mathbf{s})} = \gamma. \quad (2.14)$$

For the second case where $s_1 > s_2$, the allocation is defined as

$$g_1^{\text{E2-SR}}(s_1, s_2) = \frac{\gamma - 1}{\gamma} + \frac{1}{\gamma} \exp \left(-\frac{\gamma}{2(\gamma - 1)} \cdot \left(\frac{s_2}{s_1} \right)^2 \right)$$

and the derivative is

$$\left. \frac{\partial g_1^{\text{E2-SR}}(y, s_2)}{\partial y} \right|_{y=s_1} = \frac{s_2^2}{(\gamma - 1)s_1^3} \exp \left(-\frac{\gamma}{2(\gamma - 1)} \cdot \left(\frac{s_2}{s_1} \right)^2 \right).$$

The payment derivative is again equal to $1/s_2$ and, hence,

$$\lambda_1^{\text{E2-SR}}(\mathbf{s}) = (\gamma - 1) \left(\frac{s_1}{s_2} \right)^3 \exp \left(\frac{\gamma}{2(\gamma - 1)} \cdot \left(\frac{s_2}{s_1} \right)^2 \right).$$

By substituting $p_2^{\text{E2-SR}}(\mathbf{s})$, $\lambda_1^{\text{E2-SR}}(\mathbf{s})$, and $g_1^{\text{E2-SR}}(\mathbf{s})$, we obtain

$$\frac{p_2^{\text{E2-SR}}(\mathbf{s}) + \lambda_1^{\text{E2-SR}}(\mathbf{s})}{p_2^{\text{E2-SR}}(\mathbf{s}) + \lambda_1^{\text{E2-SR}}(\mathbf{s}) g_1^{\text{E2-SR}}(\mathbf{s})} = \gamma \frac{1 + (\gamma - 1) \left(\frac{s_1}{s_2} \right)^4 \exp \left(\frac{\gamma}{2(\gamma - 1)} \cdot \left(\frac{s_2}{s_1} \right)^2 \right)}{\gamma + (\gamma - 1)^2 \left(\frac{s_1}{s_2} \right)^4 \exp \left(\frac{\gamma}{2(\gamma - 1)} \cdot \left(\frac{s_2}{s_1} \right)^2 \right) + (\gamma - 1) \left(\frac{s_1}{s_2} \right)^4} \leq \gamma. \quad (2.15)$$

The inequality follows because the quantity at its left is decreasing in s_1/s_2 (its derivative with respect to s_1/s_2 can be shown by tedious calculations to be non-positive for $s_1/s_2 \geq 1$) and, hence, it is upper-bounded by its value for $s_1/s_2 = 1$; this is equal to γ by its definition.

The theorem follows by Lemma 2.3 using (2.14) and (2.15). \square

Interestingly, a simpler 2-player mechanism that uses the allocation function of SH and the signal-ratio payment function of E2-SR has a slightly worse LPoA of $\phi = 1.618$. In fact, this bound is tight since this particular mechanism belongs to class \mathcal{C} . The proof is left as an exercise to the reader.

2.8 Some extensions

In this section, we will shortly discuss two possible extensions of our work. In particular, we will discuss the possibility of achieving improved LPoA bounds (1) via mechanisms that have access to the player budgets, and (2) by allowing the players to be more expressive and submit signals that carry more information.

2.8.1 Budget-aware mechanisms

All of our work so far in this chapter has focused on the scenario where the budget of each player is private. Let us now discuss a bit the case of *budget-aware* mechanisms, which have access to the budget value of each player.

It is easy to verify that our analysis for mechanisms Kelly, SH, E2-PYS, and E2-SR carries over to this case. In contrast, our lower bound (Theorem 2.1) is not true anymore. The proof constructs two games, in which almost every player has different budgets. The main property that we exploited (for non-budget-aware mechanisms) is that the strategic behavior of the players results in the same set of equilibria in both games. This argument fails for budget-aware mechanisms; a small change in the budget of a single player could be enough to alter the set of equilibria. So, in principle, one might hope even for full efficiency at equilibria (LPoA equal to 1) in this case, analogously to the results of Maheswaran and Basar [2006], Johari and Tsitsiklis [2009], and Yang and Hajek [2007] in the no-budget setting. Interestingly, our next statement rules out this possibility.

Theorem 2.11. *For $n \geq 2$, every n -player budget-aware resource allocation mechanism has liquid price of anarchy at least $4/3$.*

Proof. Let M be any n -player budget-aware resource allocation mechanism that uses an allocation function g^M and a payment function p^M . Let $\mathbf{s} = (s_1, \dots, s_n)$ be an equilibrium of the game \mathcal{G}_1^M induced by M for players with valuations $v_i(x) = x$ for $i \in \{1, 2\}$ and $v_i(x) = 0$ for $i \geq 3$, and budgets $c_i = 1$ for every $i \in [n]$. Assume that the allocation returned by M at this equilibrium is $\mathbf{d} = (d_1, \dots, d_n)$. Without loss of generality, we may assume that one of the first two players (say, player 1) gets a resource share of at most $1/2$.

Recall that, for every signal vector \mathbf{y} , the utility of any player i is defined as $u_i^M(\mathbf{y}) = v_i(g_i^M(\mathbf{y})) - p_i^M(\mathbf{y})$. Now, consider the game \mathcal{G}_2^M where player 2 has the modified valuation function $\tilde{v}_2(x) = 1 + x$ while all other players are as in \mathcal{G}_1^M ; the budgets are the same in both games and are known to the mechanism. Observe that the modified utility of player 2 is now $\tilde{u}_2^M(\mathbf{y}) = \tilde{v}_2(g_2^M(\mathbf{y})) - p_2^M(\mathbf{y}) = u_2^M(\mathbf{y}) + 1$. Hence, \mathbf{s} is an equilibrium in \mathcal{G}_2^M as well and M returns the same allocation \mathbf{d} again.

Clearly, due to the definition of the valuation functions, the contribution of players $i \geq 3$ in the liquid welfare (in any state of the game) is zero. Hence, the liquid welfare at equilibrium is $\min\{\tilde{v}_1(d_1), c_1\} + \min\{\tilde{v}_2(d_2), c_2\} = d_1 + 1 \leq 3/2$, while the optimal liquid welfare is equal

to 2, achieved at the allocation according to which the whole resource is given to player 1. We conclude that the liquid price of anarchy of M is $\text{LPoA}(M) \geq \text{LPoA}(\mathcal{G}_2^M) \geq 4/3$, as desired. \square

In spite of the lower bound in Theorem 2.11, whether budget-aware resource allocation mechanisms can have an LPoA better than $2 - 1/n$ is an important open problem.

2.8.2 Higher expressiveness

Another extension of our setting could be to allow the players to declare their budget to the mechanism in addition to their scalar signal. Taking this approach to its extreme, one could imagine resource allocation mechanisms which ask the players to submit *multi-dimensional* signals. At first glance, this seems to lead to much more powerful mechanisms than the ones we have considered here. Surprisingly, this *higher level of expressiveness* has no consequences to the LPoA at all and our lower bound of $2 - 1/n$ captures such mechanisms as well. Indeed, by inspecting the two games used in the proof of Theorem 2.1, we can verify that the same signal vector (no matter whether signals are single- or multi-dimensional) leads to the same allocation by the mechanism and the same strategic behavior of the players in both games. This observation applies to the proof of Theorem 2.11 as well.

2.9 Conclusion

In this chapter, we studied the efficiency of resource allocation mechanisms for users with private concave valuation functions and budget constraints, who compete over the acquisition of a single divisible resource. Using the liquid welfare as our efficiency benchmark, we showed a completely different picture compared to the no-budget case, for which there exist fully efficient mechanisms that align the global objective of maximizing the social welfare with the strategic objectives of the players.

First, we proved a lower bound of $2 - 1/n$ on the liquid price of anarchy of any n -player resource allocation mechanism, which indicates that there exist no fully efficient mechanisms in the case where the players have budgets. Then, we characterized the worst-case games and equilibria with respect to the liquid price of anarchy, and used this characterization to prove tight bounds on the well-known Kelly and SH mechanisms (2 and 3, respectively). Further, we exploited our characterization to design the improved mechanisms E2-PYS and E2-SR for the case of two players that achieve an LPoA of approximately 1.79 and 1.53, respectively.

Chapter 3

Bounding the inefficiency of compromise in opinion formation

In this chapter, we study questions related to the existence, computational complexity, and quality of equilibria in k -compromising opinion formation (k -COF) games; see the discussion in Section 1.2 for a high-level introduction to the problem. These results have been published in [Caragiannis et al., 2017a].

3.1 Overview of contribution and techniques

We begin by proving several properties about the geometric structure of opinions and beliefs at pure Nash equilibria (states of the game where each player minimizes her individual cost assuming that the remaining players will not change their opinions).

Using these structural properties we show that there exist simple k -COF games that do not admit pure Nash equilibria. Furthermore, we prove that even in games where equilibria do exist, their quality may be suboptimal in terms of the social cost (the total cost experienced by all players), by showing that the price of stability grows linearly with k . For the special case of 1-COF games, we show that each such game admits a representation as a directed acyclic graph, in which every pure Nash equilibrium corresponds to a path between two designated nodes. Hence, the problems of computing the best or worst (in terms of the social cost) pure Nash equilibrium (or even of computing whether such an equilibrium exists) are equivalent to simple path computations that can be performed in polynomial time.

For general k -COF games, we quantify the inefficiency of the worst-case pure equilibria by bounding the price of anarchy. Specifically, we present upper and lower bounds on the price of anarchy of k -COF games (with respect to both pure and mixed Nash equilibria) that

k	PoA	MPoA	PoS	Existence/Complexity
1	3 (Thms 3.19, 3.20)	≥ 6 (Thm. 3.21)	$\geq 17/15$ (Thm. 3.8)	PNE may not exist for any k (Thm. 3.6)
2	≤ 12 (Thm. 3.12) $\geq 18/5$ (Thm. 3.22)	$\geq 24/5$ (Thm. 3.23)	$\geq 8/7$ (Thm. 3.9)	Best and worst PNE is in P (Thm. 3.10)
≥ 3	$\leq 4(k+1)$ (Thm. 3.12) $\geq k+1$ (Thm. 3.22)	$\geq k+2$ (Thm. 3.23)	$\geq \frac{k+1}{3}$ (Thm. 3.7)	Open question: Is computing a PNE in P when $k \geq 2$?

Table 3.1: Summary of our results for k -COF games. The table presents our bounds on the price of anarchy over pure Nash equilibria (PoA) and mixed Nash equilibria (MPoA), on the price of stability (PoS) as well as the existence and complexity of pure Nash equilibria (PNE). Clearly, any upper bound on the price of anarchy is also an upper bound on the price of stability. See [Caragiannis et al., 2017a].

suggest a linear dependence on k . Our upper bound on the price of anarchy exploits, in a non-trivial way, linear programming duality in order to lower-bound the optimal social cost. For the fundamental case of 1-COF games, we obtain a tight bound of 3 using a particular charging scheme in the analysis. Our contribution is summarized in Table 3.1.

3.1.1 Chapter roadmap

In the following, we begin with a discussion of the bibliography that is related to opinion formation and to our work in particular. Then, in Section 3.2, we continue with preliminary definitions, notation and examples in Section 3.3. In Section 3.4 we present several structural properties of pure Nash equilibria, while Section 3.5 is devoted to the existence and the price of stability of these equilibria. Then, in Section 3.6 we present a polynomial-time algorithm that determines whether pure Nash equilibria exist in 1-COF games, and, in addition, computes the best and worst such equilibria, when they do exist. In Sections 3.7 and 3.8 we prove upper bounds on the price of anarchy for general k -COF and 1-COF games, respectively, while Section 3.9 contains our lower bounds on the price of anarchy. We conclude with a synopsis of our results in Section 3.10.

3.2 Related work

DeGroot [1974] proposed a framework that models the opinion formation process, where each individual updates her opinion according to a weighted averaging procedure. Subsequently, Friedkin and Johnsen [1990] refined the model by assuming that each individual has a private

belief and expresses a (possibly different) public opinion that depends on her belief and the opinions of the people in her social circle. More recently, Bindel et al. [2015] studied this model and proved that, for the setting where beliefs and opinions are real numbers in the interval $[0, 1]$, the repeated averaging process leads to an opinion vector that can be thought of as the unique equilibrium in a corresponding opinion formation game.

Deviating from the assumption that opinions depend on the whole social circle, Bhawalkar et al. [2013] considered co-evolutionary opinion formation games, where as opinions evolve so does the neighborhood of each person. This model is conceptually similar to previous ones that have been studied by Hegselmann and Krause [2002], and Holme and Newman [2006]. Both Bindel et al. [2015] and Bhawalkar et al. [2013] proved constant bounds on the price of anarchy of the games that they study. In contrast, the modified cost function we used in this chapter in order to model compromise yields considerably higher price of anarchy, that depends linearly on the size of the neighborhood.

A series of recent papers in the EconCS community considered discrete opinion formation models with binary opinions. Chierichetti et al. [2018] considered discrete preference games, where beliefs and opinions are binary and study questions related to the price of stability. For these games, Auletta et al. [2015, 2017a] characterized the social networks where the belief of the minority can emerge as the opinion of the majority, while Auletta et al. [2017b] examined the robustness of such results to variants of the model. Auletta et al. [2016] generalized the class of discrete preference games so that the players are not only interested in agreeing with their neighbors, but more complex constraints can be used to represent the preferences of the players. Bilò et al. [2016] extended the class of co-evolutionary formation games to the discrete setting. Other models assume that opinion updates do not depend on the entire social circle of each individual; instead, each person consults only a small random subset of her social acquaintances; see the recent paper by Fotakis et al. [2016] as well as the survey of Mossel and Tamuz [2014].

In scenarios where there are more than one issues to be discussed, Jia et al. [2015] proposed and analyzed the DeGroot-Friedkin model for the evolution of an influence network between individuals who form opinions on a sequence of issues, while Xu et al. [2015] introduced a modification according to which each individual may recalculate the weight assigned to her opinion (her self-confidence), after the discussion of each issue with her social circle.

Another line of research has focused on how fast a system converges to a stable state. In this

spirit, Etesami and Basar [2015] considered the dynamics of the co-evolutionary Hegselmann-Krause model [2002] and focused on the termination time in finite dimensions under different settings. Similarly, Ferraioli et al. [2016] studied the convergence of decentralized dynamics in finite opinion games, where players have only a finite number of opinions available. Ferraioli and Ventre [2017] considered the role of social pressure towards consensus in opinion games and provide tight bounds on the speed of convergence for the important special case where the social network is a clique.

Das et al. [2014] performed a set of online user studies and argued that widely studied theoretical models do not completely explain the experimental results obtained. Hence, they introduced an analytical model for opinion formation and presented preliminary theoretical and simulation results on the convergence and structure of opinions when users iteratively update their respective opinions according to the new model.

Chazelle [2012] analyzed influence systems, where each individual observes the location of her neighbors and moves accordingly, and presented an algorithmic calculus for studying such systems. Kempe et al. [2016] presented a novel model of cultural dynamics and focused on the interplay between selection and influence. Among other results, they presented an almost complete characterization of stable outcomes and showed that convergence is guaranteed from all starting states. Gomez-Rodriguez et al. [2012] considered network diffusion and contagion propagation. Their goal was to infer an unknown network over which contagion propagated, tracing paths of diffusion and influence. Finally, Kempe et al. [2015] studied the optimization problem for influence maximization in a social networks, where each individual may decide to adopt an idea or an innovation depending on how many of her neighbors already do. The goal is to select an initial seed set of early adopters so that the number of adopters is maximized.

3.3 Definitions and notation

A compromising opinion formation game defined by the k nearest neighbors (henceforth, called k -COF game) is played by a set of n players whose beliefs lie on the line of real numbers. Let $\mathbf{s} = (s_1, s_2, \dots, s_n) \in \mathbb{R}^n$ be the vector containing the players' beliefs such that $s_i \leq s_{i+1}$ for each $i \in [n - 1]$. Let $\mathbf{z} = (z_1, z_2, \dots, z_n) \in \mathbb{R}^n$ be a vector containing the (deterministic or randomized) opinions expressed by the players; these opinions define a state of the game. We denote by \mathbf{z}_{-i} the opinion vector obtained by removing z_i from \mathbf{z} . In an attempt to simplify notation, we omit k from all relevant definitions.

Given vector \mathbf{z} (or a realization of it in case \mathbf{z} contains randomized opinions), we define the neighborhood $N_i(\mathbf{z}, \mathbf{s})$ of player i to be the set of k players whose opinions are the closest to the belief of player i breaking ties arbitrarily (but consistently). For each player i , we define $I_i(\mathbf{z}, \mathbf{s})$ as the shortest interval of the real line that includes the following points: the belief s_i , the opinion z_i , and the opinion z_j for each player $j \in N_i(\mathbf{z}, \mathbf{s})$. Furthermore, let $\ell_i(\mathbf{z}, \mathbf{s})$ and $r_i(\mathbf{z}, \mathbf{s})$ be the players with the leftmost and rightmost point in $I_i(\mathbf{z}, \mathbf{s})$, respectively. For example, $\ell_i(\mathbf{z}, \mathbf{s})$ can be equal to either player i or some player $j \in N_i(\mathbf{z}, \mathbf{s})$, depending on whether the leftmost point of $I_i(\mathbf{z}, \mathbf{s})$ is s_i , z_i , or z_j . To further simplify notation, we will frequently use $\ell(i)$ and $r(i)$ instead of $\ell_i(\mathbf{z}, \mathbf{s})$ and $r_i(\mathbf{z}, \mathbf{s})$ when \mathbf{z} and \mathbf{s} are clear from the context. In the following, we present the relevant definitions for the case of possibly randomized opinion vectors; clearly, these can be simplified whenever \mathbf{z} consists entirely of deterministic opinions.

Given a k -COF game with belief vector \mathbf{s} , the cost that player i experiences at the state of the game defined by an opinion vector \mathbf{z} is

$$\begin{aligned}\mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})] &= \mathbb{E} \left[\max_{j \in N_i(\mathbf{z}, \mathbf{s})} \left\{ |z_i - s_i|, |z_j - z_i| \right\} \right] \\ &= \mathbb{E} \left[\max \left\{ |z_i - s_i|, |z_{r_i(\mathbf{z}, \mathbf{s})} - z_i|, |z_i - z_{\ell_i(\mathbf{z}, \mathbf{s})}| \right\} \right].\end{aligned}\quad (3.1)$$

For the special case of 1-COF games, we denote by $\sigma_i(\mathbf{z}, \mathbf{s})$ (or $\sigma(i)$ when \mathbf{z} and \mathbf{s} are clear from the context) the player (other than i) whose opinion is closest to the belief s_i of player i ; notice that $\sigma(i)$ is the only member of $N_i(\mathbf{z}, \mathbf{s})$. In this case, the cost of player i can be simplified as

$$\mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})] = \mathbb{E} \left[\max \left\{ |z_i - s_i|, |z_{\sigma_i(\mathbf{z}, \mathbf{s})} - z_i| \right\} \right]. \quad (3.2)$$

We say that an opinion vector \mathbf{z} consisting entirely of deterministic opinions is a *pure Nash equilibrium* if no player i has an incentive to unilaterally deviate to a deterministic opinion z'_i in order to decrease her cost, i.e.,

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq \text{cost}_i((z'_i, \mathbf{z}_{-i}), \mathbf{s}),$$

where by (z'_i, \mathbf{z}_{-i}) we denote the opinion vector in which player i chooses the opinion z'_i and all other players choose the opinions they have according to vector \mathbf{z} . Similarly, a possibly randomized opinion vector \mathbf{z} is a *mixed Nash equilibrium* if for any player i and any deviating deterministic opinion z'_i we have

$$\mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})] \leq \mathbb{E}_{\mathbf{z}_{-i}}[\text{cost}_i((z'_i, \mathbf{z}_{-i}), \mathbf{s})].$$

Let $\text{PNE}(\mathbf{s})$ and $\text{MNE}(\mathbf{s})$ denote the sets of pure and mixed Nash equilibria, respectively, of the k -COF game with belief vector \mathbf{s} .

The *social cost* of an opinion vector \mathbf{z} is the total cost experienced by all players, i.e.,

$$\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})] = \sum_{i=1}^n \mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})].$$

Let $\mathbf{z}^*(\mathbf{s})$ be a deterministic opinion vector that minimizes the social cost for the given k -COF game with belief vector \mathbf{s} ; we will refer to it as an *optimal* opinion vector for \mathbf{s} .

The *price of anarchy* (PoA) over pure Nash equilibria of a particular k -COF game with belief vector \mathbf{s} is defined as the ratio between the social cost of its *worst* (in terms of the social cost) pure Nash equilibrium and the optimal social cost, i.e.,

$$\text{PoA}(\mathbf{s}) = \sup_{\mathbf{z} \in \text{PNE}(\mathbf{s})} \frac{\text{SC}(\mathbf{z}, \mathbf{s})}{\text{SC}(\mathbf{z}^*(\mathbf{s}), \mathbf{s})}.$$

The *price of stability* (PoS) over pure Nash equilibria of the k -COF game with belief vector \mathbf{s} is defined as the ratio between the social cost of the *best* pure Nash equilibrium (in terms of social cost) and the optimal social cost, i.e.,

$$\text{PoS}(\mathbf{s}) = \inf_{\mathbf{z} \in \text{PNE}(\mathbf{s})} \frac{\text{SC}(\mathbf{z}, \mathbf{s})}{\text{SC}(\mathbf{z}^*(\mathbf{s}), \mathbf{s})}.$$

Similarly, the price of anarchy and the price of stability over mixed Nash equilibria of a k -COF game with belief vector \mathbf{s} are defined as

$$\text{MPoA}(\mathbf{s}) = \sup_{\mathbf{z} \in \text{MNE}(\mathbf{s})} \frac{\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})]}{\text{SC}(\mathbf{z}^*(\mathbf{s}), \mathbf{s})}$$

and

$$\text{MPoS}(\mathbf{s}) = \inf_{\mathbf{z} \in \text{MNE}(\mathbf{s})} \frac{\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})]}{\text{SC}(\mathbf{z}^*(\mathbf{s}), \mathbf{s})},$$

respectively.

Then, the price of anarchy and the price of stability of k -COF games, for a fixed k , are defined as the supremum of $\text{PoA}(\mathbf{s})$ and $\text{PoS}(\mathbf{s})$ over all belief vectors \mathbf{s} , respectively.

Example 3.1. Consider the 1-COF game with three players and belief vector $\mathbf{s} = (-10, 2, 5)$ which is depicted in Figure 3.1(a). For simplicity, we will refer to the players as left (ℓ), middle (m), and right (r).

Let us examine the opinion vector $\mathbf{z} = (-10, -5, 4)$ which is depicted in Figure 3.1(b). We have that $\sigma(\ell) = m$ since the opinion $z_m = -5$ of the middle player is closer to the belief

$s_\ell = -10$ of the left player than the opinion $z_r = 4$ of the right player. Therefore, the cost of the left player is $\text{cost}_\ell(\mathbf{z}, \mathbf{s}) = \max\{|-10 + 10|, |-10 + 5|\} = 5$. Similarly, the neighbors of the middle and right players are $\sigma(m) = r$ and $\sigma(r) = m$, while their costs are $\text{cost}_m(\mathbf{z}, \mathbf{s}) = \max\{2 + 5, 4 + 5\} = 9$ and $\text{cost}_r(\mathbf{z}, \mathbf{s}) = \max\{5 - 4, 4 + 5\} = 9$, respectively. The social cost is $\text{SC}(\mathbf{z}, \mathbf{s}) = 23$.

Now, consider the alternative pure Nash equilibrium opinion vector $\mathbf{z}' = (-3.5, 3, 4)$ which is depicted in Figure 3.1(c). Observe that even though $\mathbf{z}' \neq \mathbf{z}$, each player has the same neighbor as in \mathbf{z} and no player has an incentive to deviate in order to decrease her cost. Indeed, let us focus on the middle player for whom it is $\sigma(m) = r$. Her opinion is in the middle of the interval defined by her belief $s_m = 3$ and the opinion $z'_r = 5$ of the right player. Hence, this opinion minimizes her cost by minimizing the maximum between the distance from her belief and the distance from the opinion of the right player. It is easy to verify that the same holds for the left and right players. The player costs are now 6.5, 1, and 1, respectively, yielding a social cost of 8.5. \square

3.4 Some properties about equilibria

We devote this section to proving several some interesting properties of pure Nash equilibria; these will be useful in the following. The first one is obvious due to the definition of the cost function.

Lemma 3.1. *In any pure Nash equilibrium \mathbf{z} of a k -COF game with belief vector \mathbf{s} , the opinion of any player i lies in the middle of the interval $I_i(\mathbf{z}, \mathbf{s})$.*

The next lemma allows us to argue about the order of player opinions in any pure Nash equilibrium \mathbf{z} .

Lemma 3.2. *In any pure Nash equilibrium \mathbf{z} of a k -COF game with belief vector \mathbf{s} , it holds that $z_i \leq z_{i+1}$ for any $i \in [n - 1]$ such that $s_i < s_{i+1}$.*

Proof. For the sake of contradiction, let us assume that $z_{i+1} < z_i$ for a pair of players i and $i + 1$ with $s_i < s_{i+1}$. Then, it cannot be the case that the leftmost endpoint of the interval $I_i(\mathbf{z}, \mathbf{s})$ of player i is at the left of (or coincides with) the leftmost endpoint of interval $I_{i+1}(\mathbf{z}, \mathbf{s})$ of player $i + 1$ and the rightmost endpoint of $I_i(\mathbf{z}, \mathbf{s})$ is at the left of (or coincides with) the rightmost endpoint of $I_{i+1}(\mathbf{z}, \mathbf{s})$. In other words, it cannot be the case that $\min\{s_i, z_{\ell(i)}\} \leq \min\{s_{i+1}, z_{\ell(i+1)}\}$ and




Figure 3.1: The game examined in Example 3.1. (a) Illustration of the belief vector $\mathbf{s} = (-10, 2, 5)$. The black squares correspond to player beliefs, while the notation $[x]$ is used to denote the number of players that have the same beliefs; here we have only one player per belief. (b) Illustration of the opinion vector $\mathbf{z} = (-10, -5, 4)$. The dots correspond to player opinions and each arrow connects the belief of a player to her opinion. (c) Illustration of the equilibrium opinion vector $\mathbf{z}' = (-3.5, 3, 4)$.

$\max\{s_i, z_{r(i)}\} \leq \max\{s_{i+1}, z_{r(i+1)}\}$ hold simultaneously. Since, by Lemma 3.1, points z_i and z_{i+1} lie in the middle of the corresponding intervals, we would have $z_i \leq z_{i+1}$, contradicting our assumption.

So, at least one of the two inequalities between the interval endpoints above must not hold. In the following, we assume that $\min\{s_i, z_{\ell(i)}\} > \min\{s_{i+1}, z_{\ell(i+1)}\}$ (the case where $\max\{s_i, z_{r(i)}\} > \max\{s_{i+1}, z_{r(i+1)}\}$ is symmetric). This assumption implies that $z_{\ell(i+1)} < s_i < s_{i+1}$ (i.e., $\min\{s_{i+1}, z_{\ell(i+1)}\} = z_{\ell(i+1)}$), and, subsequently, that $z_{\ell(i+1)} < z_{\ell(i)}$. In words, player $\ell(i+1)$ does not belong to interval $I_i(\mathbf{z}, \mathbf{s})$. Furthermore, since $z_{\ell(i+1)} < s_{i+1}$, and as (by Lemma 3.1) z_{i+1} lies in the middle of $I_{i+1}(\mathbf{z}, \mathbf{s})$, we also have that the leftmost endpoint of interval $I_{i+1}(\mathbf{z}, \mathbf{s})$ cannot belong to player $i+1$, i.e., $\ell(i+1) \neq i+1$. An example of the relative ordering of points (beliefs and opinions), after assuming that $z_{i+1} < z_i$ and $\min\{s_i, z_{\ell(i)}\} > \min\{s_{i+1}, z_{\ell(i+1)}\}$ is depicted in Figure 3.2.

Since $\ell(i+1)$ does not belong to $I_i(\mathbf{z}, \mathbf{s})$, there are at least k players different than $\ell(i+1)$ and i that have opinions at distance at most $s_i - z_{\ell(i+1)}$ from belief s_i . Since $s_i < s_{i+1}$ and




Figure 3.2: An example of the argument used in the proof of Lemma 3.2.

$z_{\ell(i+1)} < z_{\ell(i)}$, all these players are also at distance strictly less than $s_{i+1} - z_{\ell(i+1)}$ from belief s_{i+1} . This contradicts the fact that the opinion of player $\ell(i+1)$ is among the k closest opinions to s_{i+1} . \square

In the following, in any pure Nash equilibrium \mathbf{z} , we assume that $z_i \leq z_{i+1}$ for any $i \in [n-1]$. This follows by Lemma 3.2 when $s_i < s_{i+1}$ and by a convention for the identities of players with identical belief.

In addition to the ordering of opinions in a pure Nash equilibrium, we can also specify the range of neighborhoods (in Lemma 3.3) and opinions (in Lemma 3.4).

Lemma 3.3. *Let \mathbf{z} be a pure Nash equilibrium of a k -COF game with belief vector \mathbf{s} . Then, for each player i , there exists j with $i - k \leq j \leq i$ such that $I_i(\mathbf{z}, \mathbf{s})$ is the shortest interval that contains the opinions $z_j, z_{j+1}, \dots, z_{j+k}$ and belief s_i .*

Proof. If $I_i(\mathbf{z}, \mathbf{s})$ consists of a single point, the lemma follows trivially by the definition of the neighborhood and Lemma 3.2 since at least $k+1$ consecutive players including i should have opinions in $I_i(\mathbf{z}, \mathbf{s})$. Otherwise, by Lemma 3.2, the statement is true if there is at most one opinion in each of the left and the right boundary of $I_i(\mathbf{z}, \mathbf{s})$; in this case, there are exactly $k+1$ consecutive players including player i with opinions in $I_i(\mathbf{z}, \mathbf{s})$.

In the following, we will handle the subtleties that may arise due to tie-breaking on the boundaries of $I_i(\mathbf{z}, \mathbf{s})$. Let Y_ℓ and Y_r be the set of players with opinions at the leftmost and the rightmost point of $I_i(\mathbf{z}, \mathbf{s})$, respectively. From Lemma 3.1, player i belongs neither to Y_ℓ nor to Y_r . Now consider the following set of players: the $|Y_\ell \cap N_i(\mathbf{z}, \mathbf{s})|$ players with highest indices from Y_ℓ , the $|Y_r \cap N_i(\mathbf{z}, \mathbf{s})|$ players with lowest indices from Y_r and all players with opinions that lie strictly in $I_i(\mathbf{z}, \mathbf{s})$. Due to the definition of $N_i(\mathbf{z}, \mathbf{s})$ and by Lemma 3.2, there are $k+1$ players in this set, including player i , with consecutive indices. \square




Figure 3.3: An example of the argument used in the proof of Lemma 3.4.

In the following, irrespectively of how ties are actually resolved, we assume that $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$ consists of $k + 1$ players with consecutive indices. This does not affect the cost of player i at equilibrium in the proofs of our upper bounds (since, by Lemma 3.3, the interval defined is exactly the same), while our lower bound constructions are defined carefully so that the results hold no matter how ties are actually resolved.

Lemma 3.4. *Let \mathbf{z} be a pure Nash equilibrium of a k -COF game with belief vector \mathbf{s} . Then, for each player i , it holds that $s_{\ell(i)} \leq z_i \leq s_{r(i)}$.*

Proof. Since $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$ consists of $k + 1$ players with consecutive indices, we have that $s_{\ell(i)} \leq s_i \leq s_{r(i)}$. For the sake of contradiction, let us assume that $s_{\ell(i)} \leq s_{r(i)} < z_i$ for some player i (the case where z_i lies at the left of $s_{\ell(i)}$ is symmetric). Since $s_{r(i)} < s_i$ and as z_i is at the middle of $I_i(\mathbf{z}, \mathbf{s})$, it holds that $z_{r(i)} > z_i$ (i.e., $r(i) \neq i$). Also, since $z_{r(i)} > z_i > s_{r(i)}$, and because $z_{r(i)}$ is in the middle of $I_{r(i)}(\mathbf{z}, \mathbf{s})$, it holds that $z_{r(r(i))} > z_{r(i)}$ and, by Lemma 3.2, $r(r(i)) > r(i)$; see Figure 3.3 for an example of the relative ordering of points (beliefs and opinions) when assuming that $s_{r(i)} < z_i$.

We now claim that $\ell(i) \notin N_{r(i)}(\mathbf{z}, \mathbf{s})$. Assume otherwise that $\ell(i) \in N_{r(i)}(\mathbf{z}, \mathbf{s})$. By definition, $r(r(i)) \in N_{r(i)}(\mathbf{z}, \mathbf{s})$. Then, Lemma 3.2 implies that any player j , different than $r(i)$, with $\ell(i) < j < r(r(i))$ is also in $N_{r(i)}(\mathbf{z}, \mathbf{s})$. Hence, $N_{r(i)}(\mathbf{z}, \mathbf{s})$ contains at least the $k - 1$ players in $N_i(\mathbf{z}, \mathbf{s}) \setminus \{r(i)\}$, as well as players i and $r(r(i))$. This, however, contradicts the fact that $|N_{r(i)}(\mathbf{z}, \mathbf{s})| = k$. Therefore, player $\ell(i)$ is not among the k nearest neighbors of $r(i)$.

So, we obtain that

$$\begin{aligned} z_{r(r(i))} - s_{r(i)} &> z_{r(i)} - s_{r(i)} > z_{r(i)} - z_i = z_i - \min\{s_i, z_{\ell(i)}\} \\ &> s_{r(i)} - \min\{s_i, z_{\ell(i)}\} \geq s_{r(i)} - z_{\ell(i)}. \end{aligned}$$

If $z_{\ell(i)} > s_{r(i)}$ (i.e., $z_{\ell(i)}$ is at the right of $s_{r(i)}$), then since, by Lemma 3.2, $z_{\ell(i)} \leq z_{r(r(i))}$ and $r(r(i)) \in N_{r(i)}(\mathbf{z}, \mathbf{s})$, we obtain that $\ell(i) \in N_{r(i)}(\mathbf{z}, \mathbf{s})$ as well; a contradiction. Otherwise, the

above inequality yields that $z_{r(r(i))} - s_{r(i)} > s_{r(i)} - z_{\ell(i)} \geq 0$ (i.e., the distance of $s_{r(i)}$ from $z_{r(r(i))}$ is strictly higher than the distance of $s_{r(i)}$ from $z_{\ell(i)}$), and, again, we obtain a contradiction to the fact that $\ell(i) \notin N_{r(i)}(\mathbf{z}, \mathbf{s})$ and $r(r(i)) \in N_{r(i)}(\mathbf{z}, \mathbf{s})$. \square

3.5 Existence and quality of equilibria

Our first technical contribution is a negative statement: pure Nash equilibria may not exist for any k (Theorem 3.6). Then, we show that even in games that admit pure Nash equilibria, the best equilibrium may be inefficient; in other words, the price of stability is strictly greater than 1 for any value of k , and, actually, depends linearly on k . These results appear in Theorems 3.7, 3.8, and 3.9.

3.5.1 Existence of equilibria

We begin with a technical lemma. The lemma essentially presents necessary conditions so that a particular set of neighborhoods, and corresponding intervals, may coexist in a pure Nash equilibrium.

Lemma 3.5. *Consider a k -COF game and any three players a, b, c with beliefs $s_a \leq s_b \leq s_c$, respectively. For any pure Nash equilibrium \mathbf{z} where $I_a(\mathbf{z}, \mathbf{s}) = [s_a, z_b]$, $I_b(\mathbf{z}, \mathbf{s}) = [s_b, z_c]$ and $I_c(\mathbf{z}, \mathbf{s}) = [z_b, s_c]$, it must hold that $s_b \geq \frac{3s_a + 5s_c}{8}$, while for any pure Nash equilibrium \mathbf{z} where $I_a(\mathbf{z}, \mathbf{s}) = [s_a, z_b]$, $I_b(\mathbf{z}, \mathbf{s}) = [z_a, s_b]$ and $I_c(\mathbf{z}, \mathbf{s}) = [z_b, s_c]$, it must hold that $s_b \leq \frac{5s_a + 3s_c}{8}$.*

Proof. It suffices to prove the first case; the second case is symmetric. Since $I_b(\mathbf{z}, \mathbf{s}) = [s_b, z_c]$ and $I_c(\mathbf{z}, \mathbf{s}) = [z_b, s_c]$, by Lemma 3.1 it holds that $z_b = (s_b + z_c)/2$ and $z_c = (z_b + s_c)/2$ which yield that $z_b = s_b + \frac{s_c - s_b}{3}$ and $z_c = s_b + \frac{2(s_c - s_b)}{3}$. Hence, we obtain that

$$z_c - s_b = \frac{2(s_c - s_b)}{3}. \quad (3.3)$$

Similarly, since $I_a(\mathbf{z}, \mathbf{s}) = [s_a, z_b]$, it holds that $z_a = \frac{s_a + z_b}{2} = \frac{3s_a + 2s_b + s_c}{6}$ and, therefore, we obtain that

$$s_b - z_a = \frac{-3s_a + 4s_b - s_c}{6}. \quad (3.4)$$

Since $I_b(\mathbf{z}, \mathbf{s}) = [s_b, z_c]$, we have that $a \notin N_b(\mathbf{z}, \mathbf{s})$ and, subsequently, that $z_c - s_b \leq s_b - z_a$ which, together with (3.3) and (3.4), yields that $s_b \geq \frac{3s_a + 5s_c}{8}$ as desired. \square

The proof of the next theorem is inspired by a construction of Bhawalkar et al. [2013] and exploits Lemma 3.5.

Theorem 3.6. *For any k , there exists a k -COF game with no pure Nash equilibria.*

Proof. Consider a k -COF game with $2k + 1$ players partitioned into three sets called L , M , and R , where L and R each contain k players, while $M = \{m\}$ is a singleton. We set $s_i = 0$ for each $i \in L$, $s_i = 2$ for each $i \in R$, while $s_m = 1 - \epsilon$, where $\epsilon < 1/4$ is an arbitrarily small positive constant.

Let us assume that there exists a pure Nash equilibrium \mathbf{z} . Then, clearly, for any $i \in L$ it must hold that $N_i(\mathbf{z}, \mathbf{s}) = L \setminus \{i\} \cup \{m\}$, and, therefore, $I_i(\mathbf{z}, \mathbf{s}) = [0, z_m]$. Similarly, for any $i \in R$ we have $N_i(\mathbf{z}, \mathbf{s}) = R \setminus \{i\} \cup \{m\}$, and $I_i(\mathbf{z}, \mathbf{s}) = [z_m, 2]$. Now, concerning player m , if all her neighbors are in L , then, it holds that $I_m(\mathbf{z}, \mathbf{s}) = [z_i, s_m]$ for some $i \in L$. But then, observe that even though the intervals defined above exhibit the structure described in Lemma 3.5, the belief vector \mathbf{s} does not satisfy the corresponding necessary conditions of that lemma as $1 - \epsilon > 3/4$; hence, \mathbf{z} is not a pure Nash equilibrium. The same reasoning applies in case all of m 's neighbors are in R .

It remains to consider the case where m has at least one neighbor in each of L and R . By the definition of $I_i(\mathbf{z}, \mathbf{s})$ for $i \in L \cup R$, as stated above, Lemma 3.1 implies that $z_i = z_m/2$ for any $i \in L$, while $z_i = 1 + z_m/2$ for any $i \in R$. Then, Lemma 3.4 implies that $z_m/2 \leq s_m = 1 - \epsilon$ and $1 + z_m/2 \geq s_m$, and, consequently, $I_m(\mathbf{z}, \mathbf{s}) = [z_m/2, 1 + z_m/2]$. Again, by Lemma 3.1 we have that $z_m = \frac{z_m/2 + 1 + z_m/2}{2}$, i.e., $z_m = 1$. But then, we obtain $z_i = 1/2$ for any $i \in L$ and $z_i = 3/2$ for any $i \in R$, which implies that all k players in L are strictly closer to s_m than any player in R ; this contradicts the assumption that m has neighbors in both L and R . \square

An example of the construction used in the proof of Theorem 6 is presented in Figure 3.4.

3.5.2 Price of stability

We will now prove that the price of stability of k -COF games is strictly higher than 1, i.e., there exist games without any efficient pure Nash equilibria (even when they exist). In particular, for any value of k we show that there exist rather simple games with price of stability in $\Omega(k)$.

Theorem 3.7. *The price of stability of k -COF games, for $k \geq 3$, is at least $(k + 1)/3$.*

Proof. Consider a k -COF game with $k + 1$ players, where k of them have belief 0, while the remaining one has belief 1. Let $\tilde{\mathbf{z}}$ be the opinion vector where each player has opinion 0. Clearly, $\text{SC}(\tilde{\mathbf{z}}, \mathbf{s}) = 1$, and, hence the optimal social cost is at most 1.




Figure 3.4: (a) The k -COF game considered in the proof of Theorem 3.6 where the k players of set L have belief 0, player m has $s_m = 1 - \epsilon$ and the k players of set R have belief 2. (b) Lemma 3.5 implies that there is no pure Nash equilibrium where m has neighbors in strictly one of L , R . In the remaining case, it must hold that $x = 1$, but then all players in L are strictly closer to s_m than any player in R .

Now, consider any pure Nash equilibrium \mathbf{z} . Since, there are exactly $k + 1$ players, the neighborhood of each player includes all remaining ones. Let x be the opinion that the player with belief 1 expresses at \mathbf{z} . By Lemma 3.4, we have that $x \in [0, 1]$, and by Lemma 3.1, we have that all remaining players must have opinion $x/2$. Therefore, again by Lemma 3.1, x must satisfy the equation $x = (1 + x/2)/2$, i.e., $x = 2/3$. Therefore, there exists a single pure Nash equilibrium \mathbf{z} where all players with belief 0 have opinion $1/3$ and the single player with belief 1 has opinion $2/3$, and we obtain $\text{SC}(\mathbf{z}) = (k + 1)/3$ which implies the theorem. \square

Clearly, the above result states the inefficiency of the best pure Nash equilibrium only when $k \geq 3$. For the remaining cases, where $k \in \{1, 2\}$, we will present slightly more complicated instances, where the proofs rely on Lemma 3.5. Recall that, for 1-COF games, $\sigma(i)$ denotes the single neighbor of player i .

Theorem 3.8. *The price of stability of 1-COF games is at least $17/15$.*

Proof. We use the following 1-COF game with six players and belief vector

$$\mathbf{s} = (0, 5 - 3\lambda, 8, 15, 18 + 3\lambda, 23),$$

where $\lambda \in (0, 1/4)$.

Consider the opinion vector

$$\tilde{\mathbf{z}} = (3 - \lambda, 6 - 2\lambda, 7 - 6\lambda, 16 + 6\lambda, 17 + 2\lambda, 20 + \lambda).$$

It can be easily seen that it has social cost $SC(\tilde{\mathbf{z}}, \mathbf{s}) = 10 + 12\lambda$. So, clearly, $SC(\mathbf{z}^*, \mathbf{s}) \leq 10 + 12\lambda$ for any optimal opinion vector \mathbf{z}^* .

Now, consider the opinion vector

$$\mathbf{z} = \left(\frac{5 - 3\lambda}{3}, \frac{10 - 6\lambda}{3}, \frac{31}{3}, \frac{38}{3}, \frac{59 + 6\lambda}{3}, \frac{64 + 3\lambda}{3} \right)$$

with social cost $SC(\mathbf{z}, \mathbf{s}) = 34/3 - 4\lambda$. It is not hard to verify (by showing, as Lemma 3.1 requires, that each opinion lies in the middle of its player's interval) that \mathbf{z} is a pure Nash equilibrium; we argue that this equilibrium is unique.

We claim that, by Lemma 3.5, there cannot be a pure Nash equilibrium where both $\sigma(j - 1) = j$ and $\sigma(j + 1) = j$ for any $j \in \{2, 5\}$. To see this, assume otherwise and note that the corresponding intervals satisfy the conditions of the lemma. However, by observing the belief vector \mathbf{s} , it holds that $\frac{5s_{j-1} + 3s_{j+1}}{8} < s_j < \frac{3s_{j-1} + 5s_{j+1}}{8}$, for $j \in \{2, 5\}$, i.e., \mathbf{s} does not satisfy the conditions of Lemma 3.5; this contradicts our original assumption.

The above observation, together with Lemma 3.2, implies that $\sigma(1) = 2$, $\sigma(3) = 4$, $\sigma(4) = 3$ and $\sigma(6) = 5$ in any equilibrium. This leaves only $\sigma(2) \in \{1, 3\}$ and $\sigma(5) \in \{4, 6\}$ undefined.

Consider an equilibrium \mathbf{z}' with $\sigma(2) = 3$; the case $\sigma(5) = 4$ is symmetric. Since $\sigma(3) = 4$, Lemma 3.4 implies that $z'_3 > s_3 = 8$ and, hence

$$z'_3 - s_2 > 3 + 3\lambda. \quad (3.5)$$

Since $\sigma(1) = 2$, $\sigma(2) = 3$ and $z'_1 = \frac{s_1 + z'_2}{2}$, Lemma 3.4 implies that $z'_2 > s_2$ and we obtain that $z'_1 > \frac{5 - 3\lambda}{2}$ and, hence,

$$s_2 - z'_1 < \frac{5 - 3\lambda}{2}. \quad (3.6)$$

By inequalities (3.5) and (3.6), we get $z'_3 - s_2 > s_2 - z'_1$, which contradicts our assumption that $\sigma(2) = 3$. So, it must hold that $\sigma(2) = 1$ (and, respectively, $\sigma(5) = 6$) which implies that \mathbf{z} is the unique pure Nash equilibrium.

We conclude that the price of stability is lower-bounded by

$$\frac{SC(\mathbf{z}, \mathbf{s})}{SC(\mathbf{z}^*, \mathbf{s})} = \frac{34/3 - 4\lambda}{10 + 12\lambda},$$

and the theorem follows by taking λ to be arbitrarily close to 0. □

Theorem 3.9. *The price of stability of 2-COF games is at least $8/7$.*

Proof. Consider a 2-COF game with four players a, b, c , and d , with belief vector $\mathbf{s} = (0, 1, 1, 2)$. Let $\tilde{\mathbf{z}} = (1, 1, 1, 3/2)$ be an opinion vector and observe that $SC(\tilde{\mathbf{z}}, \mathbf{s}) = 3/2$; note that $\tilde{\mathbf{z}}$ is not a pure Nash equilibrium as player a has an incentive to deviate. Clearly, the optimal social cost is at most $3/2$.

Now consider any pure Nash equilibrium \mathbf{z} . By the structural properties of equilibria, $N_a(\mathbf{z}, \mathbf{s}) = N_d(\mathbf{z}, \mathbf{s}) = \{b, c\}$, while $b \in N_c(\mathbf{z}, \mathbf{s})$ and $c \in N_b(\mathbf{z}, \mathbf{s})$. It remains to argue about the second neighbor of b and c . We distinguish between two cases depending on whether b and c have a common second neighbor in $\{a, d\}$ or not.

First, let a be the common neighbor; the case where d is that neighbor is symmetric. By Lemma 3.1, we have that $z_b = z_c = (1 + z_a)/2$. Then, we have that $I_a(\mathbf{z}, \mathbf{s}) = [0, z_b]$, $I_b(\mathbf{z}, \mathbf{s}) = [z_a, 1]$, and $I_d(\mathbf{z}, \mathbf{s}) = [z_b, 2]$. Note that by applying Lemma 3.5 on players a, b , and d , we obtain a contradiction to the fact that \mathbf{z} is a pure Nash equilibrium.

Second, without loss of generality, let $N_b(\mathbf{z}, \mathbf{s}) = \{a, c\}$ and $N_c(\mathbf{z}, \mathbf{s}) = \{b, d\}$ which, by Lemma 3.4, imply that $z_b \in [0, 1]$ and $z_c \in [1, 2]$. Then, Lemma 3.1 yields $z_a = z_c/2$, $z_b = (z_a + z_c)/2$, $z_c = (z_b + z_d)/2$, and $z_d = 1 + z_b/2$. By solving this system of equations, we obtain that $\mathbf{z} = (4/7, 6/7, 8/7, 10/7)$ and, hence, $SC(\mathbf{z}) = 12/7$. \square

3.6 Complexity of equilibria

In this section we focus entirely on 1-COF games. We present a polynomial-time algorithm that determines whether such a game admits pure Nash equilibria, and, in case it does, allows us to compute the best and worst pure Nash equilibrium with respect to the social cost. We do so by establishing a correspondence between pure Nash equilibria and source-sink paths in a suitably defined directed acyclic graph. See Example 3.2 below for an instance execution of the following procedure.

Assume that we are given neighborhood information according to which each player i has either player $i - 1$ or player $i + 1$ as neighbor. From Lemma 3.3, such a neighborhood structure is necessary in a pure Nash equilibrium. We claim that this information is enough in order to decide whether there is a consistent opinion vector that is a pure Nash equilibrium or not. All we have to do is to use Lemma 3.1 and obtain n equations that relate the opinion of each player to her belief and her neighbor's opinion. These equations have a unique solution which can then be verified whether it indeed satisfies the neighborhood conditions or not. So, the main idea of our algorithm is to cleverly search among all possible neighborhood structures that are

not excluded by Lemma 3.3 for one that defines a pure Nash equilibrium.

For integers $1 \leq a \leq b < c \leq n$, let us define the *segment* $C(a, b, c)$ to be the set of players $\{a, a+1, \dots, c\}$ together with the following neighborhood information for them: $\sigma(p) = p+1$ for $p = a, \dots, b$ and $\sigma(p) = p-1$ for $p = b+1, \dots, c$. It can be easily seen that the neighborhood information for all players at a pure Nash equilibrium can always be decomposed into disjoint segments. Importantly, given the neighborhood information in segment $C(a, b, c)$ and the beliefs of its players, the opinions they could have in any pure Nash equilibrium that contains this segment are uniquely defined using Lemma 3.1. In particular, the opinions of the players within a segment $C(a, b, c)$ are computed as follows. First, we set $z_b = s_b + \frac{s_{b+1}-s_b}{3}$ and $z_{b+1} = s_b + \frac{2(s_{b+1}-s_b)}{3}$. Then, we set $z_p = \frac{s_p+z_{p+1}}{2}$ if $a \leq p < b$, and $z_p = \frac{s_p+z_{p-1}}{2}$ if $b < p \leq c$.

We remark that the opinion vector implied by a segment is not necessarily consistent to the given neighborhood structure. So, we call segment $C(a, b, c)$ *legit* if $a \neq 2$, $c \neq n-1$ (so that it can be part of a decomposition) and the uniquely defined opinions are consistent to the neighborhood information of the segment, i.e., if $|z_{\sigma(p)} - s_p| \leq |z_{p'} - s_p|$ for any pair of players p, p' (with $p \neq p'$) in $C(a, b, c)$. This process appears in Algorithm 1.

A decomposition of neighborhood information for all players will consist of consecutive segments $C(a_1, b_1, c_1), C(a_2, b_2, c_2), \dots, C(a_t, b_t, c_t)$ so that $a_1 = 1, c_t = n, a_\ell = c_{\ell-1} + 1$ for $\ell = 2, \dots, t$. Such a decomposition will yield a pure Nash equilibrium if it consists of legit segments and, furthermore, the uniquely defined opinions of players in consecutive segments are consistent to the neighborhood information.

In particular, consider the directed graph G that has two special nodes designated as the source and the sink, and a node for each legit segment $C(a, b, c)$. Note that G has $\mathcal{O}(n^3)$ nodes. The source node is connected to all segment nodes $C(1, b, c)$ while all segment nodes $C(a, b, n)$ are connected to the sink. An edge from segment node $C(a, b, c)$ to segment node $C(a', b', c')$ exists if $a' = c+1$ and the uniquely defined opinions of players in the two segments are consistent to the neighborhood information in both of them. This consistency test has to check

1. whether the leftmost opinion $z_{a'}$ in segment $C(a', b', c')$ is indeed further away from the belief s_c of player c than the opinion z_{c-1} of the designated neighbor of c in segment $C(a, b, c)$, i.e., $|z_{c-1} - s_c| \leq |z_{a'} - s_c|$, and
2. whether the rightmost opinion z_c in segment $C(a, b, c)$ is further away from the belief $s_{a'}$

of player a' than the opinion $z_{a'+1}$ of the designated neighbor of a' in segment $C(a', b', c')$, i.e., $|z_{a'+1} - s_{a'}| \leq |z_c - s_{a'}|$.

By the definition of segments and of its edges, G is acyclic. This process appears in Algorithm 2.

Based on the discussion above, there is a bijection between pure Nash equilibria and source-sink paths in G . In addition, we can assign a weight to each node of G that is equal to the total cost of the players in the corresponding segment, i.e.,

$$\text{weight}(C(a, b, c)) = \sum_{a \leq p \leq c} |z_p - s_p|.$$

Then, the total weight of a source-sink path \mathcal{P} is equal to the social cost of the corresponding pure Nash equilibrium, i.e.,

$$\text{SC}(\mathbf{z}, \mathbf{s}) = \sum_{C(a, b, c) \in \mathcal{P}} \text{weight}(C(a, b, c)).$$

Hence, standard algorithms for computing shortest or longest paths in directed acyclic graphs can be used not only to detect whether a pure Nash equilibrium exists, but also to compute the equilibrium of best or worst social cost.

Theorem 3.10. *Given a 1-COF game, deciding whether a pure Nash equilibrium exists can be done in polynomial time. Furthermore, computing a pure Nash equilibrium of highest or lowest social cost can be done in polynomial time as well.*

Example 3.2. Consider a 1-COF game with four players with belief vector $\mathbf{s} = (0, 9, 12, 21)$. According to the discussion above, there are 10 segments of the form $C(a, b, c)$ with $1 \leq a \leq b < c \leq 4$, but it can be shown that only 3 of them are legit; these are $C(1, 1, 2)$, $C(3, 3, 4)$ (see Figure 3.5a), and $C(1, 2, 4)$ (see Figure 3.5b). For example, segment $C(1, 1, 4)$, in which $\sigma(1) = 2$, $\sigma(2) = 1$, $\sigma(3) = 2$, and $\sigma(4) = 3$, corresponds to the opinion vector $(3, 6, 9, 15)$. This is not consistent to the neighborhood information $\sigma(2) = 1$ in the segment, as the belief of player 2 coincides with the opinion of player 3, while the opinion of player 1 is further away. The resulting directed acyclic graph G (see Figure 3.5c) implies that there exist two pure Nash equilibria for this 1-COF game, namely the opinion vectors $(3, 6, 15, 18)$ and $(5, 10, 11, 16)$. \square

3.7 Upper bounds on the price of anarchy

In this section we prove upper bounds on the price of anarchy of general k -COF games. In our proof, we relate the social cost of any deterministic opinion vector, including optimal ones, to a

Algorithm 1: Segment

Input: belief vector $\mathbf{s} = (s_1, \dots, s_n)$, parameters a, b , and c such that $a \leq b < c$

Output: opinion vector $\mathbf{z}_{a:c} = (z_a, \dots, z_c)$, segment weight, legit indicator

legit $\leftarrow 0$

if $a \neq 2$ or $c \neq n - 1$ **then**

 legit $\leftarrow 1$

$z_b \leftarrow s_b + \frac{1}{3}(s_{b+1} - s_b)$

$z_{b+1} \leftarrow s_b + \frac{2}{3}(s_{b+1} - s_b)$

for $p := b - 1$ **downto** a **do**

$z_p \leftarrow \frac{1}{2}(s_p + z_{p+1})$

end

for $p := b + 2$ **to** c **do**

$z_p \leftarrow \frac{1}{2}(s_p + z_{p-1})$

end

for $p := a + 1$ **to** b **do**

if $|z_{p-1} - s_p| < |z_{p+1} - s_p|$ **then**

 legit $\leftarrow 0$

end

end

for $p := b + 1$ **to** $c - 1$ **do**

if $|z_{p+1} - s_p| < |z_{p-1} - s_p|$ **then**

 legit $\leftarrow 0$

end

end

 weight $:= 0$

for $p := a$ **to** c **do**

 weight \leftarrow weight $+$ $|z_p - s_p|$

end

end

return $[\mathbf{z}_{a:c}, \text{weight}, \text{legit}]$

Algorithm 2: ConstructGraph

Input: belief vector $\mathbf{s} = (s_1, \dots, s_n)$

Output: a node-weighted directed acyclic graph G

$V \leftarrow \emptyset$

for $a \leftarrow 1$ to $n - 1$ **do**

for $b \leftarrow a$ to $n - 1$ **do**

for $c \leftarrow b + 1$ to n **do**

$[\mathbf{z}_{a:c}, \text{weight}, \text{legit}] \leftarrow \text{Segment}(\mathbf{s}, a, b, c)$

if legit = 1 **then**

$C.a \leftarrow a, C.b \leftarrow b, C.c \leftarrow c, C.\mathbf{z}_{a:c} \leftarrow \mathbf{z}_{a:c}, C.\text{weight} \leftarrow \text{weight}$

$V \leftarrow V \cup C$

end

end

end

end

$V \leftarrow V \cup \{\text{source}, \text{sink}\}$

$E \leftarrow \emptyset$

for $C \in V$ **do**

if $C.a = 1$ **then**

$E \leftarrow E \cup (\text{source}, C)$

else if $C.c = n$ **then**

$E \leftarrow E \cup (C, \text{sink})$

end

end

for all segment pairs (C, D) such that $D.a = C.c + 1$ **do**

if $|C.z_{c-1} - s_{C.c}| \leq |D.z_a - s_{C.c}|$ and $|D.z_{a+1} - s_{D.a}| \leq |C.z_c - s_{D.a}|$ **then**

$E \leftarrow E \cup (C, D)$

end

end

return $G = (V, E)$

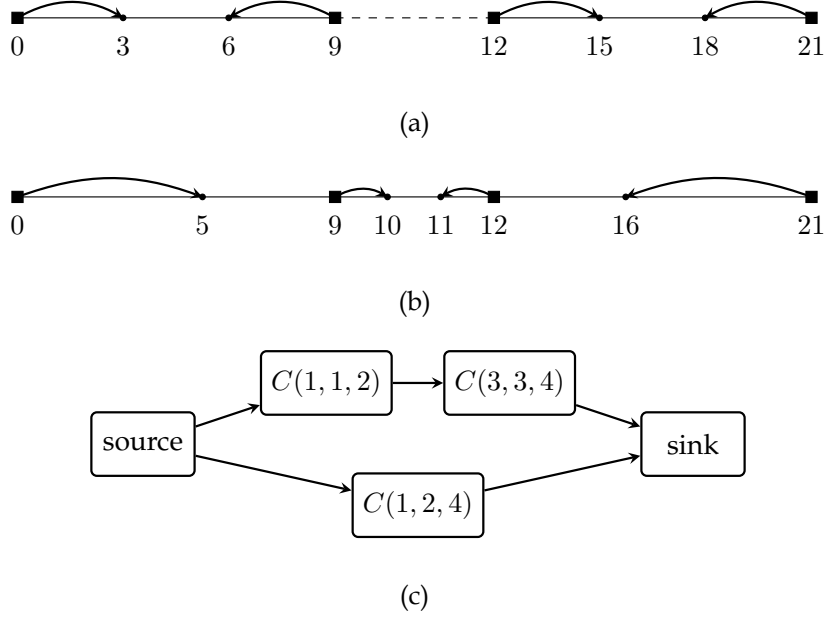


Figure 3.5: The 1-COF game considered in Example 3.2. (a) The legit segments $C(1, 1, 2)$ and $C(3, 3, 4)$ which imply the opinion vector $(3, 6, 15, 18)$. (b) The legit segment $C(1, 2, 4)$ which implies the opinion vector $(5, 10, 11, 16)$. (c) The directed acyclic graph G which shows that there exist two pure Nash equilibria in the game.

quantity that depends only on the beliefs of the players and can be thought of as the cost of the truthful opinion vector (in which the opinion of every player is equal to her belief). In particular, we prove a lower bound on the optimal social cost (in Lemmas 3.11) and an upper bound on the social cost of any pure Nash equilibrium, both expressed in terms of this quantity. The bound on the price of anarchy then follows by these relations; see the proof of Theorem 3.12.

Consider an n -player k -COF game with belief vector $\mathbf{s} = (s_1, \dots, s_n)$. For player i , we denote by $\ell^*(i)$ and $r^*(i)$ the integers in $[n]$ such that $\ell^*(i) \leq i \leq r^*(i)$, $r^*(i) - \ell^*(i) = k$, and $|s_{r^*(i)} - s_{\ell^*(i)}|$ is minimized. The proof of the next lemma exploits linear programming and duality.

Lemma 3.11. *Consider a k -COF game with belief vector $\mathbf{s} = (s_1, \dots, s_n)$ and let \mathbf{z} be any deterministic opinion vector. Then,*

$$\text{SC}(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2(k+1)} \sum_{i=1}^n |s_{r^*(i)} - s_{\ell^*(i)}|.$$

Proof. Consider any deterministic opinion vector \mathbf{z} and let π be a permutation of the players so that $z_{\pi(j)} \leq z_{\pi(j+1)}$ for each $j \in [n-1]$. We refer to player $\pi(j)$ as the player with rank j .¹ For each player i , we will identify an effective neighborhood $F_i(\mathbf{z}, \mathbf{s})$ that consists of $k+1$

¹Note that we have proved monotonicity of opinions for pure Nash equilibria only (Lemma 3.2) and it is not clear




Figure 3.6: An example of the quantities used in the proof of Lemma 3.11. Let $k = 2$ and $i = 4$. Then, the neighborhood of player 4 is $N_4(\mathbf{z}, \mathbf{s}) = \{2, 6\}$, the smallest contiguous interval containing the opinions of players in $N_4(\mathbf{z}, \mathbf{s}) \cup \{4\}$ is $J_4(\mathbf{z}, \mathbf{s}) = [z_4, z_6]$, the set of players with opinions in $J_4(\mathbf{z}, \mathbf{s})$ is $D_4(\mathbf{z}, \mathbf{s}) = \{1, 2, 3, 4, 6\}$, the effective neighborhood is $F_4(\mathbf{z}, \mathbf{s}) = \{1, 3, 4\}$, and, hence, $\tilde{\ell}(4) = 1$, and $\tilde{r}(4) = 4$.

players with consecutive ranks and includes player i . Define $\tilde{\ell}(i)$ and $\tilde{r}(i)$ to be the players in $F_i(\mathbf{z}, \mathbf{s})$ with the lowest and highest belief, respectively. In the extreme case where all players in $F_i(\mathbf{z}, \mathbf{s})$ have the same belief, we let $\tilde{\ell}(i)$ and $\tilde{r}(i)$ be the players with the lowest and highest ranks, respectively. The effective neighborhood will be defined in such a way that it satisfies the properties $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq z_{\tilde{r}(i)} - z_i$ and $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq z_i - z_{\tilde{\ell}(i)}$.

Let $N_i(\mathbf{z}, \mathbf{s})$ denote the neighborhood of player i , i.e., the set of players (not including i) with the k closest opinions to the belief s_i of player i . Let $J_i(\mathbf{z}, \mathbf{s})$ be the smallest contiguous interval containing all opinions of players in $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$ and let $D_i(\mathbf{z}, \mathbf{s})$ be the set of players with opinions in $J_i(\mathbf{z}, \mathbf{s})$. Clearly, $|D_i(\mathbf{z}, \mathbf{s})| \geq k + 1$. We define $F_i(\mathbf{z}, \mathbf{s})$ to be a subset of $D_i(\mathbf{z}, \mathbf{s})$ that consists of $k + 1$ players with consecutive ranks including player i . See Figure 3.6 for an illustrative example of all quantities defined above.

Let $\ell'(i)$ and $r'(i)$ be the players in $N_i(\mathbf{z}, \mathbf{s})$ with the leftmost and rightmost opinion. In order to show that the definition of $F_i(\mathbf{z}, \mathbf{s})$ satisfies the two desired properties, we distinguish between three different cases depending on the location of opinion z_i among the players in $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$.

- **Case I:** Player i has neither the leftmost nor the rightmost opinion in $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$, i.e., $z_{\ell'(i)} < z_i < z_{r'(i)}$.² In this case, $J_i(\mathbf{z}, \mathbf{s}) = [z_{\ell'(i)}, z_{r'(i)}]$. Then, the definition of $N_i(\mathbf{z}, \mathbf{s})$ implies that $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq z_{r'(i)} - z_i$ and $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq z_i - z_{\ell'(i)}$. Hence, $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq |z_j - z_i|$ for every $z_j \in J_i(\mathbf{z}, \mathbf{s})$ or, equivalently, $j \in D_i(\mathbf{z}, \mathbf{s})$ and, subsequently, for each $j \in$

whether such a monotonicity property holds for opinion vectors of minimum social cost. In addition, the statement of Lemma 3.11 refers to any opinion vector. This clearly includes non-monotonic ones, so we need to rank players in terms of opinions in the proof.

²Case I cannot appear when $k = 1$.

$F_i(\mathbf{z}, \mathbf{s})$. This implies the two desired properties $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq z_{\tilde{r}(i)} - z_i$ and $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq z_i - z_{\tilde{\ell}(i)}$.

- **Case II:** Player i has the leftmost opinion in $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$, i.e., $z_i \leq z_{\ell'(i)}$. Then, $J_i(\mathbf{z}, \mathbf{s}) = [z_i, z_{r'(i)}]$. Now, the definition of $N_i(\mathbf{z}, \mathbf{s})$ implies that $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq z_{r'(i)} - z_i$ and, hence, $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq |z_j - z_i|$ for every $z_j \in J_i(\mathbf{z}, \mathbf{s})$ or, equivalently, $j \in D_i(\mathbf{z}, \mathbf{s})$ and, subsequently, for each $j \in F_i(\mathbf{z}, \mathbf{s})$. Again, this implies the two desired properties.
- **Case III:** Player i has the rightmost opinion in $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$, i.e., $z_i \geq z_{r'(i)}$. Then, $J_i(\mathbf{z}, \mathbf{s}) = [z_{\ell'(i)}, z_i]$. Now, the definition of $N_i(\mathbf{z}, \mathbf{s})$ implies that $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq z_i - z_{\ell'(i)}$ and, hence, $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq |z_j - z_i|$ for every $z_j \in J_i(\mathbf{z}, \mathbf{s})$ or, equivalently, $j \in D_i(\mathbf{z}, \mathbf{s})$ and, subsequently, for every $j \in F_i(\mathbf{z}, \mathbf{s})$. Again, the two desired properties follow.

By setting the variable t_i equal to $\text{cost}_i(\mathbf{z}, \mathbf{s})$ for $i \in [n]$, the discussion above and the fact that $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq |s_i - z_i|$ imply that the opinion vector \mathbf{z} together with $\mathbf{t} = (t_1, \dots, t_n)$ is a feasible solution to the following linear program:

$$\begin{aligned}
& \text{minimize} && \sum_{i \in [n]} t_i \\
& \text{subject to} && t_i + z_i \geq s_i, \forall i \in [n] \\
& && t_i - z_i \geq -s_i, \forall i \in [n] \\
& && t_i + z_i - z_{\tilde{r}(i)} \geq 0, \forall i \in [n] \text{ such that } \tilde{r}(i) \neq i \\
& && t_i + z_{\tilde{\ell}(i)} - z_i \geq 0, \forall i \in [n] \text{ such that } \tilde{\ell}(i) \neq i \\
& && t_i, z_i \geq 0, \forall i \in [n]
\end{aligned}$$

Using the dual variables $\alpha_i, \beta_i, \gamma_i$, and δ_i associated with the four constraints of the above LP, we obtain its dual LP:

$$\begin{aligned}
& \text{maximize} && \sum_{i \in [n]} s_i \alpha_i - \sum_{i \in [n]} s_i \beta_i \\
& \text{subject to} && \alpha_i + \beta_i + \gamma_i \cdot \mathbb{1}_{\tilde{r}(i) \neq i} + \delta_i \cdot \mathbb{1}_{\tilde{\ell}(i) \neq i} \leq 1, \forall i \in [n] \\
& && \alpha_i - \beta_i + \gamma_i \cdot \mathbb{1}_{\tilde{r}(i) \neq i} - \delta_i \cdot \mathbb{1}_{\tilde{\ell}(i) \neq i} - \sum_{j \neq i: \tilde{r}(j)=i} \gamma_j + \sum_{j \neq i: \tilde{\ell}(j)=i} \delta_j \leq 0, \forall i \in [n] \\
& && \alpha_i, \beta_i, \gamma_i, \delta_i \geq 0
\end{aligned}$$

The indicator $\mathbb{1}X$ is equal to 1 when the condition X is true, and 0 otherwise. We will show that the solution defined as

$$\alpha_i = \frac{|\{j \in [n] : \tilde{r}(j) = i\}|}{2(k+1)},$$

$$\beta_i = \frac{|\{j \in [n] : \tilde{\ell}(j) = i\}|}{2(k+1)},$$

$$\gamma_i = \delta_i = \frac{1}{2(k+1)},$$

is a feasible dual solution. Indeed, to see why the first dual constraint is satisfied, first observe that player i belongs to at most $2k+1$ different effective neighborhoods. Hence, player i can have the lowest or highest belief among the players in the effective neighborhood of at most $2k+1$ players (implying that $\alpha_i + \beta_i \leq 1 - \frac{1}{2(k+1)}$) when $\tilde{r}(i) = i$ or $\tilde{\ell}(i) = i$ and of at most $2k$ players (implying that $\alpha_i + \beta_i \leq 1 - \frac{1}{k+1}$) when $\tilde{r}(i) \neq i$ and $\tilde{\ell}(i) \neq i$. The first constraint follows.

It remains to show that the second constraint is satisfied as well (with equality). We do so by distinguishing between three cases:

- When $\tilde{r}(i) \neq i$ and $\tilde{\ell}(i) \neq i$, the dual solution guarantees that $\alpha_i = \sum_{j \neq i: \tilde{r}(j)=i} \gamma_j$ and the term α_i in the left-hand side of the second constraint cancels out with the sum of γ 's. Similarly, $\beta_i = \sum_{j \neq i: \tilde{\ell}(j)=i} \delta_j$ and the term β_i cancels out with the sum of δ 's. Also, the terms γ_i and δ_i are both equal to $\frac{1}{2(k+1)}$ and cancel out as well.
- When $\tilde{r}(i) = i$ (then, clearly, $\tilde{\ell}(i) \neq i$), we have that $\alpha_i = \delta_i \cdot \mathbb{1}_{\tilde{\ell}(i) \neq i} + \sum_{j \neq i: \tilde{r}(j)=i} \gamma_j$ (cancelling out the first, fourth and fifth terms) and $\beta_i = \sum_{j \neq i: \tilde{\ell}(j)=i} \delta_j$ (cancelling out the second and sixth terms), and the second constraint is satisfied with equality as the third term is zero.
- Finally, when $\tilde{\ell}(i) = i$ (now, it is $\tilde{r}(i) \neq i$), we have that $\alpha_i = \sum_{j \neq i: \tilde{r}(j)=i} \gamma_j$ (cancelling out the first and fifth terms) and $\beta_i = \gamma_i \cdot \mathbb{1}_{\tilde{r}(i) \neq i} + \sum_{j \neq i: \tilde{\ell}(j)=i} \delta_j$ (cancelling out the second, third and sixth terms), and the second constraint is satisfied with equality as the fourth term is zero.

So, the social cost of the solution \mathbf{z} is lower-bounded by the objective value of the primal LP which, by duality, is lower-bounded by the objective value of the dual LP. Hence

$$\begin{aligned} \text{SC}(\mathbf{z}, \mathbf{s}) &\geq \sum_{i \in [n]} s_i \alpha_i - \sum_{i \in [n]} s_i \beta_i \\ &= \frac{1}{2(k+1)} \left(\sum_{i \in [n]} |\{j \in [n] : \tilde{r}(j) = i\}| s_i - \sum_{i \in [n]} |\{j \in [n] : \tilde{\ell}(j) = i\}| s_i \right) \\ &= \frac{1}{2(k+1)} \sum_{i \in [n]} (s_{\tilde{r}(i)} - s_{\tilde{\ell}(i)}) \end{aligned}$$

$$= \frac{1}{2(k+1)} \sum_{i \in [n]} |s_{\tilde{r}(i)} - s_{\tilde{\ell}(i)}|.$$

The last equality follows since $s_{\tilde{r}(i)} \geq s_{\tilde{\ell}(i)}$ by the definition of $\tilde{r}(i)$ and $\tilde{\ell}(i)$.

Note that for each player i , there are at least $k+1$ beliefs of different players with values in $[s_{\tilde{\ell}(i)}, s_{\tilde{r}(i)}]$, including player i . By the definition of $\ell^*(i)$ and $r^*(i)$ for each player i , the above inequality yields

$$\text{SC}(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2(k+1)} \sum_{i \in [n]} |s_{r^*(i)} - s_{\ell^*(i)}|,$$

as desired. \square

We are now ready to prove our upper bound on the price of anarchy for k -COF games. In our proof, we exploit the monotonicity of opinions in a pure Nash equilibrium and we associate the cost of each player in the equilibrium to the same quantity used in the statement of Lemma 3.11.

Theorem 3.12. *The price of anarchy of k -COF games over pure Nash equilibria is at most $4(k+1)$.*

Proof. Consider a k -COF game with belief vector $\mathbf{s} = (s_1, \dots, s_n)$, and let $\mathbf{z}^* = (z_1^*, \dots, z_n^*)$ be any opinion vector that minimizes the social cost. By Lemma 3.11, we have

$$\text{SC}(\mathbf{z}^*, \mathbf{s}) \geq \frac{1}{2(k+1)} \sum_{i=1}^n |s_{r^*(i)} - s_{\ell^*(i)}|. \quad (3.7)$$

Now, consider any pure Nash equilibrium \mathbf{z} of the game. We will show that

$$\text{SC}(\mathbf{z}, \mathbf{s}) \leq 2 \sum_{i=1}^n |s_{r^*(i)} - s_{\ell^*(i)}|, \quad (3.8)$$

and the theorem will then follow by inequalities (3.7) and (3.8).

The rest of this proof is, therefore, devoted to showing inequality (3.8). To this end, we will show that, for any player i , we have $\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq 2(s_{r^*(i)} - s_{\ell^*(i)})$. Then, inequality (3.8) will follow by summing over all players.

Consider an arbitrary player i and, without loss of generality, let us assume that $z_i \geq s_i$ (the case $z_i \leq s_i$ is symmetric). Recall that $\ell(i)$ and $r(i)$ denote the players in $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$ with the leftmost and rightmost point, respectively, in $I_i(\mathbf{z}, \mathbf{s})$ and note that $r(i) - \ell(i) = k$. First, observe that if $z_{r(i)} = z_i$, the assumption $z_i \geq s_i$ implies that all players in $N_i(\mathbf{z}, \mathbf{s}) \cup \{i\}$ have opinions at s_i (since, by Lemma 3.1, z_i is in the middle of interval $I_i(\mathbf{z}, \mathbf{s})$ at equilibrium). In this case,

$\text{cost}_i(\mathbf{z}, \mathbf{s}) = 0$ and the desired inequality holds trivially. So, in the following, we assume that $r(i) > i$ and $z_{r(i)} > z_i$, i.e., $z_{r(i)}$ is at the right of z_i which in turn is at the right of (or coincides with) s_i .

Recall that, for player i , $\ell^*(i)$ and $r^*(i)$ denote the integers in $[n]$ such that $\ell^*(i) \leq i \leq r^*(i)$, $r^*(i) - \ell^*(i) = k$, and $|s_{r^*(i)} - s_{\ell^*(i)}|$ is minimized. Since $r(i) - \ell(i) = r^*(i) - \ell^*(i) = k$, we distinguish between two main cases depending on the relative order of $r(i)$ and $r^*(i)$.

Case I. $r(i) > r^*(i)$ and $\ell(i) > \ell^*(i)$. Since $z_{r(i)}$ is at the right of s_i and $\ell^*(i)$ does not belong to the neighborhood of player i (while player $r(i)$ does so by definition), $z_{\ell^*(i)}$ is at the left of s_i and, furthermore, $z_{r(i)} - s_i \leq s_i - z_{\ell^*(i)}$ or, equivalently,

$$z_{r(i)} \leq 2s_i - z_{\ell^*(i)}. \quad (3.9)$$

This yields

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) = z_{r(i)} - z_i \leq 2s_i - z_{\ell^*(i)} - z_i. \quad (3.10)$$

These inequalities will be useful in several places of the proof for this case below.

If $z_{\ell^*(i)} \geq s_{\ell^*(i)}$ then, since $r^*(i) \geq i$ and $z_i \geq s_i$, inequality (3.10) becomes $\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq s_i - s_{\ell^*(i)} \leq s_{r^*(i)} - s_{\ell^*(i)}$ and the desired inequality follows. So, in the following, we assume that $z_{\ell^*(i)} < s_{\ell^*(i)}$ i.e., $z_{\ell^*(i)}$ is (strictly) at the left of $s_{\ell^*(i)}$. Hence, $\ell^*(i)$ has her leftmost neighbor with $z_{\ell(\ell^*(i))} < z_{\ell^*(i)}$ and, by Lemma 3.1,

$$z_{\ell^*(i)} = \frac{z_{\ell(\ell^*(i))} + \max\{s_{\ell^*(i)}, z_{r(\ell^*(i))}\}}{2}. \quad (3.11)$$

Since $r^*(i) - \ell^*(i) = k$ and $\ell(\ell^*(i)) < \ell^*(i)$, we have $r^*(i) - \ell(\ell^*(i)) > k$, and, therefore, $r^*(i)$ does not belong to the neighborhood of $\ell^*(i)$. Hence, $s_{\ell^*(i)} - z_{\ell(\ell^*(i))} \leq z_{r^*(i)} - s_{\ell^*(i)}$ or, equivalently

$$z_{\ell(\ell^*(i))} \geq 2s_{\ell^*(i)} - z_{r^*(i)} \geq 2s_{\ell^*(i)} - 2s_i + z_{\ell^*(i)}, \quad (3.12)$$

where the second inequality follows by our case assumption $z_{r^*(i)} \leq z_{r(i)}$ and inequality (3.9).

We now further distinguish between two cases, depending on whether player i belongs to the neighborhood of player $\ell^*(i)$ or not.

Case I.1. $i \in N_{\ell^*(i)}(\mathbf{z}, \mathbf{s})$; see also Figure 3.7a for an example of this case. Then, we have $z_i \leq z_{r(\ell^*(i))}$ and, subsequently,

$$\max\{s_{\ell^*(i)}, z_{r(\ell^*(i))}\} \geq z_{r(\ell^*(i))} \geq z_i. \quad (3.13)$$

Using inequalities (3.12) and (3.13), (3.11) yields

$$z_{\ell^*(i)} \geq s_{\ell^*(i)} - s_i + \frac{z_{\ell^*(i)}}{2} + \frac{z_i}{2},$$

which implies that $z_{\ell^*(i)} \geq 2s_{\ell^*(i)} - 2s_i + z_i$. Now, inequality (3.10) becomes

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq 4s_i - 2s_{\ell^*(i)} - 2z_i \leq 2s_i - 2s_{\ell^*(i)} \leq 2(s_{r^*(i)} - s_{\ell^*(i)})$$

as desired. The second inequality follows since $z_i \geq s_i$ and the last one follows since $r^*(i) \geq i$.

Case I.2. $i \notin N_{\ell^*(i)}(\mathbf{z}, \mathbf{s})$; see also Figure 3.7b for an example. Then, we have $s_{\ell^*(i)} - z_{\ell(\ell^*(i))} \leq z_i - s_{\ell^*(i)}$, which implies that $z_{\ell(\ell^*(i))} \geq 2s_{\ell^*(i)} - z_i$. Using this inequality together with the fact that $\max\{s_{\ell^*(i)}, z_{r(\ell^*(i))}\} \geq s_{\ell^*(i)}$, (3.11) yields

$$z_{\ell^*(i)} \geq \frac{3s_{\ell^*(i)} - z_i}{2}$$

and inequality (3.10) becomes

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq 2s_i - \frac{3}{2}s_{\ell^*(i)} - \frac{z_i}{2} \leq \frac{3}{2}s_i - \frac{3}{2}s_{\ell^*(i)} \leq 2(s_{r^*(i)} - s_{\ell^*(i)}),$$

as desired. The second last inequality follows since $z_i \geq s_i$ and the last one follows since $r^*(i) \geq i$.

Case II. $r(i) \leq r^*(i)$ and $\ell(i) \leq \ell^*(i)$. Since z_i is in the middle of the interval $I_i(\mathbf{z}, \mathbf{s})$ and $z_{r(i)}$ is the rightmost opinion in $I_i(\mathbf{z}, \mathbf{s})$, we have


$$z_i = \frac{\min\{s_i, z_{\ell(i)}\} + z_{r(i)}}{2} \leq \frac{z_{\ell(i)} + z_{r(i)}}{2} \leq \frac{z_{\ell^*(i)} + z_{r^*(i)}}{2}.$$

Since $s_i \leq z_i$, the last inequality yields


$$z_{\ell^*(i)} \geq 2s_i - z_{r^*(i)}. \quad (3.14)$$

We also have


$$\text{cost}_i(\mathbf{z}, \mathbf{s}) = z_{r(i)} - z_i \leq z_{r^*(i)} - z_i. \quad (3.15)$$




(a)



(b)



(c)



(d)

Figure 3.7: Indicative examples of the different cases in the proof of Theorem 3.12. Subfigures (a) and (b) concern Case I, as $r(i) > r^*(i)$ and $\ell(i) > \ell^*(i)$, while subfigures (c) and (d) fall under Case II, as $r(i) \leq r^*(i)$ and $\ell(i) \leq \ell^*(i)$.

If $z_{r^*(i)} \leq s_{r^*(i)}$ then, since $s_{\ell^*(i)} \leq s_i \leq z_i$, inequality (3.15) yields $\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq s_{r^*(i)} - s_i \leq s_{r^*(i)} - s_{\ell^*(i)}$, which is even stronger than the desired inequality. So, in the following we assume that $z_{r^*(i)} > s_{r^*(i)}$, i.e., $z_{r^*(i)}$ is at the right of $s_{r^*(i)}$. Since $z_{r^*(i)}$ is in the middle of the interval $I_{r^*(i)}(\mathbf{z}, \mathbf{s})$, we have that $r(r^*(i)) > r^*(i)$ and, therefore,

$$z_{r^*(i)} = \frac{\min\{s_{r^*(i)}, z_{\ell(r^*(i))}\} + z_{r(r^*(i))}}{2}. \quad (3.16)$$

Moreover, since $r(r^*(i)) - \ell^*(i) > r^*(i) - \ell^*(i) = k$, player $\ell^*(i)$ does not belong to the neighborhood of player $r^*(i)$. Hence, $z_{r(r^*(i))} - s_{r^*(i)} \leq s_{r^*(i)} - z_{\ell^*(i)}$ which, together with inequality (3.14), yields that

$$z_{r(r^*(i))} \leq 2s_{r^*(i)} - z_{\ell^*(i)} \leq 2s_{r^*(i)} - 2s_i + z_{r^*(i)}. \quad (3.17)$$

We now further distinguish between two cases, depending on whether player i belongs to the neighborhood of player $r^*(i)$ or not.

Case II.1. $i \in N_{r^*(i)}(\mathbf{z}, \mathbf{s})$; see also Figure 3.7c for an example. Then, using the fact that $\min\{s_{r^*(i)}, z_{\ell(r^*(i))}\} \leq z_{\ell(r^*(i))} \leq z_i$ and inequality (3.17), equation (3.16) becomes

$$z_{r^*(i)} \leq \frac{z_i + 2s_{r^*(i)} - 2s_i + z_{r^*(i)}}{2}$$

and, equivalently, $z_{r^*(i)} \leq z_i + 2s_{r^*(i)} - 2s_i$. Hence, inequality (3.15) yields

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq 2s_{r^*(i)} - 2s_i \leq 2(s_{r^*(i)} - s_{\ell^*(i)}),$$

as desired. The last inequality follows since $\ell^*(i) \leq i$.

Case II.2. $i \notin N_{r^*(i)}(\mathbf{z}, \mathbf{s})$; see Figure 3.7d for an example. Since i does not belong to the neighborhood of player $r^*(i)$ but player $r(r^*(i))$ does, we have that $z_{r(r^*(i))} - s_{r^*(i)} \leq s_{r^*(i)} - z_i$ or, equivalently, $z_{r(r^*(i))} \leq 2s_{r^*(i)} - z_i$. Then, using also the fact that $\min\{s_{r^*(i)}, z_{\ell(r^*(i))}\} \leq s_{r^*(i)}$, equation (3.16) becomes

$$z_{r^*(i)} \leq \frac{3s_{r^*(i)} - z_i}{2}$$

and (3.15) yields

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq \frac{3}{2}(s_{r^*(i)} - z_i) \leq \frac{3}{2}(s_{r^*(i)} - s_{\ell^*(i)}),$$

which is even stronger than the desired inequality. The last inequality follows since $z_i \geq s_i$ and $\ell^*(i) \leq i$.

So, we have shown that in the pure Nash equilibrium \mathbf{z} and for any player i , we have that $\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq 2(s_{r^*(i)} - s_{\ell^*(i)})$. By summing over all players, we obtain inequality (3.8) and the theorem follows. \square

3.8 An improved bound on the price of anarchy for 1-COF games

For the case of 1-COF games we can prove an even stronger statement following a similar proof roadmap as in the previous section, but using simpler (and shorter) arguments. We denote by $\eta(i)$ the player (other than i) that minimizes the distance $|s_i - s_{\eta(i)}|$; note that $\eta(i) \in \{i-1, i+1\}$. The proof of the next lemma (which can be thought of as a stronger version of Lemma 3.11 for 1-COF games) relies on a particular charging scheme that allows us to lower-bound the cost of each player in any deterministic opinion vector.

Lemma 3.13. *Consider a 1-COF game with belief vector $\mathbf{s} = (s_1, \dots, s_n)$ and let \mathbf{z} be any deterministic opinion vector. Then,*

$$\text{SC}(\mathbf{z}, \mathbf{s}) \geq \frac{1}{3} \sum_{i=1}^n |s_i - s_{\eta(i)}|.$$

Proof. We begin by classifying the players into groups and, subsequently, we show how the costs of different groups can be combined so that the lemma holds. We call a player i with $z_i \notin [s_{i-1}, s_{i+1}]$ a *kangaroo* player and associate the quantity excess_i with her. If $z_i \in [s_j, s_{j+1}]$ for some $j > i$, we say that the players in the set $C_i = \{i+1, \dots, j\}$ are *covered* by player i and define $\text{excess}_i = z_i - s_j$. Otherwise, if $z_i \in [s_{j-1}, s_j]$ for some $j < i$, we say that the players in the set $C_i = \{j, \dots, i-1\}$ are covered by player i and define $\text{excess}_i = s_j - z_i$.

Let \mathcal{K} be the set of kangaroo players and \mathcal{C} the set of players that are covered by a kangaroo; these need not be disjoint. We now partition the players not in $\mathcal{K} \cup \mathcal{C}$ into the set L of *large* players such that, for any $i \in L$, it holds $\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq \frac{1}{3}(|s_i - s_{\eta(i)}|)$, and the set S that contains the remaining players who we call *small*. See also Figure 3.8 for an example of these sets.

We proceed to prove five useful properties (Claims 3.14–3.18); recall that $\sigma(i)$ denotes the single neighbor of player i .

Claim 3.14. *Let $i \in \mathcal{K}$. Then, $\text{cost}_i(\mathbf{z}, \mathbf{s}) - \text{excess}_i \geq \frac{1}{3}(|s_i - s_{\eta(i)}| + \sum_{j \in C_i} |s_j - s_{\eta(j)}|)$.*

Proof. We assume that $z_i > s_i$ (the other case is symmetric). Let ℓ be the player with the rightmost belief that is covered by i . Then, $\text{excess}_i = z_i - s_\ell$. We have

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) - \text{excess}_i = \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} - (z_i - s_\ell)$$




Figure 3.8: An example with kangaroos, covered, large, and small players. In particular, $1 \in \mathcal{K}$ as $z_1 \notin [s_1, s_2]$, $2 \in \mathcal{K} \cap \mathcal{C}$ as she is covered by player 1 and, in addition, $z_2 \notin [s_1, s_3]$. Similarly, $3 \in \mathcal{C}$ as she is covered by player 2, while 4 and 5 are neither kangaroo nor covered. Since $\text{cost}_4(\mathbf{z}, \mathbf{s}) < \frac{1}{3}(s_4 - s_3)$, it is $4 \in S$, while, since $\text{cost}_5(\mathbf{z}, \mathbf{s}) \geq \frac{1}{3}(s_5 - s_4)$, we have $5 \in L$.

$$\begin{aligned} &\geq s_\ell - s_i = \sum_{j=i}^{\ell-1} (s_{j+1} - s_j) \\ &\geq \frac{1}{3}(|s_i - s_{\eta(i)}| + \sum_{j \in C_i} |s_j - s_{\eta(j)}|) \end{aligned}$$

as desired. \square

Claim 3.15. *Let $i \in S$ such that $\sigma(i) \in \mathcal{K}$. Then, $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{excess}_{\sigma(i)} \geq \frac{1}{3}|s_i - s_{\eta(i)}|$.*

Proof. We assume that $\sigma(i) > i$ (the other case is symmetric). If $z_{\sigma(i)} > s_{\sigma(i)}$, then

$$\begin{aligned} \text{cost}_i(\mathbf{z}, \mathbf{s}) &= \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} \\ &\geq \frac{1}{2}(z_{\sigma(i)} - s_i) > \frac{1}{2}(s_{\sigma(i)} - s_i) \\ &\geq \frac{1}{3}|s_i - s_{\eta(i)}|, \end{aligned}$$

which contradicts the fact that i is a small player. Hence, $z_{\sigma(i)} \in [s_i, s_{\sigma(i)}]$, otherwise player i would be covered. Let j be the player with the leftmost belief that is covered by player $\sigma(i)$. Then, $\text{excess}_{\sigma(i)} = s_j - z_{\sigma(i)}$. We have

$$\begin{aligned} \text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{excess}_{\sigma(i)} &= \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} + s_j - z_{\sigma(i)} \\ &\geq \frac{1}{2}(z_{\sigma(i)} - s_i) + \frac{1}{2}(s_j - z_{\sigma(i)}) = \frac{1}{2}(s_j - s_i) \\ &\geq \frac{1}{3}|s_i - s_{\eta(i)}| \end{aligned}$$

as desired. \square

Claim 3.16. *Let $i \in S$ such that $\sigma(i) \in L$ or $\sigma(i) \in \mathcal{C} \setminus \mathcal{K}$. Then, $\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{\sigma(i)}(\mathbf{z}, \mathbf{s}) \geq \frac{1}{3}(|s_i - s_{\eta(i)}| + |s_{\sigma(i)} - s_{\eta(\sigma(i))}|)$.*

Proof. We assume that $\sigma(i) > i$ (the other case is symmetric). If $z_{\sigma(i)} > s_{\sigma(i)}$, then

$$\begin{aligned} \text{cost}_i(\mathbf{z}, \mathbf{s}) &= \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} \\ &\geq \frac{1}{2}(z_{\sigma(i)} - s_i) > \frac{1}{2}(s_{\sigma(i)} - s_i) \\ &\geq \frac{1}{3}|s_i - s_{\eta(i)}|, \end{aligned}$$

which contradicts the fact that i is a small player. Hence, $z_{\sigma(i)} \in [s_i, s_{\sigma(i)}]$, otherwise player i would be covered. Then,

$$\begin{aligned} \text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{\sigma(i)}(\mathbf{z}, \mathbf{s}) &= \max\{|s_i - z_i|, |z_i - z_{\sigma(i)}|\} + \max\{|s_{\sigma(i)} - z_{\sigma(i)}|, |z_{\sigma(i)} - z_{\sigma(\sigma(i))}|\} \\ &\geq z_{\sigma(i)} - z_i + s_{\sigma(i)} - z_{\sigma(i)} = s_{\sigma(i)} - z_i. \end{aligned}$$

Since i is small, we have $z_i < s_i + \frac{1}{3}(s_{\sigma(i)} - s_i)$ and we get

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{\sigma(i)}(\mathbf{z}, \mathbf{s}) \geq \frac{2}{3}(s_{\sigma(i)} - s_i) \geq \frac{1}{3}|s_i - s_{\eta(i)}| + \frac{1}{3}|s_{\sigma(i)} - s_{\eta(\sigma(i))}|$$

as desired. \square

Let $N(S)$ denote the set of players j that are neighbors of players in S (i.e., $j \in N(S)$ when $\sigma(i) = j$ for some player $i \in S$).

Claim 3.17. $N(S)$ does not contain small players.

Proof. Assume otherwise that for some player $i \in S$, $\sigma(i)$ also belongs to S . Without loss of generality $\sigma(i) > i$. If $z_{\sigma(i)} \geq s_{\sigma(i)}$, then

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2}|z_{\sigma(i)} - s_i| \geq \frac{1}{2}|s_{\sigma(i)} - s_i| \geq \frac{1}{3}|s_i - s_{\eta(i)}|$$

contradicting the fact that $i \in S$. So, $z_{\sigma(i)} < s_{\sigma(i)}$. Also, $z_{\sigma(i)} \geq s_i$ (since neither i is covered nor $\sigma(i)$ is kangaroo). Since $\sigma(i)$ is small, $s_{\sigma(i)} - z_{\sigma(i)} < \frac{1}{3}|s_{\sigma(i)} - s_{\eta(\sigma(i))}| \leq \frac{1}{3}(s_{\sigma(i)} - s_i)$, i.e., $z_{\sigma(i)} > \frac{2}{3}s_{\sigma(i)} + \frac{1}{3}s_i$. Hence,

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2}(z_{\sigma(i)} - s_i) > \frac{1}{3}(s_{\sigma(i)} - s_i),$$

which contradicts $i \in S$. \square

Claim 3.18. For every two players $i, i' \in S$, $\sigma(i) \neq \sigma(i')$.

Proof. Assume otherwise and let $\sigma(i) = \sigma(i') = j$ with $i < i'$. If $z_j \notin [s_i, s_{i'}]$, then the cost of either i or i' is at least $\frac{1}{2}(s_{i'} - s_i)$, contradicting the fact that both players are small. Hence,

$z_j \in [s_i, s_{i'}]$. Notice that $s_j \in [s_i, s_{i'}]$ as well, otherwise either i or i' would be covered by j . Now the fact that i and i' are small implies that

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{i'}(\mathbf{z}, \mathbf{s}) < \frac{1}{3}|s_i - s_{\eta(i)}| + \frac{1}{3}|s_{i'} - s_{\eta(i')}| \leq \frac{1}{3}(s_j - s_i) + \frac{1}{3}(s_{i'} - s_j) = \frac{1}{3}(s_{i'} - s_i).$$

On the other hand,

$$\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{i'}(\mathbf{z}, \mathbf{s}) \geq \frac{1}{2}(z_j - s_i) + \frac{1}{2}(s_{i'} - z_j) = \frac{1}{2}(s_{i'} - s_i),$$

a contradiction. \square

We now consider the social cost of \mathbf{z} due to players of different groups and exploit the claims above so that we obtain the lemma. In particular, we have

$$\begin{aligned} \text{SC}(\mathbf{z}, \mathbf{s}) &= \sum_{i=1}^n \text{cost}_i(\mathbf{z}, \mathbf{s}) \\ &\geq \sum_{i \in S: \sigma(i) \in \mathcal{K}} (\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{excess}_{\sigma(i)}) \\ &\quad + \sum_{i \in S: \sigma(i) \in L \cup (\mathcal{C} \setminus \mathcal{K})} (\text{cost}_i(\mathbf{z}, \mathbf{s}) + \text{cost}_{\sigma(i)}(\mathbf{z}, \mathbf{s})) \\ &\quad + \sum_{i \in \mathcal{K}} (\text{cost}_i(\mathbf{z}, \mathbf{s}) - \text{excess}_i) + \sum_{i \in L \setminus N(S)} \text{cost}_i(\mathbf{z}, \mathbf{s}) \\ &\geq \frac{1}{3} \sum_{i \in S: \sigma(i) \in \mathcal{K}} |s_i - s_{\eta(i)}| \\ &\quad + \frac{1}{3} \sum_{i \in S: \sigma(i) \in L \cup (\mathcal{C} \setminus \mathcal{K})} (|s_i - s_{\eta(i)}| + |s_{\sigma(i)} - s_{\eta(\sigma(i))}|) \\ &\quad + \frac{1}{3} \sum_{i \in \mathcal{K}} \left(|s_i - s_{\eta(i)}| + \sum_{j \in C_i} |s_j - s_{\eta(j)}| \right) + \frac{1}{3} \sum_{i \in L \setminus N(S)} |s_i - s_{\eta(i)}| \\ &\geq \frac{1}{3} \sum_{i=1}^n |s_i - s_{\eta(i)}|, \end{aligned}$$

as desired. The first inequality follows by the classification of the players and due to Claims 3.17 and 3.18. The second one follows by Claims 3.15, 3.16, and 3.14, and by the definition of large players. The last one follows since the players enumerated in the first two sums at its left cover the whole set S (by Claim 3.17). \square

We are ready to present our upper bound on the price of anarchy for 1-COF games.

Theorem 3.19. *The price of anarchy of 1-COF games over pure Nash equilibria is at most 3.*

Proof. Let us consider a 1-COF game with n players and belief vector \mathbf{s} . Let \mathbf{z}^* be an optimal opinion vector and recall that $\eta(i)$ is the player that minimizes the distance $|s_i - s_{\eta(i)}|$. By Lemma 3.13, we have

$$\text{SC}(\mathbf{z}^*, \mathbf{s}) \geq \frac{1}{3} \sum_{i=1}^n |s_i - s_{\eta(i)}|. \quad (3.18)$$

Now, consider any pure Nash equilibrium \mathbf{z} of the game. We will show that

$$\text{SC}(\mathbf{z}, \mathbf{s}) \leq \sum_{i=1}^n |s_i - s_{\eta(i)}|. \quad (3.19)$$

The theorem then follows by (3.18) and (3.19).

In particular, we will show that $\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq |s_i - s_{\eta(i)}|$ for each player i . Let us assume that $\eta(i) = i - 1$; the case $\eta(i) = i + 1$ is symmetric. Recall that $\sigma(i)$ is the neighbor of player i in the pure Nash equilibrium \mathbf{z} . We distinguish between four cases.

- **Case I:** $\sigma(i) = i - 1$. By Lemma 3.4, we have $s_{i-1} \leq z_i \leq s_i$. Then, clearly, $\text{cost}_i(\mathbf{z}, \mathbf{s}) = |s_i - z_i| \leq |s_i - s_{i-1}|$ as desired.
- **Case II:** $\sigma(i) = i + 1$ and $\sigma(i - 1) = i$. By Lemmas 3.2 and 3.4, we have $s_{i-1} \leq z_{i-1} \leq s_i \leq z_i$. Since player i has player $i + 1$ as neighbor, we have $|z_{i+1} - s_i| \leq |s_i - z_{i-1}|$. Hence, $\text{cost}_i(\mathbf{z}, \mathbf{s}) = |z_i - s_i| \leq |z_{i+1} - s_i| \leq |s_i - z_{i-1}| \leq |s_i - s_{i-1}|$.
- **Case III:** $\sigma(i) = i + 1$, $\sigma(i - 1) = i - 2$, and $\text{cost}_i(\mathbf{z}, \mathbf{s}) \leq \text{cost}_{i-1}(\mathbf{z}, \mathbf{s})$. By the definition of $\sigma(\cdot)$ and Lemma 3.2, we have $z_{i-2} \leq z_{i-1} \leq s_{i-1} \leq s_i \leq z_i \leq z_{i+1}$. We have

$$\begin{aligned} \text{cost}_i(\mathbf{z}, \mathbf{s}) &\leq 2\text{cost}_{i-1}(\mathbf{z}, \mathbf{s}) - \text{cost}_i(\mathbf{z}, \mathbf{s}) \\ &= |s_{i-1} - z_{i-2}| - |z_i - s_i| \\ &\leq |z_i - s_{i-1}| - |z_i - s_i| \\ &= |s_i - s_{i-1}|. \end{aligned}$$

The second inequality follows since player $i - 2$ (instead of i) is the neighbor of player $i - 1$.

- **Case IV:** $\sigma(i) = i + 1$, $\sigma(i - 1) = i - 2$, and $\text{cost}_i(\mathbf{z}, \mathbf{s}) > \text{cost}_{i-1}(\mathbf{z}, \mathbf{s})$.

$$\begin{aligned} \text{cost}_i(\mathbf{z}, \mathbf{s}) &< 2\text{cost}_i(\mathbf{z}, \mathbf{s}) - \text{cost}_{i-1}(\mathbf{z}, \mathbf{s}) \\ &= |z_{i+1} - s_i| - |s_{i-1} - z_{i-1}| \\ &\leq |s_i - z_{i-1}| - |s_{i-1} - z_{i-1}| \end{aligned}$$

$$= |s_i - s_{i-1}|.$$

The second inequality follows since player $i + 1$ (instead of $i - 1$) is the neighbor of player i .

This completes the proof. □

3.9 Lower bounds on the price of anarchy

This section contains our lower bounds on the price of anarchy.³ We begin by considering the simpler case of 1-COF games, for which we present a tight lower bound of 3 for pure Nash equilibria (Theorem 3.20) and a lower bound of 6 for mixed Nash equilibria (Theorem 3.21). We remark that, for 1-COF games, this implies that mixed Nash equilibria are strictly worse than pure ones. Then, we study the general case of k -COF games and we show lower bounds for pure and mixed Nash equilibria (Theorems 3.22 and 3.23, respectively) that grow linearly with k .

3.9.1 The case of 1-COF games

We now present our lower bounds for the case of 1-COF games; both results rely on the same, and rather simple, instance.

Theorem 3.20. *The price of anarchy of 1-COF games over pure Nash equilibria is at least 3.*

Proof. Let $\lambda \in (0, 1)$ and consider a 1-COF game with six players and belief vector $\mathbf{s} = (-10 - \lambda, -10 - \lambda, -2 - \lambda, 2 + \lambda, 10 + \lambda, 10 + \lambda)$. This game is depicted in Figure 3.9a. We can show that the opinion vector (see Figure 3.9b)

$$\mathbf{z} = (-10 - \lambda, -10 - \lambda, -6 - \lambda, 6 + \lambda, 10 + \lambda, 10 + \lambda)$$

is a pure Nash equilibrium with social cost $\text{SC}(\mathbf{z}, \mathbf{s}) = 8$. The first two players suffer zero cost as they follow each other and their opinions coincide with their beliefs; the same holds also for the last two players. For the third player, it is $\sigma(3) \in \{1, 2\}$ since $|z_1 - s_3| = |z_2 - s_3| = 8 < |z_4 - s_3| = 8 + 2\lambda$ and z_3 is in the middle of the interval $[-10 - \lambda, -2 - \lambda]$; hence, $\text{cost}_3(\mathbf{z}, \mathbf{s}) = 4$. Similarly, we have $\sigma(4) \in \{5, 6\}$, z_4 lies in the middle of the interval $[2 + \lambda, 10 + \lambda]$ and $\text{cost}_4(\mathbf{z}, \mathbf{s}) = 4$. Hence, \mathbf{z} is indeed a pure Nash equilibrium.

³We remark that our lower bounds on the price of stability in Section 3.5 are also lower bounds on the price of anarchy. However, the lower bounds presented in this section are much stronger.




Figure 3.9: (a) The 1-COF game considered in the proofs of Theorems 3.20 and 3.21. (b) The pure Nash equilibrium vector \mathbf{z} (see the proof of Theorem 3.20) with social cost 8. (c) The opinion vector $\tilde{\mathbf{z}}$ with social cost $\frac{8+4\lambda}{3}$.

Now, consider the opinion vector (see Figure 3.9c)

$$\tilde{\mathbf{z}} = \left(-10 - \lambda, -10 - \lambda, \frac{-2 - \lambda}{3}, \frac{2 + \lambda}{3}, 10 + \lambda, 10 + \lambda \right)$$

which yields a social cost of $SC(\tilde{\mathbf{z}}, \mathbf{s}) = \frac{8+4\lambda}{3}$; here, again, the first and last two players have zero cost, but players 3 and 4 now each have cost $\frac{4+2\lambda}{3}$ since they follow each other. The optimal social cost is upper bounded by $SC(\tilde{\mathbf{z}})$ and, hence, the price of anarchy is at least

$$\frac{SC(\mathbf{z}, \mathbf{s})}{SC(\tilde{\mathbf{z}}, \mathbf{s})} = \frac{3}{1 + \lambda/2},$$

and the theorem follows by setting λ arbitrarily close to 0. \square

Our next theorem gives a lower bound on the price of anarchy over mixed Nash equilibria for 1-COF games; we remark that this lower bound is greater than the upper bound of Theorem 3.19 for the price of anarchy over pure Nash equilibria.

Theorem 3.21. *The price of anarchy of 1-COF games over mixed Nash equilibria is at least 6.*

Proof. Consider again the 1-COF game depicted in Figure 3.9a with six players and beliefs $\mathbf{s} = (-10 - \lambda, -10 - \lambda, -2 - \lambda, 2 + \lambda, 10 + \lambda, 10 + \lambda)$, where $\lambda \in (0, 1)$. To simplify the following

discussion, we will refer to the first two players as the L players, the third player as player ℓ , the fourth player as player r , and the last two players as the R players.

Let \mathbf{z} be a randomized opinion vector according to which $z_i = s_i$ for every $i \in L \cup R$, z_ℓ is chosen equiprobably from $\{-6 - \lambda, -6 + 3\lambda\}$, and z_r is chosen equiprobably from $\{6 + \lambda, 6 - 3\lambda\}$. Observe that $\sigma(\ell) \in L$ whenever $z_r = 6 + \lambda$, and $\sigma(\ell) = r$ whenever $z_r = 6 - 3\lambda$; each of these events occurs with probability $1/2$. Hence, we obtain

$$\mathbb{E}[\text{cost}_\ell(\mathbf{z}, \mathbf{s})] = \mathbb{E}[\text{cost}_r(\mathbf{z}, \mathbf{s})] = \frac{1}{2} \left(\frac{4}{2} + \frac{4 + 4\lambda}{2} \right) + \frac{1}{2} \left(\frac{12 - 2\lambda}{2} + \frac{12 - 6\lambda}{2} \right) = 8 - \lambda,$$

and, thus, $\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})] = 16 - 2\lambda$. In the following, we will prove that \mathbf{z} is a mixed Nash equilibrium. First, observe that all players in sets L and R have no incentive to deviate since they follow each other and have zero cost. We will now argue about player ℓ ; due to symmetry, our findings will apply to player r as well.

Consider a deterministic deviating opinion y for player ℓ . We will show that $\mathbb{E}[\text{cost}_\ell(\mathbf{z}, \mathbf{s})] \leq \mathbb{E}_{\mathbf{z}_{-\ell}}[\text{cost}_\ell(y, \mathbf{z}_{-\ell}), \mathbf{s}]$ for any y , which implies that player ℓ has no incentive to deviate from the randomized opinion z_ℓ . Indeed, we have that

$$\begin{aligned} & \mathbb{E}_{\mathbf{z}_{-\ell}}[\text{cost}_\ell((y, \mathbf{z}_{-\ell}), \mathbf{s})] \\ &= \frac{1}{2} \max\{|-2 - \lambda - y|, |y + 10 + \lambda|\} + \frac{1}{2} \max\{|-2 - \lambda - y|, |6 - 3\lambda - y|\} \\ &\geq \frac{1}{2}(y + 10 + \lambda) + \frac{1}{2}(6 - 3\lambda - y) \\ &= 8 - \lambda, \end{aligned}$$

where the inequality holds since $\max\{|a|, |b|\} \geq a$ for any a and b . Hence, player ℓ has no incentive to deviate from her strategy in \mathbf{z} , and neither has player r due to symmetry. Therefore, \mathbf{z} is a mixed Nash equilibrium.

Now, consider the opinion vector

$$\tilde{\mathbf{z}} = \left(-10 - \lambda, -10 - \lambda, \frac{-2 - \lambda}{3}, \frac{2 + \lambda}{3}, 10 + \lambda, 10 + \lambda \right)$$

which, as in Theorem 3.20, yields a social cost of $\text{SC}(\tilde{\mathbf{z}}, \mathbf{s}) = \frac{8+4\lambda}{3}$. Hence, the optimal social cost is upper bounded by $\text{SC}(\tilde{\mathbf{z}}, \mathbf{s})$, and the price of anarchy over mixed equilibria is at least

$$\frac{\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})]}{\text{SC}(\tilde{\mathbf{z}}, \mathbf{s})} = 3 \frac{16 - 2\lambda}{8 + 4\lambda},$$

and the theorem follows by setting λ arbitrarily close to 0. □




Figure 3.10: (a) The k -COF game considered in the proofs of Theorems 3.22 and 3.23, for $k \geq 2$. (b) The pure Nash equilibrium opinion vector \mathbf{z} (see the proof of Theorem 3.22). (c) The optimal opinion vector $\tilde{\mathbf{z}}$ for $k \geq 3$. (d) The optimal opinion vector $\tilde{\mathbf{z}}$ for $k = 2$. Observe that the optimal opinion vector changes at $k = 2$ due to the neighborhood size.

3.9.2 The general case of k -COF games with $k \geq 2$

We will now present lower bounds on the price of anarchy for k -COF games, with $k \geq 2$. We start with the case of pure Nash equilibria and continue with the more general case of mixed equilibria. As in the case of 1-COF games, a particular game will be used in order to derive the lower bounds both for pure and mixed Nash equilibria.

Theorem 3.22. *The price of anarchy of k -COF games over pure Nash equilibria is at least $k + 1$ for $k \geq 3$, and at least $18/5$ for $k = 2$.*

Proof. Let $\lambda \in (0, 1)$ and consider a k -COF game with $3k + 3$ players, for $k \geq 2$, that are partitioned into the following five sets. The first set L consists of $k+1$ players with $s_i = -16 - 2\lambda$ for any $i \in L$, the second set consists of a single player ℓ with $s_\ell = -4 - \lambda$, the third set M has

$k - 1$ players with $s_i = 0$ for any $i \in M$, the fourth set is a single player r with $s_r = 4 + \lambda$, and the last set R consists of $k + 1$ players with $s_i = 16 + 2\lambda$ for any $i \in R$. This instance is depicted in Figure 3.10a.

Let \mathbf{z} be the following opinion vector: $z_i = -16 - 2\lambda$ for any $i \in R$, $z_\ell = -8 - \lambda$, $z_i = 0$ for any $i \in M$, $z_r = 8 + \lambda$, and $z_i = 16 + 2\lambda$ for any $i \in R$; see Figure 3.10b. It is not hard to verify that this opinion vector is a pure Nash equilibrium with social cost $\text{SC}(\mathbf{z}, \mathbf{s}) = (8 + \lambda)(k + 1)$. First, observe that all players in sets L and R have zero cost, and, hence, have no incentive to deviate to another opinion. Furthermore, no player $i \in M$ has an incentive to deviate either since z_i lies in the middle of the interval $[-8 - \lambda, 8 + \lambda]$ which is defined by the opinions of players ℓ and r who, together with the remaining players of M , constitute the neighborhood $N_i(\mathbf{z}, \mathbf{s})$ of player i . The cost experienced by such a player i is $8 + \lambda$. Finally, the neighborhood $N_\ell(\mathbf{z}, \mathbf{s})$ of player ℓ consists of all players in M (who have opinions that are closest to s_ℓ) and some player $i \in L$; note that player r does not belong to $N_\ell(\mathbf{z}, \mathbf{s})$ since $z_r - s_\ell = 12 + 2\lambda > 12 - \lambda = s_\ell - z_i$ for all $i \in L$. Hence, player ℓ has no incentive to deviate to another opinion since z_ℓ lies in the middle of the interval $[-16 - 2\lambda, 0]$ and she experiences cost equal to $8 + \lambda$. Due to symmetry, player r does not have incentive to deviate as well. Hence, \mathbf{z} is indeed a pure Nash equilibrium with $\text{SC}(\mathbf{z}, \mathbf{s}) = (8 + \lambda)(k + 1)$.

We now present an opinion vector $\tilde{\mathbf{z}}$ with social cost $\text{SC}(\tilde{\mathbf{z}}, \mathbf{s}) = 8 + 2\lambda$ for $k \geq 3$ and $\text{cost}(\tilde{\mathbf{z}}, \mathbf{s}) = \frac{5}{3}(4 + \lambda)$ for $k = 2$. In particular, for $k \geq 3$, $\tilde{\mathbf{z}}$ is defined as follows: $\tilde{z}_i = -16 - 2\lambda$ for any $i \in L$, $\tilde{z}_\ell = \tilde{z}_i = \tilde{z}_r = 0$ for any $i \in M$, and $\tilde{z}_i = 16 + 2\lambda$ for any $i \in R$; see Figure 3.10c. Observe that all players in L , M , and R have zero cost, while players ℓ and r have cost equal to $4 + \lambda$ each. For $k = 2$, $\tilde{\mathbf{z}}$ is defined as follows: $\tilde{z}_i = -16 - 2\lambda$ for any $i \in L$, $\tilde{z}_\ell = -\frac{1}{3}(4 + \lambda)$, $\tilde{z}_i = 0$ for any $i \in M$, $\tilde{z}_r = \frac{1}{3}(4 + \lambda)$, and $\tilde{z}_i = 16 + 2\lambda$ for any $i \in R$; see Figure 3.10d. Again, all players in L and R have zero cost. However, players ℓ and r now each have cost $\frac{2}{3}(4 + \lambda)$ and the unique player in M has cost $\frac{1}{3}(4 + \lambda)$.

Clearly, since $\text{SC}(\tilde{\mathbf{z}}, \mathbf{s})$ is an upper bound on the optimal social cost, we conclude that the price of anarchy over pure Nash equilibria is at least $\frac{(8+\lambda)(k+1)}{8+2\lambda}$ for $k \geq 3$ and $\frac{9(8+\lambda)}{5(4+\lambda)}$ for $k = 2$, and the theorem follows by setting λ arbitrarily close to 0. \square

We now consider the case of mixed Nash equilibria; we remark that, in this case, our lower bounds for $k \geq 2$ are smaller than the corresponding upper bounds for pure Nash equilibria.

Theorem 3.23. *The price of anarchy of k -COF games over mixed Nash equilibria is at least $k + 2$ for*

$k \geq 3$, and at least $24/5$ for $k = 2$.

Proof. As in the proof of Theorem 3.22, let $\lambda \in (0, 1)$ and consider the k -COF game depicted in Figure 3.10a with $3k + 3$ players that form 5 sets. Again, the first set L consists of $k + 1$ players where $s_i = -16 - 2\lambda$ for all $i \in L$, the second set consists of a single player ℓ with $s_\ell = -4 - \lambda$, the third set M has $k - 1$ players with $s_i = 0$ for all $i \in M$, the fourth set is a single player r with $s_r = 4 + \lambda$, and the last set R consists of $k + 1$ players with $s_i = 16 + 2\lambda$ for all $i \in R$.

Consider the following (randomized) opinion vector \mathbf{z} : $z_i = s_i$ for every $i \in L \cup M \cup R$, while z_ℓ is chosen uniformly at random among $\{-8 - \lambda, -8 + 3\lambda\}$ and z_r is chosen uniformly at random among $\{8 - 3\lambda, 8 + \lambda\}$. We will show that the opinion vector \mathbf{z} is a mixed Nash equilibrium with $\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})] = 8k + 16 - \lambda$.

First, observe that the players in sets L and the R constitute local neighborhoods, that is, $N_i(\mathbf{z}, \mathbf{s}) = L \setminus \{i\}$ for any player $i \in L$, and $N_i(\mathbf{z}, \mathbf{s}) = R \setminus \{i\}$ for any player $i \in R$. Hence, all these players have zero cost and no incentive to deviate.

Next, let us focus on a player $i \in M$. Clearly, the neighborhood of player i consists of the remaining $k - 2$ players in M as well as players ℓ and r . The expected cost of player i in \mathbf{z} is $\mathbb{E}[\text{cost}_i(\mathbf{z}, \mathbf{s})] = \frac{3}{4}(8 + \lambda) + \frac{1}{4}(8 - 3\lambda) = 8$ since at least one of players ℓ and r is at distance $8 + \lambda$ with probability $3/4$ and both of them are at distance $8 - 3\lambda$ with probability $1/4$. Hence, these $k - 1$ players contribute $8(k - 1)$ to the expected social cost of \mathbf{z} . We now argue that if player $i \in M$ deviates to a deterministic opinion y , her expected cost does not decrease. Clearly, if $y \geq 3\lambda$, then this trivially holds as the expected cost of i is at least $y - z_\ell$ which is at least $y + 8 - 3\lambda$; the case where $y \leq -3\lambda$ is symmetric. Hence, it suffices to consider the case where $|y| < 3\lambda$. The expected cost of i when deviating to y is

$$\begin{aligned} & \mathbb{E}_{\mathbf{z}_{-i}}[\text{cost}_i((y, \mathbf{z}_{-i}), \mathbf{s})] \\ &= \frac{1}{4} \max\{8 + \lambda - y, y + 8 + \lambda\} + \frac{1}{4} \max\{8 + \lambda - y, y + 8 - 3\lambda\} \\ & \quad + \frac{1}{4} \max\{8 - 3\lambda - y, y + 8 + \lambda\} + \frac{1}{4} \max\{8 - 3\lambda - y, y + 8 - 3\lambda\} \\ & \geq \frac{1}{4}(8 + \lambda - y) + \frac{1}{4}(8 + \lambda - y) + \frac{1}{4}(y + 8 + \lambda) + \frac{1}{4}(y + 8 - 3\lambda) \\ &= 8, \end{aligned}$$

where the inequality holds since $\max\{a, b\} \geq a$ for any a and b .

Now, let us examine player r ; the case of player ℓ is symmetric. Observe that the $k - 1$ players in M always belong to the neighborhood $N_r(\mathbf{z}, \mathbf{s})$ of player r and it remains to argue

about the identity of the last player in $N_r(\mathbf{z}, \mathbf{s})$. Whenever $z_\ell = -8 + 3\lambda$, then $\ell \in N_r(\mathbf{z}, \mathbf{s})$, otherwise, if $z_\ell = -8 - \lambda$, one of the players in set R belongs to $N_r(\mathbf{z}, \mathbf{s})$. The expected cost of player r is $\mathbb{E}[\text{cost}_r(\mathbf{z}, \mathbf{s})] = \frac{1}{4}(8 + \lambda) + \frac{1}{4}(8 + 5\lambda) + \frac{1}{4}(16 - 2\lambda) + \frac{1}{4}(16 - 6\lambda) = 12 - \lambda/2$, and, hence, players ℓ and r contribute $24 - \lambda$ to the expected social cost of \mathbf{z} . It remains to show that player r cannot decrease her expected cost by deviating to another opinion y . The expected cost of player r when deviating to y is

$$\begin{aligned}\mathbb{E}_{\mathbf{z}_{-r}}[\text{cost}_r((y, \mathbf{z}_{-r}), \mathbf{s})] &= \frac{1}{2} \max\{|16 + 2\lambda - y|, |y|\} + \frac{1}{2} \max\{|y + 8 - 3\lambda|, |4 + \lambda - y|\} \\ &\geq \frac{1}{2}(16 + 2\lambda - y) + \frac{1}{2}(y + 8 - 3\lambda) \\ &= 12 - \lambda/2,\end{aligned}$$

where the inequality holds since $\max\{|a|, |b|\} \geq a$ for any a and b . Hence, we conclude that \mathbf{z} is a mixed Nash equilibrium with expected social cost $\mathbb{E}[\text{SC}(\mathbf{z}, \mathbf{s})] = 8k + 16 - \lambda$.

As in the proof of Theorem 3.22, there exists an opinion vector $\tilde{\mathbf{z}}$ with social cost $\text{SC}(\tilde{\mathbf{z}}, \mathbf{s}) = 8 + 2\lambda$ for $k \geq 3$ and $\text{SC}(\tilde{\mathbf{z}}, \mathbf{s}) = \frac{5}{3}(4 + \lambda)$ for $k = 2$. Since $\text{SC}(\tilde{\mathbf{z}}, \mathbf{s})$ is an upper bound on the optimal social cost, we have that the price of anarchy over mixed equilibria is at least $\frac{8k+16-\lambda}{8+2\lambda}$ for $k \geq 3$ and $\frac{3(32-\lambda)}{5(4+\lambda)}$ for $k = 2$, and the theorem follows, by setting λ arbitrarily close to 0. \square

3.10 Conclusion

In this chapter, we focused on the efficiency and complexity of a simple class of compromising opinion formation games, which we call k -COF games. In such a game, there exists a set of players with personal beliefs over some issue, but each of them expresses a public opinion in order to minimize an explicit cost that is defined as the maximum between the distance of her opinion from her belief and the distance of her opinion from every opinion expressed in her neighborhood, which dynamically changes depending on the other players that express opinion chose to her belief.

In particular, we first proved several structural properties about pure Nash equilibria as well as that pure equilibria may not exist for any value of k . Then, we proved that the price of stability and anarchy of general k -COF games grows linearly in terms of k . For the special case of $k = 1$, we showed a tight bound of 3 on the price of anarchy, and designed an efficient algorithm for computing the best and worst equilibrium (in terms of the social cost) by reducing the corresponding problems to the problems of computing minimum and maximum paths in particular directed acyclic graphs.

Chapter 4

Truthful mechanisms for ownership transfer with expert advice

In this chapter we focus on the design and analysis of near-optimal truthful mechanisms for ownership transfer; see the discussion in Section 1.3 for an introduction to the problem and motivating examples. The results presented here can be found in [Caragiannis et al., 2018].

4.1 Overview of contribution and techniques

We focus on ownership transfer and study the very simple but fundamental setting of two competing agents A and B , and a single expert with cardinal preferences over the three options of selling to agent A , selling to agent B , or not selling at all (in which case the ownership transfer does not take place). A mechanism takes as input the bids of the agents and the expert's preferences, and decides one of the three options as outcome. In general, mechanisms are *randomized*. For a given input, they select the outcome using a probability distribution (or *lottery*) over the three options.

We consider mechanisms that can be *truthfully implemented* as follows. First, the outcome of the mechanism is complemented with payments that are imposed to the agents. Then, the lottery and the payments should be such that the expert is incentivized to report her true preferences in order to maximize her (expected) value for the outcome and the agents are incentivized to report their true values as bids in order to maximize their utility, i.e., their expected value for the outcome minus their payment to the mechanism. In the following, we refer to mechanisms with such implementations as *truthful mechanisms*.

Interestingly, the theory of mechanism design allows us to abstract away from payments and view truthful mechanisms simply as lotteries. Well-known characterizations for single-

parameter mechanism design with money from the literature, as well as new characterizations that we prove here for lotteries that guarantee truthfulness from the expert's side, are the main tools we use in order to constrain the design space of truthful mechanisms in our setting.

We make additional informational restrictions that can further divide truthful mechanisms into the following classes:

- *ordinal* mechanisms, which ignore the exact bids and the expert's preference values and instead take into account only their relative order,
- *bid-independent* mechanisms, which ignore the bids and base their decision solely on the expert's cardinal preferences,
- *expert-independent* mechanisms, which ignore the expert's preferences and base their decision solely on the bids, and
- general truthful mechanisms, which may take *both* the bids and the expert's preference values into account.

We measure the quality of truthful mechanisms in terms of the *social welfare*, the aggregate value of the agents and the expert for the outcome. Unfortunately, our setting does not allow for a truthful implementation of the social welfare-maximizing outcome. Therefore, we resort to near-optimal truthful mechanisms and use the notion of the *approximation ratio* to measure their quality.

For the classes of ordinal, bid-independent, and expert-independent mechanisms, we prove lower bounds on the approximation ratio of truthful mechanisms in the class and identify the best possible among them, with approximation ratios of 1.5, 1.377, and 1.343, respectively. Furthermore, by slightly enhancing expert-independent mechanisms and allowing them to utilize a single bit of information about the expert's preferences, we define a template for the design of new truthful mechanisms. The template defines *always-sell* mechanisms that select either agent A or agent B as the outcome. We present two mechanisms that follow our template, one deterministic and one randomized, with approximation ratios 1.618 and 1.25, respectively. The former is best-possible among all deterministic truthful mechanisms. The latter is best-possible among all always-sell truthful mechanisms. We also present an unconditional lower bound of 1.141 on the approximation ratio of any truthful mechanism. These results are summarized in Table 4.1.

Class of mechanisms	apx. ratio	Comment
ordinal	1.5	mechanisms EOM, BOM (Theorem 4.3) best possible (see Theorem 4.4)
bid-independent	1.377	mechanism BIM (Theorem 4.7) best possible (Theorem 4.8)
expert-independent	1.343	mechanism EIM (Theorem 4.11) best possible (Theorem 4.11)
our template	1.25	randomized mechanism R (Theorem 4.14) best possible, always-sell (Theorem 4.15)
	1.618	deterministic mechanism D (Theorem 4.14) best possible, deterministic (Theorem 4.17)
	1.14	lower bound (Theorem 4.16)
all mechanisms		

Table 4.1: Overview of our results; see [Caragiannis et al., 2018].

Both our positive and negative results have been possible by narrowing the design space using the truthfulness characterizations, the particular structure in each class of mechanisms, as well as the goal of low approximation ratio. In most cases, the design of new mechanisms turns out to be as simple as drawing a curve in a restricted area of a 2-dimensional plot (for instance, see Figures 4.2 and 4.3).

4.1.1 Chapter roadmap

We begin with a discussion of related work in Section 4.2. Then, we continue with preliminary definitions, notation and examples in Section 4.3. Then, Sections 4.4, 4.5, and 4.6 are devoted to ordinal, bid-independent and expert-independent mechanisms, respectively. Our template and the corresponding best possible deterministic and randomized mechanisms are presented in Section 4.7, while our unconditional lower bounds are presented in Section 4.8. Finally, we conclude in Section 4.9

4.2 Related work

Our setting can be viewed as an instance of *approximate mechanism design*, with [Nisan and Ronen, 2001] and without money [Procaccia and Tennenholtz, 2013], which was proposed for problems where the goal is to optimize an objective under strict truthfulness requirements. Myerson [1981] proved necessary and sufficient conditions for (deterministic or randomized) truthful mechanisms with money. This characterization allowed us to abstract away from the payment functions (which are uniquely determined given the winning probabilities) on the

agents' side, and provided us with tools to argue about the structure of truthful mechanisms without money on the expert's side as well.

For settings with money, the VCG mechanism [Clarke, 1971, Groves, 1973, Vickrey, 1961] is deterministic, truthful, and maximizes the social welfare. However, as we pointed out in Section 1.3, in our hybrid mechanism design setting we need to take the values of the expert into account as well, and therefore VCG is no longer truthful nor optimal (see Example 4.1). On the expert's side, truthful mechanisms can be thought of as truthful voting rules; any positive results for deterministic such rules are impaired by impossibility theorems [Gibbard, 1973, Satterthwaite, 1975] which limit this class to only *dictatorial mechanisms*.

In contrast, the class of randomized truthful voting rules is much richer and includes many reasonable truthful rules that are not dictatorial. In fact, Gibbard [1977] characterized the class of all such ordinal rules; a general characterization for all cardinal rules is still elusive. To this end, a notable amount of work in the classical economics literature as well as in computer science has been devoted towards designing such rules and proving structural properties for restricted classes. Gibbard [1978] provided a characterization which only holds for discrete strategy spaces, and later Hylland [1980]¹ proved that the class of truthful rules that are Pareto-efficient reduces to random dictatorships.

Freixas [1984] used the differential approach to mechanism design proposed by Laffont and Maskin [1980] to design a class of truthful mechanisms which actually characterize the class of twice differentiable mechanisms over subintervals of the valuation space; the best possible truthful bid-independent mechanism that we propose in this chapter can be seen as a mechanism in this class. Barbera et al. [1998] showed that there are many interesting truthful mechanisms that do not fall into the classes considered by Freixas [1984]. In the computer science literature, Feige and Tennenholtz [2010] designed a class of one-voter cardinal truthful mechanisms, where the election probabilities are given by certain polynomials.

Social welfare maximization without payments has been studied in a plethora of papers, in general social choice settings [Bhaskar et al., 2018, Filos-Ratsikas and Miltersen, 2014] as well as in restricted domains such as matching and allocation problems [Cheng, 2016, Filos-Ratsikas et al., 2014, Guo and Conitzer, 2010]. Similarly to our work, Filos-Ratsikas and Miltersen [2014] use one-voter cardinal truthful mechanisms to achieve improved welfare guarantees. The presence of the agents significantly differentiates our setting from theirs (as well as other

¹Quite remarkably, this paper is unpublished – the result was revisited by Dutta et al. [2007].

related paper). Another relevant notion is that of the *distortion* of (non-truthful) mechanisms which operate under limited (ordinal) information [Anshelevich et al., 2015, Boutilier et al., 2015, Caragiannis et al., 2017b, Caragiannis and Procaccia, 2011, Caragiannis et al., 2016]. While the lack of information has also been a restrictive factor for some of our results (in conjunction with truthfulness), we mainly focused on cardinal mechanisms for which truthfulness is the limiting constraint.

4.3 Definitions and notation

Our setting consists of two agents A and B who compete for an item (to be thought of as an abstraction of a merger or acquisition) and an expert E . The agents have valuations w_A and w_B denoting the amount of money that they would be willing to spend for the item, and the expert has a valuation function $v : \mathcal{O} \rightarrow \mathbb{R}$ over the following three *options*: agent A is selected to get the item, or agent B is selected, or no agent is selected to get the item. We use \emptyset to denote this last option; hence, $\mathcal{O} = \{A, B, \emptyset\}$. We use $\mathbf{w} = (w_A, w_B)$ to denote an *agent profile* and let \mathcal{W} be the set of all such profiles. Similarly, we use $\mathbf{v} = (v(A), v(B), v(\emptyset))$ to denote an *expert profile* and let \mathcal{V} be the set of all such profiles. The domain of our setting is $\mathcal{D} = \mathcal{V} \times \mathcal{W}$. From now on, we use the term *profile* to refer to elements of \mathcal{D} .

A mechanism M takes as input from the expert and the agents a profile (\mathbf{v}, \mathbf{w}) and decides, according to a probability distribution (or *lottery*) P^M , the pair (o, \mathbf{p}) consisting of an option $o \in \mathcal{O}$ and a vector $\mathbf{p} = (p_A, p_B)$ indicating the payments that are imposed to the agents. The execution of the mechanism yields a utility to the expert and the agents. Given an outcome (o, \mathbf{p}) of the mechanism, the utility of the expert is

$$u_E(o, \mathbf{p}) = v(o)$$

and the utility of agent $i \in \{A, B\}$ is

$$u_i(o, \mathbf{p}) = \begin{cases} w_i - p_i, & \text{if } i = o \\ -p_i, & \text{otherwise.} \end{cases}$$

The expert and the agents submit an *expert's report* and *bids* to the mechanism representing their corresponding profiles, but may have incentives to misreport their true values in order to maximize their utility. We are interested in mechanisms that do not allow such strategic manipulations. We say that a mechanism M is *truthful for agent* $i \in \{A, B\}$ if for any agent value w_i and any profile $(\mathbf{v}', \mathbf{w}')$,

$$\mathbb{E}[u_i(M(\mathbf{v}', (w_i, w'_{-i})))] \geq \mathbb{E}[u_i(M(\mathbf{v}', \mathbf{w}'))],$$

where the expectation is taken with respect to the lottery P^M . This means that bidding her true value w_i is a utility-maximizing strategy for the agent, no matter what the other agent and the expert's report are. Mechanism M is *truthful for the expert* if for any expert profile \mathbf{v} and any profile $(\mathbf{v}', \mathbf{w}')$,

$$\mathbb{E}[u_E(M(\mathbf{v}, \mathbf{w}))] \geq \mathbb{E}[u_E(M(\mathbf{v}', \mathbf{w}))].$$

Again, this means that reporting her true valuation profile is a utility-maximizing strategy for the expert, no matter what the agents bid. A mechanism M is *truthful* if it is truthful for the agents and truthful for the expert.

Our goal is to design truthful mechanisms that achieve high social welfare, which is the total value of the agents and the expert for the outcome. For a meaningful definition of the social welfare that weighs equally the expert's and the agents' valuations, we adopt a canonical representation of profiles. The expert has *von Neumann-Morgenstern valuations*, i.e., she has valuations of 0 and 1 for two of the options and a value in $[0, 1]$ for the third one. The agent values are normalized in the definition of the social welfare by dividing with the maximum of them. Then, the social welfare of an option $o \in \mathcal{O}$ is

$$SW(o, \mathbf{v}, \mathbf{w}) = \begin{cases} v(o) + \frac{w_o}{\max\{w_A, w_B\}}, & \text{if } o \in \{A, B\} \\ v(\emptyset), & \text{otherwise.} \end{cases}$$

We measure the quality of a truthful mechanism M by its *approximation ratio*, which (by abusing notation a bit and interpreting $M(\mathbf{v}, \mathbf{w})$ as the option decided by the mechanism) is defined as

$$\rho(M) = \sup_{(\mathbf{v}, \mathbf{w}) \in \mathcal{D}} \frac{\max_{o \in \mathcal{O}} SW(o, \mathbf{v}, \mathbf{w})}{\mathbb{E}[SW(M(\mathbf{v}, \mathbf{w}), \mathbf{v}, \mathbf{w})]}.$$

Low values of $\rho(M)$, as close as possible to 1, are most desirable.

Before we continue with the discussion of alternative representation of profiles, we present an example demonstrating the reason why the mechanism that simply selects the option that maximizes the social welfare based on the reported profile provided by the expert and the agents is not truthful.

Example 4.1. Let α and β be two parameters in $(0, 1)$ such that $\alpha > \beta$. Consider a profile in which the expert has values $\mathbf{v} = (v(A), v(B), v(\emptyset)) = (1, \alpha, 0)$ and the agents have values $\mathbf{w} = (w_A, w_B) = (\beta, 1)$.

If the expert and the agents were truthful, then since $\alpha > \beta$, the mechanism that chooses the option that maximizes the social welfare would select agent B . From the agent's side, it is well-understood how such a mechanism can be implemented; a simple second-price auction

would incentivize both agents to be truthful and, of course, would choose the one with the highest value. However, at the same time, we want the expert to be truthful as well, which is not possible in this particular example. The expert has strong incentive to misreport her value for agent B and decrease it from α to zero. This way, the output of the mechanism is agent A for whom the expert has value 1, as opposed to agent B for whom her value is $\alpha < 1$. \square

4.3.1 An alternative view of profiles

In order to simplify the exposition in the following sections, we devote some space here to introduce two alternative ways of representing profiles, which in turn will showcase more intuitive ways of realizing truthfulness and will help us in the design of efficient mechanisms.

Without restricting the space of mechanisms that can achieve good approximation ratios according to our definition of the social welfare, we focus on mechanisms that base their decisions on the *normalized bid* values $\frac{w_A}{\max\{w_A, w_B\}}$ and $\frac{w_B}{\max\{w_A, w_B\}}$. It will be convenient to use the following two alternative ways

$$\begin{pmatrix} 1 & x & 0 \\ h & \ell & z \end{pmatrix} \text{ and } \begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$$

to represent profiles. These representations are the *expert's* and *agents' view* of the profile, respectively. Each column corresponds to an option. According to the expert's view at the left, the columns are ordered in terms of the expert's values, which appear in the first row. The quantities h , ℓ , and z hold the normalized agent bids for the corresponding option and 0 for option \emptyset . Essentially, h is the value that the expert's favourite option has, which can be equal to 1 if it corresponds to the value of the agent with the highest value (*high-bidder*), equal to some value $y \in [0, 1]$ if it corresponds to the value of the agent with the lowest value (*low-bidder*), or 0 if it corresponds to the no-sale option \emptyset . Similarly, ℓ and z are the values that expert's second and third favourite options have, respectively. According to the agents' view at the right, the columns are ordered in terms of the bids, which appear in the second row. The quantities h , ℓ , and n now hold the expert valuations for the corresponding options. Now, h is the value that the expert has for the high-bidder, ℓ is the value of the expert for the low-bidder, and z is the value that the expert has for the no-sale option. All of them can take values in the interval $[0, 1]$.

These representations yield a crisper way to argue about truthfulness for the expert and the agents in our main results. Specifically, in Section 4.5, we will study bid-independent mechanisms, and therefore it makes sense to use the expert's view of profiles, whereas in Section 4.6, it will be easier to argue about our expert-independent mechanisms based on the

agents' view instead. The agents' view will also be used in Section 4.7, where, the mechanisms we present use the expert's opinion only to appropriately partition the input profiles into categories, and it is therefore easier to argue about their properties using the agent's view.

Similarly, we use two different representations of the lottery P^M , depending on whether we represent profiles according to the expert's or the agents' view. From the expert's viewpoint, P^M is represented by three functions g^M , f^M , and η^M , which correspond to the probability of selecting the first, second, and third favourite option of the expert, respectively. Similarly, from the agents' viewpoint, P^M is represented by three functions d^M , c^M , and e^M , which correspond to the probability of selecting the agent with the highest bid (or *high-bidder*), the other agent (or *low-bidder*), or option \emptyset .

Example 4.2. Consider a profile with expert valuations 1 for option \emptyset , 0.3 for option A , and 0 for option B and normalized bids of 1 and 0.9 from agents A and B , respectively. Consider a lottery which, for the particular profile, uses probabilities 0.4, 0.1, and 0.5 for options A , B , and \emptyset , respectively. The expert's and agents' views of the profile are

$$\begin{pmatrix} 1 & 0.3 & 0 \\ 0 & 1 & 0.9 \end{pmatrix} \text{ and } \begin{bmatrix} 0.3 & 0 & 1 \\ 1 & 0.9 & 0 \end{bmatrix},$$

respectively. The functions g^M , f^M , and η^M are defined over the 4-tuple of arguments $(x, h, \ell, z) = (0.3, 0, 1, 0.9)$ following the expert's view of the profile and take values 0.5, 0.4, and 0.1, respectively. Similarly, the functions d^M , c^M , and e^M are defined over the 4-tuple of arguments $(y, h, \ell, n) = (0.9, 0.3, 0, 1)$ following the agents' view of the profile and take values 0.4, 0.1, and 0.5, respectively. \square

In order to handle situations of equal values (e.g., equal bids), we adopt the convention to resolve ties using the fixed priority $A \succ B \succ \emptyset$ in order to identify the high- and low-bidder as well as the highest and/or lowest expert valuation. For example, if the expert has valuations of 1 for options \emptyset and B , we interpret this as option B being her most favourite one. Similarly, agent A is always the high-bidder and agent B is the low-bidder when their bids are equal. This is used in the definition of our mechanisms only; lower bound arguments do not depend on such assumptions in order to be as general as possible.

4.3.2 Reasoning about truthfulness

Let us now explain the truthfulness requirements having these profile representations in mind. There are two different kinds of possible misreports by the expert. In particular, the expert can attempt to make

- a *level change in the reported valuation* (ECh) by changing her second highest valuation without affecting the order of her valuations for the options, or
- a *reported valuation swap* (ESw), i.e., change the order of her valuations for the options as well as the particular values.

Example 4.3. The profile

$$\begin{pmatrix} 1 & 0.6 & 0 \\ 0.9 & 0 & 1 \end{pmatrix}$$

is the result of a reported valuation swap by the expert who changes her valuations from $(1, 0.3, 0)$ to $(0.6, 0, 1)$ for the options (\oslash, A, B) . \square

There are also two different kinds of possible misreports by each agent, who can attempt to make

- a *level change in the reported bid* (BCh) by changing her bid without affecting the order of bids, or
- a *bid swap* (BSw) by changing both the bid order and the corresponding values.

Example 4.4. The profile

$$\begin{bmatrix} 0 & 0.3 & 1 \\ 1 & 0.25 & 0 \end{bmatrix}$$

is the result of a bid swap deviation by the low-bidder, who increases her bid in the profile above to a new bid that is four times the bid of the other agent. \square

A truthful mechanism never incentivizes (i.e., it is *incentive compatible* with respect to) such misreportings. We use the terms ECh-IC, ESw-IC, BCh-IC, and BSw-IC to refer to incentive compatibility with respect to the misreporting attempts mentioned above. Therefore, a *truthful* mechanism satisfies all these IC conditions. Before we proceed, we provide a few examples of truthful mechanisms.

Example 4.5 (A bid-independent ordinal mechanism). Consider the following mechanism that ignores the bids reported by the agents. With probability $2/3$, output the expert's most preferred option and with probability $1/3$, output the expert's second most preferred option. Adopting the expert's view of profiles and the corresponding representation of the lottery P^M , the mechanism can be written as

$$g^M(x, h, \ell, z) = \frac{2}{3}, \quad f^M(x, h, \ell, z) = \frac{1}{3} \quad \text{and} \quad \eta^M(x, h, \ell, z) = 0.$$

The mechanism can be seen to be truthful by the fact that (a) ignores the bids of the agents and (b) it always assigns higher probability to the most-preferred outcome for the expert and 0 probability to the least-preferred outcome. Note that using the terminology above, any ordinal mechanism is ECh by construction, since changing the level in the reported valuation does not change the outcome. \square

Example 4.6 (A bid-independent non-ordinal mechanism). Consider the following mechanism that ignores the bids reported by the agents. Again, we adopt the expert's view of profiles and the corresponding representation of the lottery P^M ; recall that x is the value of the expert for her second most-preferred outcome. Let P^M be given by

$$g^M(x, h, \ell, z) = \frac{4 - x^2}{6}, \quad f^M(x, h, \ell, z) = \frac{1 + 2x}{6} \quad \text{and} \quad \eta^M(x, h, \ell, z) = \frac{1 - 2x + x^2}{6}.$$

Note that the mechanism uses the cardinal information of the expert's report and therefore it is not ordinal. This mechanism has been referred to in the literature as the *quadratic lottery* and has been proven to be truthful [Feige and Tennenholtz, 2010, Freixas, 1984]. \square

Example 4.7 (An expert-independent mechanism). Consider the following mechanism that ignores the expert's values for the different outcomes. Among the two agents, output the agent with the highest bid (breaking ties arbitrarily) and charge this agent a payment that is equal to the bid of the other agent. Charge the other agent a payment of 0. In terms of the agents' view, the outcome of the mechanism can be written as

$$d^M(y, h, \ell, n) = 1, \quad c^M(y, h, \ell, n) = 0 \quad \text{and} \quad e^M(y, h, \ell, n) = 0.$$

This mechanism is the well-known *second-price auction* [Vickrey, 1961], which is known (and easily seen) to be truthful. \square

It is not hard to observe that none of the mechanisms presented in Examples 4.5, 4.6 and 4.7 can achieve very strong approximation ratios. As we will see in Section 4.4, the mechanism of Example 4.5 is actually the best possible among the restricted class of ordinal mechanisms; later on, the use of cardinal information will allow us to decisively outperform it. We also note that while the second-price auction in Example 4.7 is welfare-optimal for the agents, which is a well-known fact, it can only provide a 2-approximation when it comes to our objective of the combined welfare of the agents and the expert.

We continue with important conditions that are necessary and sufficient for BCh-IC and ECh-IC. The next lemma is essentially the well-known characterization of [Myerson, 1981] for single-parameter domains.

Lemma 4.1 (Myerson, 1981). *A mechanism M is BCh-IC if and only if the functions d^M and c^M are non-increasing and non-decreasing in terms of their first argument, respectively.*

The correct interpretation of the lemma is that, as long as the output of a mechanism satisfies the monotonicity condition above, one can always find payments for the agents that will make the mechanism BCh-IC. In fact, when the mechanisms are required to charge a payment of zero to an agent with a zero bid, then these payments are uniquely defined, and are given by the following formula

$$p_i(w_i, w_{-i}) = w_i \cdot q_i(w_i, w_{-i}) - \int_0^{w_i} q_i(t, w_{-i}) dt,$$

where q_i is the probability that agent $i \in \{A, B\}$ gets selected as the outcome, p_i is the payment function, w_i is the bid of agent i and w_{-i} is the bid of the other agent. Therefore, we can avoid referring to the payment function when designing our mechanisms, as we can choose the above payment function, provided that the outcome probabilities satisfy the monotonicity conditions of Lemma 4.1. On the other hand, our lower bounds apply to all mechanisms, regardless of the payment function, as they only use the monotonicity condition.

Next, we provide a similar proof to that of Myerson [1981] for characterizing ECh-IC in our setting.

Lemma 4.2. *A mechanism M is ECh-IC if and only if the function f^M is non-decreasing in terms of its first argument and the function g^M satisfies*

$$g^M(x, h, \ell, z) = g^M(0, h, \ell, z) - x f^M(x, h, \ell, z) + \int_0^x f^M(t, h, \ell, z) dt, \quad (4.1)$$

for every 4-tuple (x, h, ℓ, z) representing a profile as seen by the expert.

As a corollary, functions g^M and h^M are non-increasing in terms of the first argument.

Proof. To shorten notation, we use $\mathbf{b} = (h, \ell, z)$ as an abbreviation of the information in the second row of a profile in expert's view and (x, \mathbf{b}) as an abbreviation of (x, h, ℓ, z) . Also, we drop M from notation (hence, $f(x, \mathbf{b})$ is used instead of $f^M(x, h, \ell, z)$) since it is clear from context. Due to ECh-IC, the expert has no incentive to attempt a level change of her utility for her second favourite option from x to x' . This means that

$$g(x, \mathbf{b}) + x f(x, \mathbf{b}) \geq g(x', \mathbf{b}) + x f(x', \mathbf{b}). \quad (4.2)$$

Similarly, she has no incentive to attempt a level change of her utility for her second favourite option from x' to x . This means that

$$g(x', \mathbf{b}) + x'f(x', \mathbf{b}) \geq g(x, \mathbf{b}) + x'f(x, \mathbf{b}). \quad (4.3)$$

By summing (4.2) and (4.3), we obtain that

$$(x - x')(f(x, \mathbf{b}) - f(x', \mathbf{b})) \geq 0,$$

which implies that f is non-decreasing in terms of its first argument.

To prove equation (4.1), we observe that inequality (4.2) yields

$$g(x, \mathbf{b}) + xf(x, \mathbf{b}) \geq g(x', \mathbf{b}) + x'f(x', \mathbf{b}) + (x - x')f(x', \mathbf{b}). \quad (4.4)$$

This means that function $g(x, \mathbf{b}) + xf(x, \mathbf{b})$ is convex with respect to its first argument and has f as its subgradient [Rockafellar, 2015]. Hence, from the standard results of convex analysis we get

$$g(x, \mathbf{b}) + xf(x, \mathbf{b}) = g(0, \mathbf{b}) + \int_0^x f(t, \mathbf{b}) dt,$$

which is equivalent to (4.1). □

Before we conclude the section, we remark here that while Lemma 4.2 will be fundamental for our proofs, it does not provide a characterization of all truthful one-voter mechanisms in the unrestricted social choice setting (such mechanisms are referred to as *unilateral* in the literature). The reason is that (a) it applies only to changes in the intensity of the preferences and not swaps in the ordering of alternatives and (b) it only provides conditions for three alternatives, as opposed to many alternatives in the general setting.

4.4 Ordinal mechanisms

We will consider several classes of truthful mechanisms depending on the level of information that they use. Let us warm up with some easy results on ordinal mechanisms, which do not use the exact values of the expert's report and the bids but only their relative order. It turns out that the best possible approximation ratio of such mechanisms is $3/2$ and is achieved by two symmetric mechanisms, one depending only on the ordinal information provided by the expert (expert-ordinal), while the other depends only on the relation between the bids (bid-ordinal).

The expert-ordinal mechanism EOM selects the expert's favourite and second best option with probabilities $2/3$ and $1/3$, respectively. Symmetrically, the bid-ordinal mechanism BOM selects the high- and low-bidder with probabilities $2/3$ and $1/3$, respectively.

Theorem 4.3. *Mechanisms EOM and BOM are truthful mechanisms that have approximation ratio at most $3/2$.*

Proof. Mechanism EOM is clearly truthful for the agents since it ignores their bids. It is also clearly truthful for the expert since the probabilities of selecting the options follow the order of the expert's valuations for them. BOM is clearly truthful for the expert (since her input is ignored); truthfulness for the agents follows by observing that the probability of selecting an agent is non-decreasing in terms of her bid.

We prove the approximation ratio for mechanism BOM only; the proof for the case of EOM is completely symmetric. Consider the profile $\begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$ in agents' view. We distinguish between two cases. If $1 + h \geq y + \ell$, the optimal welfare is $1 + h$ and the approximation ratio is

$$\frac{1 + h}{\frac{2}{3}(1 + h) + \frac{1}{3}(y + \ell)} \leq \frac{3}{2}$$

since $y + \ell \geq 0$. If $1 + h \leq y + \ell$, the optimal welfare is $y + \ell$ and the approximation ratio is

$$\frac{y + \ell}{\frac{2}{3}(1 + h) + \frac{1}{3}(y + \ell)} = \frac{1}{\frac{2}{3}\frac{1+h}{y+\ell} + \frac{1}{3}} \leq \frac{3}{2}$$

since $\frac{1+h}{y+\ell} \geq \frac{1}{2}$. □

We conclude this section by showing that both EOM and BOM are best possible among all ordinal mechanisms.

Theorem 4.4. *The approximation ratio of any ordinal mechanism is at least $3/2$.*

Proof. Let $\epsilon \in (0, 1/2)$ and consider the following two profiles:

$$\begin{pmatrix} 1 & \epsilon & 0 \\ 0 & \epsilon & 1 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & 1 - \epsilon & 0 \\ 0 & 1 - \epsilon & 1 \end{pmatrix}.$$

Since the order of the expert utilities and the bids is the same in both profiles, an ordinal mechanism behaves identically in all these profiles for every $\epsilon \in (0, 1/2)$. Assume that such a mechanism selects the middle option with probability p . Then, the approximation ratio of this mechanism is at least the maximum between its approximation ratio for these two profiles. Considering all profiles for $\epsilon \in (0, 1/2)$, we get an approximation ratio of at least

$$\sup_{\epsilon \in (0, 1/2)} \left\{ \frac{1}{1 - p + 2\epsilon p}, \frac{2(1 - \epsilon)}{1 - p + 2(1 - \epsilon)p} \right\} = \max \left\{ \frac{1}{1 - p}, \frac{2}{1 + p} \right\}.$$

This is minimized to $3/2$ for $p = 1/3$. □

4.5 Bid-independent mechanisms

In this section, we consider cardinal mechanisms but restrict our attention to ones that ignore the bids and base their decisions only on the expert's report. It is convenient to use the expert's view of profiles $\begin{pmatrix} 1 & x & 0 \\ h & \ell & z \end{pmatrix}$. Then, a *bid-independent* mechanism can be thought of as using univariate functions g^M , f^M , and η^M which indicate the probability of selecting the expert's first, second, and third favourite option when she has value $x \in [0, 1]$ for the second favourite option. We drop M from notation since the mechanism will be clear from context. The next lemma provides sufficient and necessary conditions for bid-independent mechanisms with good approximation ratio.

Lemma 4.5. *Let M be a bid-independent mechanism that uses functions g , f and η . Then M has approximation ratio at most ρ if and only if the inequalities*

$$2g(x) + xf(x) \geq 2/\rho \quad (4.5)$$

$$g(x) + (1+x)f(x) \geq (1+x)/\rho \quad (4.6)$$

hold for every $x \in [0, 1]$.

Proof. Consider the application of M on the profile $\begin{pmatrix} 1 & x & 0 \\ h & \ell & z \end{pmatrix}$. If $1+h \geq x+\ell$ the optimal welfare is $1+h$ and the approximation ratio is

$$\begin{aligned} \frac{1+h}{(1+h)g(x) + (x+\ell)f(x) + z\eta(x)} &\leq \frac{1+h}{(1+h)g(x) + (x+\ell)f(x)} \\ &\leq \frac{2}{2g(x) + xf(x)}. \end{aligned}$$

The first inequality follows since $z, \eta(x) \geq 0$ and the second one follows since the expression at the RHS is non-increasing in ℓ and non-decreasing in h . Then, the first inequality of the statement follows as a sufficient condition so that M has approximation ratio at most ρ . To see why it is also necessary, observe that the inequalities in the derivation above are tight for $h = 1$, $\ell = 0$, and $z = 0$.

If $1+h \leq x+\ell$ the optimal welfare is $x+\ell$ and the approximation ratio is

$$\begin{aligned} \frac{x+\ell}{(1+h)g(x) + (x+\ell)f(x) + z\eta(x)} &\leq \frac{x+\ell}{(1+h)g(x) + (x+\ell)f(x)} \\ &\leq \frac{1+x}{g(x) + (1+x)f(x)}. \end{aligned}$$

The first inequality follows since $z, \eta(x) \geq 0$ and the second one follows since the expression at the RHS is non-increasing in ℓ and non-decreasing in h . Then, the second inequality of the

statement follows as a sufficient condition so that M has approximation ratio at most ρ . To see why it is also necessary, observe that the two inequalities in the derivation above are tight for $h = 0$, $\ell = 1$, and $z = 0$. \square

Truthfulness of bid-independent mechanisms in terms of the agents follows trivially (since the bids are ignored). In order to guarantee truthfulness from the expert's side, we will use the characterization of ECh-IC from Lemma 4.2 together with additional conditions that will guarantee ESw-IC. These are provided by the next lemma.

Lemma 4.6. *An ECh-IC bid-independent mechanism is truthful if and only if the functions g , f , and η it uses satisfy $g(x) \geq f(x')$ and $f(x) \geq \eta(x')$ for every pair $x, x' \in (0, 1)$.*

Proof. We first show that the first condition is necessary. Assume that the first condition is violated, i.e., $f(x_1) > g(x_2)$ for two points $x_1, x_2 \in (0, 1)$. If $x_1 > x_2$, by the monotonicity of g we have $g(x_1) \leq g(x_2)$ and $f(x_1) > g(x_1)$. Otherwise, by the monotonicity of f , we have $f(x_2) \geq f(x_1)$ and $f(x_2) > g(x_2)$. In any case, there must exist $x^* \in (0, 1)$ such that $f(x^*) > g(x^*)$. Now consider the swap from expert valuation profile $(1, x^*, 0)$ to the profile $(x^*, 1, 0)$. The utility of the expert in the initial true profile is $g(x^*) + x^*f(x^*)$ while her utility at the new profile becomes $f(x^*) + x^*g(x^*)$, which is strictly higher.

Now, we show that the second condition is necessary. Again, assuming that the second condition is violated, we obtain that there is a point $x^* \in (0, 1)$ such that $\eta(x^*) > f(x^*)$. Now, the swap from expert's valuation profile $(1, x^*, 0)$ to the profile $(1, 0, x^*)$ increases the utility of the expert from $g(x^*) + x^*f(x^*)$ to $g(x^*) + x^*\eta(x^*)$, which is again strictly higher.

To show that the condition is sufficient for ECh-IC, we need to distinguish between five possible attempts for valuation swap by the expert.

Case 1. Consider the swap from the valuation profile $(1, x, 0)$ to the profile $(1, 0, x')$. The utility of the expert at the new profile is $g(x') + x\eta(x') \leq g(0) + \int_0^x f(t) dt = g(x) + xf(x)$, where the inequality holds due to the fact that $\eta(x') \leq f(t)$, for every $t \in [0, x]$. Observe that the RHS of the derivation is the expert's utility at the initial true profile.

Case 2. Consider the swap from the valuation profile $(1, x, 0)$ to the profile $(x', 1, 0)$. The utility of the expert at the new profile is $f(x') + xg(x') \leq g(x') + xf(x') = g(x') + x'f(x') + (x - x')f(x') \leq g(x) + xf(x)$, which is her utility at the initial true profile. The first inequality follows by the condition $g(x') \geq f(x)$ of the lemma and the second one is due to the convexity of function

$g(x) + xf(x)$. See also the proof of Lemma 4.2.

Case 3. Consider the swap from the valuation profile $(1, x, 0)$ to the profile $(x', 0, 1)$. The utility of the expert at the new profile is $f(x') + x\eta(x')$, which is at most $g(x) + xf(x)$ due to the conditions of the lemma.

Case 4. Consider the swap from the valuation profile $(1, x, 0)$ to the profile $(0, x', 1)$. The utility of the expert at the new profile is $\eta(x') + xf(x') \leq f(x) + xg(x) \leq g(x) + xf(x)$, which is her utility at the initial true profile.

Case 5. Consider the swap from the valuation profile $(1, x, 0)$ to the profile $(0, 1, x')$. The utility of the expert at the new profile is $\eta(x') + xg(x') \leq f(x') + xg(x')$ and the proof proceeds as in Case 2 above. \square

We are now ready to propose our mechanism BIM. Let $\tau = -W\left(-\frac{1}{2e}\right)$, where W is the Lambert function, i.e., τ is the solution of the equation $2\tau = e^{\tau-1}$. Mechanism BIM is defined as follows:

$$f(x) = \begin{cases} \frac{\tau}{1+3\tau}, & x \in [0, \tau] \\ \frac{1+\tau-2\tau e^{1-x}}{1+3\tau}, & x \in [\tau, 1] \end{cases}$$

$$g(x) = \begin{cases} \frac{1+\tau}{1+3\tau}, & x \in [0, \tau] \\ \frac{2\tau(1+x)e^{1-x}}{1+3\tau}, & x \in [\tau, 1] \end{cases}$$

$$\eta(x) = \begin{cases} \frac{\tau}{1+3\tau}, & x \in [0, \tau] \\ \frac{2\tau(1-x)e^{1-x}}{1+3\tau}, & x \in [\tau, 1] \end{cases}$$

BIM is depicted in Figure 4.1. All functions are constant in $[0, \tau]$ and have (admittedly, counter-intuitive at first glance) exponential terms in $[\tau, 1]$. Interestingly, as we will show later, this is the unique best possible bid-independent truthful mechanism. Its properties are proved in the next statement.

Theorem 4.7. *Mechanism BIM is truthful and has approximation ratio at most*

$$\frac{1 - 3W\left(-\frac{1}{2e}\right)}{1 - W\left(-\frac{1}{2e}\right)} \approx 1.37657,$$

where W is the Lambert function.




Figure 4.1: The lottery used by mechanism BIM.

Proof. Tedious calculations can verify that BIM is truthful. The function f is non-decreasing in x and g is defined exactly as in equation (4.1); hence, ECh-IC follows by Lemma 4.2. ESw-IC follows since f , g , and h satisfy the conditions of Lemma 4.6.

Now, let $\rho = \frac{1+3\tau}{1+\tau}$. We use the definition of BIM and Lemma 4.5 to show the bound on the approximation ratio. If $x \in [0, \tau]$, inequalities (4.5) and (4.6) are clearly satisfied since $x \geq 0$ and $x \leq \tau$, respectively. If $x \in [\tau, 1]$, we have

$$2g(x) + xf(x) = 2 \frac{2\alpha(1+x)e^{1-x}}{1+3\alpha} + x \frac{1+\alpha-2\alpha e^{1-x}}{1+3\alpha},$$

which is minimized for $x = \tau$ (recall that $2\tau = e^{\tau-1}$) at $\frac{2+2\tau+\tau^2}{1+3\tau} \geq 2/\rho$. Hence, inequality (4.5) holds. Also, inequality (4.6) can be easily seen to hold with equality. \square

We now show that BIM is optimal among all bid-independent truthful mechanisms. The proof of the next theorem exploits the characterization of ECh-IC mechanisms from Lemma 4.2, the characterization of ESw-IC bid-independent mechanisms from Lemma 4.6, and Lemma 4.5.

Theorem 4.8. *The approximation ratio of any truthful bid-independent mechanism is at least*

$$\frac{1 - 3W\left(-\frac{1}{2e}\right)}{1 - W\left(-\frac{1}{2e}\right)} \approx 1.37657,$$

where W is the Lambert function.

Proof. Let M be a bid-independent mechanism that uses functions g , f , and h to define the probability of selecting the expert's first, second, and third favourite option and has approximation ratio $\rho \geq 1$. By the necessary condition (4.1) for ECh-IC in Lemma 4.2, we know

that

$$g(x) = g(0) - xf(x) + \int_0^x f(t) dt. \quad (4.7)$$

Let α be any value in $[0, 1]$.

Due to the fact that $f(1) + g(1) \leq 1$, we have

$$g(0) + \int_0^1 f(t) dt \leq 1. \quad (4.8)$$

By the necessary condition for ESw-IC in Lemma 4.6 and since g is non-increasing (by Lemma 4.2), we also have $f(x) \geq \eta(x) = 1 - f(x) - g(x) \geq 1 - f(x) - g(0)$, i.e., $g(0) + 2f(x) \geq 1$, for $x \in (0, 1)$. Integrating in the interval $(0, \alpha]$, we get

$$\alpha g(0) + 2 \int_0^\alpha f(t) dt \geq \alpha. \quad (4.9)$$

Since, the mechanism is ρ -approximate, Lemma 4.5 yields

$$g(0) \geq 1/\rho \quad (4.10)$$

(by applying inequality (4.5) with $x = 0$) and

$$g(x) + (1+x)f(x) \geq (1+x)/\rho, \forall x \in [\alpha, 1].$$

Using (4.7), this last inequality becomes

$$g(0) + f(x) + \int_0^x f(t) dt \geq (1+x)/\rho, \forall x \in [\alpha, 1].$$

Now, let λ be a continuous function with $\lambda(x) \leq f(x)$ in $[\alpha, 1]$ such that

$$g(0) + \int_0^\alpha f(t) dt + \int_\alpha^x \lambda(t) dt + \lambda(x) = (1+x)/\rho.$$

Setting $\Lambda(x) = \int_\alpha^x \lambda(t) dt$ (clearly, Λ is differentiable due to the continuity of λ in $[0, 1]$), we get the differential equation

$$g(0) + \int_0^\alpha f(t) dt + \Lambda(x) + \Lambda'(x) = (1+x)/\rho$$

which, given that $\Lambda(\alpha) = 0$, has the solution

$$\Lambda(x) = \frac{x}{\rho} - g(0) - \int_0^\alpha f(t) dt + \left(g(0) - \frac{\alpha}{\rho} + \int_0^\alpha f(t) dt \right) \exp(\alpha - x)$$

for $x \in [\alpha, 1]$. Hence,

$$\int_\alpha^1 f(t) dt \geq \Lambda(1) = \frac{1 - \alpha e^{\alpha-1}}{\rho} - (1 - e^{\alpha-1}) g(0) - (1 - e^{\alpha-1}) \int_0^\alpha f(t) dt. \quad (4.11)$$

Now, by multiplying inequalities (4.8), (4.9), (4.10), and (4.11) by coefficients 2 , $e^{\alpha-1}$, $(2 - \alpha)e^{\alpha-1}$, and 2 , respectively, and then summing them, we obtain

$$\rho \geq \frac{2 - \alpha e^{\alpha-1}}{2 - 3\alpha e^{\alpha-1} + 2e^{\alpha-1}}.$$

Picking $\alpha = -W\left(-\frac{1}{2e}\right)$ (i.e., α is the solution of the equation $e^{\alpha-1} = 2\alpha$), we get that

$$\rho \geq \frac{1 - 3W\left(-\frac{1}{2e}\right)}{1 - W\left(-\frac{1}{2e}\right)}.$$

This completes the proof. \square

4.6 Expert-independent mechanisms

Here, we consider mechanisms that depend only on the bids. Now, it is convenient to use the agents' view of profiles $\begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$. Then, an expert-independent mechanism can be thought of as using univariate functions d^M , c^M , and e^M which indicate the probability of selecting the high-bidder, the low-bidder, and the option \oslash in terms of the normalized low-bid y . Again, we drop M from notation. Following the same roadmap as in the previous section, the next lemma provides sufficient and necessary conditions for expert-independent mechanisms with good approximation ratio.

Lemma 4.9. *Let M be an expert-independent mechanism that uses functions d , c , and e with $d(y) = 1 - c(y)$ and $e(y) = 0$ for $y \in [0, 1]$. If*

$$\frac{1}{\rho} - \frac{1 - 1/\rho}{y} \leq c(y) \leq \frac{2(1 - 1/\rho)}{2 - y} \quad (4.12)$$

for every $y \in [0, 1]$, then M has approximation ratio at most ρ . Condition (4.12) is necessary for every ρ -approximate expert-independent mechanism.

Proof. Consider the application of M on the profile $\begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$. We distinguish between two cases. If $1 + h \geq y + \ell$, assuming that condition (4.12) is true, the approximation ratio of M is

$$\begin{aligned} \frac{1 + h}{(y + \ell)c(y) + (1 + h)(1 - c(y))} &= \frac{1}{\frac{y + \ell}{1 + h}c(y) + 1 - c(y)} \\ &\leq \frac{1}{1 - (1 - y/2)c(y)} \leq \rho. \end{aligned}$$

The first inequality follows since $\frac{y + \ell}{1 + h} \geq y/2$ when $y \in [0, 1]$, while the second one is essentially the right inequality in condition (4.12).




Figure 4.2: A pictorial view of the statement in Lemma 4.9.

Otherwise, if $1 + h \leq y + \ell$, the approximation ratio of M is

$$\begin{aligned} \frac{y + \ell}{(y + \ell)c(y) + (1 + h)(1 - c(y))} &= \frac{1}{c(y) + \frac{1+h}{y+\ell}(1 - c(y))} \\ &\leq \frac{1 + y}{1 + yc(y)} \leq \rho. \end{aligned}$$

The first inequality follows since $\frac{1+h}{y+\ell} \geq \frac{1}{1+y}$ when $y \in [0, 1]$; again, the second one is essentially the left inequality in condition (4.12).

To see that condition (4.12) is necessary for every mechanism, first consider a mechanism M' that uses functions \bar{c} , \bar{d} , and \bar{e} such that the function \bar{c} violates the left inequality in (4.12), i.e., $\bar{c}(y^*) < \frac{1}{\rho} - \frac{1-1/\rho}{y^*}$ for some $y^* \in [0, 1]$. Then, using this inequality and the fact that $\bar{d}(y^*) \leq 1 - \bar{c}(y^*)$, the approximation ratio of M' at profile $\begin{bmatrix} 0 & 1 & 0 \\ 1 & y^* & 0 \end{bmatrix}$ is

$$\frac{y^* + 1}{(y^* + 1)\bar{c}(y^*) + \bar{d}(y^*)} \geq \frac{1 + y^*}{1 + y^*\bar{c}(y^*)} > \rho.$$

Now, assume that function \bar{c} violates the right inequality in (4.12), i.e., $\bar{c}(y^*) > \frac{2(1-1/\rho)}{2-y^*}$. Then, using this inequality and the fact that $\bar{d}(y^*) \leq 1 - \bar{c}(y^*)$, the approximation ratio of M' at profile $\begin{bmatrix} 1 & 0 & 0 \\ 1 & y^* & 0 \end{bmatrix}$ is

$$\frac{2}{2\bar{d}(y^*) + y^*\bar{c}(y^*)} \geq \frac{2}{2 - (2 - y^*)\bar{c}(y^*)} > \rho$$

as desired. \square

Figure 4.2 shows the available space (grey area) for the definition of function $c(y)$, so that the corresponding mechanism has an approximation ratio of at most $\rho = 7 - 4\sqrt{2}$. It can be easily verified that this is the minimum value for which the LHS of condition (4.12) in Lemma 4.9 is smaller than or equal to the RHS so that a function satisfying (4.12) does exist.

Our aim now is to define an expert-independent truthful mechanism achieving the best possible approximation ratio of $\rho = 7 - 4\sqrt{2}$. Since the expert's report is ignored, truthfulness for the expert follows trivially. We restrict our attention to the design of a mechanism that never selects option \emptyset , i.e., it has $d(y) = 1 - c(y)$ for every $y \in [0, 1]$. Lemmas 4.1 and 4.9 guide this design as follows. In order to be BCh-IC and ρ -approximate, our mechanism should use a non-decreasing function $c(y)$ in the space available by condition (4.12). Still, we need to guarantee BSw-IC; the next lemma gives us the additional sufficient (and necessary) condition.

Lemma 4.10. *A BCh-IC expert-independent mechanism is truthful if and only if $d(1) \geq c(1)$.*

Proof. Consider an attempted bid swap according to which the low-bidder increases her normalized bid of y so that it becomes the high-bidder and the normalized bid of the other agent is y' . Essentially, this attempted bid swap modifies the initial profile $\begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$ to $\begin{bmatrix} \ell & h & n \\ 1 & y' & 0 \end{bmatrix}$. The deviating agent corresponds to the middle column in the initial profile and has probability $c(y)$ of being selected. In the new profile, she corresponds to the first column, and has probability $d(y')$ of being selected. So, the necessary and sufficient condition so that BSw-IC is guaranteed is $c(y) \leq d(y')$ for every $y, y' \in [0, 1]$. Since, by Lemma 4.1, c and d are non-decreasing and non-increasing, respectively, this condition boils down to $d(1) \geq c(1)$.

The case in which the high-bidder decreases her bid so that it gets a normalized value of y' is symmetric. \square

We are ready to propose our mechanism EIM, which uses functions

$$c(y) = \begin{cases} \frac{2(1-1/\rho)}{2-y}, & y \in [0, \frac{3-\rho}{2}] \\ \frac{1}{\rho} - \frac{1-1/\rho}{y}, & y \in [\frac{3-\rho}{2}, 1] \end{cases}$$

for $\rho = 7 - 4\sqrt{2}$ and $d(y) = 1 - c(y)$ for $y \in [0, 1]$.

Essentially, EIM uses the blue line in the upper right plot of Figure 4.2, which consists of the curve that upper bounds the grey area up to point $\frac{3-\rho}{2} = 2\sqrt{2} - 2$ and the curve that lower-bounds the grey area after that point. The properties of mechanism EIM are summarized in the next statement. It should be clear though that the statement holds for every mechanism that uses a non-decreasing function in the grey area that is below $1/2$ (together with the restriction $d(y) = 1 - c(y)$, this is necessary and sufficient for BSw-IC). Given the discussion about the optimality of $\rho = 7 - 4\sqrt{2}$ above, all these mechanisms are optimal within the class of expert-independent mechanisms.

Theorem 4.11. *Mechanism EIM is truthful and has approximation ratio at most $7 - 4\sqrt{2} \approx 1.3431$. This ratio is optimal among all truthful expert-independent mechanisms.*

4.7 Beyond expert-independent mechanisms

In this section, we present a template for the design of better truthful mechanisms, compared to those presented in the previous sections. The template strengthens expert-independent mechanisms by exploiting a single additional bit of information that allows to distinguish between profiles that have the same (normalized) bid values.

We denote by \mathcal{T} the set of mechanisms that are produced according to our template. In order to define a mechanism $M \in \mathcal{T}$, it is convenient to use the agents' view of a profile as $\begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$. We partition the profiles of \mathcal{D} into two categories. Category $T1$ contains all profiles with $\ell > h$ or with $\ell = h$ such that the tie between the expert valuations ℓ and h is resolved in favour of the low-bidder. All other profiles belong to category $T2$.

For each profile in category $T1$, mechanism M selects the low-bidder with probability $c(y, T1)$ that is non-decreasing in y and the high-bidder with probability $1 - c(y, T1)$. For each profile in category $T2$, mechanism M selects the low-bidder with probability 0, and the high-bidder with probability 1. Different mechanisms following our template are defined using different functions $c(y, T1)$. The mechanisms of the template ignore neither the bids nor the expert's report; still, it is not hard to show that they are truthful.

Lemma 4.12. *Every mechanism $M \in \mathcal{T}$ is truthful.*

Proof. We first show that M is truthful for the agents. BCh-IC follows easily by Lemma 4.1, since $c(y, T1)$ and $c(y, T2)$ are non-decreasing in y . To show BSw-IC, notice that a bid swap attempt in a profile of category $T1$ yields a profile of category $T2$, and vice versa. This involves either a high-bidder who decreases her bid to become the low-bidder in the new profile, or the low-bidder who increases her bid to become the high-bidder in the new profile. In both cases, the increase or decrease in the selection probability according to M follows the increase or decrease of the deviating bid.

To show that M is truthful for the expert, first observe that according to the expert's view, the lottery uses constant functions f , g , and h in terms of her valuation for her second favourite option. Hence, Lemma 4.2 implies ECh-IC. To show ESw-IC, observe again that an expert's report swap attempt from a profile of category $T1$ creates a profile of category $T2$ and vice versa.

The expected utility that M yields to the expert in the initial profile is $\ell c(y, T1) + h(1 - c(y, T1)) = h + (\ell - h)c(y, T1) \geq h$ if it is of category $T1$ and $h + (\ell - h)c(y, T2) = h$ if it is of category $T2$. After the deviation, the utility of the expert becomes $\ell c(y, T1) + h(1 - c(y, T1)) = h + (\ell - h)c(y, T1) \leq h$ if the new profile is of category $T1$ and $h + (\ell - h)c(y, T2) = h$ if it is of category $T2$. Hence, such a swap attempt is never profitable for the expert. \square

The next lemma is useful for proving bounds on the approximation ratio of mechanisms in the template class \mathcal{T} .

Lemma 4.13. *Let M be a mechanism of \mathcal{T} and $\rho \geq 1$ be such that the function $c(y, T1)$ used by M satisfies*

$$\frac{1}{\rho} - \frac{1 - 1/\rho}{y} \leq c(y, T1) \leq \frac{1 - 1/\rho}{1 - y}.$$

Then, M has approximation ratio at most ρ .

Proof. Clearly, the approximation ratio of M in profiles of category $T2$ is always 1 since the mechanism takes the optimal decision of selecting the high-bidder with probability 1.

Now, consider a profile $\begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$ of category $T1$, i.e., $\ell \geq h$. We distinguish between two cases. If $1 + h \geq y + \ell$, then the approximation ratio of M is

$$\begin{aligned} \frac{1 + h}{(y + \ell)c(y, T1) + (1 + h)(1 - c(y, T1))} &= \frac{1}{1 - c(y, T1) + \frac{y + \ell}{1 + h}c(y, T1)} \\ &\leq \frac{1}{1 - (1 - y)c(y, T1)} \leq \rho. \end{aligned}$$

The first inequality follows since $\frac{y + \ell}{1 + h} \geq y$ when $y \in [0, 1]$ and $\ell \geq h \geq 0$ while the second one is due to the right inequality in the condition of the lemma.

Otherwise, if $1 + h \leq y + \ell$, the approximation ratio of M is

$$\begin{aligned} \frac{y + \ell}{(y + \ell)c(y, T1) + (1 + h)(1 - c(y, T1))} &= \frac{1}{c(y, T1) + \frac{1 + h}{y + \ell}(1 - c(y, T1))} \\ &\leq \frac{1 + y}{1 + yc(y, T1)} \leq \rho. \end{aligned}$$

The first inequality follows since $\frac{1 + h}{y + \ell} \geq \frac{1}{1 + y}$ when $y \in [0, 1]$ and $h \geq \ell \geq 0$; again, the second one is due to the left inequality in the condition of the lemma. \square

The conditions of Lemma 4.13 are depicted in the two plots of Figure 4.3 (for $\rho = 5/4$ and $\rho = \phi$, respectively). The grey area represents the available space for the definition of the (non-decreasing) function $c(y, T1)$ that a mechanism of \mathcal{T} should use on profiles of category $T1$ so that its approximation ratio is at most ρ .




Figure 4.3: Pictorial views of the statement in Lemma 4.13 for $\rho = 5/4$ (left) and $\rho = \phi$ (right).

These plots explain the definition of the next two mechanisms that follow our template: the randomized mechanism R and the deterministic mechanism D . For each profile of category $T1$, mechanism R uses

$$c^R(y, T1) = \begin{cases} \frac{1}{5(1-y)}, & y \in [0, 4/5] \\ 1, & y \in [4/5, 1] \end{cases}$$

(i.e., the blue line in the lower left plot of Figure 4.2) and mechanism D uses

$$c^D(y, T1) = \begin{cases} 0, & y \in [0, 1/\phi] \\ 1, & y \in [1/\phi, 1] \end{cases}$$

(i.e., the blue line in the lower right plot of Figure 4.2), where $\phi = \frac{1+\sqrt{5}}{2} \approx 1.618$ is the golden ratio. Their properties are as follows.

Theorem 4.14. *Mechanisms R and D are $5/4$ - and ϕ -approximate truthful mechanisms, respectively.*

Proof. Since $R, D \in \mathcal{T}$, truthfulness follows by Lemma 4.12. Their approximation ratios follow by Lemma 4.13 for $\rho = 5/4$ and $\rho = \phi$, respectively. \square

We remark that the condition of Lemma 4.13 can be proved to be not only sufficient but also necessary for achieving a ρ -approximation. Then, it can be easily seen that the value of $5/4$ is the lowest value for which the condition of the lemma is feasible. Hence, mechanism R is best possible among mechanisms that use our template. More interestingly, $5/4$ turns out to be the lower bound of any mechanism that always sells the item, as we prove in the next theorem. Mechanism D will be proved to be optimal among all deterministic truthful mechanisms in the next section.

Theorem 4.15. *The approximation ratio of any mechanism that always sells the item is at least $5/4$.*

Proof. Consider preference profiles in agents' view $\begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$, and let M be any truthful always-sell mechanism. Then, M can be thought of as using functions $d(y, h, \ell, n)$, $c(y, h, \ell, n)$ and $e(y, h, \ell, n)$ to assign probabilities to the high-bidder, the low-bidder and the no-sale option, respectively, such that $d(y, h, \ell, n) = 1 - c(y, h, \ell, n)$ and $e(y, h, \ell, n) = 0$.

Since M is truthful for the expert, the expert does not have any incentive to misreport her valuations from (h, ℓ, n) to (h', ℓ', n') , for any $\ell > h$ and $\ell' > h'$. This means that

$$h \cdot (1 - c(y, h, \ell, n)) + \ell \cdot c(y, h, \ell, n) \geq h' \cdot (1 - c(y, h', \ell', n')) + \ell' \cdot c(y, h', \ell', n')$$

or, equivalently, since $\ell > h$,

$$c(y, h, \ell, n) \geq c(y, h', \ell', n') \quad (4.13)$$

Similarly, the expert does not have incentive to misreport her valuations from (h', ℓ', n') to (h, ℓ, n) , for any $\ell > h$ and $\ell' > h'$. This gives us that

$$h' \cdot (1 - c(y, h', \ell', n')) + \ell' \cdot c(y, h', \ell', n') \geq h' \cdot (1 - c(y, h, \ell, n)) + \ell' \cdot c(y, h, \ell, n)$$

or, equivalently, since $\ell' > h'$,

$$c(y, h', \ell', n') \geq c(y, h, \ell, n) \quad (4.14)$$

Therefore, by (4.13) and (4.14), we have that $c(y, h, \ell, n)$ is constant in all profiles $\begin{bmatrix} h & \ell & n \\ 1 & y & 0 \end{bmatrix}$ with $\ell > h$.

Now, let $\epsilon \in (0, 1/2)$ and consider the following two profiles:

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & 1/2 & 0 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & \epsilon & 1 \\ 1 & 1/2 & 0 \end{bmatrix}$$

Since $\ell > h$ in both profiles, any truthful mechanism M that always sells the item behaves identically in all such profiles, for any $\epsilon \in (0, 1/2)$. Hence, assume that such a mechanism M selects the low-bidder with probability p (and the high-bidder with probability $1 - p$). Then, the approximation ratio of M is at least the maximum between its approximation ratio for these profiles, i.e.,

$$\sup_{\epsilon \in (0, 1/2)} \left\{ \frac{\frac{3}{2}}{1 - p + \frac{3}{2}p}, \frac{1}{1 - p + (\epsilon + \frac{1}{2})p} \right\} = \max \left\{ \frac{3}{2 + p}, \frac{2}{2 - p} \right\}.$$

This is minimized to $5/4$ for $p = 2/5$. □

4.8 Unconditional lower bounds

In the previous sections, we presented (or informally discussed) lower bounds on the approximation ratio of truthful mechanisms belonging to particular classes. Here, we present our most general lower bound that holds for every truthful mechanism. The proof exploits the ECh-IC characterization from Lemma 4.2.

Theorem 4.16. *The approximation ratio of any truthful mechanism is at least 1.14078.*

Proof. Let $\gamma \in [0, 1]$ be such that $1 - 2\gamma - 4\gamma^2 - 2\gamma^3 = 0$ and $\beta = (1 + \gamma)^{-1}$. The corresponding values are $\beta \approx 0.7709$ and $\gamma \approx 0.29716$.

Consider any ρ -approximate truthful mechanism and the profiles

$$\begin{pmatrix} 1 & \beta & 0 \\ \gamma & 1 & 0 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & 0 & 0 \\ \gamma & 1 & 0 \end{pmatrix}.$$

Since the bids are identical in both profiles, we can assume that the functions f and g are univariate (depending only on the expert's second highest utility). Since the mechanism is ρ -approximate in both profiles, we have

$$(1 + \gamma)g(\beta) + (1 + \beta)f(\beta) \geq \frac{1 + \beta}{\rho} \quad (4.15)$$

and

$$(1 + \gamma)g(0) + f(0) \geq \frac{1 + \gamma}{\rho}. \quad (4.16)$$

By condition (4.1) in Lemma 4.2, $g(x) = g(0) - xf(x) + \int_0^x f(t) dt$ which, due to the fact that f is non-decreasing (again by Lemma 4.2), yields $\int_0^\beta f(t) dt \geq \beta f(0)$. Hence,

$$g(x) \geq g(0) - \beta f(\beta) + \beta f(0). \quad (4.17)$$

Also, clearly,

$$1 \geq g(\beta) + f(\beta). \quad (4.18)$$

Now, multiplying inequalities (4.15), (4.16), (4.17), and (4.18) by the coefficients $\frac{\gamma}{\beta + 2\beta\gamma - \gamma^2}$, $\frac{\beta - \gamma}{\beta + 2\beta\gamma - \gamma^2}$, $\frac{(\beta - \gamma)(1 + \gamma)}{\beta + 2\beta\gamma - \gamma^2}$, and $\frac{\beta(1 + \gamma)}{\beta + 2\beta\gamma - \gamma^2}$, and by summing them, we get

$$\rho \geq \frac{\beta + 2\beta\gamma - \gamma^2}{\beta(1 + \gamma)}.$$

Substituting β and γ , we obtain that $\rho \geq 1.14078$ as desired. \square

Our last statement shows that mechanism D , which was presented in Section 4.7, is optimal among all deterministic truthful mechanisms.

Theorem 4.17. *No truthful deterministic mechanism has approximation ratio better than ϕ .*

Proof. Let M be a deterministic truthful mechanism. Consider a profile $\begin{pmatrix} 1 & x & 0 \\ h & \ell & z \end{pmatrix}$ for some combination of values for h , ℓ , and z . We first show that M selects the same option for every value of $x \in (0, 1)$. Indeed, assume otherwise; due to Lemma 4.2, f must be non-decreasing in x and, hence, $f(x_1, h, \ell, z) = 0$ and $f(x_2, h, \ell, z) = 1$ for two different values x_1 and x_2 in $(0, 1)$ with $x_1 < x_2$. Let $x_3 \in (x_2, 1)$, i.e., $f(x_3, h, \ell, z) = 1$ due to monotonicity. Property (4.1) in Lemma 4.2 requires that

$$g(x_3, h, \ell, z) = g(0, h, \ell, z) - x_3 + \int_0^{x_3} f(t, h, \ell, z) dt.$$

By our assumptions on f (and due to its monotonicity), we also have that

$$x_3 - x_2 \leq \int_0^{x_3} f(t, h, \ell, z) dt \leq x_3 - x_1.$$

These last two (in)equalities imply that $g(0, h, \ell, z) - g(x_3, h, \ell, z)$ lies between x_2 and x_3 , i.e., it is non-integer. This contradicts the fact that M is deterministic.

Now let $\epsilon > 0$ be negligibly small and consider the two profiles

$$\begin{pmatrix} 1 & 1 - \epsilon & 0 \\ 0 & 1/\phi & 1 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & \epsilon/\phi^2 & 0 \\ 0 & 1/\phi & 1 \end{pmatrix}.$$

If M selects the low-bidder in both profiles, its approximation ratio at the right one is $\frac{1}{\epsilon/\phi^2 + 1/\phi} \geq \phi - \epsilon$. Otherwise, its approximation ratio at the left profile is $1 + 1/\phi - \epsilon$. In any case, the approximation ratio is at least $\phi - \epsilon$, and the proof is complete. \square

Of course, Theorem 4.17 is meaningful for cardinal mechanisms. Deterministic ordinal mechanisms can be easily seen to be at least 2-approximate.

4.9 Conclusion

In this chapter we focused on designing truthful mechanisms for a simple ownership transfer scenario with two agents and one expert. The agents have monetary values for an item that is up for sale, while the expert has values over the different options of selling to some agent or not selling at all. The goal was to design mechanisms to incentivize all parties to truthfully report their values, while at the same time maximize the social welfare that takes into account the values of the agents and the expert for the outcome.

We considered several different classes of truthful mechanisms depending on the level of information that they used. For each such class, we identified the best possible mechanism in terms of the approximation ratio of the optimal social welfare. Indicatively, we showed an unconditional lower bound of 1.14 on the approximation ratio of all mechanisms, and designed a particular randomized mechanism with approximation ratio of 1.25, that uses the cardinal information provided by the agents as well as a single extra bit of information by the expert that allows for a classification of the possible valuation profiles.

Chapter 5

Near-optimal asymmetric binary matrix partitions

In this chapter, we study the asymmetric binary matrix partition problem and design simple algorithms with improved approximation guarantees; for a short introduction to the problem and possible applications for revenue maximization in take-it-or-leave-it sales by exploiting information asymmetry see the discussion in Section 1.4. The results that we present here have been published in [Abed et al., 2018].

5.1 Problem definition and overview of contribution

Consider an $n \times m$ matrix A with non-negative entries and a probability distribution p over its columns; p_j denotes the probability associated with column j . We distinguish between two cases for the probability distribution over the columns of the given matrix, depending on whether it is uniform or non-uniform. A partition scheme $B = (B_1, \dots, B_n)$ for matrix A consists of a partition B_i of $[m]$ for each row i of A . More specifically, B_i is a collection of k_i pairwise disjoint subsets $B_{ik} \subseteq [m]$ (with $1 \leq k \leq k_i$) such that $\bigcup_{k=1}^{k_i} B_{ik} = [m]$. We can think of each partition B_i as a smoothing operator, which acts on the entries of row i and changes their value to the expected value of the partition subset they belong to. Formally, the *smooth value* of an entry (i, j) such that $j \in B_{ik}$ is defined as

$$A_{ij}^B = \frac{\sum_{\ell \in B_{ik}} p_\ell \cdot A_{i\ell}}{\sum_{\ell \in B_{ik}} p_\ell}. \quad (5.1)$$

Notice that all entries (i, j) such that $j \in B_{ik}$ have the same smooth value. Given a scheme B that induces the *smooth matrix* A^B , the resulting *partition value* is the expected maximum column

entry of A^B , namely,

$$v^B(A, p) = \sum_{j \in [m]} p_j \cdot \max_i A_{ij}^B. \quad (5.2)$$

The objective of the *asymmetric binary matrix partition problem* is to find a partition scheme B such that the resulting partition value $v^B(A, p)$ is maximized.

The problem was first introduced by Alon et al. [2013] who proved it to be APX-hard even for input matrices containing binary values and uniform probability distributions. Further, for the binary version, they presented a 0.563- and a $1/13$ -approximation algorithms for the cases where the probability distribution over the columns of the input matrix is uniform and non-uniform, respectively. We significantly improve both of these results, by designing a $9/10$ -approximation algorithm for the uniform case and a $(1 - 1/e)$ -approximation algorithm for non-uniform distributions.

For the uniform case, the algorithm of Alon et al. [2013] use several interesting phases. We borrow two of them, namely a *covering* and a *greedy completion* phase, which we put together into an intuitive greedy algorithm. Despite the purely combinatorial nature of this algorithm, to analyze it and prove its approximation ratio, we exploit linear programming techniques and duality.

For non-uniform distributions, we exploit a nice relation of asymmetric matrix partition to submodular welfare maximization, and use well-known algorithms from the literature. First, we discuss the application of a simple greedy $1/2$ -approximation algorithm that has been studied by Lehmann et al. [2006]. Then, we apply the smooth greedy algorithm of Vondrák [2008] to achieve a $(1 - 1/e)$ -approximation for our problem, which is optimal in the value query model due to Khot et al. [2008]. In a more powerful model where it is assumed that demand queries can be answered efficiently, Feige and Vondrák [2010] proved that $(1 - 1/e + \epsilon)$ -approximation algorithms are possible, where ϵ is a small positive constant. We briefly discuss the possibility/difficulty of applying such algorithms to asymmetric binary matrix partition and observe that the corresponding demand query problems are, in general, NP-hard.

5.1.1 Chapter roadmap

In the following, shortly discuss other related work in Section 5.2. Then, we give preliminary definitions and examples, and prove several important structural observations in Section 5.3. We present our $9/10$ -approximation algorithm for the case where the probability distribution over the columns of the matrix is uniform in Section 5.4. Our $(1 - 1/e)$ -approximation algorithm

for the non-uniform case is analyzed in Section 5.5. We conclude with a short synopsis in Section 5.6.

5.2 Related work

Apart from the binary version of the asymmetric matrix partition problem, Alon et al. [2013] considered also the more general case of input matrices with non-binary entries, for which they presented a $1/2$ - and an $\Omega(1/\log m)$ -approximation algorithm for uniform and non-uniform distributions, respectively. A common idea underlying these results is that they try to identify a set of high-value entries that can be bundled together with other entries in order to increase the total contribution.

The possible application of asymmetric matrix partition to revenue maximization in take-it-or-leave-it sales (see Section 1.4) falls within the line of research that studies the impact of information asymmetry to the quality of markets. Akerlof [1970] was the first to introduce a formal analysis of the *markets for lemons*, where the seller has much more accurate information than the buyers regarding the quality of the products.

The particular approach of partitioning in take-it-or-leave-it sales is closer in spirit to the strategic information transmission that was initiated in the work of Crawford and Sobel [1982], where the seller has information about the valuations of the buyers, and strategically aims to exploit this advantage in order to maximize her revenue. In order for such an approach to work, an additional constraint is that the potential buyers need to be unaware of each other as well as of details of the underlying mechanism which could be used to extract information about the quality of the items. If this not possible, then the *linkage principle* of Milgrom and Weber [1982] suggests that the seller should reveal all possible information to the buyers in order to maximize her revenue.

Levin and Milgrom [2010] as well as Milgrom [2010] suggest that careful bundling of the items is the best way to exploit information asymmetry. Many different frameworks that reveal different kinds of information to the bidders have been proposed in the literature over the years. For instance, Ghosh et al. [2007] considered full information and proposed a clustering scheme according to which, the items are partitioned into bundles and, then, for each such bundle, a separate second-price auction takes place. This way, the potential buyers cannot bid only for the items that they actually want, but have to also compete for items that they do not value much. Hence, the demand for each item is increased and the revenue of the seller gets increased. Emek

et al. [2012] and Dughmi [2014] presented several complexity results in similar settings, while Miltersen and Sheffet [2012] considered fractional bundling schemes for signaling.

Finally, it is worth mentioning that exploiting linear programming for the analysis of purely combinatorial algorithms, like we did in Section 5.4, is a now well-established technique that has already been used in many different settings, such as facility location [Jain et al., 2003], variants of set cover [Athanasopoulos et al., 2009a,b, Caragiannis et al., 2013], online matching [Mahdian and Yan, 2011], maximum directed cut [Feige and Jozeph, 2015], and wavelength routing [Caragiannis, 2009].

5.3 Definitions, examples and structural observations

An algorithm for the asymmetric matrix partition that computes a partition scheme with value that is at least $\rho \in [0, 1]$ times the partition value of the best possible partition scheme is called a ρ -approximation algorithm. Henceforth, we focus our attention on the case where the input matrix A consists only of *binary* values.

Let $A^+ = \{j \in [m] : \text{there exists a row } i \text{ such that } A_{ij} = 1\}$ denote the set of columns of A that contain at least one 1-value entry, and $A^0 = [m] \setminus A^+$ denote the set of columns of A that contain only 0-value entries. In the next sections, we usually refer to the sets A^+ and A^0 as the sets of *one-columns* and *zero-columns*, respectively. Furthermore, let $A_i^+ = \{j \in [m] : A_{ij} = 1\}$ and $A_i^0 = \{j \in [m] : A_{ij} = 0\}$ denote the sets of columns that intersect with row i at a 1- and 0-value entry, respectively. All columns in A_i^+ are one-columns and, furthermore, $A^+ = \cup_{i=1}^n A_i^+$. The columns of A_i^0 can be either one- or zero-columns and, thus, $A^0 \subseteq \cup_{i=1}^n A_i^0$. Also, denote by $r = \sum_{j \in A^+} p_j$ the total probability of the one-columns. As an example, consider the 3×6 matrix

$$A = \begin{pmatrix} 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \end{pmatrix}$$

and a uniform probability distribution over its columns. We have $A^+ = \{2, 3, 5\}$ and $A^0 = \{1, 4, 6\}$. In the first two rows, the sets A_i^+ and A_i^0 are identical to A^+ and A^0 , respectively. In the third row, the sets A_3^+ and A_3^0 are $\{2, 3\}$ and $\{1, 4, 5, 6\}$. Finally, the total probability of the one-columns r is $1/2$.

A partition scheme B can be thought of as consisting of n partitions B_1, B_2, \dots, B_n of the set of columns $[m]$. We use the term *bundle* to refer to the elements of a partition B_i ; a bundle is just a non-empty set of columns. For a bundle b of partition B_i corresponding to row i , we

say that b is an *all-zero* bundle if $b \subseteq A_i^0$ and an *all-one* bundle if $b \subseteq A_i^+$. A singleton all-one bundle of partition B_i is called *column-covering bundle in row i* . A bundle that is neither all-zero nor all-one is called *mixed*. A mixed bundle corresponds to a set of columns that intersects with row i at both 1- and 0-value entries.

Let us examine the following partition scheme B for matrix A that defines the smooth matrix A^B according to equation (5.1).

B_1	$\{1, 2, 3, 4\}, \{5, 6\}$	A^B	1/2	1/2	1/2	1/2	1/2	1/2
B_2	$\{1, 2\}, \{3\}, \{4, 6\}, \{5\}$		1/2	1/2	1	0	1	0
B_3	$\{1, 4, 6\}, \{2, 3, 5\}$		0	2/3	2/3	0	2/3	0
		$\max_i A_{ij}^B$	1/2	2/3	1	1/2	1	1/2

Here, the bundle $\{1, 2, 3, 4\}$ of (the partition B_1 of) the first row is a mixed one. The bundle $\{3\}$ of B_2 is all-one and, in particular, column-covering in row 2. The bundle $\{1, 4, 6\}$ of B_3 is all-zero.

By equation (5.2), the partition value is $25/36$ and it can be further improved. First, observe that the leftmost zero-column is included in two mixed bundles (in the first two rows). Also, the mixed bundle in the third row contains a one-column that has been covered through a column-covering bundle in the second row and intersects with the third row at a 0-value entry. Let us modify these two bundles.

B'_1	$\{1\}, \{2, 3, 4\}, \{5, 6\}$	$A^{B'}$	0	2/3	2/3	2/3	1/2	1/2
B'_2	$\{1, 2\}, \{3\}, \{4, 6\}, \{5\}$		1/2	1/2	1	0	1	0
B'_3	$\{1, 4, 5, 6\}, \{2, 3\}$		0	1	1	0	0	0
		$\max_i A_{ij}^{B'}$	1/2	1	1	2/3	1	1/2

The partition value becomes $7/9 > 25/36$. Now, by merging the two mixed bundles $\{2, 3, 4\}$ and $\{5, 6\}$ in the first row, we obtain the smooth matrix below with partition value $47/60 > 7/9$. Observe that the contribution of column 4 to the partition value decreases but, overall, we have an increase in the partition value due to the increase in the contribution of column 6. Actually, such merges never decrease the partition value (see Lemma 5.1 below).

B''_1	$\{1\}, \{2, 3, 4, 5, 6\}$	$A^{B''}$	0	3/5	3/5	3/5	3/5	3/5
B''_2	$\{1, 2\}, \{3\}, \{4, 6\}, \{5\}$		1/2	1/2	1	0	1	0
B''_3	$\{1, 4, 5, 6\}, \{2, 3\}$		0	1	1	0	0	0
		$\max_i A_{ij}^{B''}$	1/2	1	1	3/5	1	3/5

Finally, by merging the bundles $\{1, 2\}$ and $\{3\}$ in the second row and decomposing the bundle $\{2, 3\}$ in the last row into two singletons, the partition value becomes $73/90 > 47/60$ which can

be verified to be optimal.

B_1'''	$\{1\}, \{2, 3, 4, 5, 6\}$	$A^{B'''}$	0	3/5	3/5	3/5	3/5	3/5
B_2'''	$\{1, 2, 3\}, \{4, 6\}, \{5\}$		2/3	2/3	2/3	0	1	0
B_3'''	$\{1, 4, 5, 6\}, \{2\}, \{3\}$		0	1	1	0	0	0
		$\max_i A_{ij}^{B'''}$	2/3	1	1	3/5	1	3/5

We will now give some more definitions that will be useful in the following. We say that a one-column j is *covered* by a partition scheme B if there is at least one row i in which $\{j\}$ is column-covering. For example, in B''' , the singleton $\{5\}$ is a column-covering bundle in the second row and the singletons $\{2\}$ and $\{3\}$ are column-covering in the third row. We say that a partition scheme *fully covers* the set A^+ of one-columns if all of them are covered. In this case, we use the term *full cover* to refer to the pairs of indices (i, j) of the 1-value entries A_{ij} such that $\{j\}$ is a column-covering bundle in row i . For example, the partition scheme B''' has the full cover $(2, 5), (3, 2), (3, 3)$.

It turns out that optimal partition schemes always have a special structure like the one of B''' . Alon et al. [2013] formalized observations like the above into the following statement.

Lemma 5.1 (Alon et al. [2013]). *Given a uniform instance of the asymmetric binary matrix partition problem with a matrix A , there is an optimal partition scheme B with the following properties:*

- P1. B fully covers the set A^+ of one-columns.
- P2. For each row i , B_i has at most one bundle containing all columns of A_i^+ that are not included in column-covering bundles in row i (if any). This bundle can be either all-one (if it does not contain zero-columns) or the unique mixed bundle of row i .
- P3. For each zero-column j , there exists at most one row i such that j is contained in the mixed bundle of B_i (and j is contained in the all-zero bundles of the remaining rows).
- P4. For each row i , the zero-columns that are not contained in the mixed bundle of B_i form an all-zero bundle.

Properties P1 and P3 imply that we can think of the partition value as the sum of the contributions of the column-covering bundles and the contributions of the zero-columns in mixed bundles. Property P2 comes from the following more general statement that has been proved by Alon et al. [2013]; we give an alternative more direct proof here using Milne inequality [Hardy et al., 1934, page 61]. Lemma 5.2 will be very useful several times in our analysis in both the uniform and the non-uniform case.

Lemma 5.2 (Alon et al. [2013]). Consider $t \geq 2$ mixed bundles. For $i = 1, \dots, t$, bundle i contains 1-value entries of total probability x_i and zero-columns of probability y_i . The total contribution of the zero-columns in these mixed bundles to the partition value is upper bounded by the contribution of zero-columns of probability $\sum_{i=1}^t y_i$ that form a single mixed bundle together with 1-value entries of probability $\sum_{i=1}^t x_i$.

Proof. By the definitions, the smooth value of the i -th bundle is $\frac{x_i}{x_i + y_i}$ and the contribution of its zero-columns to the the partition value is $\frac{x_i y_i}{x_i + y_i}$. The proof follows by Milne inequality which states that

$$\sum_{i=1}^t \frac{x_i y_i}{x_i + y_i} \leq \frac{\sum_{i=1}^t x_i \cdot \sum_{i=1}^t y_i}{\sum_{i=1}^t x_i + \sum_{i=1}^t y_i},$$

where the right-hand side expression is the contribution of the zero-columns in the partition value of the single mixed bundle. \square

Now, property P2 should be apparent; the columns of A_i^+ that do not form column-covering bundles in row i are bundled together with zero-columns (if possible) in order to increase the contribution of the latter to the partition value. Property P4 makes B consistent to the definition of a partition scheme where the disjoint union of all the partition subsets in a row should be $[m]$. Clearly, the contribution of the all-zero bundles to the partition value is 0. Also, the non-column-covering all-one bundles do not contribute to the partition value either.

Unfortunately, as we will see later in Section 5.5, Lemma 5.1 does not hold for non-uniform instances. This is due only to property P1 which requires a uniform probability distribution over columns. Luckily, it turns out that non-uniform instances also exhibit some structure (recall that the crucial Lemma 5.2 applies to the non-uniform case as well), which allows us to consider the problem of computing an optimal partition scheme as a *welfare maximization problem*. In welfare maximization, there are m items and n agents; agent i has a valuation function $v_i : 2^{[m]} \rightarrow \mathbb{R}^+$ that specifies her value for each subset of the items. I.e., for a set S of items, $v_i(S)$ represents the value of agent i for S . Given a disjoint partition (or allocation) $S = (S_1, S_2, \dots, S_n)$ of the items to the agents, where S_i denotes the set of items allocated to agent i , the social welfare is the sum of values of the agents for the sets of items allocated to them, i.e., $\text{SW}(S) = \sum_i v_i(S_i)$. The term welfare maximization refers to the problem of computing an allocation of maximum social welfare. We will discuss only the variant of the problem where the valuations are monotone and submodular; following the literature, we use the term submodular welfare maximization to refer to it.

Definition 5.1. A valuation function v is monotone if $v(S) \leq v(T)$ for any pair of sets S, T such that $S \subseteq T$. A valuation function v is submodular if $v(S \cup \{x\}) - v(S) \geq v(T \cup \{x\}) - v(T)$ for any pair of sets S, T such that $S \subseteq T$ and for any item x .

An important issue in (submodular) welfare maximization arises with the representation of valuation functions. A valuation function can be described in detail by listing explicitly the values for each of the possible subsets of items. Unfortunately, this is clearly inefficient due to the necessity for exponential input size. A solution that has been proposed in the literature is to assume access to these functions by queries of a particular form. The simplest such form of queries reads as

What is the value of agent i for the set of items S ?

These are known as *valuation queries*. Another type of queries, known as *demand queries*, are phrased as follows:

Given a non-negative price for each item, compute a set S of items for which the difference of the valuation of agent i minus the sum of prices for the items in S is maximized

Approximation algorithms that use a polynomial number of valuation or demand queries and obtain solutions to submodular welfare maximization with a constant approximation ratio are well-known in the literature (e.g. see the papers of Feige and Vondrák [2010], Lehmann et al. [2006], Vondrák [2008]). Our improved approximation algorithm for the non-uniform case of asymmetric binary matrix partition exploits such algorithms.

5.4 The uniform case

In this section, we focus on the case where the probability distribution p over the columns of the given matrix is uniform and present the analysis of a greedy approximation algorithm. Our algorithm uses a *greedy completion procedure* that was also considered by Alon et al. [2013]. This procedure starts from a full cover of the matrix, i.e., from column-covering bundles in some rows so that all one-columns are covered (by exactly one column-covering bundle). Once this initial full cover is given, the set of columns from A_i^+ that are not included in column-covering bundles in row i can form a mixed bundle together with some zero-columns in order to increase the contribution of the latter to the partition value. Greedy completion proceeds as follows. It

goes over the zero-columns, one by one, and adds a zero-column to the mixed bundle of the row that maximizes the marginal contribution of the zero-column. The marginal contribution of a zero-column to the partition value when it is added to a mixed bundle that consists of x zero-columns and y one-columns is proportional (due to the uniform distribution over columns) to the quantity

$$\Delta(x, y) = (x + 1) \frac{y}{x + y + 1} - x \frac{y}{x + y} = \frac{y^2}{(x + y)(x + y + 1)}.$$

The right-hand side of the first equality is simply the difference between the contribution of $x + 1$ and x zero-columns to the partition value when they form a mixed bundle with y one-columns. Note that $\Delta(0, y)$ indicates the marginal contribution of a zero-column when put together with y one-columns to form a (new) mixed bundle. Alon et al. [2013] made the next extremely important observation. We extensively use it below, as well as the fact that $\Delta(x, y)$ is non-decreasing with respect to y .

Lemma 5.3 (Alon et al. [2013]). *Among all partition schemes that include a given full cover, the greedy completion procedure yields the maximum contribution from the zero-columns to the partition value.*

Our algorithm consists of two phases. In the first phase, called the *cover phase*, the algorithm computes an arbitrary full cover for set A^+ . In the second phase, called the *greedy phase*, it simply runs the greedy completion procedure mentioned above. Note that, intentionally, we have not used much detail in the description of the algorithm and there are three issues that might seem to cause ambiguity at first glance. First, we have not described any particular way the full cover is constructed. Second, we have not defined some particular order in which the zero-columns are examined during the greedy phase. And, third, we have not discussed how ties are broken when there are multiple rows that maximize the marginal contribution of a zero-column. So, our description essentially defines a family of greedy algorithms; a different greedy algorithm is defined, depending on how the above three issues are implemented. In the rest of this section, we will show that *any* greedy algorithm has an approximation ratio of at least $9/10$; actually, the three issues do not affect the analysis at all. We will also show that our analysis is tight by presenting a simple instance for which some greedy algorithm is at most $9/10$ -approximate. Even though greedy algorithms are purely combinatorial, our analysis exploits linear programming duality. In the following, unless otherwise specified, the term greedy algorithm refers to any member of the family of greedy algorithms.

Overall, the partition value obtained by the algorithm can be thought of as the sum of contributions from column-covering bundles (this is exactly r) plus the contribution from the

mixed bundles created during the greedy phase (i.e., the contribution from the zero-columns). Denote by ρ the ratio between the total number of appearances of one-columns in the mixed bundles of the optimal partition scheme (so, the number of times each one-column is counted equals the number of mixed bundles that contain it) and the number of zero-columns. For example, in the partition scheme B''' in the example of the previous section, the two mixed bundles are $\{2, 3, 4, 5, 6\}$ in the first row and $\{1, 2, 3\}$ in the second row. So, the one-columns 2 and 3 appear twice while the one-column 5 appears once in these mixed bundles. Since we have three zero-columns, the value of ρ is $5/3$. We can use the quantity ρ to upper-bound the optimal partition value as follows.

Lemma 5.4. *The optimal partition value is at most $r + (1 - r)\frac{\rho}{\rho+1}$.*

Proof. The first term in the above expression represents the contribution of the one-columns in the full cover of the optimal partition scheme. To reason about the second term, recall that our definitions imply that the total probability of one-columns in the mixed bundles of an optimal partition scheme is $\rho(1 - r)$, while the total probability of zero-columns in these mixed bundles is $1 - r$. By Lemma 5.2, the second term upper-bounds the total contribution of the zero-columns to the optimal partition value. \square

In our analysis, we distinguish between two main cases depending on the value of ρ . The first case is when $\rho < 1$; in this case, we show that the additional partition value which is obtained during the greedy phase of the algorithm (i.e., the contribution of the zero-columns; recall that the greedy algorithm maximizes this quantity) is lower-bounded by the additional partition value we would have by creating bundles containing exactly one one-column and an almost equal number of zero-columns each.

Lemma 5.5. *If $\rho < 1$, then the partition value obtained by the algorithm is at least 0.97 times the optimal one.*

Proof. Using the definition of ρ , we can lower-bound the number of 1-value entries in the input matrix A by the sum of the mr column-covering bundles that form the full cover of the optimal partition scheme and the at least $\rho m(1 - r)$ appearances of one-columns in the mixed bundles.

Now, consider a selection of the full cover during the cover phase of the greedy algorithm (this can, of course be different than the full cover of the optimal partition scheme) and let X

be a set of (exactly) $\rho m(1-r)$ 1-value entries in the matrix A among those that are not included in the cover.

Using Lemma 5.3, we will lower-bound the partition value returned by the algorithm by considering the following formation of mixed bundles as an alternative to the greedy completion procedure used in the greedy phase. If $1/\rho$ is an integer, for each 1-value entry of X , we create a mixed bundle that contains the corresponding one-column together with $1/\rho$ distinct zero-columns. Hence, the smooth value of each zero-column is $\frac{1}{1+1/\rho}$ and the total partition value of this scheme is $r + (1-r)\frac{\rho}{\rho+1}$; by Lemma 5.4, this is optimal.

If instead $1/\rho$ is not an integer, let $k = \lfloor 1/\rho \rfloor$. For each 1-value entry of X , we create a mixed bundle that contains the corresponding one-column together with k or $k+1$ distinct zero-columns. In particular, $m(1-r)(1-\rho k)$ of these mixed bundles contain one one-column and $k+1$ zero-columns and the remaining $m(1-r)(\rho(k+1)-1)$ mixed bundles contain one one-column and k zero-columns. Observe that the smooth value of a zero-column is $\frac{1}{k+2}$ in the first case and $\frac{1}{k+1}$ in the second case. Hence, we can bound the partition value obtained by the algorithm as follows:

$$\begin{aligned} \text{ALG} &\geq r + (1-r)(1-\rho k)\frac{k+1}{k+2} + (1-r)(\rho(k+1)-1)\frac{k}{k+1} \\ &= r + (1-r)\frac{1+\rho k(k+1)}{(k+1)(k+2)}. \end{aligned}$$

Using Lemma 5.4, we have

$$\frac{\text{ALG}}{\text{OPT}} \geq \frac{r + (1-r)\frac{1+\rho k(k+1)}{(k+1)(k+2)}}{r + (1-r)\frac{\rho}{\rho+1}} \geq \frac{\frac{1+\rho k(k+1)}{(k+1)(k+2)}}{\frac{\rho}{\rho+1}} = \frac{(1+1/\rho)(1+\rho k(k+1))}{(k+1)(k+2)}.$$

This last expression is minimized (with respect to ρ) for $1/\rho = \sqrt{k(k+1)}$. Hence,

$$\frac{\text{ALG}}{\text{OPT}} \geq \frac{\left(1 + \sqrt{k(k+1)}\right)^2}{(k+1)(k+2)},$$

which is minimized for $k = 1$ to approximately 0.97. □

For the case $\rho \geq 1$, we use completely different arguments. Of course, we assume that $r < 1$, i.e., the input matrix contains some zero-columns since, otherwise, any full cover computed during the cover phase of the greedy algorithm would give an optimal partition value. We will reason about the partition value of the solution produced by the algorithm by considering a particular decomposition of the set of mixed bundles computed in the greedy phase. Then, using Lemmas 5.2 and 5.3, the contribution of the zero-columns to the partition value in the

solution computed by the algorithm is lower-bounded by their contribution to the partition value when they are part of the mixed bundles obtained after the decomposition. To justify the correctness of the decomposition, we will use the following observation.

Lemma 5.6. *If $\rho \geq 1$, no mixed bundle computed by the greedy algorithm has more zero-columns than one-columns.*

Proof. First observe that the total number of appearances of one-columns in mixed and column-covering bundles in the optimal partition scheme is at least $rm + (1 - r)\rho m$, which includes rm appearances of one-columns in column-covering bundles and $(1 - r)\rho m$ appearances of one-columns in mixed bundles (there may be additional 1-value entries included in all-one bundles). So, after the end of the cover phase, there are at least $(1 - r)\rho m \geq (1 - r)m$ 1-value entries that can be included in mixed bundles together with the $(1 - r)m$ zero-columns.

Assume, for the sake of contradiction, that some zero-column Z is included as the $(x + 1)$ -th zero-column in a mixed bundle b together with x 1-value entries for $x \geq 1$ at some step of the greedy phase. Prior to that step, there is either some 1-value entry not included in any mixed bundle which could be used to form a mixed bundle together with Z for a marginal contribution of $\Delta(0, 1) = 1/2$ or some mixed bundle with $y \geq 1$ zero-columns and $y + \alpha$ 1-value entries (with $\alpha \geq 1$) in which case the marginal contribution would be $\Delta(y, y + \alpha) > 1/4$. This contradicts the definition of the greedy algorithm since the marginal contribution of Z was $\Delta(x, x) < 1/4$ when included in b . \square

Now, the decomposition procedure is defined as follows. It takes as input a mixed bundle with y zero-columns and x one-columns (by Lemma 5.6, it must be $x \geq y$) and decomposes it into y bundles as follows. If x/y is an integer, each bundle has one zero-column and x/y one-columns. Otherwise, $x - y\lfloor x/y \rfloor$ bundles have one zero-column and $\lceil x/y \rceil$ one-columns and $y\lceil x/y \rceil - x$ bundles have one zero-column and $\lfloor x/y \rfloor$ one-columns. Clearly, this process does not alter bundles with a single zero-column. The solution obtained after the decomposition of the solution returned by the algorithm has a very special structure as our next lemma suggests.

Lemma 5.7. *There exists an integer $s \geq 1$ such that each bundle in the decomposition has at least s and at most $3s$ one-columns.*

Proof. Consider the application of the decomposition procedure to the mixed bundles that are computed by the algorithm and let s be the minimum number of one-columns among the

decomposed mixed bundles. This implies that one of the mixed bundles, say b_1 , computed by the algorithm has μ zero-columns and at most $(s+1)\mu - 1$ one-columns. Denoting by ν the number of one-columns in this bundle, we have that the marginal partition value when the last zero-column Z is included in b_1 is exactly

$$\Delta(\mu, \nu) = \frac{\nu^2}{(\nu + \mu)(\nu + \mu - 1)} \leq \frac{((s+1)\mu - 1)^2}{((s+2)\mu - 1)((s+2)\mu - 2)}$$

since $\Delta(\mu, \nu)$ is increasing in ν and $\nu \leq (s+1)\mu - 1$. The rightmost expression is decreasing in μ and $\mu \geq 1$; hence, the marginal partition value of Z is at most $\frac{s}{s+1}$.

Now assume for the sake of contradiction that one of the mixed bundles obtained after the decomposition has at least $3s+1$ one-columns. Clearly, this must have been obtained by the decomposition of a mixed bundle b_2 (returned by the algorithm) with λ zero-columns and at least $(3s+1)\lambda$ one-columns. Denote by ν' the number of one-columns in this bundle and let us compute the marginal partition value if the zero-column Z would be included in b_2 . This would be

$$\Delta(\lambda+1, \nu') = \frac{\nu'^2}{(\nu' + \lambda + 1)(\nu' + \lambda)} \geq \frac{(3s+1)^2\lambda}{((3s+2)\lambda + 1)(3s+2)} \geq \frac{(3s+1)^2}{(3s+3)(3s+2)}.$$

The first inequality follows since the marginal partition value function is increasing in ν' and $\nu' \geq (3s+1)\lambda$, and the second one follows since $\lambda \geq 1$. Now, the last quantity can be easily verified to be strictly higher than $\frac{s}{s+1}$ and the algorithm should have included Z in b_2 instead. We have reached the desired contradiction that proves the lemma. \square

Now, our analysis proceeds as follows. For every triplet $r \in [0, 1], \rho \geq 1$ and integer $s \geq 1$, we will prove that any solution consisting of an arbitrary cover of the rm one-columns and the decomposed set of bundles containing at least s and at most $3s$ one-columns yields a 9/10-approximation of the optimal partition value. By the discussion above (in particular, by Lemmas 5.2 and 5.3), this will also be the case for the solution returned by the algorithm. In order to account for the worst-case contribution of zero-columns to the partition value for a given triplet of parameters, we will use the following linear program, which we denote by $\text{LP}(r, \rho, s)$:

$$\begin{aligned} & \text{minimize} && \sum_{k=s}^{3s} \frac{k}{k+1} \theta_k \\ & \text{subject to:} && \sum_{k=s}^{3s} \theta_k = 1 - r \end{aligned}$$

$$\begin{aligned} \sum_{k=s}^{3s} k\theta_k &\geq \rho(1-r) - r \\ \theta_k &\geq 0, k = s, \dots, 3s \end{aligned}$$

The variable θ_k corresponds to the total probability of the zero-columns that participate in decomposed mixed bundles with k one-columns. The objective is to minimize the contribution of the zero-columns to the partition value. The equality constraint means that all zero-columns have to participate in bundles, while the inequality constraint requires that the total number of appearances of one-columns in bundles used by the algorithm is at least the total number of appearances of one-columns in mixed bundles of the optimal partition scheme minus one appearance for each one-column, since for every selection of the cover, the algorithm will have the same number of (appearances of) one-columns available to form mixed bundles.

Informally, the linear program answers (pessimistically) to the question of how inefficient the algorithm can be. In particular, given an instance with parameters r and ρ , the quantity $\min_{\text{int } s \geq 1} \text{LP}(r, \rho, s)$ lower-bounds the contribution of the zero-columns to the partition value and $r + \min_{\text{int } s \geq 1} \text{LP}(r, \rho, s)$ is a lower bound on the partition value. The next lemma completes the analysis of the greedy algorithm for the case $\rho \geq 1$.

Lemma 5.8. *For every $r \in [0, 1]$ and $\rho \geq 1$,*

$$r + \min_{\text{int } s \geq 1} \text{LP}(r, \rho, s) \geq \frac{9}{10} \text{OPT}.$$

Proof. We will prove the lemma using LP-duality. The dual of $\text{LP}(r, \rho, s)$ is:

$$\begin{aligned} &\text{maximize} \quad (1-r)\alpha + ((1-r)\rho - r)\beta \\ &\text{subject to:} \quad k\beta + \alpha \leq \frac{k}{k+1}, k = s, \dots, 3s \\ &\quad \beta \geq 0 \end{aligned}$$

Using Lemma 5.4, we bound the optimal partition value as

$$\text{OPT} \leq r + (1-r) \frac{\rho}{\rho+1} = \frac{\rho+r}{\rho+1}.$$

Hence, it suffices to show that, for every triplet of parameters (r, ρ, s) , there is a feasible dual solution of objective value $D(r, \rho, s)$ that satisfies

$$r + D(r, \rho, s) - \frac{9}{10} \frac{\rho+r}{\rho+1} \geq 0. \tag{5.3}$$

The feasible region of the dual is defined by the lines $\beta = 0$, $\alpha = \frac{s}{s+1} - s\beta$ and $\alpha = \frac{3s}{3s+1} - 3s\beta$; the remaining constraints can be easily seen to be redundant. The two important intersections of those lines are the points

$$(\alpha, \beta) = \left(\frac{s}{s+1}, 0 \right) \text{ and } (\alpha, \beta) = \left(\frac{3s^2}{(s+1)(3s+1)}, \frac{1}{(s+1)(3s+1)} \right)$$

with objective values

$$D_1(r, \rho, s) = \frac{s}{s+1}(1-r) \text{ and } D_2(r, \rho, s) = \frac{3s^2}{(s+1)(3s+1)}(1-r) + \frac{\rho(1-r) - r}{(s+1)(3s+1)},$$

respectively. We will show that one of these two points can always be used as a feasible dual solution in order to prove inequality (5.3). We distinguish between two cases.

Case I. $r \geq \frac{\rho-1}{\rho}$. We will show that the point with dual objective value $D_1(r, \rho, s)$ satisfies inequality (5.3), i.e.,

$$r + \frac{s}{s+1}(1-r) - \frac{9}{10} \frac{\rho+r}{\rho+1} \geq 0. \quad (5.4)$$

Since $s \geq 1$, we have that the left hand side of inequality (5.4) is at least

$$\frac{1+r}{2} - \frac{9}{10} \frac{\rho+r}{\rho+1} = \frac{1}{2} - \frac{9\rho}{10(\rho+1)} + r \left(\frac{1}{2} - \frac{9}{10(\rho+1)} \right).$$

Since $\rho \geq 1$, we have that $\frac{1}{2} - \frac{9}{10(\rho+1)} \geq 0$, and we can lower-bound the above quantity using the assumption $r \geq \frac{\rho-1}{\rho}$, as follows:

$$\frac{1+r}{2} - \frac{9}{10} \frac{\rho+r}{\rho+1} \geq \frac{1}{2} - \frac{9\rho}{10(\rho+1)} + \frac{\rho-1}{\rho} \left(\frac{1}{2} - \frac{9}{10(\rho+1)} \right) = \frac{(\rho-2)^2}{10\rho(\rho+1)} \geq 0,$$

and inequality (5.4) follows.

Case II. $r < \frac{\rho-1}{\rho}$. We will now show that the point with dual objective value $D_2(r, \rho, s)$ satisfies inequality (5.3), i.e.,

$$r + \frac{3s^2}{(s+1)(3s+1)}(1-r) + \frac{\rho(1-r) - r}{(s+1)(3s+1)} - \frac{9}{10} \frac{\rho+r}{\rho+1} \geq 0. \quad (5.5)$$

Let us denote by F the left hand side of inequality (5.5). With simple calculations, we obtain

$$F = \frac{10\rho^2 - (-3s^2 + 36s - 1)\rho + 30s^2}{10(3s+1)(s+1)(\rho+1)} - r \cdot \frac{10\rho^2 - (40s - 10)\rho + 27s^2 - 4s + 9}{10(3s+1)(s+1)(\rho+1)}. \quad (5.6)$$

Observe that the numerator of the left fraction in (5.6) is a quadratic function with respect to ρ with positive coefficient in the leading term. Its discriminant is $-1191s^4 - 216s^3 + 1296s^2 - 72s + 7$

which is clearly negative for every integer $s \geq 1$. Hence, the numerator of the left fraction is always positive. Now, if the numerator of the rightmost fraction is negative, then inequality (5.5) is obviously satisfied. Otherwise, using the assumption $r < \frac{\rho-1}{\rho}$, we have

$$\begin{aligned} F &\geq \frac{10\rho^2 - (-3s^2 + 36s - 1)\rho + 30s^2}{10(3s+1)(s+1)(\rho+1)} - \frac{\rho-1}{\rho} \cdot \frac{10\rho^2 - (40s-10)\rho + 27s^2 - 4s + 9}{10(3s+1)(s+1)(\rho+1)} \\ &= \frac{(3s^2 + 4s + 1)\rho^2 + (3s^2 - 36s + 1)\rho + 27s^2 - 4s + 9}{10\rho(3s+1)(s+1)(\rho+1)}. \end{aligned}$$

Now, the numerator of the last fraction is again a quadratic function in terms of ρ with positive coefficient in the leading term and discriminant equal to

$$-315s^4 - 600s^3 + 1150s^2 - 200s - 35 = (-315s^3 - 915s^2 + 235s - 35)(s-1) \leq 0,$$

for every integer $s \geq 1$. Hence, $F \geq 0$ and the proof is complete. \square

The next statement summarizes the discussion above.

Theorem 5.9. *The greedy algorithm always yields a 9/10-approximation of the optimal partition value in the uniform case.*

Our analysis is tight as our next counter-example suggests.

Theorem 5.10. *There exists an instance of the uniform asymmetric binary matrix partition problem for which a greedy algorithm computes a partition scheme with value (at most) 9/10 of the optimal one.*

Proof. Consider the instance of the asymmetric binary matrix partition problem that consists of the matrix

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{pmatrix}$$

with $p_i = 1/4$ for $i = 1, 2, 3, 4$. The optimal partition value is obtained by covering the one-columns in the first two rows and then bundling each of the two zero-columns with a pair of one-columns in the third and fourth row, respectively. This yields a partition value of 5/6. A greedy algorithm may select to cover the one-columns using the 1-value entries A_{31} and A_{42} . This is possible since the greedy algorithm has no particular criterion for breaking ties when selecting the full cover. Given this full cover, the greedy completion procedure will assign each of the two zero-columns with only one one-column. The partition value is then 3/4, i.e., 9/10 times the optimal partition value. \square

5.5 Asymmetric binary matrix partition as welfare maximization

We now consider the more general non-uniform case. Interestingly, property P1 of Lemma 5.1 does not hold any more as the following statement shows.

Lemma 5.11. *For every $\epsilon > 0$, there exists an instance of the asymmetric binary matrix partition problem in which any partition scheme containing a full cover of the columns in A^+ yields a partition value that is at most $8/9 + \epsilon$ times the optimal one.*

Proof. Consider the instance of the asymmetric binary matrix partition problem consisting of the matrix

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \end{pmatrix}$$

with column probabilities $p_j = \frac{1}{\beta+3}$ for $j = 1, 2, 3$ and $p_4 = \frac{\beta}{\beta+3}$ for $\beta > 2$. We will first prove an upper bound on the partition value of any partition scheme containing a full cover. Then, we will present a partition scheme without a full cover, which has a strictly higher partition value. The desired ratio of $8/9 + \epsilon$ will then follow by setting the parameter β appropriately.

Observe that there are four partition schemes containing a full cover (depending on the rows that contain the column-covering bundle of the first two columns). In each of them, there are two 1-value entries in different rows that are not included in the full cover, and only one of them can be bundled together with the zero-column. By making calculations, we obtain that the partition value in these cases is $\frac{4\beta+3}{(\beta+1)(\beta+3)}$. Here is one of these partition schemes:

B_1	$\{1\}, \{2, 3, 4\}$	1	0	0	0
B_2	$\{2\}, \{1, 3, 4\}$	0	1	0	0
B_3	$\{1, 3\}, \{2, 4\}$	0	$\frac{1}{\beta+1}$	0	$\frac{1}{\beta+1}$
B_4	$\{1\}, \{3\}, \{2, 4\}$	1	0	1	0
$p_j \cdot \max_i A_{ij}^B$		$\frac{1}{\beta+3}$	$\frac{1}{\beta+3}$	$\frac{1}{\beta+3}$	$\frac{\beta}{(\beta+1)(\beta+3)}$

In contrast, consider the partition scheme B' in which the 1-value entries A_{11} and A_{22} form column-covering bundles in rows 1 and 2, the entries A_{32} and A_{33} are bundled together in row 3 and the entries A_{41} , A_{43} , and A_{44} are bundled together in row 4. As it can be seen from the tables below (recall that $\beta > 2$), the partition value now becomes $\frac{4.5\beta+5}{(\beta+2)(\beta+3)}$.

Clearly, the ratio of the two partition values approaches $8/9$ from above as β tends to infinity. Hence, the theorem follows by selecting β sufficiently large for any given $\epsilon > 0$. \square

B'_1	$\{1\}, \{2, 3, 4\}$	$A^{B'}$	1	0	0	0
B'_2	$\{2\}, \{1, 3, 4\}$		0	1	0	0
B'_3	$\{1, 4\}, \{2, 3\}$		0	1/2	1/2	0
B'_4	$\{2\}, \{1, 3, 4\}$		$\frac{2}{\beta+2}$	0	$\frac{2}{\beta+2}$	$\frac{2}{\beta+2}$
$p_j \cdot \max_i A_{ij}^{B'}$			$\frac{1}{\beta+3}$	$\frac{1}{\beta+3}$	$\frac{1}{2(\beta+3)}$	$\frac{2\beta}{(\beta+2)(\beta+3)}$

Still, as the next statement indicates, the optimal partition scheme has some structure which we will exploit later.

Lemma 5.12. *Consider an instance of the asymmetric binary matrix partition problem consisting of a matrix A and a probability distribution p over its columns. There is an optimal partition scheme B that satisfies properties P2, P3, P4 (from Lemma 5.1) as well as the new property P5:*

- P2. *For each row i , B_i has at most one bundle containing all columns of A_i^+ that are not included in column-covering bundles in row i (if any). This bundle can be either all-one (if it does not contain zero-columns) or the unique mixed bundle of row i .*
- P3. *For each zero-column j , there exists at most one row i such that j is contained in the mixed bundle of B_i (and j is contained in the all-zero bundles of the remaining rows).*
- P4. *For each row i , the zero-columns that are not contained in the mixed bundle of B_i form an all-zero bundle.*
- P5. *Given any column j , denote by $H_j = \arg \max_i A_{ij}^B$ the set of rows through which column j contributes to the partition value $v^B(A, p)$. For every $i \in H_j$ such that $A_{ij} = 1$, the bundle of partition B_i that contains column j is not mixed.*

Proof. We first focus on property P5. Consider an optimal partition scheme B that does not satisfy property P5, and let j^* be a column such that $A_{i^*j^*} = 1$ for some $i^* \in H_{j^*}$. Furthermore, assume that the mixed bundle b of partition B_{i^*} that contains column j^* , also contains the columns of a (possibly empty) set $b_1 \subseteq A_{i^*}^+ \setminus \{j^*\}$ and the columns of a non-empty set $b_0 \subseteq A_{i^*}^0$. Let $p^+ \geq 0$ and $p^0 > 0$ be the sum of probabilities of the columns in b_1 and b_0 , respectively.

Let B' be the partition scheme that is obtained from B when splitting bundle b into two bundles $\{j^*\}$ and $b \setminus \{j^*\}$; we will show that B' must be optimal as well. Observe that $A_{i^*j}^B = \frac{p_{j^*} + p^+}{p_{j^*} + p^+ + p^0}$ and $A_{i^*j}^{B'} = \frac{p^+}{p^+ + p^0}$ for every $j \in b \setminus \{j^*\}$; hence, $A_{i^*j}^B > A_{i^*j}^{B'}$. Since, this is the only difference between B and B' , the difference $\max_i A_{ij}^B - \max_i A_{ij}^{B'}$ is at most $A_{i^*j}^B - A_{i^*j}^{B'}$ for every

$j \in b \setminus \{j^*\}$, and for j^* , $\max_i A_{ij^*}^B - \max_i A_{ij^*}^{B'} = A_{i^*j^*}^B - A_{i^*j^*}^{B'} = \frac{p_{j^*} + p^+}{p_{j^*} + p^+ + p^0} - 1$. Hence, we have

$$\begin{aligned}
v^B(A, p) - v^{B'}(A, p) &= \sum_{j \in [m]} p_j \cdot \max_i A_{ij}^B - \sum_{j \in [m]} p_j \cdot \max_i A_{ij}^{B'} \\
&= \sum_{j \in b} p_j \left(\max_i A_{ij}^B - \max_i A_{ij}^{B'} \right) \\
&\leq \sum_{j \in b} p_j \left(A_{i^*j}^B - A_{i^*j}^{B'} \right) \\
&= p_{j^*} \left(\frac{p_{j^*} + p^+}{p_{j^*} + p^+ + p^0} - 1 \right) + \sum_{j \in b \setminus \{j^*\}} p_j \left(\frac{p_{j^*} + p^+}{p_{j^*} + p^+ + p^0} - \frac{p^+}{p^+ + p^0} \right) \\
&= \frac{p_{j^*} + p^+}{p_{j^*} + p^+ + p^0} \left(p_{j^*} + \sum_{j \in b \setminus \{j^*\}} p_j \right) - p_{j^*} - \frac{p^+}{p^+ + p^0} \sum_{j \in b \setminus \{j^*\}} p_j \\
&= 0,
\end{aligned}$$

where the second last equality is just a rearrangement of terms and the last one follows from the fact that $\sum_{j \in b \setminus \{j^*\}} p_j = p^+ + p^0$. Hence, the partition value does not decrease. By repeating this argument, we will reach an optimal partition scheme that satisfies property P5. Then, using arguments similar to the ones used in the proof of Alon et al. [2013] for Lemma 5.1¹ is we can prove that the resulting partition scheme can be transformed in such a way so that it satisfies properties P2, P3, and P4. \square

What Lemma 5.12 says is that the contribution of column $j \in A^+$ to the partition value comes from a row i such that either $j \in A_i^+$ and $\{j\}$ forms a column-covering bundle (and, hence, its smooth value is 1) or $j \in A_i^0$ and j belongs to the mixed bundle of row i (and the smooth value of its entries is strictly smaller than 1). A non-zero contribution of a column $j \in A^0$ to the partition value always comes from a row i where j belongs to the mixed bundle. A column $j \in A^0$ can have a contribution of zero to the optimal partition value when no mixed bundle exists². Hence, the problem of computing the partition scheme of optimal partition value is equivalent to deciding the row from which each column contributes to the partition value, either as a one-column that is part of a (not necessarily full) cover or as a zero-column that is part of a mixed bundle.

Let B be a partition scheme and S be a set of columns whose contribution to the partition value of B comes from row i (i.e., i is a row that maximizes the smooth value A_{ij}^B for each

¹Invoking Lemma 5.2 in order to prove property P2 is crucial here; verifying properties P3 and P4 is much easier.

²As an example of such an extreme case, consider an instance with a $k \times (k+1)$ matrix that consists of the identity $k \times k$ matrix and an extra zero-column, and has a uniform probability distribution over the columns. The optimal partition scheme contains a full cover and all-zero bundles only, and the zero-column has no contribution to the partition value.

column j in S). Denoting the sum of these contributions by $R_i(S) = \sum_{j \in S} p_j \cdot A_{ij}^B$, we can equivalently express $R_i(S)$ as

$$R_i(S) = \sum_{j \in S \cap A_i^+} p_j + \frac{\sum_{j \in S \cap A_i^0} p_j \sum_{j \in A_i^+ \setminus S} p_j}{\sum_{j \in S \cap A_i^0} p_j + \sum_{j \in A_i^+ \setminus S} p_j}.$$

The first sum represents the contribution of columns of $S \cap A_i^+$ to the partition value (through column-covering bundles) while the second sum represents the contribution of the columns in $S \cap A_i^0$ which are bundled together with all 1-value entries in $A_i^+ \setminus S$ in the mixed bundle of row i . Then, the partition scheme B can be thought of as a collection of disjoint sets S_i (with one set per row) such that S_i contains those columns whose entries achieve their maximum smooth value in row i . Hence, the partition value of B is $v^B(A, p) = \sum_{i \in [n]} R_i(S_i)$ and the problem is essentially equivalent to welfare maximization where the rows act as the agents who will be allocated bundles of items (corresponding to columns).

Lemma 5.13. *For every row i , the function R_i is non-decreasing and submodular.*

Proof. We will show that the function R_i is non-decreasing and has decreasing marginal utilities, i.e.,

- (monotonicity) for every set S and item $x \notin S$, it holds that $R_i(S) \leq R_i(S \cup \{x\})$;
- (decreasing marginal utilities) for every pair of sets S, T such that $S \subseteq T$ and every item $x \notin T$, it holds that $R_i(S \cup \{x\}) - R_i(S) \geq R_i(T \cup \{x\}) - R_i(T)$.

In order to simplify notation, let us define the functions $\alpha(S) = \sum_{j \in S \cap A_i^+} p_j$, $\beta(S) = \sum_{j \in S \cap A_i^0} p_j$ and $\gamma(S) = \sum_{j \in A_i^+ \setminus S} p_j$. We can rewrite the function R_i as

$$R_i(S) = \alpha(S) + \frac{\beta(S) \cdot \gamma(S)}{\beta(S) + \gamma(S)}.$$

Let $S, T \subseteq [m]$ be two sets of columns such that $S \subseteq T$ and let x be a column that does not belong to set T . We distinguish between two cases depending on x . If $x \in A_i^+$, observe that

- $\alpha(S \cup \{x\}) = \alpha(S) + p_x$ and $\alpha(T \cup \{x\}) = \alpha(T) + p_x$;
- $\beta(S \cup \{x\}) = \beta(S)$ and $\beta(T \cup \{x\}) = \beta(T)$;
- $\gamma(S \cup \{x\}) = \gamma(S) - p_x$ and $\gamma(T \cup \{x\}) = \gamma(T) - p_x$.

Using the definition of function R_i , we have

$$\begin{aligned}
R_i(S \cup \{x\}) - R_i(S) &= p_x + \beta(S) \left(\frac{\gamma(S) - p_x}{\beta(S) + \gamma(S) - p_x} - \frac{\gamma(S)}{\beta(S) + \gamma(S)} \right) \\
&= p_x - \frac{p_x \beta(S)^2}{(\beta(S) + \gamma(S))(\beta(S) + \gamma(S) - p_x)} \\
&\geq p_x - \frac{p_x \beta(S)^2}{(\beta(S) + \gamma(T))(\beta(S) + \gamma(T) - p_x)} \\
&\geq p_x - \frac{p_x \beta(T)^2}{(\beta(T) + \gamma(T))(\beta(T) + \gamma(T) - p_x)} \\
&= R_i(T \cup \{x\}) - R_i(T).
\end{aligned}$$

The first inequality follows since γ is non-increasing and $S \subseteq T$ and the second inequality follows by applying twice (with $a = \gamma(T)$ and $a = \gamma(T) - p_x$, respectively) the fact that the function $f(z) = \frac{z}{z+a}$ for $a \geq 0$ is non-decreasing.

If instead $x \in A_i^0$, observe that

- $\alpha(S \cup \{x\}) = \alpha(S)$ and $\alpha(T \cup \{x\}) = \alpha(T)$;
- $\beta(S \cup \{x\}) = \beta(S) + p_x$ and $\beta(T \cup \{x\}) = \beta(T) + p_x$;
- $\gamma(S \cup \{x\}) = \gamma(S)$ and $\gamma(T \cup \{x\}) = \gamma(T)$.

Hence, we have

$$\begin{aligned}
R_i(S \cup \{x\}) - R_i(S) &= \gamma(S) \left(\frac{\beta(S) + p_x}{\beta(S) + \gamma(S) + p_x} - \frac{\beta(S)}{\beta(S) + \gamma(S)} \right) \\
&= \frac{p_x \gamma(S)^2}{(\beta(S) + \gamma(S))(\beta(S) + \gamma(S) + p_x)} \\
&\geq \frac{p_x \gamma(S)^2}{(\beta(T) + \gamma(S))(\beta(T) + \gamma(S) + p_x)} \\
&\geq \frac{p_x \gamma(T)^2}{(\beta(T) + \gamma(T))(\beta(T) + \gamma(T) + p_x)} \\
&= R_i(T \cup \{x\}) - R_i(T).
\end{aligned}$$

Again, the first inequality follows since β is clearly non-decreasing and $S \subseteq T$ and the second inequality follows by applying twice (with $a = \beta(T)$ and $a = \beta(T) + p_x$, respectively) the fact that the function $f(z) = \frac{z}{z+a}$ with $a \geq 0$ is non-decreasing.

We have completed the proof that R_i has decreasing marginal utilities. In order to establish monotonicity, it suffices to observe that the quantity at the right-hand side of the second equality in each of the above two derivations starting with $R_i(S \cup \{x\}) - R_i(S)$ is non-negative. □

Lehmann et al. [2006] presented a simple greedy algorithm that uses value queries and yields a $1/2$ -approximation of the optimal welfare for the submodular welfare maximization problem. This algorithm considers the items one by one in arbitrary order and assigns item j to an agent that maximizes the marginal valuation (the additional value from the allocation of item j). In our setting, this algorithm considers the one-columns first and the zero-columns afterwards. Whenever considering a one-column j , a column-covering bundle $\{j\}$ is formed at an arbitrary row i with $j \in A_i^+$ (such a decision definitely maximizes the increase in the partition value). Once all one-columns have been processed, the remaining 1-value entries (that did not form column-covering bundles) in each row are grouped into a bundle. All these bundles are available to host zero-columns (that will be processed next) and evolve into mixed ones. Afterwards, whenever considering a zero-column, the algorithm includes it to a mixed bundle that maximizes the increase in the partition value. Using the terminology we used in Section 5.4, the algorithm essentially starts with an arbitrary cover of the one-columns and then it runs the greedy completion procedure.

Again, we use the term *greedy algorithm* to refer to the whole family of algorithms that are defined by different implementations of the several missing details in the above description, such as the order in which the one-columns are processed, the particular way the column-covering bundles are selected, the order in which the zero-columns are processed, and the way ties are broken between different mixed bundles to which a zero-column can be added. Our analysis below holds for any member of this family.

Theorem 5.14. *The greedy algorithm for the asymmetric binary matrix partition problem has approximation ratio at least $1/2$. This bound is tight.*

Proof. The lower bound holds by the equivalence of the greedy algorithm with the algorithm studied by Lehmann et al. [2006]. Below, we prove the upper bound. In particular, we show that for every $\epsilon > 0$, there exists an instance of the problem in which the greedy algorithm obtains a partition scheme whose value is at most $1/2 + \epsilon$ of the optimal one.

Let $k > 0$ be a positive integer and α significantly higher than k . Consider the instance of the asymmetric binary matrix partition that consists of the following $(k + 1) \times (k + 1)$ matrix

$$A = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \\ 1 & 1 & \cdots & 1 & 0 \end{pmatrix}$$

where $p_j = \frac{1}{k+\alpha}$ for $j \in [k]$ and $p_{k+1} = \frac{\alpha}{k+\alpha}$. So, the first k columns and rows of A form an identity matrix, the last column has only 0-value entries and the last row consists of k 1-value entries in the first k columns. In order to lower-bound the optimal partition value, consider the partition scheme consisting of a full cover that contains the 1-value entries (i, i) for $i \leq k$, and a bundle containing the whole $(k+1)$ -th row. The optimal partition value is lower-bounded by the value of this partition scheme. By simple calculations, we obtain

$$\text{OPT} \geq \frac{k^2 + 2\alpha k}{(k + \alpha)^2}.$$

On the other hand, the greedy algorithm may select first to cover the k one-columns using the 1-value entries $(k+1, j)$ for $j \leq k$ and, then, bundle the zero-column together with only one 1-value entry in some of the first k rows. The partition value of the greedy algorithm is then

$$\text{GREEDY} = \frac{k + (k+1)\alpha}{(k + \alpha)(\alpha + 1)}.$$

Hence, the ratio between the two partition values is

$$\frac{\text{GREEDY}}{\text{OPT}} \leq \frac{(k + \alpha)(k + (k+1)\alpha)}{(k^2 + 2\alpha k)(\alpha + 1)}.$$

Pick an arbitrarily small $\delta > 0$; then, there exist a value for α (significantly higher than k) so that the above ratio satisfies $\frac{\text{GREEDY}}{\text{OPT}} \leq \frac{k+1}{2k} + \delta$. The theorem follows by picking k sufficiently large and δ sufficiently small. \square

We can use the more sophisticated smooth greedy algorithm of Vondrák [2008], which uses value queries to obtain the following.

Corollary 5.15. *There exists a $(1 - 1/e)$ -approximation algorithm for the asymmetric binary matrix partition problem.*

One might hope that due to the particular form of the functions R_i , better approximation guarantees could be possible using the $(1 - 1/e + \epsilon)$ -approximation algorithm of Feige and Vondrák [2010] which requires that demand queries of the form

given agent i and a price q_j for every item $j \in [m]$, select the bundle S that maximizes the difference $R_i(S) - \sum_{j \in S} q_j$

can be answered in polynomial time. Unfortunately, in our setting, this is not the case in spite of the very specific form of the function R_i .

Lemma 5.16. *Answering demand queries associated with the asymmetric binary matrix partition problem are NP-hard.*

Proof. We use reduction from Partition to show that the following (very restricted) decision version DQ of a demand query is NP-hard.

DQ: Given a $1 \times m$ binary matrix A , probabilities p_j and prices q_j for column $j \in [m]$, is there a set $S \subseteq [m]$ such that $R_i(S) - \sum_{j \in S} q_j \geq 5/18$?

We start from an instance of Partition consisting of a collection C of t items of integer size w_1, w_2, \dots, w_t and the question of whether there exists a subset $Y \subseteq C$ of items such that

$$\sum_{j \in Y} w_j = \sum_{j \in C \setminus Y} w_j = \frac{1}{2} \sum_{j \in C} w_j.$$

Define $W = \sum_{j \in C} w_j$. Given this instance, we construct an instance of DQ with $m = t + 1$ as follows. The binary matrix A consists of a single row that contains t 1-value entries with associated probabilities $\frac{w_1}{2W}, \frac{w_2}{2W}, \dots, \frac{w_t}{2W}$ and a 0-value entry with associated probability $1/2$. Set the prices as $q_j = \frac{5w_j}{18W}$ for $j = 1, \dots, t$ and $q_{t+1} = 0$.

By the definition of the function R_i , given a set $S \subseteq [t + 1]$, we have

$$\begin{aligned} R_i(S) - \sum_{j \in S} q_j &= \frac{1}{2W} \sum_{j \in S \setminus \{t+1\}} w_j + \frac{\frac{1}{4W} \sum_{j \in [t] \setminus S} w_j}{\frac{1}{2} + \frac{1}{2W} \sum_{j \in [t] \setminus S} w_j} - \frac{5}{18W} \sum_{j \in S \setminus \{t+1\}} w_j \\ &= \frac{2}{9} - \frac{2}{9W} \sum_{j \in [t] \setminus S} w_j + \frac{\sum_{j \in [t] \setminus S} w_j}{2W + 2 \sum_{j \in [t] \setminus S} w_j}. \end{aligned}$$

Now, consider the function $f(z) = \frac{2}{9} - \frac{2z}{9W} + \frac{z}{2W + 2z}$; the equality above implies that

$$R_i(S) - \sum_{j \in S} q_j = f\left(\sum_{j \in [t] \setminus S} w_j\right).$$

By nullifying the derivative of function f , we obtain that it has a unique maximum at $z = W/2$. Since $f(W/2) = 5/18$, the instance of DQ is equivalent to asking whether there exists a set S such that $\sum_{j \in [t] \setminus S} w_j = W/2$, which is equivalent to asking whether there exists a set of items of total size $W/2$ in the instance of Partition. \square

5.6 Conclusion

In this chapter, we studied the asymmetric matrix partition problem that is related to revenue maximization in take-it-or-leave-it sales, and focused on its binary version. In short, an instance

of the problem consists of a matrix of non-negative real values, and a probability distribution over its columns that can either be uniform or non-uniform. The goal is to find a partition of every row of the matrix into asymmetric bundles so that the expected value of each column is maximized.

For the case where the probability distribution over the matrix columns is uniform, we designed a simple greedy $9/10$ -approximation algorithm, whose analysis was heavily based on dual fitting techniques. For the case where the probability distribution is non-uniform, we showed that there exists a $(1 - 1/e)$ -approximation algorithm, by reducing the problem to the problem of submodular welfare maximization. Both of these results significantly improve upon the corresponding results presented in the previous work of Alon et al. [2013].

Chapter 6

Conclusions and open problems

In the previous four chapters of this thesis, we focused on the presentation of the results that we were able to obtain for the different problems that we studied. In particular, we designed and analyzed simple resource allocation mechanisms for budget-constrained users in Chapter 2, we bounded the price of anarchy and stability of compromising opinion formation games in Chapter 3, we designed truthful mechanisms for ownership transfer using expert advice in Chapter 4, and, finally, we designed efficient approximation algorithms for the asymmetric binary matrix partition problem in Chapter 5. However, in each of these problems, our work inevitably leaves open several interesting and important questions as well as reveals new ones. In this concluding chapter of the thesis, we discuss several of these possible directions for future research.

6.1 Resource allocation and auctions for budget-constrained users

Even though we have revealed an almost complete picture on the liquid price of anarchy of resource allocation mechanisms in Chapter 2, the gap between the lower bound of $2 - 1/n$ for all mechanisms and the bound of 2 that the Kelly mechanism is able to achieve leaves the following interesting open question:

Open question 6.1. Is the $2 - 1/n$ bound achievable, preferably by a simple mechanism?

In particular, is there a mechanism with proportional allocation function and appropriate non-pay-your-signal payments that achieves this LPoA bound? This question seems technically challenging even for the case of two players only, where our best 2-player mechanism E2-SR (presented in Section 2.7.2) can achieve an LPoA bound of approximately 1.53, while the lower bound is 1.5.

Regarding the liquid price of anarchy over more general equilibrium concepts (like mixed and correlated equilibria) or settings with incomplete information (and Bayes-Nash equilibria), our results lead to the following natural open question:

Open question 6.2. Is the Kelly mechanism still optimal within low-order terms for general equilibrium concepts?

Caragiannis and Voudouris [2016] showed that the set of mixed Nash equilibria induced by the Kelly mechanism coincides with that of pure Nash equilibria, even when the users have budget constraints. Therefore, it turns out that Kelly is indeed optimal within low-order terms for mixed Nash equilibria. However, for even more general equilibrium concepts, we are far from answering this question. The papers by Caragiannis and Voudouris [2016] and by Christodoulou et al. [2016b] present such LPoA bounds for Kelly over coarse-correlated and pure Bayes-Nash equilibria, but these are not known to be tight. We conjecture that the proof of tight LPoA bounds over more general equilibrium concepts for any resource allocation mechanism should exploit the structure of worst-case games and equilibria as we did in Chapter 2 for pure Nash equilibria. Unfortunately, extending our characterization from Section 2.5 to more general equilibrium concepts seems elusive at this point.

In Section 2.8.1, we proved a slightly weaker lower bound of $4/3$ on the liquid price of anarchy of any budget-aware resource allocation mechanism. This leaves open the possibility of finding such a mechanism that could beat the bound of $2 - 1/n$.

Open question 6.3. Which is the best budget-aware resource allocation mechanism?

Unfortunately, our characterization of worst-case games and equilibria from Section 2.5 does not seem to extend to the case of known budgets. Therefore, in order to be able to prove tight bounds and pinpoint the best budget-aware mechanism, we need to obtain a different characterization, which is an extremely challenging and technically non-trivial task.

Finally, in general, we believe that the liquid welfare is an appropriate efficiency benchmark for auctions with budget-constrained players. The recent paper by Azar et al. [2017] studies the LPoA of simultaneous first-price auctions over Bayes-Nash equilibria, while the paper by Voudouris [2018] focuses on the LPoA of position mechanisms over pure Nash equilibria. Obtaining similar results for other auction formats is certainly an important future research direction; see the recent survey of Roughgarden et al. [2017] on the price of anarchy of auction mechanisms.

Open question 6.4. Which is the best auction mechanism for budget-constrained players with respect to the liquid welfare efficiency benchmark?

Needless to say, we do not expect that the liquid welfare is unique as a measure of efficiency in settings with budgets. Defining alternative efficiency benchmarks and studying the price of anarchy with respect to them would shed extra light to the strengths and weaknesses of auction mechanisms.

6.2 Compromising opinion formation

In Chapter 3, we introduced the class of compromising opinion formation (k -COF) games by enriching that of co-evolutionary opinion games with a cost function that urges players to essentially *meet halfway*. Our findings indicate that the quality of their equilibria grows linearly with the neighborhood size k , but there exists a gap between our lower and upper bounds for $k \geq 2$; closing this gap seems to be a challenging technical task and may require different analysis techniques.

Open question 6.5. What is the tight bound on the price of anarchy and stability of k -COF games for $k \geq 2$?

Furthermore, for 1-COF games, due to the tight bound of 3 for pure equilibria and the lower bound of 6 for mixed equilibria, we know the equality of mixed equilibria is strictly worse than that of pure ones. However, we were not able to prove any upper bounds on their price of anarchy.

Open question 6.6. Is the price of anarchy over mixed equilibria still linear?

Another natural question is about the complexity of pure equilibria in k -COF games. For $k = 1$, we managed to show that computing the best and worst pure equilibria can be done by searching for paths of minimum and maximum total weight in directed acyclic graphs where the node correspond to partial segments of the game.

Open question 6.7. Can we efficiently compute pure Nash equilibria for $k \geq 2$?

Driven by our positive results for 1-COF games, we conjecture that there exists a polynomial time algorithm for computing equilibria in more general k -COF games, but finding such an algorithm remains elusive at this point. Similarly, one could also focus on the complexity of computing optimal opinion vectors, even for $k = 1$.

Finally, our modeling assumption that the number of neighbors is equal for all players is rather restrictive. It would be interesting to investigate whether our results (qualitative and algorithmic) can be extended to more general scenarios.

Open question 6.8. Do our results extend to compromising opinion formation games with players of different neighborhood sizes?

One possible such generalization is to combine our approach with the Hegselmann-Krause model so that the neighborhood of each player i consists solely the players $j \neq i$ with opinions that are sufficiently close to the i 's belief.

6.3 Ownership transfer

In Chapter 4, we presented a series of positive and negative results for a simple mechanism design model with and without monetary transfers, which we believe that captures the main challenges in the implementation of ownership transfer. Still, closing the gap between the approximation ratio of $5/4$ of the template mechanism R (see Section 4.7) and our general unconditional lower bound of approximately 1.14 for any truthful mechanism (see Section 4.8) is an important and definitely non-trivial challenge.

Open question 6.9. Which is the best possible achievable approximation ratio?

A possible direction towards answering the above question could be to consider extensions of the template mechanisms by exploiting a few more bits of information about the preferences of the expert. One could also consider the alternative of using bid-independent mechanisms embedded with extra bits of information that could be distilled by the values reported by the bidders.

Besides the aforementioned concrete open problem that is directly related to our results in this thesis, there are many natural extensions of the model that are worth studying. For example, we have weighed equally the contribution of the expert and the agents to the social welfare. We can generalize the definition of the welfare by introducing a factor of $\alpha > 0$, by which the contribution of the expert will be multiplied.

Open question 6.10. Can we design near-optimal truthful mechanisms for the different values of the parameter α ?

In cases where the parameter α is very large or very small, we expect that bid-independent and expert-independent mechanisms will be almost optimal, respectively. However, we suspect

that there are values of the parameter α (close to 1) that make the mechanism design problem even more interesting.

Another extension could be to consider a different optimization objective; for example, by mixing the welfare of the expert with the revenue that can be extracted by the bidders.

Open question 6.11. Can we design truthful mechanisms that maximize the sum of expert welfare and revenue?

We remark that in order for the revenue to be (part of) a meaningful objective, one would have to restrict attention to *individually rational* mechanisms that guarantee non-negative utility to the agents for participating. This is an important property, since otherwise a truthful mechanism could simply ignore their bids and charge them the maximum possible amount. In fact, the related literature on revenue-maximization focuses on mechanisms which are individually rational for this reason. However, in our setting, it is not hard to see that bid-independent, individually rational mechanisms must always extract zero revenue. It is also well-documented that revenue maximization is a less meaningful objective in the absence of prior knowledge of the values of the agents [Hartline, 2013], and it is commonly assumed that these values are drawn from some known distributions [Myerson, 1981, Nisan et al., 2007]. Therefore, designing efficient truthful mechanisms for such an optimization objective requires radically different ideas, or perhaps even the migration to a Bayesian setting.

Our model of one expert and two competing bidders can be thought of as the simplest possible non-trivial ownership transfer scenario. There are many important generalizations that one could consider for future research. Indicatively, these could include larger populations of experts and agents, more than one assets to be transferred with combinatorial constraints governing their acquisition, or even dynamic expert preferences that depend on the bidding information. All of these lead to the following abstract open question:

Open question 6.12. Can we design near-optimal truthful mechanisms for generalizations of ownership transfer?

Finally, we believe that the combination of mechanism design with and without money can be exploited in different contexts as well, especially in settings where the agents are partitioned into groups depending on whether they value money or not. Our setting of ownership transfer is such a setting, but it is definitely not unique in its kind.

6.4 Asymmetric matrix partition

In Chapter 5, we focused on the binary version of the asymmetric matrix partition problem and presented improved approximation algorithms for uniform and non-uniform probability distributions, compared to the previous work of Alon et al. [2013]. Designing algorithms with even better approximation guarantees or proving stronger inapproximability results for this version of this problem is a first obvious open problem.

Open question 6.13. What are the limits of approximation for the asymmetric (binary) matrix partition problem?

Recall (see Section 1.4) that the motivation behind the definition of the asymmetric matrix partition problem comes from revenue maximization in take-it-or-leave-it sales, where the goal is to exploit possible asymmetries in the information of the seller and of the potential buyers. Admittedly, in the (uniform) binary case of the problem, the fact that the greedy partition schemes contain column-covering bundles makes it possible for a buyer to distinguish between cases in which she is actually offered an item that she values as 1 (a singleton bundle with smooth value of 1) or 0 (a mixed bundle). This is clearly a drawback and asymmetric binary matrix partition should not be used to model such simple take-it-or-leave-it sales. One possible remedy could be to lower-bound the size of any bundle with non-zero value or require some symmetry among the bundles that contain any given zero-column, so that no information about the item selected by nature is revealed to the buyer by the seller.

Open question 6.14. Given additional constraints that guarantee no information revelation, can we design near-optimal approximation algorithms?

Still, we believe that asymmetric binary matrix partition is important as an algorithmically challenging problem and can provide insights to efficient solutions for revenue maximization. In this direction, the above issue does not seem to be as severe in the general asymmetric matrix partition. This is justified by the assumption that buyers do not know each other and information about the particular item that is selected to be sold is not as easy to be inferred.

Bibliography

- F. Abed, I. Caragiannis, and A. A. Voudouris. 2018. Near-optimal asymmetric binary matrix partitions. *Algorithmica* 80, 1 (2018), 48–72.
- G. Aggarwal, A. Goel, and R. Motwani. 2006. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC)*. 1–7.
- G. A. Akerlof. 1970. The market for ‘lemons’: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84 (1970), 488–500.
- N. Alon, M. Feldman, I. Gamzu, and M. Tennenholtz. 2013. The asymmetric matrix partition problem. In *Proceedings of the 9th Conference on Web and Internet Economics (WINE)*. 1–14.
- E. Anshelevich, O. Bhardwaj, and J. Postl. 2015. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*. 777–783.
- E. Anshelevich, A. Dasgupta, J. M. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. 2008. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing* 38, 4 (2008), 1602–1623.
- A. Archer and E. Tardos. 2001. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS)*. 482–491.
- I. Ashlagi, F. A. Fischer, I. A. Kash, and A. D. Procaccia. 2015. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior* 91 (2015), 284–296.
- S. Athanassopoulos, I. Caragiannis, and C. Kaklamanis. 2009a. Analysis of approximation algorithms for k-set cover using factor-revealing linear programs. *Theory of Computing Systems* 45, 3 (2009), 555–576.

- S. Athanassopoulos, I. Caragiannis, C. Kaklamanis, and M. Kyropoulou. 2009b. An improved approximation bound for spanning star forest and color saving. In *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science (MFCS)*. 90–101.
- A. J. Auerbach. 2008. *Mergers and acquisitions*. University of Chicago Press.
- V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, and G. Persiano. 2015. Minority becomes majority in social networks. In *Proceedings of the 11th International Conference on Web and Internet Economics (WINE)*. 74–88.
- V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, and G. Persiano. 2016. Generalized discrete preference games. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*. 53–59.
- V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, and G. Persiano. 2017a. Information retention in heterogeneous majority dynamics. In *Proceedings of the 13th International Conference on Web and Internet Economics (WINE)*. 30–43.
- V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, and G. Persiano. 2017b. Robustness in discrete preference games. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1314–1322.
- Y. Azar, M. Feldman, N. Gravin, and A. Roytman. 2017. Liquid price of anarchy. In *Proceedings of the 10th International Symposium on Algorithmic Game Theory (SAGT)*. 3–15.
- S. Barbera, A. Bogomolnaia, and H. van der Stel. 1998. Strategy-proof probabilistic rules for expected utility maximizers. *Mathematical Social Sciences* 35, 2 (1998), 89–103.
- U. Bhaskar, V. Dani, and A. Ghosh. 2018. Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*. 925–932.
- K. Bhawalkar, S. Gollapudi, and K. Munagala. 2013. Coevolutionary opinion formation games. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*. 41–50.
- V. Bilò, A. Fanelli, and L. Moscardelli. 2016. Opinion formation games with dynamic social influences. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*. 444–458.

- D. Bindel, J. M. Kleinberg, and S. Oren. 2015. How bad is forming your own opinion? *Games and Economic Behavior* 92 (2015), 248–265.
- C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. 2015. Optimal social choice functions: a utilitarian view. *Artificial Intelligence* 227 (2015), 190–213.
- F. Brandt, V. Conitzer, U. Endriss, A. D. Procaccia, and J. Lang. 2016. *Handbook of computational social choice*. Cambridge University Press.
- P. Briest, P. Krysta, and B. Vöcking. 2011. Approximation techniques for utilitarian mechanism design. *SIAM J. Comput.* 40, 6 (2011), 1587–1622.
- I. Caragiannis. 2009. Wavelength management in WDM rings to maximize the number of connections. *SIAM Journal on Discrete Mathematics* 23, 2 (2009), 959–978.
- I. Caragiannis, X. Chatzigeorgiou, G. A. Krimpas, and A. A. Voudouris. 2017. Optimizing positional scoring rules for rank aggregation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 430–436. Full version available as CoRR abs/1609.07460.
- I. Caragiannis, A. Filos-Ratsikas, S. Nath, and A. A. Voudouris. 2018. Truthful mechanisms for ownership transfer with expert advice. *CoRR abs/1802.01308* (2018).
- I. Caragiannis, A. Filos-Ratsikas, and A. D. Procaccia. 2015. An improved 2-agent kidney exchange mechanism. *Theoretical Computer Science* 589 (2015), 53–60.
- I. Caragiannis, C. Kaklamanis, and M. Kyropoulou. 2013. Tight approximation bounds for combinatorial frugal coverage algorithms. *Journal of Combinatorial Optimization* 26, 2 (2013), 292–309.
- I. Caragiannis, P. Kanellopoulos, and A. A. Voudouris. 2017a. Bounding the inefficiency of compromise. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. 142–148. Full version available as CoRR abs/1702.07309.
- I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. 2017b. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58 (2017), 123–152.
- I. Caragiannis and A. D. Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9–10 (2011), 1655–1671.

- I. Caragiannis, A. D. Procaccia, and N. Shah. 2016. When do noisy votes reveal the truth? *ACM Transactions on Economics and Computation* 4, 3 (2016), article 15.
- I. Caragiannis and A. A. Voudouris. 2016. Welfare guarantees for proportional allocations. *Theory of Computing Systems* 59, 4 (2016), 581–599.
- I. Caragiannis and A. A. Voudouris. 2018. The efficiency of resource allocation mechanisms for budget-constrained users. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 681–698. Full version available as CoRR abs/1707.03551.
- G. Chalkiadakis, E. Elkind, and M. Wooldridge. 2011. *Computational aspects of cooperative game theory*. Morgan & Claypool Publishers.
- B. Chazelle. 2012. Natural algorithms and influence systems. *Communications of the ACM* 55, 12 (2012), 101–110.
- X. Chen, X. Deng, and S. H. Teng. 2009. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM* 56, 3 (2009), 14:1–14:57.
- Y. K. Cheng. 2016. Better strategyproof mechanisms without payments or prior – an analytic approach. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*. 194–200.
- F. Chierichetti, J. M. Kleinberg, and S. Oren. 2018. On discrete preferences and coordination. *Journal of Computer and System Sciences* 93 (2018), 11–29.
- G. Christodoulou and E. Koutsoupias. 2005. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*. 67–73.
- G. Christodoulou, A. Kovács, and M. Schapira. 2016a. Bayesian combinatorial auctions. *Journal of the ACM* 63, 2 (2016), 11:1–11:19.
- G. Christodoulou, A. Sgouritsa, and B. Tang. 2016b. On the efficiency of the proportional allocation mechanism for divisible resources. *Theory of Computing Systems* 59, 4 (2016), 600–618.
- E.H. Clarke. 1971. Multipart pricing of public goods. *Public Choice* 11, 1 (1971), 17–33.
- R. Cole, V. Gkatzelis, and G. Goel. 2013. Mechanism design for fair division: allocating divisible items without payments. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*. 251–268.

- V. Crawford and J. Sobel. 1982. Strategic information transmission. *Econometrica* 50 (1982), 1431–1451.
- A. Das, S. Gollapudi, and K. Munagala. 2014. Modeling opinion dynamics in social networks. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining (WSDM)*. 403–412.
- C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *Commun. ACM* 52, 2 (2009), 89–97.
- M. H. DeGroot. 1974. Reaching a consensus. *Journal of the American Statistical Association* 69, 345 (1974), 118–121.
- S. Dobzinski, N. Nisan, and M. Schapira. 2012. Truthful randomized mechanisms for combinatorial auctions. *J. Comput. System Sci.* 78, 1 (2012), 15–25.
- S. Dobzinski and R. Paes Leme. 2014. Efficiency guarantees in auctions with budgets. In *Proceedings of the 41st International Colloquium on Automata, Languages and Programming (ICALP)*. 392–404.
- S. Dughmi. 2014. On the hardness of signaling. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. 354–363.
- S. Dughmi, A. Eden, M. Feldman, A. Fiat, and S. Leonardi. 2016. Lottery pricing equilibria. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. 401–418.
- B. Dutta, H. Peters, and A. Sen. 2007. Strategy-proof cardinal decision schemes. *Social Choice and Welfare* 28, 1 (2007), 163–179.
- Y. Emek, M. Feldman, I. Gamzu, R. Paes Leme, and M. Tennenholtz. 2012. Signaling schemes for partition value maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 514–531.
- S. Rasoul Etesami and T. Basar. 2015. Game-theoretic analysis of the Hegselmann-Krause model for opinion dynamics in finite dimensions. *IEEE Transactions on Automatic Control* 60, 7 (2015), 1886–1897.
- European Commission. 2018. Merger control procedures. Online. (2018). http://ec.europa.eu/competition/mergers/procedures_en.html.

- U. Feige and S. Jozeph. 2015. Oblivious algorithms for the maximum directed cut problem. *Algorithmica* 71, 2 (2015), 404–428.
- U. Feige and M. Tennenholtz. 2010. Responsive lotteries. In *Proceeding of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*. 150–161.
- U. Feige and J. Vondrák. 2010. The submodular welfare problem with demand queries. *Theory of Computing* 6 (2010), 247–290.
- D. Ferraioli, P. W. Goldberg, and C. Ventre. 2016. Decentralized dynamics for finite opinion games. *Theoretical Computer Science* 648 (2016), 96–115.
- D. Ferraioli and C. Ventre. 2017. Social pressure in opinion games. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*. 3661–3667.
- FIFA.com. 2018. Bidding process. Online. (2018). <http://www.fifa.com/governance/competition-organisation/bidding-process.html>.
- A. Filos-Ratsikas, S. K. S. Frederiksen, and J. Zhang. 2014. Social welfare in one-sided matchings: Random priority and beyond. In *Proceedings of the 7th Symposium of Algorithmic Game Theory (SAGT)*. Springer, 1–12.
- A. Filos-Ratsikas and P. B. Miltersen. 2014. Truthful approximations to range voting. In *Proceedings of the 10th International Conference on Web and Internet Economics (WINE)*. Springer, 175–188.
- D. Fotakis, D. Palyvos-Giannas, and S. Skoulakis. 2016. Opinion dynamics with local interactions. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*. 279–285.
- X. Freixas. 1984. A cardinal approach to straightforward probabilistic mechanisms. *Journal of Economic Theory* 34, 2 (1984), 227–251.
- N. E. Friedkin and E. C. Johnsen. 1990. Social influence and opinions. *Journal of Mathematical Sociology* 15, 3-4 (1990), 193–205.
- A. Ghosh, H. Nazerzadeh, and M. Sundararajan. 2007. Computing optimal bundles for sponsored search. In *Proceedings of the 3rd International Workshop on Web and Internet Economics (WINE)*. 576–583.

- A. Gibbard. 1973. Manipulation of voting schemes. *Econometrica* 41 (1973), 587–602.
- A. Gibbard. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45 (1977), 665–681.
- A. Gibbard. 1978. Straightforwardness of game forms with lotteries as outcomes. *Econometrica* 46, 3 (1978), 595–614.
- M. Gomez-Rodriguez, J. Leskovec, and A. Krause. 2012. Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data* 5, 4 (2012), 21:1–21:37.
- T. Groves. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society* (1973), 617–631.
- M. Guo and V. Conitzer. 2010. 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)*. 881–888.
- M. Guo and A. Deligkas. 2013. Revenue maximization via hiding item attributes. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*. 157–163.
- B. Hajek and G. Gopalakrishnan. 2002. Do greedy autonomous systems make for a sensible Internet? *Unpublished manuscript* (2002).
- G. H. Hardy, J. E. Littlewood, and G. Pólya. 1934. *Inequalities*. Cambridge University Press.
- J. D. Hartline. 2013. Mechanism design and approximation, Chapter 6. *Book draft* 122 (2013).
- R. Hegselmann and U. Krause. 2002. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation* 5, 3 (2002).
- M. A. Hitt, R. D. Ireland, and J. S. J. S. Harrison. 2001. Mergers and acquisitions. *The Blackwell Handbook of Strategic Management* (2001), 377–402.
- P. Holme and M. E. J. Newman. 2006. Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E* 74 (2006).
- A. Hylland. 1980. Strategy proofness of voting procedures with lotteries as outcomes and infinite sets of strategies. *Unpublished manuscript* (1980).

- K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. V. Vazirani. 2003. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *Journal of the ACM* 50, 6 (2003), 795–824.
- M. C. W. Janssen (Ed.). 2004. *Auctioning public assets: analysis and alternatives*. Cambridge University Press.
- P. Jia, A. MirTabatabaei, N. E. Friedkin, and F. Bullo. 2015. Opinion dynamics and the evolution of social power in influence networks. *SIAM Review* 57, 3 (2015), 367–397.
- R. Johari. 2007. The price of anarchy and the design of scalable resource allocation algorithms. *Chapter 21 in Algorithmic Game Theory* (2007), 543–568.
- R. Johari and J. N. Tsitsiklis. 2004. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research* 29, 3 (2004), 407–435.
- R. Johari and J. N. Tsitsiklis. 2009. Efficiency of scalar-parameterized mechanisms. *Operations Research* 57, 4 (2009), 823–839.
- F. P. Kelly. 1997. Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8 (1997), 33–37.
- F. P. Kelly, A. Maulloo, and D. Tan. 1998. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* 49 (1998), 237–252.
- D. Kempe, J. M. Kleinberg, S. Oren, and A. Slivkins. 2016. Selection and influence in cultural dynamics. *Network Science* 4, 1 (2016), 1–27.
- D. Kempe, J. M. Kleinberg, and E. Tardos. 2015. Maximizing the spread of influence through a social network. *Theory of Computing* 11 (2015), 105–147.
- S. Khot, R. Lipton, E. Markakis, and A. Mehta. 2008. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica* 52 (2008), 3–18.
- E. Koutsoupias and C. H. Papadimitriou. 1999. Worst-case equilibria. In *Proceedings of the 16th Symposium on Theoretical Aspects of Computer Science (STACS)*. 404–413.
- R. J. La and V. Anantharam. 2000. Charge-sensitive TCP and rate control in the Internet. In *Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. 1166–1175.

- J. Laffont and E. Maskin. 1980. A differential approach to dominant strategy mechanisms. *Econometrica: Journal of the Econometric Society* (1980), 1507–1520.
- B. Lehmann, D. J. Lehmann, and N. Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2 (2006), 270–296.
- J. Levin and P. Milgrom. 2010. Online advertising: Heterogeneity and conflation in market design. *American Economic Review* 100 (2010), 603–607.
- P. Lu and T. Xiao. 2015. Improved efficiency guarantees in auctions with budgets. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*. 397–413.
- P. Lu and T. Xiao. 2017. Liquid welfare maximization in auctions with multiple items. In *Proceedings of the 10th International Symposium on Algorithmic Game Theory (SAGT)*. 41–52.
- M. Mahdian and Q. Yan. 2011. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*. 597–606.
- R. T. Maheswaran and T. Basar. 2003. Nash equilibrium and decentralized negotiation in auctioning divisible resources. *Group Decision and Negotiation* 12, 5 (2003), 361–395.
- R. T. Maheswaran and T. Basar. 2006. Efficient signal proportional allocation (ESPA) mechanisms: decentralized social welfare maximization for divisible resources. *IEEE Journal on Selected Areas in Communications* 24, 5 (2006), 1000–1009.
- P. Milgrom. 2010. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior* 70 (2010), 62–70.
- P. R. Milgrom and R. J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* 50 (1982), 1089–1122.
- P. B. Miltersen and O. Sheffet. 2012. Send mixed signals - Earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 234–247.
- E. Mossel and O. Tamuz. 2014. Opinion exchange dynamics. arXiv:1401.4770. (2014).
- A. Mualem and N. Nisan. 2008. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior* 64, 2 (2008), 612–631.

- R. B. Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- J. F. Nash. 1951. Non-cooperative games. *The Annals of Mathematics* 54, 2 (1951), 286–295.
- N. Nisan and A. Ronen. 2001. Algorithmic mechanism design. *Games and Economic Behavior* 35, 1-2 (2001), 166–196.
- N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- A. Olshevsky and J. N. Tsitsiklis. 2009. Convergence speed in distributed consensus and averaging. *SIAM Journal on Control and Optimization* 48, 1 (2009), 33–55.
- A. D. Procaccia, S. J. Reddi, and N. Shah. 2012. A Maximum likelihood approach for selecting sets of alternatives. In *Proceedings of the 28th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 695–704.
- A. D. Procaccia and M. Tennenholtz. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation* 1, 4 (2013), 18:1–18:26.
- R. T. Rockafellar. 2015. *Convex analysis*. Princeton university press.
- T. Roughgarden. 2015. Intrinsic robustness of the price of anarchy. *Journal of the ACM* 62, 5 (2015), 32:1–42.
- T. Roughgarden, V. Syrgkanis, and E. Tardos. 2017. The price of anarchy in auctions. *Journal of Artificial Intelligence Research* 59 (2017), 59–101.
- T. Roughgarden and E. Tardos. 2002. How bad is selfish routing? *Journal of the ACM* 49, 2 (2002), 236–259.
- S. Sanghavi and B. Hajek. 2004. Optimal allocation of a divisible good to strategic buyers. In *Proceedings of the 43rd IEEE Conference on Decision and Control (CDC)*.
- M. A. Satterthwaite. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10, 2 (1975), 187–217.
- N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. 2015. Belief revision games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*. 1590–1596.

- P. Skowron, P. Faliszewski, and J. Lang. 2016. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence* 241 (2016), 191–216.
- V. Syrgkanis and E. Tardos. 2013. Composable and efficient mechanisms. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*. 211–220.
- A. Tsang and K. Larson. 2014. Opinion dynamics of skeptical agents. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 277–284.
- V. V. Vazirani. 2001. *Approximation algorithms*. Springer.
- W. Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16, 1 (1961), 8–37.
- J. Vondrák. 2008. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th ACM Symposium on Theory of Computing (STOC)*. 67–74.
- A. A. Voudouris. 2018. A note on the efficiency of position mechanisms with budget constraints. *CoRR* abs/1804.09649 (2018).
- D. P. Williamson and D. B. Shmoys. 2011. *The design of approximation algorithms*. Cambridge University Press.
- Z. Xu, J. Liu, and T. Basar. 2015. On a modified DeGroot-Friedkin model of opinion dynamics. In *Proceedings of the 2015 American Control Conference (ACC)*. 1047–1052.
- S. Yang and B. E. Hajek. 2007. VCG-Kelly mechanisms for allocation of divisible goods: adapting VCG mechanisms to one-dimensional signals. *IEEE Journal on Selected Areas in Communications* 25, 6 (2007), 1237–1243.

Appendix A

Extended abstract in Greek

Σχεδιασμός και ανάλυση αλγορίθμων για μη συνεργατικά περιβάλλοντα

Αλέξανδρος Ανδρέας Βουδούρης

Τις τελευταίες δύο δεκαετίες, η ταχύτατη και συνεχώς αυξανόμενη ανάπτυξη του Διαδικτύου και των κοινωνικών δικτύων, έχει οδηγήσει στην υλοποίηση μη συνεργατικών περιβαλλόντων, όπου πολλαπλές εγωκεντρικές οντότητες ανταγωνίζονται η μία την άλλη. Για παράδειγμα, οι οντότητες μπορεί να είναι χρήστες ενός τηλεπικοινωνιακού καναλιού που ανταγωνίζονται για το περιορισμένο διαθέσιμο εύρος ζώνης, διαφημιστές που ανταγωνίζονται για τον διαθέσιμο χώρο διαφήμισης σε ιστοσελίδες αποτελεσμάτων αναζήτησης, εργολάβοι που ανταγωνίζονται για συμμετοχή σε δημόσια έργα, ή ακόμη και απλοί άνθρωποι οι οποίοι συζητάνε με τα άτομα του κοινωνικού τους περίγυρου για πολιτικά θέματα εκφράζοντας απόψεις. Σε όλα αυτά τα σενάρια, οι οντότητες είναι συνήθως εγωκεντρικές και κάθε μία από αυτές έχει ως σκοπό το να επιλέξει την καλύτερη δυνατή στρατηγική για να βελτιστοποιήσει τους προσωπικούς της στόχους, οι οποίοι δεν επηρεάζονται μόνο από την υποκείμενη δομή του περιβάλλοντος, αλλά και από τις άλλες οντότητες (και τις στρατηγικές που αυτές επιλέγουν).

Υπάρχουν πολλές σημαντικές υπολογιστικές ερωτήσεις που αφορούν την ευστάθεια και την απόδοση των συστημάτων που αναδύονται σε μη συνεργατικά περιβάλλοντα. Ποια είναι η ποιότητα (τόσο καλύτερης όσο και χειρότερης περίπτωσης) των καταστάσεων ισορροπίας που καταλήγουν τα στρατηγικά παιχνίδια (τα οποία ανακύπτουν από τη στρατηγική συμπεριφορά των οντοτήτων, οι οποίες λειτουργούν ως παίκτες); Μπορούμε να σχεδιάσουμε βελτιωμένους μηχανισμούς οι οποίοι επιπλέον να παρέχουν και τα κατάλληλα κίνητρα στους παίκτες ώστε να

αναφέρουν πάντα την αλήθεια για τις προτιμήσεις τους; Σε αυτήν την Διατριβή, απαντάμε σε τέτοιες ερωτήσεις για τέσσερα προβλήματα που προκύπτουν σε μη συνεργατικά περιβάλλοντα κατανομής διαιρέσιμων πόρων, διαμόρφωσης απόψεων, μεταφοράς ιδιοκτησίας, και μεγιστοποίησης εσόδων σε συνδυαστικές πωλήσεις. Στη συνέχεια, δίνουμε μια συνοπτική περιγραφή αυτών των προβλημάτων, καθώς και των αποτελεσμάτων μας.

A.1 Κατανομή πόρων με περιορισμούς προϋπολογισμού

Η κατανομή πόρων είναι ένα από τα βασικά προβλήματα που προκύπτουν αναπόφευκτα σε όλα τα υπολογιστικά συστήματα, σε διάφορες μορφές. Μάλιστα, τις περισσότερες φορές ο σχεδιασμός αποδοτικών λύσεων για κατανομή πόρων δημιουργεί μη-τετριμμένες αλγοριθμικές προκλήσεις. Ως εκ τούτου, η σχετική αλγοριθμική ερευνητική κοινότητα έχει απασχοληθεί με την σχεδίαση και ανάλυση αποδοτικών αλγορίθμων για προβλήματα κατανομής πόρων εδώ και δεκαετίες. Η πρόσφατη ανάπτυξη κατανεμημένων συστημάτων μεγάλης κλίμακας με μη συνεργατικούς χρήστες οι οποίοι ανταγωνίζονται για πρόσβαση σε περιορισμένους πόρους, έχει οδηγήσει στην ανάλυση σχετικών προβλημάτων κατανομής πόρων με χρήση εννοιών και εργαλείων της Θεωρίας Παιγνίων.

Μελετάμε μια συγκεκριμένη κλάση *μηχανισμών κατανομής πόρων* οι οποίοι μοιράζουν έναν διαιρέσιμο πόρο (όπως το εύρος ζώνης ενός τηλεπικοινωνιακού καναλιού, ο υπολογιστικός χρόνος μιας CPU, ο αποθηκευτικός χώρος ενός cloud κτλ.) στους χρήστες ως εξής. Κάθε χρήστης υποβάλλει ένα βαθμωτό *σήμα* (έναν μη-αρνητικό πραγματικό αριθμό). Δεδομένων αυτών των σημάτων, ο μηχανισμός αποφασίζει το μέρος του πόρου που θα πάρει κάθε χρήστης, καθώς και το ποσό των χρημάτων που θα πρέπει να πληρώσει για αυτό. Ένα κλασικό παράδειγμα είναι ο αναλογικός μηχανισμός που προτάθηκε από τον Kelly [1997] (δείτε επίσης την εργασία των Kelly et al. [1998]), σύμφωνα με τον οποίο το μέρος του πόρου που παίρνει κάθε χρήστης είναι ανάλογο του σήματος που υποβάλλει, ενώ το σήμα είναι η πληρωμή του.

Ακολουθώντας τις τυπικές υποθέσεις στη σχετική βιβλιογραφία, θεωρούμε ότι η αποτίμηση κάθε χρήστη για τα διάφορα μέρη του πόρου υπολογίζεται μέσω μιας ιδιωτικής *συνάρτησης αποτίμησης*. Ο παραπάνω ορισμός των μηχανισμών κατανομής πόρων επιτρέπει στους χρήστες να συμπεριφερθούν στρατηγικά υπό την έννοια ότι το σήμα που επιλέγουν να υποβάλλουν είναι τέτοιο ώστε η ωφέλεια τους (αποτίμηση για το μέρος του πόρου που παίρνουν μείον την πληρωμή τους) μεγιστοποιείται. Φυσικά, αυτή η συμπεριφορά ορίζει ένα στρατηγικό παιχνίδι μεταξύ των χρηστών, οι οποίοι δρουν ως παίκτες. Μετά τον ορισμό του μηχανισμού του Kelly,

μια σειρά από εργασίες μελέτησαν την ύπαρξη και τη μοναδικότητα των *αμιγών ισορροπιών κατά Nash* (καταστάσεις του παιχνιδιού όπου το σήμα κάθε παίκτη μεγιστοποιεί την προσωπική του ωφέλεια) των υποκείμενων στρατηγικών παιχνιδιών [Hajek and Gopalakrishnan, 2002, La and Anantharam, 2000, Maheswaran and Basar, 2003] και ποσοτικοποίησαν την αποδοτικότητα τους [Johari and Tsitsiklis, 2004] φράσσοντας το *κόστος της αναρχίας* [Koutsourias and Papadimitriou, 1999].

Πιο συγκεκριμένα, οι Johari and Tsitsiklis [2004] χρησιμοποίησαν ως μέτρο απόδοσης το *κοινωνικό όφελος* (συνολική αποτίμηση των παικτών για τα μέρη του πόρου που παίρνουν) και απέδειξαν ότι το κοινωνικό όφελος οποιασδήποτε ισορροπίας είναι τουλάχιστον $3/4$ φορές το βέλτιστο κοινωνικό όφελος. Αυτό μεταφράζεται σε ένα φράγμα $4/3$ για το κόστος της αναρχίας, το οποίο είναι αυστηρό. Η εργασία των Johari and Tsitsiklis [2004] πυροδότησε μεταγενέστερη έρευνα για άλλους μηχανισμούς κατανομής πόρων, οι οποίοι χρησιμοποιούν διαφορετικούς κανόνες κατανομής και πληρωμών.

Μια πρώτη προφανή ερώτηση αφορούσε το αν είναι δυνατόν να αποδείξουμε βελτιωμένα φράγματα για το κόστος της αναρχίας αλλάζοντας τον αναλογικό κανόνα κατανομής, αλλά διατηρώντας τον απλό κανόνα πληρωμών σύμφωνα με τον οποίο κάθε χρήστης πληρώνει το σήμα του (για ευκολία, θα αναφέρομαστε σε τέτοιου είδους μηχανισμό ως PYS). Οι Sanghavi and Hajek [2004] απέδειξαν ότι κανένας PYS μηχανισμός δεν μπορεί να πετύχει κόστος της αναρχίας καλύτερο του $8/7$, σχεδίασαν μια συνάρτηση κατανομής η οποία πετυχαίνει αυτό το φράγμα για την περίπτωση των δυο παικτών, και έδωσαν ισχυρές ενδείξεις ότι ένα λίγο χειρότερο φράγμα ισχύει για αυθαίρετα πολλούς παίκτες. Ίσως ένα από τα πιο απρόσμενα αποτελέσματα είναι το ότι υπάρχουν πλήρως αποδοτικοί μηχανισμοί με κόστος της αναρχίας 1. Αυτή η ανακάλυψη έγινε σε τρεις ανεξάρτητες εργασίες από τους Maheswaran and Basar [2006], Yang and Hajek [2007], και Johari and Tsitsiklis [2009]. Ο μηχανισμός των Maheswaran and Basar [2006] χρησιμοποιεί αναλογική κατανομή αλλά διαφορετικές πληρωμές, ενώ οι μηχανισμοί των Johari and Tsitsiklis [2009] και Yang and Hajek [2007] προσαρμόζουν το γνωστό VCG υπόδειγμα στην περίπτωση βαθμωτών σημάτων και διαιρέσιμων πόρων (δείτε επίσης το survey του Johari [2007] πάνω σε αυτά τα αποτελέσματα).

Εδώ, επικεντρωνόμαστε στο πιο ρεαλιστικό σενάριο όπου κάθε παίκτης έχει έναν ιδιωτικό προϋπολογισμό ο οποίος περιορίζει το ποσό των χρημάτων που μπορεί να πληρώσει και, άρα, περιορίζει και τον χώρο των πιθανών στρατηγικών του. Καθώς οι μηχανισμοί κατανομής πόρων δεν έχουν άμεση πρόσβαση στους προϋπολογισμούς, το σύνολο των ισορροπιών μπορεί

να αλλάξει δραστικά και το κοινωνικό τους όφελος ενδέχεται να είναι εξαιρετικά μικρό σε σχέση με το βέλτιστο δυνατό, το οποίο δεν συσχετίζεται με τις στρατηγικές των παικτών, τις πληρωμές τους, ή τους προϋπολογισμούς που μπορεί να έχουν. Ένα μέτρο απόδοσης το οποίο είναι πιο κατάλληλο για αυτό το σενάριο είναι γνωστό ως *ρευστό όφελος* (παρουσιάστηκε από τους Dobzinski and Paes Leme [2014] και, ανεξάρτητα, από τους Syrgkanis and Tardos [2013]) και προκύπτει αλλάζοντας λίγο τον ορισμό του κοινωνικού οφέλους. Συγκεκριμένα, για κάθε παίκτη, το ρευστό όφελος λαμβάνει υπόψη το ελάχιστο μεταξύ της αποτίμησης του παίκτη για το μέρος του πόρου που παίρνει και του προϋπολογισμού του. Ακολουθώντας την πρόσφατη εργασία των Azar et al. [2017], χρησιμοποιούμε τον όρο *ρευστό κόστος της αναρχίας* (LPoA, για συντομία) για να αναφερθούμε στο κόστος της αναρχίας ως προς το ρευστό όφελος, δηλαδή, τον λόγο του μέγιστου δυνατού ρευστού οφέλους σε οποιαδήποτε κατάσταση του παιχνιδιού προς το ελάχιστο ρευστό όφελος σε κατάσταση ισορροπίας.

A.1.1 Αποτελέσματα και τεχνικές

Στόχος μας είναι να μελετήσουμε *όλους* τους μηχανισμούς κατανομής πόρων και να βρούμε εκείνον με το καλύτερο δυνατό LPoA. Τα αποτελέσματα μας υποδεικνύουν μια ολοκληρωτικά διαφορετική εικόνα σε σχέση με την περίπτωση όπου οι παίκτες δεν έχουν προϋπολογισμούς. Αρχικά, δείχνουμε ένα κάτω φράγμα $2 - 1/n$ για το LPoA κάθε μηχανισμού κατανομής πόρων για n παίκτες (υπό τυπικές υποθέσεις σχετικά με τις συναρτήσεις αποτίμησης των παικτών και τα χαρακτηριστικά των μηχανισμών) το οποίο αποδεικνύει ότι δεν υπάρχουν πλήρως αποδοτικοί μηχανισμοί. Έπειτα, δείχνουμε ότι ο μηχανισμός του Kelly έχει LPoA ακριβώς 2 το οποίο είναι σχεδόν το καλύτερο δυνατό, ενώ ο μηχανισμός των Sanghavi and Hajek (SH) έχει LPoA ίσο με 3. Βελτιωμένα φράγματα για το LPoA είναι δυνατά για την περίπτωση των δυο παικτών. Σχεδιάζουμε τον PYS μηχανισμό E2-PYS για δυο παίκτες ο οποίος έχει LPoA 1.792. Αυτό το φράγμα είναι μάλιστα βέλτιστο για μια ευρεία κλάση μηχανισμών. Επίσης, σχεδιάζουμε τον μηχανισμό E2-SR για δυο παίκτες ο οποίος έχει LPoA το πολύ 1.529 (σχεδόν καλύτερο δυνατό με βάση το κάτω φράγμα 1.5 για δυο παίκτες) και χρησιμοποιεί πληρωμές που ορίζονται από τον λόγο των σημάτων των παικτών. Στον Πίνακα A.1 μπορείτε να βρείτε μια σύνοψη των αποτελεσμάτων μας.

Τα αποτελέσματα μας εκμεταλλεύονται την ιδιαίτερη δομή των παιχνιδιών και ισορροπιών χειρότερης περίπτωσης (ως προς το LPoA). Αποδεικνύουμε ότι για κάθε μηχανισμό κατανομής πόρων, το χειρότερο LPoA παρουσιάζεται σε στιγμιότυπα όπου οι παίκτες έχουν *γραμμικές με μετατόπιση* συναρτήσεις αποτίμησης. Επιπλέον, όλοι οι παίκτες εκτός ενός έχουν πεπερασμένους

Μηχανισμός	LPOA	Σχόλιο
όλοι	$\geq 2 - 1/n$	Δεν υπάρχουν πλήρως αποδοτικοί μηχανισμοί
Kelly	2	Αυστηρό φράγμα. Σχεδόν καλύτερος δυνατός μηχανισμός για n παίκτες
SH	3	Αυστηρό φράγμα. Πλεον δεν είναι καλύτερος του Kelly
E2-PYS	1.792	Αυστηρό φράγμα. Καλύτερος δυνατός PYS μηχανισμός με κοίλη συνάρτηση κατανομής για 2 παίκτες
E2-SR	1.529	Σχεδόν καλύτερος δυνατός μηχανισμός για 2 παίκτες

Table A.1: Μηχανισμοί και σχετικά φράγματα για το ρευστό κόστος της αναρχίας. Τα αποτελέσματα αυτά δημοσιεύτηκαν στην εργασία [Caragiannis and Voudouris, 2018].

προϋπολογισμούς και επιλέγουν στρατηγικές οι οποίες συνεπάγονται είτε μηδενικές πληρωμές είτε πληρωμές οι οποίες είναι ίσες με τους προϋπολογισμούς τους, ενώ ο μοναδικός παίκτης με απεριόριστο προϋπολογισμό υποβάλλει το σήμα που μηδενίζει την παράγωγο της ωφέλειας του.

Για το σενάριο όπου οι παίκτες δεν έχουν προϋπολογισμούς, οι Johari and Tsitsiklis [2004] απέδειξαν έναν ανάλογο χαρακτηρισμό χειρότερης περίπτωσης για τον μηχανισμό του Kelly, ο οποίος έπειτα γενικεύτηκε για όλους τους μηχανισμούς κατανομής πόρων. Συγκεκριμένα, η χειρότερη περίπτωση είναι όταν όλοι οι παίκτες έχουν γραμμικές συναρτήσεις αποτίμησης (με μηδενική μετατόπιση) και υποβάλλουν σήματα που μηδενίζουν τις παραγώγους των ωφελειών τους. Συγκριτικά, ο δικός μας χαρακτηρισμός χειρότερης περίπτωσης είναι πιο πλούσιος σε δομή, και η απόδειξη του είναι αρκετά πιο πολύπλοκη.

Ο χαρακτηρισμός περιέχει τόση πληροφορία που τα φράγματα για το LPOA ακολουθούν σχετικά εύκολα. Το πιο ακραίο παράδειγμα είναι η απόδειξη του καλύτερου φράγματος μας για τον μηχανισμό του Kelly, η οποία είναι μόλις μερικές γραμμές. Επίσης, ο χαρακτηριστικός μπορεί να χρησιμοποιηθεί για τη σχεδίαση νέων μηχανισμών. Για παράδειγμα, ο σχεδιασμός και η ανάλυση των μηχανισμών E2-PYS και E2-SR για δυο παίκτες προκύπτουν από απλές διαφορικές εξισώσεις πρώτου βαθμού, τις οποίες δεν θα μπορούσαμε να αναγνωρίσουμε χωρίς τον χαρακτηρισμό μας. Ακόμη, υπό ορισμένες προϋποθέσεις (όπως, για παράδειγμα, κοίλες συναρτήσεις κατανομής και κυρτές συναρτήσεις πληρωμών), μπορούμε να δείξουμε αυτόματα ότι τα φράγματα για το LPOA είναι αυστηρά, χωρίς να παρουσιάσουμε κάποιο ειδικό κάτω φράγμα (αντι-παράδειγμα).

Σχετική βιβλιογραφία

Το ρευστό όφελος έχει μελετηθεί πρόσφατα ως μέτρο απόδοσης για την σχεδίαση φιλαληθών μηχανισμών [Dobzinski and Paes Leme, 2014, Lu and Xiao, 2015, 2017] και για την ανάλυση συνδυαστικών Walrasian ισορροπιών με προϋπολογισμούς [Dughmi et al., 2016]. Επίσης, έχει χρησιμοποιηθεί από τους Azar et al. [2017] για την ανάλυση του ρευστού κόστους της αναρχίας σε ταυτόχρονες δημοπρασίες πρώτης τιμής, και από τον Voudouris [2018] για δημοπρασίες θέσεων διαφήμισης.

Οι Caragiannis and Voudouris [2016] ήταν οι πρώτοι που απέδειξαν ότι το ρευστό κόστος της αναρχίας του μηχανισμού του Kelly είναι σταθερό. Συγκεκριμένα, έδειξαν ένα άνω φράγμα 2.78 και ένα κάτω φράγμα 2. Ουσιαστικά, εδώ αποδεικνύουμε και πάλι το ίδιο κάτω φράγμα με διαφορετικό και πολύ πιο ενδιαφέρον τρόπο. Οι Christodoulou et al. [2016b] βελτίωσαν το άνω φράγμα σε 2.618 και επέκτειναν τα αποτελέσματα σε ένα πιο γενικό μοντέλο με περισσότερους πόρους. Πριν από αυτές τις εργασίες, οι Syrgkanis and Tardos [2013] έδειξαν ότι το κοινωνικό όφελος σε κατάσταση ισορροπίας είναι το πολύ ένας σταθερός παράγοντας του βέλτιστου ρευστού οφέλους.

Σε αντίθεση με τις τεχνικές ανάλυσης που χρησιμοποιούμε εδώ, η ανάλυση του μηχανισμού του Kelly από τους Caragiannis and Voudouris [2016], Christodoulou et al. [2016b] και Syrgkanis and Tardos [2013] αξιοποιεί το πρότυπο smoothness [Roughgarden, 2015, Roughgarden et al., 2017] και βασίζεται στη φραγή της ωφέλειας κάθε παίκτη από την ωφέλεια που θα μπορούσε να έχει αν απέκλινε μονομερώς και υπέβαλλε κάποιο άλλο σήμα. Τα αποτελέσματα αυτά ισχύουν και για πιο γενικές έννοιες ισορροπίας όπως οι συσχετιζόμενες ισορροπίες και οι ισορροπίες κατά Bayes-Nash. Τα LPoA φράγματα μας ισχύουν μόνο για αμιγείς ισορροπίες κατά Nash, αλλά είναι πιο ισχυρά και αυστηρά.

A.2 Διαμόρφωση απόψεων

Εδώ και αιώνες, η *διαμόρφωση απόψεων* έχει αποτελέσει αντικείμενο έρευνας σε επιστήμες όπως η κοινωνιολογία, τα οικονομικά, η φυσική καθώς και η επιδημιολογία. Η διάδοση και υιοθέτηση του Διαδικτύου έχει επιτρέψει την πρόσφατη άνθιση των κοινωνικών δικτύων, τα οποία έχουν λειτουργούν ως μέσα διάδοσης πληροφοριών οι οποίες είναι σε πολλές περιπτώσεις ευεργητικές για τους χρήστες, αλλά συχνά χρησιμοποιούνται και στρατηγικά από συγκεκριμένα μέρη τα οποία θέλουν με αυτόν τον τρόπο να πετύχουν τους προσωπικούς τους στόχους. Αυτές οι ιδιότητες έχουν πρόσφατα προσελκύσει το ενδιαφέρον των ερευνητών της τεχνητής νοημοσύνης

[Auletta et al., 2016, Schwind et al., 2015, Tsang and Larson, 2014] καθώς και της επιστήμης των υπολογιστών γενικότερα [Bindel et al., 2015, Mossel and Tamuz, 2014, Olshevsky and Tsitsiklis, 2009], και έχει οδηγήσει σε αναθεωρήσεις των κλασικών μοντέλων για διαμόρφωση απόψεων, χρησιμοποιώντας έννοιες και εργαλεία της Θεωρίας Παιγνίων.

Οι Friedkin and Johnsen [1990] προσπάθησαν να μοντελοποιήσουν την διάδοση απόψεων μεταξύ ατόμων που αλληλεπιδρούν μεταξύ τους. Σύμφωνα με το μοντέλο τους, κάθε άτομο έχει μια προσωπική πεποίθηση για κάποιο θέμα συζήτησης και εκφράζει δημοσίως κάποια, ενδεχομένως να είναι διαφορετική από την πεποίθηση, άποψη. Οι πεποιθήσεις και οι απόψεις αναπαριστώνται ως πραγματικοί αριθμοί. Συγκεκριμένα, η άποψη ενός ατόμου προκύπτει από τον μέσο όρο της προσωπικής του πεποίθησης και των απόψεων που εκφράζουν τα άτομα στον κοινωνικό του κύκλο (ο οποίος θεωρείται σταθερός).

Πρόσφατα, οι Bindel et al. [2015] έδειξαν ότι αυτή η συμπεριφορά μπορεί να ερμηνευθεί παιγνιο-θεωρητικά ως εξής: ο μέσος όρος μεταξύ της πεποίθησης του ατόμου και των απόψεων που εκφράζονται στον κοινωνικό του κύκλο είναι απλώς μια *στρατηγική* η οποία ελαχιστοποιεί ένα συγκεκριμένο κόστος. Αυτό το κόστος ορίζεται ως μια τετραγωνική συνάρτηση η οποία είναι ίση με την ολική απόσταση της άποψης του ατόμου από την πεποίθηση του αλλά και από τις απόψεις που εκφράζονται στον κοινωνικό του κύκλο, στο τετράγωνο. Κατά μια έννοια, η στρατηγική αυτή συμπεριφορά έχει ως αποτέλεσμα απόψεις που ακολουθούν την πλειοψηφία του κοινωνικού κύκλου.

Οι Bindel et al. [2015] θεώρησαν στατικά στιγμιότυπα του υποκείμενου κοινωνικού δικτύου και υπέθεσαν ότι η άποψη κάθε ατόμου επηρεάζεται από όλο τον κοινωνικό του κύκλο. Ωστόσο, στη πραγματικότητα, καθώς οι απόψεις εξελίσσονται, οι άνθρωποι τείνουν να παραβλέπουν τις απόψεις που είναι μακριά από την προσωπική τους πεποίθηση, ακόμη και αν εκφράζονται από τους καλύτερους τους φίλους. Ακολουθώντας αυτή τη λογική, οι Bhawalkar et al. [2013] υπέθεσαν ότι η άποψη ενός ατόμου εξαρτάται μόνο από ένα *μικρό* μέρος των ανθρώπων στον κοινωνικό του κύκλο, τους οποίους καλούμε *γείτονες*. Επομένως, στο μοντέλο των Bhawalkar et al. [2013], η διαμόρφωση απόψεων συν-εξελίσσεται με τη γειτονιά κάθε ατόμου, η οποία αποτελείται από εκείνους τους ανθρώπους που έχουν απόψεις αρκετά κοντά στην πεποίθηση του ατόμου. Τώρα, η άποψη που εκφράζει ένα άτομο είναι μια στρατηγική η οποία ελαχιστοποιεί και πάλι την ίδια τετραγωνική συνάρτηση κόστους, αλλά λαμβάνοντας υπόψη τη γειτονιά και όχι ολόκληρο τον κοινωνικό κύκλο.

Τόσο οι Bindel et al. [2015] όσο και οι Bhawalkar et al. [2013] απέδειξαν μικρά σταθερά

φράγματα (9/8 και 14, αντίστοιχα) για το κόστος της αναρχίας των στρατηγικών παιχνιδιών που προκύπτουν από τις υποθέσεις των μοντέλων τους. Ουσιαστικά, αυτά τα φράγματα υποδεικνύουν ότι ένα μη φυσιολογικά μεγάλο μέρος του πληθυσμού των ανθρώπων εκφράζουν απόψεις οι οποίες είναι πάρα πολύ κοντά στις πεποιθήσεις τους. Δυστυχώς, αυτό είναι δύσκολο να το πιστέψει κανείς δεδομένων των τόσο διαφορετικών και, σε πολλές περιπτώσεις, ακραίων απόψεων που εκφράζονται στα κοινωνικά δίκτυα, για παράδειγμα, σε συζητήσεις σχετικά με πολιτική και θρησκεία.

Ακολουθούμε το μοντέλο συν-εξέλιξης, και υποθέτουμε ότι η γειτονιά κάθε ατόμου ορίζεται από τα k άτομα που εκφράζουν απόψεις οι οποίες είναι οι πιο κοντινές στην πεποίθηση του. Ωστόσο, αποκλίνουμε από τον ορισμό της τετραγωνικής συνάρτησης κόστος και, αντιθέτως, θεωρούμε ότι τα άτομα προσπαθούν να *συμβιβαστούν* περισσότερο με τους γείτονες τους. Έτσι, υποθέτουμε ότι κάθε άτομο προσπαθεί να ελαχιστοποιήσει τη *μέγιστη* απόσταση της άποψης τους από την προσωπική του πεποίθηση και κάθε άποψης που εκφράζεται στη γειτονιά του. Φυσικά, αυτές οι υποθέσεις οδηγούν στον ορισμό στρατηγικών παιχνιδιών, τα οποία καλούμε k -COF παιχνίδια, όπου κάθε άτομο λειτουργεί ως παίκτης.

A.2.1 Αποτελέσματα και τεχνικές

Αρχικά αποδεικνύουμε διάφορες ιδιότητες σχετικά με τη γεωμετρική δομή των απόψεων και των πεποιθήσεων σε αμιγείς ισορροπίες κατά Nash (καταστάσεις τους παιχνιδιού όπου κάθε παίκτης ελαχιστοποιεί το προσωπικό τους κόστος, υποθέτοντας ότι οι υπόλοιποι παίκτες δεν θα αλλάξουν τις απόψεις τους).

Χρησιμοποιώντας τις δομικές αυτές ιδιότητες, δείχνουμε ότι υπάρχουν k -COF παιχνίδια τα οποία δεν επιδέχονται αμιγείς ισορροπίες κατά Nash. Επίσης, αποδεικνύουμε ότι ακόμη και για παιχνίδια όπου υπάρχουν ισορροπίες, η ποιότητα τους μπορεί να μην είναι βέλτιστη ως προς το κοινωνικό κόστος (το συνολικό κόστος των παικτών), δείχνοντας ότι το κόστος της ευστάθειας μεγαλώνει γραμμικά ως προς το k . Για την ειδική περίπτωση των 1-COF παιχνιδιών, δείχνουμε ότι κάθε τέτοιο παιχνίδι μπορεί να αναπαρασταθεί ως ένα διευθυνόμενο άκυκλο γράφημα, στο οποίο κάθε αμιγής ισορροπία κατά Nash αντιστοιχεί σε ένα μονοπάτι μεταξύ δυο συγκεκριμένων κόμβων. Επομένως, το πρόβλημα ύπαρξης αμιγών ισορροπιών κατά Nash καθώς και τα προβλήματα υπολογισμού της καλύτερης και της χειρότερης ισορροπίας (ως προς το κοινωνικό κόστος) είναι ισοδύναμα με απλούς υπολογισμούς ελάχιστων και μέγιστων μονοπατιών σε άκυκλα γραφήματα οι οποίοι μπορούν να γίνουν σε πολυωνυμικό χρόνο.

k	PoA	MPoA	PoS	Υπαρξη/Πολυπλοκότητα
1	3	≥ 6	$\geq 17/15$	Δεν υπάρχουν πάντα ισορροπίες
2	$\in [18/5, 12]$	$\geq 24/5$	$\geq 8/7$	Καλύτερη/χειρότερη ισορροπία στο P
≥ 3	$\in [k+1, 4(k+1)]$	$\geq k+2$	$\geq (k+1)/3$	Ανοικτό: πολυπλοκότητα για $k \geq 2$

Table A.2: Τα αποτελέσματα μας για k -COF παιχνίδια. Ο πίνακας περιέχει φράγματα για το κόστος της αναρχίας ως προς αμιγείς (PoA) και μικτές ισορροπίες (MPoA), για το κόστος της ευστάθειας (PoS) καθώς και για την ύπαρξη και πολυπλοκότητα ως προς αμιγείς ισορροπίες. Προφανώς, κάθε άνω φράγμα για το κόστος της αναρχίας είναι επίσης άνω φράγμα και για το κόστος της ευστάθειας. Τα αποτελέσματα αυτά έχουν δημοσιευτεί στην εργασία [Caragiannis et al., 2017a].

Για γενικά k -COF παιχνίδια, ποσοτικοποιούμε την ποιότητα των αμιγών ισορροπιών κατά Nash (ως προς το κοινωνικό κόστος) στη χειρότερη περίπτωση, φράσσοντας το κόστος της αναρχίας. Συγκεκριμένα, παρουσιάζουμε άνω και κάτω φράγματα για το κόστος της αναρχίας των k -COF παιχνιδιών (ως προς αμιγείς και μικτές ισορροπίες) που εξαρτώνται γραμμικά από το k . Στις αποδείξεις των άνω φραγμάτων μας εκμεταλλευόμαστε, με μη τετριμμένο τρόπο, τεχνικές γραμμικού προγραμματισμού για να φράξουμε από κάτω το βέλτιστο κοινωνικό κόστος. Για την θεμελιώδη περίπτωση των 1-COF παιχνιδιών, αποδεικνύουμε ένα αυστηρό φράγμα 3 χρησιμοποιώντας ένα συγκεκριμένο σχήμα τιμολόγησης στην ανάλυση μας. Τα αποτελέσματα μας συνοψίζονται στον Πίνακα A.2.

A.2.2 Σχετική βιβλιογραφία

Ο DeGroot [1974] πρότεινε ένα μοντέλο για την μοντελοποίηση της διαμόρφωσης απόψεων. Σύμφωνα με αυτό το μοντέλο, κάθε άτομο ενημερώνει την άποψη του μέσω μιας διαδικασίας βεβαρυμένου μέσου όρου. Έπειτα, οι Friedkin and Johnsen [1990] αναθεώρησαν το μοντέλο του DeGroot υποθέτοντας ότι κάθε άτομο έχει μια ιδιωτική πεποίθηση και εκφράζει μια (πιθανώς διαφορετική) δημόσια άποψη, η οποία εξαρτάται τόσο από την πεποίθηση του όσο και από τις απόψεις που εκφράζουν οι άνθρωποι του κοινωνικού του κύκλου. Πιο πρόσφατα, οι Bindel et al. [2015] μελέτησαν αυτό το μοντέλο και απέδειξαν ότι η επαναληπτική αυτή διαδικασία μέσου όρου οδηγεί σε απόψεις οι οποίες μπορούν να ερμηνευτούν ως η μοναδική ισορροπία ενός παιχνιδιού διαμόρφωσης απόψεων.

Οι Bhawalkar et al. [2013] απέκλιναν από την υπόθεση ότι οι απόψεις επηρεάζονται από όλο τον κοινωνικό κύκλο, και θεώρησαν τα παιχνίδια συν-εξέλιξης που συζητήσαμε παραπάνω,

όπου οι απόψεις αλλάζουν καθώς οι γειτονιές των ατόμων αλλάζουν. Αυτό το μοντέλο είναι εννοιολογικά όμοιο με προηγούμενα μοντέλα που μελετηθεί από τους Hegselmann and Krause [2002] και Holme and Newman [2006], σύμφωνα με τα οποία κάθε άτομο έχει μια παράμετρο εμπιστοσύνης ε και λαμβάνει υπόψη μόνο τους ανθρώπους που εκφράζουν απόψεις οι οποίες είναι ε κοντά στην πεποίθηση του.

Μια σειρά από πρόσφατες εργασίες θεωρούν διακριτά μοντέλα διαμόρφωσης απόψεων με δυαδικές απόψεις. Οι Chierichetti et al. [2018] μελέτησαν διακριτά παιχνίδια προτίμησης, όπου οι πεπεισμένοι και οι απόψεις είναι δυαδικές και απάντησαν ερωτήσεις σχετικές με το κόστος της ευστάθειας. Για τέτοιου είδους παιχνίδια, οι Auletta et al. [2015, 2017a] χαρακτήρισαν τα κοινωνικά δίκτυα όπου η πεποίθηση της μειονότητας μπορεί να αναδειχτεί ως η άποψη της πλειονότητας, ενώ οι Auletta et al. [2017b] εξέτασαν αν αυτά τα αποτελέσματα μπορούν να επεκταθούν σε άλλα μοντέλα. Οι Auletta et al. [2016] γενίκευσαν την κλάση των διακριτών παιχνιδιών προτίμησης στην περίπτωση όπου οι προτιμήσεις των παικτών είναι αρκετά πιο πολύπλοκες. Τέλος, οι Bilò et al. [2016] επέκτειναν την κλάση των παιχνιδιών συν-εξέλιξης στο διακριτό μοντέλο.

Στην περίπτωση όπου υπάρχουν παραπάνω από ένα θέματα συζήτησης, οι Jia et al. [2015] πρότειναν και ανέλυσαν το λεγόμενο DeGroot-Friedkin μοντέλο για την εξέλιξη ενός δικτύου επιρροής μεταξύ των ατόμων οι οποίοι εκφράζουν απόψεις για μια σειρά από θέματα, ενώ οι Xu et al. [2015] παρουσίασαν μια παραλλαγή σύμφωνα με την οποία κάθε άτομο μπορεί να επαν-υπολογίσει το βάρος που θέτει στην άποψη του, μετά από συζήτηση κάθε θέματος με τους ανθρώπους του κοινωνικού του κύκλου.

Μια άλλη γραμμή έρευνας έχει επικεντρωθεί στη μελέτη της ταχύτητας με την οποία ένα σύστημα συγκλίνει σε μια σταθερή κατάσταση. Σε αυτά τα πλαίσια, οι Etesami and Basar [2015] μελέτησαν την δυναμική του Hegselmann-Krause μοντέλου συν-εξέλιξης και εστίασαν στον χρόνο τερματισμού για διάφορες περιπτώσεις. Όμοια, οι Ferraioli et al. [2016] μελέτησαν την σύγκλιση αποκεντρωμένων δυναμικών σε πεπερασμένα παιχνίδια διαμόρφωσης απόψεων, με παίκτες που έχουν μόνο έναν πεπερασμένο αριθμό από διαθέσιμες απόψεις. Οι Ferraioli and Ventre [2017] μελέτησαν τον ρόλο που παίζει η κοινωνική πίεση και έδειξαν αυστηρά φράγματα για τον χρόνο τερματισμού στην πολύ σημαντική ειδική περίπτωση όπου το κοινωνικό δίκτυο είναι κλίκια.

Οι Das et al. [2014], ύστερα από μια σειρά διαδικτυακών ερωτηματολογίων, κατέληξαν ότι τα πιο γνωστά θεωρητικά μοντέλα δεν εξηγούν εντελώς τα πειραματικά τους αποτελέσματα.

Έτσι, παρουσίασαν ένα νέο αναλυτικό μοντέλο για διαμόρφωση απόψεων, και έδειξαν τόσο θεωρητικά όσο και πειραματικά προκαταρκτικά αποτελέσματα για την σύγκλιση και την δομή των απόψεων όταν οι χρήστες ενημερώνουν τις απόψεις του επαναληπτικά σύμφωνα με το μοντέλο τους.

Ο Chazelle [2012] μελέτησε και ανέλυσε συστήματα επιρροής, όπου κάθε άτομο παρατηρεί τις τοποθεσίες των γειτόνων του και κινείται αναλόγως. Οι Kempe et al. [2016] παρουσίασαν ένα μοντέλο για ανάλυση πολιτιστικής δυναμικής, και εστίασαν στην αλληλεπίδραση μεταξύ επιλογής και επιρροής. Μεταξύ άλλων αποτελεσμάτων, παρουσίασαν έναν χαρακτηρισμό των σταθερών καταστάσεων και έδειξαν ότι το σύστημα συγκλίνει πάντα χωρίς να επηρεάζεται από την αρχική κατάσταση. Οι Gomez-Rodriguez et al. [2012] μελέτησαν φαινόμενα διάχυσης και μεταδοτικής διάδοσης πληροφοριών σε δίκτυα. Τέλος, οι Kempe et al. [2015] μελέτησαν ένα πρόβλημα βελτιστοποίησης για την μεγιστοποίηση της επιρροής σε κοινωνικά δίκτυα, όπου κάθε άτομο μπορεί να αποφασίσει να υιοθετήσει μια ιδέα ανάλογα με το πόσοι από τους γείτονες του την έχουν ήδη υιοθετήσει. Ο στόχος είναι η επιλογή ενός αρχικού συνόλου ατόμων που υιοθετούν την ιδέα με σκοπό να μεγιστοποιηθεί ο τελικός αριθμός των ατόμων που τους ακολουθούν.

A.3 Μεταφορά ιδιοκτησίας

Τα πιο γνωστά προβλήματα της *υπολογιστικής κοινωνικής επιλογής* [Brandt et al., 2016] αφορούν την αποδοτική συγχώνευση ατομικών προτιμήσεων επί των εναλλακτικών επιλογών (οι οποίες εκφράζονται συνήθως υπό τη μορφή *κατατάξεων*) σε μια συλλογική απόφαση [Caragiannis et al., 2017,b, Procaccia et al., 2012, Skowron et al., 2016]. Οι περισσότεροι μηχανισμοί που υλοποιούν τη διαδικασία συγχώνευσης είναι *θεσιακοί* και, άρα, δεν εκμεταλλεύονται τις συγκεκριμένες αριθμητικές αποτιμήσεις των ατόμων. Επιπλέον, λόγω διάφορων γνωστών αποτελεσμάτων αδυναμίας [Gibbard, 1973, Satterthwaite, 1975], τέτοιου είδους μηχανισμοί *δεν είναι φιλαλήθεις*. Δηλαδή, κάποιοι από τους συμμετέχοντες ενδέχεται να έχουν ισχυρά κίνητρα να πουν ψέματα σχετικά με τις προτιμήσεις τους ώστε να χειραγωγήσουν τον μηχανισμό και να τον οδηγήσουν στο να επιλέξει μια εναλλακτική επιλογή την οποία προτιμούν περισσότερο (από αυτή που θα επέλεγε αν έλεγαν την αλήθεια).

Αντίθετα, η κλάση των φιλαληθών *αριθμητικών* μηχανισμών έχει αποδειχτεί ότι είναι πολύ πιο πλούσια [Barbera et al., 1998, Feige and Tennenholtz, 2010, Freixas, 1984]. Μπορούμε να βελτιώσουμε την συνολική ευημερία της κοινωνίας, αξιοποιώντας τις επιπλέον πληροφορίες

που περιέχουν οι αριθμητικές ατομικές αποτιμήσεις [Cheng, 2016, Filos-Ratsikas and Miltersen, 2014, Guo and Conitzer, 2010]. Οι φιλαλήθεις μηχανισμοί που χρησιμοποιούν *χρήματα* έχουν αναλυθεί εκτενώς στη σχετική βιβλιογραφία, και οι μηχανισμοί που μεγιστοποιούν το κοινωνικό όφελος έχουν ήδη σχεδιαστεί για διάφορες κλάσεις προβλημάτων [Nisan et al., 2007]. Ένα διάσημο τέτοιο παράδειγμα είναι η οικογένεια των VCG μηχανισμών [Clarke, 1971, Groves, 1973, Vickrey, 1961].

Ωστόσο, υπάρχει ένα αρκετά μεγάλο σύνολο προβλημάτων *υβριδικής κοινωνικής επιλογής*, όπου η μεταφορά χρημάτων δεν είναι δυνατή για κάποιο μέρος του πληθυσμού. Επομένως, η σχεδίαση *φιλαληθών, αριθμητικών* μηχανισμών είναι μια αρκετά πολύπλοκη διαδικασία και πρέπει να συνδυαστούν στοιχεία σχεδίασης μηχανισμών με χρήματα καθώς και κοινωνικής επιλογής. Μελετάμε τέτοιου είδους σενάρια στα πλαίσια της μεταφοράς ιδιοκτησίας, όπου έχουμε ένα σύνολο πιθανών αγοραστών με χρηματικές αποτιμήσεις για μια εταιρεία, ενώ υπάρχει και ένα σύνολο ειδικών (π.χ., το διοικητικό συμβούλιο της εταιρείας) οι οποίοι έχουν απόψεις για το σε ποιον θα πρέπει τελικά να πουληθεί η εταιρεία. Ο στόχος είναι να πάρουμε την απόφαση που μεγιστοποιεί το κοινωνικό όφελος, το οποίο λαμβάνει υπόψη τις αριθμητικές αποτιμήσεις τόσο των αγοραστών όσο και των ειδικών.

Αυτό το σενάριο μοντελοποιεί διάφορες ενδιαφέρουσες πραγματικές περιπτώσεις. Μια πρώτη εφαρμογή αφορά την αποκρατικοποίηση κρατικών περιουσιακών στοιχείων, όπου ένα σύνολο πιθανών αγοραστών ενδιαφέρεται να τα αποκτήσει, ενώ διάφοροι οργανισμοί πολιτών στοχεύουν στο να εγγυηθούν ότι η επιλογή θα είναι υπέρ των πολιτών. Παρόμοιες περιπτώσεις μεταφοράς ιδιοκτησίας προκύπτουν κατά την ανάθεση αθλητικών διοργανώσεων όπως το Παγκόσμιο Κύπελλο ποδοσφαίρου και οι Ολυμπιακοί αγώνες, όπου λαμβάνονται υπόψη τόσο οι προσφορές των ενδιαφερόμενων διοργανωτών (χώρες) όσο και οι συστάσεις των αντίστοιχων διοικητικών αρχών (π.χ. FIFA.com, 2018).

A.3.1 Αποτελέσματα και τεχνικές

Στην Διατριβή αυτή εστιάζουμε στο θεμελιώδες σενάριο όπου έχουμε δυο πιθανούς αγοραστές A και B , και έναν ειδικό με αριθμητικές αποτιμήσεις για τις τρεις εναλλακτικές επιλογές του να πουλήσουμε στον αγοραστή A , να πουλήσουμε στον B , ή να μην πουλήσουμε καθόλου (στην οποία περίπτωση η μεταφορά ιδιοκτησίας δεν πραγματοποιείται). Ένα μηχανισμός δέχεται ως είσοδο τις προσφορές των αγοραστών καθώς και τις αποτιμήσεις του ειδικού, και αποφασίζει μία εναλλακτική ως το αποτέλεσμα. Γενικά, οι μηχανισμοί είναι *πιθανοτικοί* και το αποτέλεσμα επιλέγεται σύμφωνα με μια πιθανοτική κατανομή (ή *λοταρία*) επί των τριών εναλλακτικών.

Κλάση μηχανισμών	Λόγος προσέγγισης	Σχόλιο
Θεσιακοί	1.5	μηχανισμοί EOM, BOM καλύτερο δυνατό φράγμα
Ανεξάρτητοι των προσφορών	1.377	μηχανισμός BIM καλύτερος δυνατός
Ανεξάρτητοι του ειδικού	1.343	μηχανισμός EIM καλύτερος δυνατός
Πρότυπο	1.25	πιθανοτικός μηχανισμός R καλύτερος always-sell
	1.618	ντετερμινιστικός μηχανισμός D καλύτερος ντετερμινιστικός
Όλοι	1.14	κάτω φράγμα

Table A.3: Περίληψη των αποτελεσμάτων μας. Δείτε την εργασία [Caragiannis et al., 2018].

Θεωρούμε *φιλαλήθεις* μηχανισμούς οι οποίοι αξιοποιούν διαφορετικά επίπεδα γνώσης. Πιο συγκεκριμένα, μελετάμε *θεσιακούς*, *ανεξάρτητους των προσφορών*, *ανεξάρτητους του ειδικού*, καθώς και γενικούς *φιλαλήθεις* μηχανισμούς οι οποίοι λαμβάνουν υπόψη *τις αποτιμήσεις τόσο των αγοραστών όσο και του ειδικού*. Υποθέτουμε *ντετερμινιστικούς* και *πιθανοτικούς* μηχανισμούς, ενώ ως μέτρο απόδοσης χρησιμοποιούμε την έννοια του λόγου προσέγγισης ως προς το κοινωνικό όφελος, το οποίο λαμβάνει υπόψη *τις αποτιμήσεις τόσο των αγοραστών όσο και του ειδικού*. Για κάθε κλάση μηχανισμών, αποδεικνύουμε κάτω φράγματα για τον λόγο προσέγγισης όλων των σχετικών μηχανισμών και εντοπίζουμε τον καλύτερο μεταξύ αυτών. Τα αποτελέσματα μας συνοψίζονται στον Πίνακα A.3.

Οι τεχνικές μας στηρίζονται στο γεγονός ότι κάθε μηχανισμός μπορεί να θεωρηθεί ως μια λοταρία η οποία αναθέτει πιθανότητες στις εναλλακτικές επιλογές (A , B , ή κανένας από τους δυο) που ορίζονται από την είσοδο που παρέχουν ο ειδικός και οι αγοραστές. Η γενική μας στρατηγική για να εντοπίσουμε τον καλύτερο δυνατό φιλαλήθη μηχανισμό για κάθε κλάση που θεωρούμε είναι η εξής. Για κάθε πιθανό λόγο προσέγγισης, πρώτα χαρακτηρίζουμε τον χώρο των φιλαληθών μηχανισμών υπολογίζοντας κατάλληλα φράγματα για τις πιθανότητες που μπορεί να ορίσει η αντίστοιχη λοταρία. Έπειτα, υπολογίζουμε τον ελάχιστο δυνατό λόγο προσέγγισης που ορίζει έναν εφικτό (όχι κενό) χώρο φιλαληθών μηχανισμών.

A.3.2 Σχετική βιβλιογραφία

Σε μια αγορά, οι συγχωνεύσεις και οι αγορές εταιρειών παίζουν κεντρικό ρόλο στον ανταγωνισμό μεταξύ δημοσίων και ιδιωτικών οργανισμών. Υπάρχουν άπλετες ενδείξεις ότι η μεταφορά της ιδιοκτησίας μια εταιρείας επηρεάζει σημαντικά την οικονομία τόσο των εργαζομένων όσο και των καταναλωτών της [Auerbach, 2008, Hitt et al., 2001]. Σύμφωνα με δεδομένα της Ευ-

ρωπαικής Ένωσης, έχουν πραγματοποιηθεί περισσότερες από 6500 συγχωνεύσεις από το 1990, ενώ έχουν οριστεί αυστηροί κανόνες που διέπουν το πώς τέτοιου είδους συγχωνεύσεις πρέπει να γίνονται.

Το μοντέλο μας είναι ένα στιγμιότυπο *προσεγγιστικής σχεδίασης μηχανισμών* με χρήματα [Nisan and Ronen, 2001] και χωρίς χρήματα [Procaccia and Tennenholtz, 2013], η οποία έχει προταθεί για προβλήματα βελτιστοποίησης υπό αυστηρούς περιορισμούς φιλαλήθειας. Ο Myerson [1981] απέδειξε ικανούς και αναγκαίους περιορισμούς για φιλαλήθεις μηχανισμούς με χρήματα (ντετερμινιστικούς ή πιθανοτικούς). Αυτός ο χαρακτηρισμός μας επιτρέπει να μην ασχοληθούμε με συναρτήσεις πληρωμών για την πλευρά των αγοραστών, καθώς αυτές είναι καλώς ορισμένες δεδομένων των πιθανοτήτων επιλογής. Επίσης, μας έδωσε όλα τα απαραίτητα εργαλεία ώστε να επιχειρηματολογήσουμε σχετικά με την δομή των φιλαληθών μηχανισμών χωρίς χρήματα από την πλευρά του ειδικού.

Για προβλήματα όπου μπορούν να χρησιμοποιηθούν χρήματα, ο γνωστός μηχανισμός VCG [Clarke, 1971, Groves, 1973, Vickrey, 1961] έχει μερικές πολύ σημαντικές ιδιότητες: είναι ντετερμινιστικός, φιλαλήθης, και μεγιστοποιεί το κοινωνικό όφελος. Ωστόσο, όπως είδαμε και προηγουμένως, στο υβριδικό μοντέλο σχεδίασης μηχανισμών που θεωρούμε εδώ, θα πρέπει να λάβουμε υπόψη και τις αποτιμήσεις του ειδικού. Επομένως, ο μηχανισμός VCG δεν είναι πλέον φιλαλήθης ούτε βέλτιστος. Από την πλευρά του ειδικού, οι φιλαλήθεις μηχανισμοί μπορούν να θεωρηθούν ως φιλαλήθεις κανόνες ψηφοφορίας. Δυστυχώς, πολύ γνωστά αποτελέσματα αδυναμίας [Gibbard, 1973, Satterthwaite, 1975] περιορίζουν την κλάση των ντετερμινιστικών φιλαληθών κανόνων ψηφοφορίας μόνο σε *δικτατορικούς* μηχανισμούς.

Αντίθετα, η κλάση των πιθανοτικών κανόνων ψηφοφορίας είναι πλουσιότερη και περιέχει λογικούς φιλαλήθεις κανόνες οι οποίοι δεν είναι δικτατορικοί. Ο Gibbard [1977] χαρακτήρισε την κλάση των θεσιακών πιθανοτικών κανόνων, αλλά δεν είναι ακόμη γνωστό αν υπάρχει ένας πιο γενικός χαρακτηρισμός για όλους τους αριθμητικούς κανόνες. Πάρα πολλές εργασίες στην κλασική οικονομική βιβλιογραφία καθώς και στην αντίστοιχη της επιστήμης των υπολογιστών έχουν αφιερωθεί στην σχεδίαση φιλαληθών αριθμητικών κανόνων καθώς και στην απόδειξη δομικών ιδιοτήτων για περιορισμένες κλάσεις. Ο Gibbard [1978] απέδειξε έναν χαρακτηρισμό ο οποίος ισχύει μόνο για χώρους διακριτών στρατηγικών, ενώ αργότερα ο Hylland [1980] απέδειξε ότι η κλάση των φιλαληθών μηχανισμών που είναι αποδοτικοί κατά Pareto ανάγεται στην κλάση των πιθανοτικών δικτατορικών κανόνων.

Ο Freixas [1984] χρησιμοποίησε την μέθοδο των διαφορών [Laffont and Maskin, 1980] για

να σχεδιάσει μια κλάση φιλαληθών μηχανισμών η οποία ουσιαστικά χαρακτηρίζει την κλάση των (δυο φορές) παραγωγίσιμων (επί υποδιαστημάτων του χώρου αποτιμήσεων) μηχανισμών. Ο καλύτερος φιλαλήθης ανεξάρτητος-των-προσφορών μηχανισμός που προτείνουμε μπορεί να θεωρηθεί ως ένας μηχανισμός αυτής της κλάσης. Έπειτα, οι Barbera et al. [1998] έδειξαν ότι υπάρχουν πολλοί φιλαλήθεις μηχανισμοί οι οποίοι δεν ανήκουν στις κλάσεις που μελέτησε ο Freixas [1984]. Ακόμη, οι Feige and Tennenholtz [2010] σχεδίασαν μια κλάση αριθμητικών φιλαληθών μηχανισμών για ένα ψηφοφόρο μόνο, όπου οι πιθανότητες επιλογής ορίζονται από συγκεκριμένα πολυώνυμα.

Το πρόβλημα της μεγιστοποίησης του κοινωνικού οφέλους χωρίς χρήματα έχει μελετηθεί σε πολλές εργασίες, τόσο για γενικά μοντελά κοινωνικής επιλογής [Bhaskar et al., 2018, Filos-Ratsikas and Miltersen, 2014] όσο και για ειδικές περιπτώσεις όπως προβλήματα ταιριασμάτων και κατανομής πόρων [Cheng, 2016, Filos-Ratsikas et al., 2014, Guo and Conitzer, 2010]. Όμοια με εμάς, οι Filos-Ratsikas and Miltersen [2014] χρησιμοποιούν φιλαλήθεις μηχανισμούς για ένα ψηφοφόρο ώστε να πετύχουν βελτιωμένες εγγυήσεις απόδοσης. Ωστόσο, η παρουσία των αγοραστών διακρίνει σημαντικά το μοντέλο μας από το δικό τους (όπως και από τα μοντέλα άλλων σχετικών εργασιών).

Μια άλλη σχετική έννοια είναι αυτή της *παραμόρφωσης* των (μη φιλαληθών) μηχανισμών οι οποίοι λειτουργούν υπό περιορισμένη(θεσιακή) πληροφόρηση [Anshelevich et al., 2015, Boutilier et al., 2015, Caragiannis et al., 2017b, Caragiannis and Procaccia, 2011, Caragiannis et al., 2016]. Αν και η έλλειψη πληροφόρησης έχει αποτελέσει έναν περιοριστικό παράγοντα για κάποια από τα αποτελέσματα μας, εστιάζουμε κυρίως σε αριθμητικούς μηχανισμούς για τους οποίους η φιλαλήθεια είναι ο βασικός περιορισμός.

A.4 Ασυμμετρία πληροφορίας για μεγιστοποίηση εσόδων

Η εκμετάλλευση ασυμμετριών στην πληροφόρηση είναι ένα αντικείμενο έρευνας που ξεκίνησε από την πρωτοποριακή εργασία του Akerlof [1970] ο οποίος μελέτησε τέτοια φαινόμενα στην λεγόμενη *αγορά των λεμονιών*. Υποθέστε μια αγορά αυτοκινήτων η οποία περιέχει αυτοκίνητα υψηλής ποιότητας (τα οποία είναι γνωστά ως ροδάκινα) καθώς και χαμηλής ποιότητας που εμφανίζουν προβλήματα μετά την τελική τους αγορά (τα οποία είναι γνωστά ως λεμόνια). Σε μια τέτοια αγορά, ο πωλητής έχει αρκετά πιο ακριβή πληροφόρηση σχετικά με την ποιότητα των αυτοκινήτων, ενώ οι αγοραστές δεν μπορούν να ξεχωρίσουν τα ροδάκινα από τα λεμόνια. Επομένως, προκύπτει ένα ενδιαφέρον πρόβλημα στρατηγικής απόφασης από την πλευρά του

πωλητή με σκοπό να βρει τον καλύτερο δυνατό τρόπο ώστε να εκμεταλλευτεί την κατάσταση και να θέσει υψηλότερες τιμές για τα αυτοκίνητα. Αναμενόμενα, η εργασία του Akerlof σχετικά με την ασυμμετρική πληροφόρηση έδωσε το έναυσμα για περαιτέρω έρευνα στην περιοχή της οικονομικής θεωρίας [Crawford and Sobel, 1982, Levin and Milgrom, 2010, Milgrom, 2010, Milgrom and Weber, 1982] και, πρόσφατα, στην επιστήμη των υπολογιστών [Dughmi, 2014, Emek et al., 2012, Ghosh et al., 2007, Guo and Deligkas, 2013, Miltersen and Sheffet, 2012].

Ακολουθούμε την εργασία των Alon et al. [2013] και εστιάζουμε σε πιθανοτικές πωλήσεις τύπου αποδέξου ή απέρριψε, όπου υπάρχουν m αντικείμενα και n πιθανοί αγοραστές. Κάθε αγοραστής έχει μια αποτίμηση για κάθε αντικείμενο, και υποθέτουμε ότι γενικά δεν γνωρίζει την ύπαρξη των άλλων αγοραστών και των αποτιμήσεων τους. Σύμφωνα με μια πιθανοτική κατανομή, η φύση επιλέγει τυχαία ένα μοναδικό αντικείμενο προς πώληση. Έπειτα, ο πωλητής προσεγγίζει τον αγοραστή με τη μεγαλύτερη αποτίμηση και του προσφέρει το αντικείμενο σε τιμή ίση με την αποτίμηση του για το αντικείμενο. Ένα συγκεκριμένο στιγμιότυπο αυτού του σεναρίου θα μπορούσε να είναι το εξής: τα αντικείμενα αντιστοιχούν σε λέξεις κλειδιά και οι πιθανοί αγοραστές αντιστοιχούν σε διαφημιστές. Κάθε διαφημιστής έχει μια αποτίμηση για κάθε λέξη κλειδί η οποία αναπαριστά το μέγιστο ποσό χρημάτων που είναι διαθέσιμος να πληρώσει προκειμένου να δεσμεύσει τον χώρο διαφήμισης που δίνεται όταν γίνεται αναζήτηση της συγκεκριμένης λέξης κλειδί. Η φύση αντιστοιχεί στους χρήστες που κάνουν αναζητήσεις, ενώ ο πωλητής αντιστοιχεί στην μηχανή αναζήτησης, η οποία δεσμεύει τον χώρο διαφήμισης ανάλογα με την λέξη κλειδί και στοχεύει στη μεγιστοποίηση των εσόδων της.

Μπορεί να ο πωλητής να εκμεταλλευτεί το γεγονός ότι διαθέτει πιο ακριβή πληροφόρηση σχετικά με τα αντικείμενα προς πώληση σε σχέση με τους πιθανούς αγοραστές; Συγκεκριμένα, η ασυμμετρία πληροφορίας προκύπτει από το γεγονός ότι ο πωλητής γνωρίζει το αντικείμενο που επιλέγει η φύση τυχαία, ενώ οι αγοραστές δεν το γνωρίζουν. Μια πιθανή προσέγγιση είναι ο πωλητής να μπορεί να ορίσει ένα σχήμα σημάτων ανάλογα με τον αγοραστή. Δηλαδή, για κάθε αγοραστή, ο πωλητής μπορεί να χωρίσει τα αντικείμενα σε ανα δυο ανεξάρτητα σύνολα και να δηλώσει αυτήν την διαμέριση στον αγοραστή. Για παράδειγμα, η μηχανή αναζήτησης θα μπορούσε να ομαδοποιήσει μαζί συγκεκριμένες λέξεις κλειδιά οι οποίες συσχετίζονται η μια με την άλλη. Μετά την τυχαία επιλογή της φύσης, ο πωλητής μπορεί να αποκαλύψει στον κάθε αγοραστή την ομάδα που περιέχει το αντικείμενο, έτσι ώστε να του επιτρέψει να επαν-υπολογίσει την αποτίμηση του για την ομάδα αυτή.

Οι Alon et al. [2013] παρουσίασαν το πρόβλημα *μη-συμμετρικής διαμέρισης πίνακα* με σκοπό

να μοντελοποιήσουν το πρόβλημα της μεγιστοποίησης εσόδων σε πωλήσεις τύπου αποδέξου ή απέρριψε. Τα στιγμιότυπα του προβλήματος αποτελούνται από έναν $n \times m$ πίνακα A με μη-αρνητικές εγγραφές και μια πιθανοτική κατανομή p επί των στηλών του. Διακρίνουμε δύο περιπτώσεις για την πιθανοτική κατανομή επί των στηλών του πίνακα εισόδου, ανάλογα με το αν είναι ομοιόμορφη ή μη-ομοιόμορφη. Ένα σχήμα διαμέρισης $B = (B_1, \dots, B_n)$ για τον πίνακα A αποτελείται από μια διαμέριση B_i του συνόλου $[m]$ για κάθε γραμμή i του A . Συγκεκριμένα, το B_i είναι μια συλλογή η οποία αποτελείται από k_i , ανά δύο ανεξάρτητα, υποσυνόλα $B_{ik} \subseteq [m]$ (με $1 \leq k \leq k_i$) έτσι ώστε $\bigcup_{k=1}^{k_i} B_{ik} = [m]$. Μπορούμε να φανταστούμε κάθε διαμέριση B_i ως έναν τελεστή ομαλότητας ο οποίος δρα πάνω στις εγγραφές της γραμμής i και αλλάζει τις τιμές τους στην αναμενόμενη τιμή του υποσυνόλου διαμέρισης στο οποίο ανήκουν. Τυπικά, η ομαλή τιμή μιας εγγραφής (i, j) τέτοια ώστε $j \in B_{ik}$ ορίζεται ως

$$A_{ij}^B = \frac{\sum_{\ell \in B_{ik}} p_\ell \cdot A_{i\ell}}{\sum_{\ell \in B_{ik}} p_\ell}.$$

Παρατηρήστε ότι όλες οι εγγραφές (i, j) με $j \in B_{ik}$ έχουν την ίδια ομαλή τιμή. Δεδομένου ενός σχήματος διαμέρισης B το οποίο συνεπάγεται έναν ομαλό πίνακα A^B , η τιμή διαμέρισης ορίζεται ως η αναμενόμενη μέγιστη εγγραφή στις στήλες του A^B , δηλαδή,

$$v^B(A, p) = \sum_{j \in [m]} p_j \cdot \max_i A_{ij}^B.$$

Σκοπός του προβλήματος είναι ο υπολογισμός ενός σχήματος διαμέρισης B έτσι ώστε η τιμή διαμέρισης $v^B(A, p)$ να μεγιστοποιείται.

Η σχέση του προβλήματος μη-συμμετρικής διαμέρισης με το πρόβλημα μεγιστοποίησης εσόδων σε πωλήσεις τύπου αποδέξου ή απέρριψε είναι η εξής: οι στήλες του πίνακα εισόδου αντιστοιχούν σε αντικείμενα, οι γραμμές αντιστοιχούν στους πιθανούς αγοραστές, και η τιμή της εγγραφής (i, j) αντιστοιχεί στην αποτίμηση του αγοραστή i για το αντικείμενο j . Μετά τη διαμέριση των αντικειμένων σε υποσύνολα για ένα συγκεκριμένο αγοραστή, η ομαλή τιμή ενός υποσυνόλου αντιστοιχεί στην αναμενόμενη αποτίμηση του αγοραστή για κάθε αντικείμενο του συγκεκριμένου υποσυνόλου. Τέλος, η τιμή διαμέρισης αντιστοιχεί στα έσοδα που αναμένεται να έχει ο πωλητής.

A.4.1 Αποτελέσματα και τεχνικές

Μεταξύ άλλων αποτελεσμάτων, οι Alon et al. [2013] απέδειξαν ότι το πρόβλημα είναι APX-hard ακόμη και για την περίπτωση όπου ο πίνακας περιέχει δυαδικές τιμές, και σχεδίασαν

έναν 0.563- και έναν $1/13$ -προσεγγιστικό αλγόριθμο για τις περιπτώσεις όπου η πιθανοτική κατανομή επί των στηλών του πίνακα είναι ομοιόμορφη και μη-ομοιόμορφη, αντίστοιχα. Βελτιώνουμε σημαντικά και τα δυο αυτά αποτελέσματα. Παρουσιάζουμε έναν $9/10$ -προσεγγιστικό αλγόριθμο για ομοιόμορφες πιθανοτικές κατανομές, καθώς και έναν $(1 - 1/e)$ -προσεγγιστικό αλγόριθμο για μη-ομοιόμορφες πιθανοτικές κατανομές. Το σύνολο αυτών των αποτελεσμάτων έχουν δημοσιευτεί στην εργασία [Abed et al., 2018].

Για την ομοιόμορφη περίπτωση, ο αλγόριθμος μας πρώτα καλύπτει τις στήλες που έχουν τουλάχιστον έναν άσσο, και έπειτα ταιριάζει με άπληστο τρόπο τις στήλες που περιέχουν μόνο μηδενικά με άσσους σε συγκεκριμένες γραμμές. Η ανάλυση αυτού του αλγορίθμου είναι εξαιρετικά ενδιαφέρουσα καθώς, παρά το γεγονός ότι ο αλγόριθμος είναι αμιγώς συνδυαστικός, εκμεταλλεύεται τεχνικές γραμμικού προγραμματισμού και δυσκότητας.

Για την γενική μη-ομοιόμορφη περίπτωση, εκμεταλλευόμαστε τη σχέση του προβλήματος μη-συμμετρικής διαμέρισης δυαδικού πίνακα με το πρόβλημα μεγιστοποίησης κοινωνικού οφέλους με κοίλες συναρτήσεις αποτίμησης, και χρησιμοποιούμε γνωστούς αλγορίθμους από τη σχετική βιβλιογραφία. Πρώτα συζητάμε την πιθανή εφαρμογή ενός απλού άπληστου $1/2$ -προσεγγιστικού αλγορίθμου, ο οποίος έχει μελετηθεί από τους Lehmann et al. [2006]. Έπειτα, εφαρμόζουμε τον ομαλό άπληστο $(1 - 1/e)$ -προσεγγιστικό αλγόριθμο του Vondrák [2008].

Στο πρόβλημα μεγιστοποίησης κοινωνικού οφέλους με κοίλες συναρτήσεις, ο αλγόριθμος του Vondrák είναι ο καλύτερος δυνατός στο μοντέλο ερωτημάτων αποτίμησης, όπου έχουμε πρόσβαση σε ένα oracle που απαντάει γρήγορα σε ερωτήσεις σχετικά με τις αποτιμήσεις των αγοραστών για συγκεκριμένα σύνολα αντικειμένων [Khot et al., 2008]. Οι Feige and Vondrák [2010] έδειξαν ότι υπάρχουν βελτιωμένοι $(1 - 1/e + \epsilon)$ -προσεγγιστικοί αλγόριθμοι για το πιο ισχυρό μοντέλο ερωτημάτων απαίτησης, όπου έχουμε πρόσβαση σε ένα oracle που απαντάει γρήγορα σε ερωτήσεις σχετικά με το ποιο σύνολο αντικειμένων πετυχαίνει συγκεκριμένη αποτίμηση για έναν αγοραστή. Συζητάμε την δυνατότητα εφαρμογής τέτοιων αλγορίθμων για το πρόβλημα μη-συμμετρικής διαμέρισης πίνακα και παρατηρούμε ότι η απάντηση ερωτημάτων απαίτησης είναι NP-hard γενικά.

A.4.2 Σχετική βιβλιογραφία

Πέρα από την δυαδική περίπτωση, οι Alon et al. [2013] μελέτησαν επίσης και την πιο γενική περίπτωση του προβλήματος μη-συμμετρικής διαμέρισης πίνακα όπου ο πίνακας αποτελείται από μη-αρνητικούς πραγματικούς αριθμούς, και παρουσίασαν έναν $1/2$ - και έναν $\Omega(1/\log m)$ -

προσεγγιστικό αλγόριθμο για ομοιόμορφες και μη-ομοιόμορφες πιθανοτικές κατανομές. Η κοινή ιδέα των αλγορίθμων είναι ο εντοπισμός ενός συνόλου εγγραφών με μεγάλες τιμές το οποίο μπορεί να ομαδοποιηθεί μαζί με άλλες εγγραφές που περιέχουν αρκετά μικρές τιμές έτσι ώστε να αυξηθεί η συνολική τους προσφορά στην τιμή διαμέρισης.

Η πιθανή μοντελοποίηση του προβλήματος μεγιστοποίησης εσόδων σε πωλήσεις τύπου αποδέξου ή απέρριψε από το πρόβλημα μη-συμμετρικής διαμέρισης πίνακα, αποτελεί μέρος μιας γραμμής έρευνας η οποία μελετά την επίπτωση της μη συμμετρικής πληροφόρησης στην ποιότητα των αγορών. Όπως συζητήσαμε και παραπάνω, ο Akerlof [1970] ήταν ο πρώτος που παρουσίασε μια τυπική ανάλυση για την αγορά των λεμονιών, όπου ο πωλητής έχει πιο ακριβή πληροφόρηση για την ποιότητα των προϊόντων σε σχέση με τους πιθανούς αγοραστές.

Η ιδέα του να διαμερίσουμε το σύνολο των αντικειμένων σε διαφορετικές ομάδες για κάθε αγοραστή και ύστερα να ενημερώσουμε σε κάθε αγοραστή την ομάδα που περιέχει το τυχαίο αντικείμενο, προέρχεται από την μέθοδο στρατηγικής μετάδοσης πληροφοριών των Crawford and Sobel [1982], όπου ο πωλητής έχει πληροφορίες σχετικά με τις αποτιμήσεις των αγοραστών, και στρατηγικά στοχεύει στο να εκμεταλλευτεί αυτό το πλεονέκτημα ώστε να μεγιστοποιήσει τα έσοδα του. Για να δουλέψει μια τέτοια προσέγγιση όμως, πρέπει να υποθέσουμε τον επιπλέον περιορισμό ότι οι αγοραστές δεν γνωρίζουν ο ένας τον άλλο και δεν γνωρίζουν λεπτομέρειες σχετικά με τον υποκείμενο μηχανισμό, καθώς διαφορετικά θα μπορούσαν να μάθουν ποιο είναι πραγματικά το αντικείμενο προς πώληση. Αν αυτό δεν είναι δυνατό, τότε η αρχή σύνδεσης των Milgrom and Weber [1982] υποδεικνύει ότι η καλύτερη στρατηγική του πωλητή είναι να αποκαλύψει όλες τις πληροφορίες που έχει στους αγοραστές.

Οι Levin and Milgrom [2010] καθώς και ο Milgrom [2010] έδειξαν ότι ο καλύτερος τρόπος για να εκμεταλλευτεί κανείς πιθανές ασυμμετρίες στην πληροφόρηση είναι μέσω κατάλληλης διαμέρισης των αντικειμένων σε ομάδες. Στη βιβλιογραφία έχουν προταθεί διάφορα μοντέλα για αποκάλυψη κατάλληλων πληροφοριών στους αγοραστές. Για παράδειγμα, οι Ghosh et al. [2007] θεώρησαν πλήρη πληροφόρηση και πρότειναν ένα σχήμα ομαδοποίησης σύμφωνα με το οποίο, τα αντικείμενα διαμερίζονται σε υποσύνολα και, έπειτα, για κάθε ένα από αυτά τα υποσύνολα, εκτελείται μια ξεχωριστή δημοπρασία δεύτερης τιμής. Αυτό έχει ως αποτέλεσμα οι πιθανοί αγοραστές να μην μπορούν να κάνουν προσφορές μόνο για τα αντικείμενα που πραγματικά θέλουν να αποκτήσουν, αλλά να πρέπει να ανταγωνιστούν και για αντικείμενα τα οποία δεν επιθυμούν να αγοράσουν. Δηλαδή, η ζήτηση κάθε αντικειμένου αυξάνεται το οποίο συνεπάγεται μεγαλύτερο κέρδος για τον πωλητή. Οι Emek et al. [2012] παρουσίασαν

αποτελέσματα πολυπλοκότητας σε ένα παρόμοιο πιθανοτικό μοντέλο, ενώ οι Miltersen and Sheffet [2012] μελέτησαν κλασματικά σχήματα διαμέρισης για αποστολή σημάτων.

Τέλος, αξίζει να αναφέρουμε ότι η χρήση γραμμικού προγραμματισμού για την ανάλυση αμιγώς συνδυαστικών αλγορίθμων είναι πλέον μια πολύ καλά ορισμένη τεχνική και έχει ήδη αξιοποιηθεί σε πολλά διαφορετικά προβλήματα σχετικά με τοποθέτηση εγκαταστάσεων [Jain et al., 2003], κάλυψη συνόλων [Athanassopoulos et al., 2009a,b, Caragiannis et al., 2013], ταιριάσματα άμεσης απόκρισης [Mahdian and Yan, 2011], μέγιστες διευθυνόμενες τομές [Feige and Jozeph, 2015], και δρομολόγηση μήκους κύματος [Caragiannis, 2009].