



UNIVERSITY OF

LIVERPOOL

THESIS FOR DOCTOR OF PHILOSOPHY DEGREE

Algorithms for Game-Theoretic Environments

ALKMINI SGOURITSA

SUPERVISOR: DR. GEORGE CHRISTODOULOU

To Dimitris and Ifigenia

Abstract

Game Theory constitutes an appropriate way for approaching the Internet and modelling situations where participants interact with each other, such as networking, online auctions and search engine's page ranking. Mechanism Design deals with the design of private-information games and attempts implementing desired social choices in a strategic setting. This thesis studies how the efficiency of a system degrades due to the selfish behaviour of its agents, expressed in terms of the *Price of Anarchy* (PoA). Our objective is to design mechanisms with improved PoA, or to determine the exact value of the PoA for existing mechanisms for two well-known problems, Auctions and Network Cost-Sharing Design.

We study three different settings of auctions, combinatorial auction, multi-unit auction and bandwidth allocation. The combinatorial auction constitutes a fundamental resource allocation problem that involves the interaction of selfish agents in competition for indivisible goods. Although it is well-known that by using the VCG mechanism the selfishness of the agents does not affect the efficiency of the system, i.e. the social welfare is maximised, this mechanism cannot generally be applied in computationally tractable time. In practice, several simple auctions (lacking some nice properties of the VCG) are used, such as the generalised second price auction on AdWords, the simultaneous ascending price auction for spectrum allocation, and the independent second-price auction on eBay. The latter auction is of particular interest in this thesis. Precisely, we give *tight* bounds on the PoA when the goods are sold in independent and simultaneous *first-price* auctions, where the highest bidder gets the item and pays her own bid. Then, we generalise our results to a class of auctions that we call *bid-dependent* auctions, where the goods are also sold in independent and simultaneous auctions and further the payment of each bidder is a function of her bid, even if she doesn't get the item. Overall, we show that the first-price auction is optimal among all bid-dependent auctions.

The multi-unit auction is a special case of combinatorial auction where all

items are identical. There are many variations: the discriminatory auction, the uniform price auction and the Vickrey multi-unit auction. In all those auctions, the goods are allocated to the highest marginal bids, and their difference lies on the pricing scheme. Our focus is on the *discriminatory auction*, which can be seen as the variant of the first-price auction adjusted to multi-unit auctions.

The bandwidth allocation is equivalent to auctioning *divisible* resources. Allocating network resources, like bandwidth, among agents is a canonical problem in the network optimisation literature. A traditional model for this problem was proposed by Kelly [94], where each agent receives a fraction of the resource proportional to her bid and pays her own bid. We complement the PoA bounds known in the literature and give tight bounds for a more general case. We further show that this mechanism is optimal among a wider class of mechanisms.

We further study *design* issues for network games: given a rooted undirected graph with nonnegative edge costs, a set of players with terminal vertices need to establish connectivity with the root. Each player selects a path and the global objective is to minimise the cost of the used edges. The cost of an edge may represent infrastructure cost for establishing connectivity or renting expense, and needs to be covered by the users. There are several ways to split the edge cost among its users and this is dictated by a *cost-sharing protocol*. Naturally, it is in the players best interest to choose paths that charge them with small cost.

The seminal work of Chen et al. [36] was the first to address design questions for this game. They thoroughly studied the PoA for the following informational assumptions. i) The designer has full knowledge of the instance, that is, she knows both the network topology and the players' terminals. ii) The designer has *no knowledge of the underlying graph*. Arguably, there are situations where the former assumption is too optimistic while the latter is too pessimistic. We propose a model that lies in the middle-ground; the designer has prior knowledge of the underlying metric, but *knows nothing* about the positions of the terminals. Her goal is to process the graph and choose a *universal* cost-sharing protocol that has low PoA against *all possible* requested subsets. The main question is to what extent prior knowledge of the underlying metric can help in the design.

We first demonstrate that there exist graph metrics where knowledge of the underlying metric can dramatically improve the performance of good network cost-sharing design. However, in our main technical result, we show that there exist graph metrics for which knowing the underlying metric does not help and *any universal* protocol matches the bound of [36] which ignores the graph metric.

We further study the stochastic and Bayesian games where the players choose their terminals according to a probability distribution. We showed that in the stochastic setting there exists a priority protocol that achieves constant PoA, whereas the PoA under the the Bayesian setting can be very high for *any* cost-sharing protocol satisfying some natural properties.

Contents

Abstract	i
Contents	iv
List of Figures	vii
List of Tables	vii
Acknowledgments	ix
1 Introduction	1
1.1 Auctions	2
1.1.1 Single-item Auction	2
1.1.2 Combinatorial Auctions	4
1.1.3 Bandwidth Allocation	4
1.1.4 Information Models	5
1.2 Network Games	6
1.2.1 Congestion Games	7
1.2.2 Network Design	10
1.3 Thesis Contribution	12
1.3.1 Auctions	12
1.3.2 Network Games	14
I Simple Auctions	17
2 Overview	19
2.1 Results	21
2.2 Literature Review	24
3 Preliminaries	29
3.1 Mechanisms and Valuations	29

3.2	Item-bidding Auctions	30
3.3	Bidding Strategies and Equilibria	32
3.4	Price of Anarchy	32
3.5	Types of Valuations	33
3.5.1	Valuations for Combinatorial Auctions	33
3.5.2	Valuations for Multi-Unit Auctions	34
3.5.3	Valuations for Bandwidth Allocation	34
4	First Price Auction	37
4.1	Fractionally Subadditive Valuations	37
4.1.1	Upper Bound	38
4.1.2	Tight Lower Bound	44
4.2	Subadditive Valuations	48
5	Bid-Dependent Auctions	51
5.1	Fractionally Subadditive Valuations	51
5.1.1	Upper Bounds	52
5.1.2	Lower Bound	56
5.2	Subadditive Valuations	59
5.2.1	Upper Bound	60
5.2.2	Lower Bound	62
6	Discriminatory auctions	67
6.1	Preliminaries	67
6.2	Submodular Valuations	69
6.2.1	Upper Bound	69
6.2.2	Lower Bound	71
6.3	Subadditive Valuations	73
7	Proportional Allocation Mechanism	75
7.1	Concave Valuations	75
7.2	Subadditive Valuations	77
7.2.1	Upper Bound	78
7.2.2	Lower Bound for Simple Mechanisms	80
7.2.3	Lower Bound for Scale-free Mechanisms	82

II	Cost-Sharing Networks	85
8	Overview	87
8.1	Results	91
8.2	Techniques	93
8.3	Related Work	94
9	Models and Preliminaries	97
9.1	Information Models	97
9.2	Universal Cost-Sharing Protocols	98
9.3	Generalised Weighted Shapley Protocol	100
9.4	Pure Nash Equilibrium	100
9.5	Price of Anarchy	101
9.6	Graph Notation	101
10	Design Against the Adversary	103
10.1	Outerplanar Graphs	103
10.2	Lower Bound of Ordered Protocols	106
10.2.1	Description of \mathbf{G}_m	111
10.2.2	Ramsey-type Theorem	112
10.2.3	Colouring Based on the Labels	114
10.2.4	Lower Bound Construction	115
10.3	Lower Bound for All Universal Protocols	117
10.3.1	Cost-Sharing Preliminaries	118
10.3.2	Lower Bound	119
11	Stochastic Design	127
11.1	Randomised Protocol	128
11.2	Deterministic Protocol	130
12	Bayesian Design	137
12.1	Fully Budget-Balance	138
12.2	Budget-Balance in Equilibrium	139
12.2.1	BBiE Protocols	140
12.2.2	BBiE Posted Prices	142
	Conclusion	145
	Bibliography	147

List of Figures

Figure 1.1:	Pigou’s example.	9
Figure 1.2:	“Prim” Protocol.	11
Figure 4.1:	Illustration for the first-price auction upper bound.	39
Figure 4.2:	Illustration for the first-price auction lower bound.	46
Figure 8.1:	Example of how knowing the network helps in the design.	89
Figure 8.2:	Ordered and Shapley protocols may not be the best.	91
Figure 10.1:	Example of improved design in outerplanar graphs.	104
Figure 10.2:	Example of zig-zag pattern.	107
Figure 10.3:	Extract a zig-zag pattern by using [5].	108
Figure 10.4:	Replacing hypercube’s edges by paths.	114
Figure 10.5:	Lower bound construction against the adversary.	121
Figure 12.1:	Lower bound under the Bayesian setting.	139

List of Tables

2.1	Price of Anarchy for the simultaneous first-price auction.	25
10.1	Example of unfolding the bit mapping.	116

Acknowledgments

First and foremost, I would like to express my sincere gratitude to my supervisor Dr. George Christodoulou for his insightful mentoring and guidance. I was totally inspired by his enthusiasm for pure research and gained a valuable experience throughout seeking and answering interesting research questions together. I would also like to thank him for his help and support in overcoming obstacles I have been facing through my research.

I would further like to thank the University of Liverpool for giving me the opportunity to work in such a stimulated environment. I am grateful to the staff members and my fellow doctoral students for their feedback, cooperation and friendship. Special thanks to Annamária Kovács, Bo Tang and Stefano Leonardi for our collaborations, which appear to be a great experience and pleasure.

I am also grateful to the members of my committee, Dr. Martin Gairing and Prof. Elias Koutsoupias, for their insightful comments and encouragement and also for the interesting questions which incentivised me to widen my research from various perspectives. A very special gratitude goes out to Prof. Elias Koutsoupias who inspired me during my undergraduate studies to follow an academic career in Algorithmic Game Theory.

I am also thankful to my family and friends for their belief in my abilities and interest in my work. Last but most importantly, I would like to especially thank my partner, Dimitris, for his encouragement and patience along the way. He provides me with moral, emotional and spiritual support in my life and I feel very fortunate to have him by my side!

CHAPTER 1

Introduction

Game Theory constitutes an appropriate way for approaching and modelling situations where participants interact with each other. Mechanism Design is a subarea of both Game Theory and Economic Theory, which attempts implementing desired social choices in a strategic setting, by assuming that the different members of the society act rationally in a game theoretic sense. The emergence of the Internet motivates the development of both Algorithmic Game Theory and Mechanism Design which are applicable in many topics, such as networking, peering, online auctions and exchanges, online advertising, and search engine's page ranking. The relation between Game Theory, Mechanism Design, Economic Theory, Theoretical Computer Science and the Internet is described in more details in [118].

The focus of this thesis is to examine the efficiency/inefficiency of equilibria expressed as the *Price of Anarchy* (PoA). The PoA was introduced in 1999 by Koutsoupias and Papadimitriou [98, 99] and is the concept in Game Theory that measures how the efficiency of a system degrades due to the selfish behaviour of its agents. The game designer's objective is to maximise the *social welfare* which is the total sum of agents' payoffs. On the other hand, the agents may act strategically under the incentive of maximising their own payoff. When agents' strategies are such that no agent can benefit by unilaterally deviating from their strategy, the strategy profile is called *Nash equilibrium*. Price of Anarchy (PoA) [98] and Price of Stability (PoS) [6] are tools provided by the Algorithmic Game Theory in order to measure the quality of the equilibrium solutions; PoA (or PoS) is the ratio between the worst-case (or the best-case) social welfare in a Nash equilibrium and the maximum social welfare. Our objective is to design mecha-

nisms with improved PoA, or to determine the exact value of the PoA for existing mechanisms for the following well-known problems: Auctions of indivisible and divisible resources and Network Cost-Sharing Design.

The thesis is divided into two parts, one for each problem. The background and the related work of each problem is discussed in the related part. For the rest of the introduction we review in a more general way auctions and network games.

1.1 Auctions

Auctions can be modelled as games of incomplete information between many selfish agents/bidders in competition for one or more items or resources. The preference of each bidder is expressed via a valuation function over the different allocations of the items/resources to the bidders¹. The game designer should decide a mechanism that asks the agents to provide information (bids) about their valuations, based on which it computes the allocation and bidders' payments. Each agent's utility (payoff) is defined as the difference between her valuation for the allocation and her payment. The agents' valuations are usually unknown to the auctioneer and the other participants. Designing *truthful mechanisms*, where there is no way that the bidders can increase their utilities by lying, is one way to deal with this lack of information. The appeal of truthful mechanisms is twofold: they relieve the decision-making burden of the participants and their outcome is predictable. However, the main challenge is to design truthful mechanisms that allocate the items in an *efficient* way, i.e., so that they maximise the *social welfare* which equals the sum of bidders' valuations.

1.1.1 Single-item Auction

The most simple auction is that of a single item, where the main question to address is to decide who should get the item. Suppose that the objective of the auctioneer is to assign the item to the player who values it the most, maximising that way the social welfare. She should decide it, though, without knowing participants' valuations. A way to access this information is to ask the participants directly. The participants, on the other side, would declare a very high valuation

¹The valuation may be considered as the agent's willingness to pay.

in their attempt to obtain the item and then the auctioneer could not decide who truly values the item the most.

To overcome this problem, prices have been introduced in order to force the participants to declare their true valuation, or at least a value close to it. Obviously, the existence of prices serves another, very important purpose: they provide a revenue for the seller. For the scope of this thesis, we focus on the former reason.

Auctions that sell a single indivisible item are well-understood in terms of both truthfulness and efficiency (i.e., allocating the item to the agent that values it the most). The *English* (or ascending price) auction is the oldest form of auction and satisfies the above two properties. This auction is an iterative auction, meaning that it is conducted in steps. By starting from a zero or a small price for the item, the participants repeatedly respond to the current price by announcing a bid higher than that price; this implies that they are interested in purchasing at that bid. The price is then updated and the auction ends when no participant is willing to increase the price. The higher bidder receives the item and pays her (last) bid. A variant of the English auction is conducted by an auctioneer who increases the price at each step by a small amount, as long as there are at least two participants interested in purchasing at the current price. The auction ends when there is only one participant left², who buys the item by paying the current price.

The celebrated *Vickrey* (or *sealed-bid second-price*) auction [138] is a counterpart of the latter variant of the English auction, conducted in a single step. Each bidder simultaneously submits a sealed bid and then the bidder with the highest bid receives the item and pays the *second* highest bid. The Vickrey auction is both truthful and efficient.

To check the truthfulness of the Vickrey auction, we need to examine two cases: i) a participant receives the item by bidding her true valuation and ii) a participant doesn't receive the item by bidding truthfully. In the first case her bid is the highest one and she should pay at most her valuation which leads to nonnegative utility. If she bids any value that remains the highest bid, she still receives the item at the same price leading to the same utility and if she bids lower than the second highest bid she loses the item resulting in zero utility. Therefore, in the first case she cannot increase her utility by lying. In the second case her utility is zero and she can only receive the item by bidding higher than

²The increments of the price should be sufficiently small such that there is always one participant left.

the highest bid. Note though that the highest bid is higher or equal to her valuation and she should pay that value as it would become the second highest bid. This results in nonpositive utility and hence, in this case she cannot increase her utility by lying, meaning that overall the auction is truthful. Given that the participants act in a truthful way, the Vickrey auction allocates by definition the item to the participant with the highest bid/valuation satisfying efficiency.

1.1.2 Combinatorial Auctions

Combinatorial auctions are a natural generalisation of the single-item auction that involve the interaction of n selfish agents in competition for m *indivisible* items. The preferences of each player for different bundles of items are expressed via a valuation *set* function. It is well-known that in this setting truthfulness and efficiency are preserved by the *Vickrey-Clarke-Groves (VCG) mechanism* [50, 79, 138]. Unfortunately, computing the allocation and the payments of the VCG mechanism may take exponential time in m and n [114, 115] (unless P=NP). This fact led in dropping the truthfulness requirement and designing simple (in terms of computation) mechanisms, under the objective of maximising the social welfare in the equilibria.

In practice, several simple *non-truthful* mechanisms are used. The most notable examples are the generalised second price (GSP) auctions used by AdWords [61, 137], the simultaneous ascending price auctions for wireless spectrum allocation [109] and the independent second price auctions on eBay. Furthermore, in these auctions the expressive power of the buyers is heavily restricted by the bidding language, so that they are not able to represent their complex preferences precisely. In light of the above, Christodoulou, Kovács and Schapira [40] proposed the study of simple, non-truthful auctions, called *simultaneous (item-bidding) auctions*, using the PoA as a measure of inefficiency of such auctions. In such an auction, the auctioneer sells each item by running *simultaneously* m independent single-item auctions.

1.1.3 Bandwidth Allocation

Auctioning *divisible* resources is also of particular interest. Allocating network resources, like bandwidth, among agents is a canonical problem in the network optimisation literature. In such situations, it is sometimes difficult to talk about

truthful mechanism, since even expressing the valuation function for every fractional allocation becomes extremely complex.

A traditional model for this problem was proposed by Kelly [94], where allocating these infinitely divisible resources is treated as a market with prices. More precisely, agents in the system submit bids on resources to express their willingness to pay. After soliciting the bids, the system manager prices each resource and then agents buy portions of resources by paying a proportional amount of the prices. The users act as price takers, trying to maximise their utility, i.e. the difference between their valuations and payments, and they do not anticipate the effect of their actions on the prices. Kelly [94] showed that, under certain assumptions, the aggregate utility of the users is maximised when the players receive portions of the resources that are proportional to their bids. In the case of a single resource, each user receives a fraction of the resource equal to the ratio of their bid over the sum of all bids; additionally, they should pay an amount equal to their own bid. This is known as the *proportional allocation mechanism* or *Kelly's mechanism* in the literature.

Johari and Tsitsiklis [91] relaxed the assumption that the users act as price takers and instead they can anticipate the effects of their actions on the prices of the resources. They observed that this strategic bidding in the proportional allocation mechanism leads to inefficient allocations that do not maximise the social welfare.

The proportional allocation mechanism has also been used in the *trading post* game proposed by Shapley and Shubik [129]. In the trading post game, each good is sold in a separate trading post and each trader makes a monetary bid on each trading post and receives an allocation based on the proportional allocation mechanism. The trading post model differentiates from Kelly's mechanism in the sense that traders also receive payments for the fraction of goods that they sell.

Additionally, the proportional allocation mechanism is widely used in network pricing and has been implemented for allocating computing resources in several distributed systems [49], for time-sharing of resources [131] and in resource allocation in Capacity-constrained Clouds [135].

1.1.4 Information Models

Regarding non-truthful mechanisms, there are two basic *information models* with respect to agents' valuations: the *full information* and the *Bayesian*. In the full

information setting the valuation function of each player is fixed and known by all other players. The Bayesian setting was introduced by Harsanyi [84], and is an elegant way of modelling partial-information settings. In this setting, the valuation function of each player is drawn from some known probability distribution that, in a sense, represents the players' beliefs. Clearly, the full information model is a special Bayesian one, in which each player has some valuation function with probability 1. Accordingly, there are different concepts of Nash equilibria with respect to the information models. The pure and mixed Nash equilibria refer to the full information model, where players choose a single (pure) strategy or a probability distribution over pure strategies (mixed strategy), respectively. In the Bayesian model, the Bayesian Nash equilibrium is defined, in which each player cannot increase their *expected* utility by unilaterally deviating from their strategy, where the expectation is taken over the valuations of the other players.

In Part I of the thesis, we provide bounds for the PoA of several auction with respect to both information models. We first study simultaneous auctions: the *first-price auction* (Chapter 4) and a wider class of auctions that we call *bid-dependent* (Chapter 5). We further study the counterpart of the first-price auction, called *discriminatory auction* (Chapter 6), which is applied in the special combinatorial auction where all items are identical. Finally, regarding divisible items (bandwidth allocation) we study the *simultaneous proportional allocation mechanism* (Chapter 7).

1.2 Network Games

Recall that in auctions with indivisible items the goal was to allocate the items in such a way that no item is assigned to more than one bidder. As for the case of auctioning divisible resources, we allocate portions of the resources to the bidders, under the restriction that those portions sum up to 1 for each resource. In this section, still resources are assigned to players, but in a different notion: the players now choose some resources to use and it is common that the same resource is used by more than one player. Then, the players experience some cost that depends on their choices and the congestion on the resources.

The Internet is full of such applications: communication networks, peer-to-peer networks and job scheduling are only few of those application. The size of the Internet makes it impossible to use a central authority in order to impose

optimal solutions, i.e. solutions that minimise the social cost. Therefore, it is essential for players to interact with each other, without imposing them any additional control (except possibly from some underlying *protocol*), and result in equilibrium solutions that guarantee good approximations of the optimum solution, i.e. designing mechanisms with low PoA guarantees.

We focus on special cases of such games that are called *network games*. In a network game, the resources are edges of a network and players want to establish connectivity between two vertices, a source and a destination (possibly different for each player). Then, the strategy space of the players are all possible paths connecting those two vertices. The inefficiency of equilibria has been observed in several network games such as congestion in parallel links [99], selfish load balancing [132], selfish routing games [9, 38, 127], network design games [6, 36]. Next we explore such games in more details.

1.2.1 Congestion Games

The congestion games were first defined by Rosenthal [123]. In these games players choose a subset of resources and their overall cost (or latency) is additive over the resources that they chose and depends on how many players chose the same resources. A *latency function* for each resource takes as input the number of players using that resource and determines the exact latency that each player experiences for using it.

An interesting connection between congestion games and potential games is known due to Rosenthal [123] and Monderer and Shapley [110]. A game is called *exact potential* [110] if there exists a function (called *exact potential function*) over players' strategies with the following property: the difference of the function's outcome when a player's strategy unilaterally changes equals the difference of that player's utility or cost. Rosenthal [123] proved that every congestion game admits a pure Nash equilibrium, by providing an exact potential function. Then, Monderer and Shapley [110] showed that every finite exact potential game is (isomorphic to) a congestion game, and as a result, the two classes of games coincide.

We next review two special cases of congestion games that lie in the class of network games, namely *atomic* and *nonatomic selfish routing*. Routing games (see also Chapter 18 of [116]) deal with problems of how to route traffic in a network. The ground difference between atomic and nonatomic routing games is that in the former game each player controls a significant amount of traffic

that she should route unsplittably via a path, whereas in the latter game the populations is large enough, such that each player controls only an infinitesimal amount of traffic.

Atomic selfish routing

The atomic selfish routing game is defined by an underlying (directed) graph, whose edges are associated with cost functions, and a finite set of players each of whom wants to route a *unit* amount of traffic from a source vertex to a destination vertex. Each player should choose a single path and the congestion on the edges that she used induces a latency computed via the edge cost functions. The objective here is again to maximise the social welfare, or alternatively, to *minimise the social cost* which is equal to the aggregate latency that players experience.

The players route their traffic via their chosen paths which can be described by using a (multicommodity) flow. A flow is an *equilibrium flow* (or a pure Nash equilibrium) when no player can decrease their latency by unilaterally changing their path. The PoA compares the worst-case social cost of an equilibrium flow with the minimum social cost. If only affine cost functions are allowed, the PoA is exactly $5/2$, a result due to Christodoulou and Koutsoupias [38] and Awerbuch, Azar and Epstein [9]. We refer the reader to [125] for extension of this PoA bound to other equilibria concepts, via the so-called *smoothness technique*. For cost functions that are bounded-degree polynomials the exact PoA is known due to Aland et al. [3].

A more general setting is the *weighted* atomic routing games, where each player may route many units of unsplittable traffic via a single path. Fotakis, Kontogiannis and Spirakis [69] proved that a weighted atomic routing game with affine cost functions always admits a pure Nash equilibrium and the PoA is $\frac{3+\sqrt{5}}{2} \approx 2.618$. However, if we consider even quadratic cost function, the property of Nash equilibrium existence vanishes and there exists an instance [74] with no pure Nash equilibrium. Nevertheless, for bounded-degree polynomial cost functions, Aland et al. [3] provided tight PoA bounds for more general equilibria concepts which also hold for the pure Nash equilibrium whenever it exists. Bhawalkar, Gairing and Roughgarden [18] characterised the PoA for general cost functions (under mild conditions) and extend their results to other equilibria concepts by using the smoothness technique, similar to [125].

Nonatomic selfish routing

The nonatomic selfish routing game and its equilibria were formally defined by Wardrop [140] and Beckmann, McGuire and Winsten [11]. In those games we assume that each unit of traffic is controlled by a large population of players, each of whom routes an infinitesimal amount. Therefore, the unit traffic may be routed from many different paths, in contrast with the atomic games.

Example 1. (Pigou’s example [119]) A notable example of the nonatomic selfish routing is the Pigou’s example illustrated in Figure 1.1. A unit amount of traffic should be routed from s to t via two alternative paths with cost functions $c_1(x) = 1$ and $c_2(x) = x$, respectively, that express the latency/cost that players experience if they use the paths, where x is the congestion on them. It is in players’ best interest to choose the second path as it is cheaper; the latency in the second path cannot exceed 1, which is the latency of the first path. This is the unique Nash equilibrium (or Wardrop equilibrium) of the game. The cost, that each player experiences, is defined by $c_2(1)$ multiplied by the (tiny) amount of traffic they control and the social cost is then the aggregate cost which equals $c_2(1) = 1$.

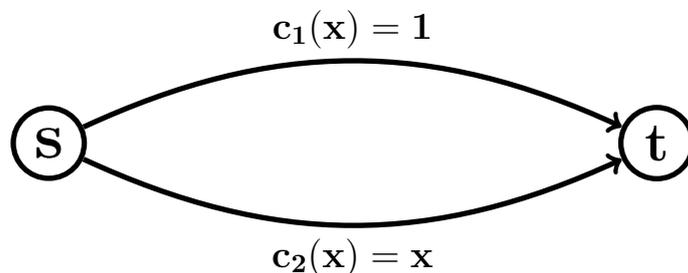


Figure 1.1: Pigou’s example.

However, there is a less costly (optimum) solution which is to split the traffic equally between the two paths. The upper path then is used by a half unit of traffic which experience a latency of $c_1\left(\frac{1}{2}\right) = \frac{1}{2}$ and the lower path is used by the other half unit with latency of $c_2\left(\frac{1}{2}\right) = 1$, resulting in social cost of $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 1 = 3/4$. Therefore, the PoA of the game is $4/3$.

Roughgarden and Tardos [127] were the first to study the PoA of the nonatomic selfish routing games and they showed that Pigou’s example is the worst case instance for affine latency functions, by proving that the PoA is at most $4/3$ for this case. They further showed the first upper bound on the PoA for polyno-

mial latency functions. For special cases, Dumrauf and Gairing [60] provided improved bounds.

1.2.2 Network Design

In a different network game we consider a graph with a nonnegative *constant* cost on each edge and a set of players that need to establish connectivity between their pair of source-destination vertices by selecting a path connecting them. The cost of an edge may represent infrastructure cost for establishing connectivity or renting expense, and needs to be covered by the players that use that edge in the solution. There are several ways to split the edge costs among the users and this is dictated by a cost-sharing protocol. Naturally, it is in the players best interest to choose paths that charge them with small cost, and therefore it is reasonable to assume that the solution will be a Nash equilibrium. The global objective is to minimise the total cost of the used edges in the graph, which is the Minimum Steiner Forest. The PoA can be used here in order to analyse the quality of the equilibrium solutions. This is a fundamental network design game that was originated by Anshelevich et al. [6] and has been extensively studied since. In [6], they studied the Shapley cost-sharing protocol, where the cost of each edge is equally split among its users. They showed that the quality of equilibria can be really poor; even for simple networks the PoA grows linearly with the number of players³.

However, different cost-sharing protocols may result in different quality of equilibria. Chen, Roughgarden and Valiant [36] were the first to address design questions for network cost-sharing games by identifying the cost-sharing protocols with the best PoA (and Price of Stability). For the case of undirected graphs they showed that the *priority/ordered* protocols improve the PoA in comparison to the Shapley protocol. In an ordered protocol, a global order of the players is first defined and then whenever a subset of players is using an edge, its cost is covered by the player who is first in the order. For the case of multicast games, where all players have a common source and therefore the optimum coincides with the minimum Steiner tree, Chen, Roughgarden and Valiant [36] noticed that a simple ordered protocol that simulates Prim’s algorithm [121] results in the best PoA. In fact they showed that there exists a global order of the play-

³The Price of Stability of this game is not well-understood. For undirected graphs it is a big open question to determine its exact value that is between constant and $O(\log k / \log \log k)$ [103], where k is the number of players.

ers resulting in PoA of 2, as the following example illustrates, which is a great improvement upon the linear PoA of Shapley protocol.

Example 2. (The “Prim” Protocol) Suppose we are given an undirected graph with a designated root r , $G(V, E, r)$, nonnegative edge costs and a set of k players with terminals, t_1, \dots, t_k . Consider now the weighted graph, H , with vertices $\{r, t_1, \dots, t_k\}$ where the edge weights are induced by the shortest weighted paths in G between the corresponding vertices. We order the terminals in the same way that Prim’s algorithm processes them in H , starting from r . Figure 1.2 shows such an example. The left graph is the original network and the right graph is the complete weighted graph on which we run Prim’s algorithm. The order that Prim’s algorithm processes the terminals, denoted by bold integers, is used in the order protocol on the left graph. The Nash equilibrium and the outcome of Prim’s algorithm (minimum spanning tree) are shown by the thicker edges. It is not hard to verify that the cost of the path that each player selects is upper bounded by the weight of the edge used in Prim’s algorithm in order to connect her terminal. Hence, the approximation ratio of 2 of the Prim’s algorithm immediately provides an upper bound on the PoA of the ordered protocol.

A lower bound construction of [36] shows that under some natural axioms (discussed in details in Chapter 9) no cost-sharing protocol can result in better PoA.

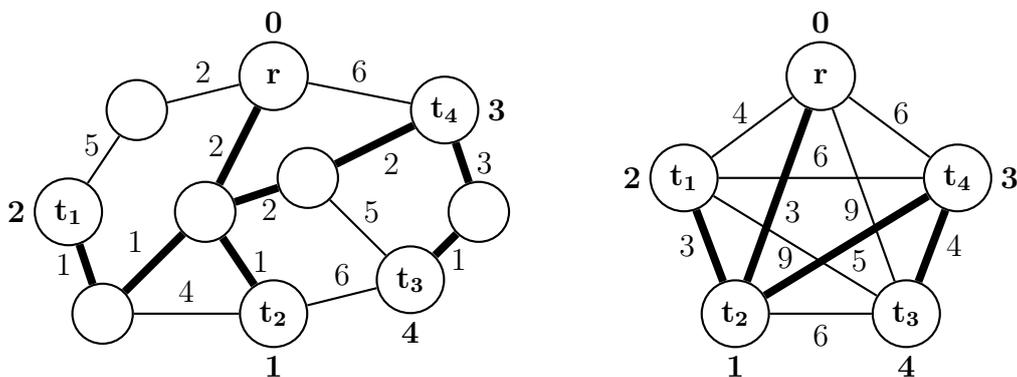


Figure 1.2: An illustration of “Prim” Protocol.

In Part II of the thesis, we study designing issues of the multicast game with respect to the PoA. We consider different informational assumptions, first from the perspective only of the designer, namely games under uncertainty (Chapter 10) and stochastic games (Chapter 11), and then from the perspective of

both the designer and the players, which are the Bayesian games (Chapter 12). For the latter case, we show negative results that can be dramatically improved by relaxing some designing requirement in a reasonable way; we discuss this in details in Chapter 12.

1.3 Thesis Contribution

In this section, we make an overview of the contribution of this thesis with respect to the current literature.

1.3.1 Auctions

The first part of the thesis is devoted to combinatorial auctions, multi-unit auctions and bandwidth allocation. In the combinatorial auctions, selfish agents compete for the allocation of a bundle of items. The allocation of the items in a way that the social welfare is maximised has been proved to be computationally inefficient. Christodoulou, Kovács and Schapira [40] proposed the study of simple auctions, called *simultaneous (item-bidding)* auctions, under the objective of maximising the social welfare. In such auctions, the items are sold separately in simultaneous and independent single-item auctions. Many variants have been studied and the most notable ones are the *first-price*, *second-price* and *all-pay* auctions. In all those auctions, the bidders are asked to submit a bid for each item. Then each item is assigned to the highest bidder. In the first-price auction the winner is charged his own bid, while in the second-price auction the winner pays the second highest bid. In both auctions, the rest of the bidders pay zero. On the contrary, in all-pay auction all bidders (winner and losers) pay their bids. The first-price auction appears to outperform the other auctions, so far, in terms of the PoA.

Regarding the first-price and all-pay auctions we provide PoA bounds which are, most of the cases, *tight*. Precisely, in [41] (co-authored with Christodoulou, Kovács and Tang), our results complement the current knowledge about simultaneous first-price auctions [64, 133]; for two important classes of valuation functions, namely submodular and subadditive, we provide matching lower bounds to the upper bounds of 1.58 by Syrgkanis and Tardos [133] and 2 by Feldman et al. [64], respectively. For subadditive valuations (valuations without complementarities), the tight bounds hold for a more general class of auctions that includes

the all-pay auctions, as well, and it is further extended to multi-unit auctions [41], where the items are identical, and to divisible resources, like bandwidth, as we show in [46] (co-authored by Christodoulou and Tang). Regarding submodular valuations, i.e. functions with decreasing marginals, the lower bounds are tight only for the simultaneous first-price auction. Instead, in [45] the upper bound of the simultaneous all-pay auctions is improved for such valuations by using several structural theorems that characterise the Nash equilibria.

Independently, Roughgarden [124] presented a very elegant methodology to provide PoA lower bounds via a reduction from communication or computational complexity lower bounds for the underlying optimisation problem. One consequence is the indication that simultaneous first-price auction is the most efficient (i.e. has the lowest PoA) simple auctions for some classes of valuations, like subadditive valuations. However, a combination of [124] and our work in [41] indicates that the question of the most efficient auction remains, regarding more specialised, nevertheless significant, classes of valuations such as submodular valuations; either a different approach than the one in [124] is needed in order to prove optimality of the first-price auction for those valuations or there is another auction that improves the PoA.

We further study multi-unit auctions as a special case of combinatorial auctions, where all items are identical. Our focus is on the discriminatory auctions, where the highest bidders receive the items and pay their bid for the obtained items. In [41] (co-authored with Christodoulou, Kovács and Tang), we complement the results of de Keijzer et al. [56] for the case of subadditive valuations, by providing a matching lower bound of 2. For the case of submodular valuations, we show a lower bound of 1.099 improving upon the previous lower bound of 1.0004 [56] which holds only for the Bayesian setting.

Additionally, auctions with divisible resources is in our scope of interest. A traditional model for allocating network resources, like bandwidth was proposed by Kelly [94], where allocating these infinitely divisible resources is treated as a market with prices. Johari and Tsitsiklis [91] relaxed the assumption that the users act as price takers and instead they can anticipate the effects of their actions on the prices of the resources. They considered many divisible resources that are sold simultaneously by using the *proportional allocation mechanism* (or *Kelly's mechanism*); each agent submits a bid for each divisible resource and then receives a fraction proportional to their bids and pay their own bids. Johari and Tsitsiklis [91] showed tight bounds on the PoA under the full-information model

when the valuation functions are concave and Caragiannis and Voudouris [30] showed an upper bound of 2 on the PoA for the single resource case and under the Bayesian model. In [46] (co-authored with Christodoulou and Tang) we provide a lower bound of $\sqrt{m}/2$ for the case of m resources and for the Bayesian setting. We then considered subadditive valuations⁴ and we provided a *tight* bound of 2 on the PoA which generalises and improves the previous bound of 3.73 for lattice-submodular⁵ valuations by Syrgkanis and Tardos [133]. We further showed optimality of the proportional allocation mechanism among any simple mechanism, as defined in the framework of Roughgarden [124].

1.3.2 Network Games

The second part of the thesis examines design questions in network cost-sharing games. Given a rooted undirected graph with nonnegative edge costs, players with terminal vertices need to establish connectivity with the root. Each player selects a path and the global objective is to minimise the cost of the used edges. This cost needs to be covered by the users and there are several ways to split the edge cost among its users which is decided based on a *cost-sharing protocol*.

Different cost-sharing protocols result in different values of the PoA. The seminal work of Chen, Roughgarden and Valiant [36] was the first to address design questions for this game. They gave a characterisation of protocols that satisfy some natural axioms and they thoroughly studied their PoA for the following two classes of protocols that use different informational assumptions from the perspective of the designer.

Non-uniform protocols: the designer has full knowledge of the instance, that is, she knows both the network topology and in addition the players' terminals; they showed that a simple priority (ordered) protocol has a constant PoA.

Uniform protocols: the designer has *no knowledge of the underlying graph*; they showed that the PoA is logarithmic.

There are situations where the former assumption is too optimistic while the latter is too pessimistic. In [43] (co-authored with Christodoulou) we propose a model that lies in the middle-ground; the designer has prior knowledge of the underlying metric, but is *uncertain* about the positions of the terminals. We

⁴The subadditive functions constitute a superclass of the concave functions *only* in the case of a single resource. Otherwise the two classes of functions are incomparable.

⁵The class of lattice-submodular functions is a subclass of subadditive functions and coincides with the concave functions in the case of a single resource.

consider three different models, the *adversarial*, the *stochastic* and the *Bayesian* model.

In the adversarial model, the designer *knows nothing* about the positions of the terminals and needs to process the graph and choose a single, *universal* cost-sharing protocol that has low PoA against *all possible* requested subsets (no distributional assumptions are made for the requested subset of players). The main question we address is: *to what extent can prior knowledge of the underlying metric help in the design?* We first demonstrate that there exist graph metrics, the *outerplanar* graphs, where knowledge of the underlying metric can dramatically improve the performance of good network cost-sharing design. For *outerplanar* graph metrics, we provide a universal cost-sharing protocol with PoA of at most 2, in contrast to protocols that, by ignoring the graph metric, cannot achieve better than a logarithmic PoA. However, in our main technical result, we show that there exist graph metrics, namely the hypercube, for which knowing the underlying metric does not help; that is, *any universal* protocol has logarithmic PoA which is tight and matches the bound of [36] that ignores the graph metric. We attack this problem by developing new techniques that employ powerful tools from extremal combinatorics, and more specifically Ramsey Theory in high dimensional hypercubes.

Those results partially answer our questions, however, the question remains for other significant metrics such as the Euclidean metric and planar graphs. Furthermore, very few is known with respect to randomised protocols. The only known result is a logarithmic lower bound of the protocol that chooses an order uniformly at random [72].

In [43] we further study the stochastic model, where the players' terminals are drawn from some probability distribution which is given to the designer. The goal is now to choose a universal protocol where the *expected* worst-case cost in the Nash equilibrium is not far from the *expected* optimal cost. We show that there exists a randomised ordered protocol that achieves constant PoA.

In [42] (co-authored with Christodoulou and Leonardi) we study the Bayesian design, which is similar to the stochastic, with the following ground difference: the players have also incomplete information, i.e. they are only aware of the probability distribution over the players' terminals and they choose their paths so that they minimise their *expected* cost-share. We show that the PoA under the Bayesian setting can be very high ($\Omega(\sqrt{n})$, where n is the number of players) for *any* cost-sharing protocol satisfying some natural properties. One of them is

the budget-balance, where players cost-shares should cover exactly the occurring usage cost. However, by relaxing this assumption and requiring budget-balance only in all equilibria solutions, we design a cost-sharing protocol with constant PoA. The protocol is derived after showing an interesting connection between algorithms for oblivious stochastic optimisation problems and cost-sharing design. We further show how to derive the same constant bound on the PoA by using *anonymous* posted prices.

Part I

Simple Auctions

CHAPTER 2

Overview

This part is based on joint works with George Christodoulou, Annamária Kovács and Bo Tang. Chapters 4, 5, 6 are based on the paper [41], co-authored with George Christodoulou, Annamária Kovács and Bo Tang, which was published in the ACM Transaction on Economics and Computation in 2016. Chapter 7 is based on the paper [46], co-authored with George Christodoulou and Bo Tang, which appeared in the Proceedings of the 8th International Symposium on Algorithmic Game Theory in 2015 and was invited to Special Issue of Theory of Computing Systems [47] in 2016.

Allocating many resources to many players with combinatorial valuation functions over allocations constitutes a fundamental optimisation problem. In this thesis we consider both indivisible and divisible resources. We study the former case as a *combinatorial auction* and the latter in the context of *bandwidth allocation*.

A very common objective is to maximise the social welfare, i.e. the aggregation of players' valuation. Under this objective, a mechanism designer should decide the allocation and possible payments for the players, without having access to their valuations, and they should only declare their (maybe non-true) preferences.

This is a well-studied problem for which it is known that the celebrated VCG mechanism is truthful, meaning that declaring their true valuation is players' (weakly) *dominant* strategy, i.e. it is in their best interest to declare the truth. Furthermore, the VCG mechanism allocates the resources in such a way that the aggregation of players' declared valuation is maximised. Therefore, under the assumption that the participants are rational and hence, they declare their true

preference, the VCG mechanism provides the desired outcome, i.e. maximises the social welfare. In spite of VCG's nice properties, it is rarely used in practice due to two main drawbacks: a) it is hard for the players to describe fully their valuations and b) it is computationally inefficient, meaning that computing the allocation and the payments requires time possibly exponential with respect to the number of resources and players. As a consequence many simpler mechanisms have been recruited.

Simultaneous (item-bidding) Auctions. Of particular interest are the so-called *simultaneous auctions*, (also known as *item-bidding* auctions) from both practical and theoretical aspects. In such an auction, the resources are sold *simultaneously* in single-item auctions. Bikhchandani [21] was the first who studied the simultaneous sealed bid auctions in full information settings and observed the inefficiency of their equilibria.

Depending on the type of single-item auctions used, the two main variants that have been studied for combinatorial auctions are *simultaneous second-price auctions* [19, 40, 64] and *simultaneous first-price auctions* [64, 86, 133]. In both cases, the bidders are asked to submit a bid for each item. Then each item is assigned to the highest bidder. The main difference is that in the former a winner is charged an amount equal to the second highest bid while in the latter a winner pays his own bid.

Another interesting variant is the *simultaneous all-pay (first-price) auction* [45, 133], where each item is still assigned to the highest bidder, however, all players should pay their bids. It is a common economic phenomenon in competitions that agents make irreversible investments without knowing the outcome. All-pay auctions are widely used in economics to capture such situations, where all players, even the losers, pay their bids.

Simultaneous first-price auctions have been shown to be more efficient than second-price and all-pay auctions. For general valuation functions, Hassidim et al. [86] showed that pure equilibria of first-price auctions are efficient whenever they exist, but mixed and Bayesian Nash equilibria of first-price auctions can be highly inefficient in settings with *complementarities*. For two important classes of valuation functions, namely *fractionally subadditive* and *subadditive*⁶, for mixed and Bayesian Nash equilibria, Hassidim et al. [86], Syrgkanis and Tardos [133] and Feldman et al. [64] showed that simultaneous first-price auctions have lower

⁶Fractionally subadditive valuations are also known as XOS valuations. For definitions of these valuation functions we refer the reader to Chapter 3.

or the same Price of Anarchy than the respective bounds obtained for second-price auctions, due to Bhawalkar and Roughgarden [19], Christodoulou, Kovács and Schapira [40] and Feldman et al. [64], and for the all-pay auctions, due to Christodoulou et al. [41] and Christodoulou, Sgouritsa and Tang [45].

Simultaneous auctions can be also used for allocating many divisible resources like bandwidth. Many resources are auctioned in simultaneous single-resource auctions. In the literature [91, 46, 133], the proportional allocation mechanism has been used as the single-resource auction, where each bidder submits a single bid and receives a portion of the resource proportional to her bid by paying her own bid.

Multi-Unit Auctions. The multi-unit auction is a special case of combinatorial auction where all items are identical (or alternatively, they are units of the same item). Hence, the bidders do not discriminate among items and their valuation becomes simpler since it depends only on the number of items received and not on the specific set. The natural bidding scheme now is to submit one bid for each *number* of items. In this thesis we consider the *standard bidding format* [56] of non-increasing marginal bids; that is the willingness to pay for an additional unit decreases with the number of units already obtained.

There are many variations for auctioning m units of the same item: the discriminatory auction, the uniform price auction and the Vickrey multi-unit auction [56, 100, 108]. In all those auctions, the units are allocated to the m highest marginal bids, and their difference lies on the pricing scheme. Our focus is on the *discriminatory auction*, which can be seen as the variant of the first-price auction adjusted to multi-unit auctions, where each bidder pays her *winning* marginal bids.

Even though we do not sell each item separately any more, we consider this auction along with item-bidding auctions due to the simplicity of its bidding scheme.

2.1 Results

Combinatorial Auction. Following the work of [86, 64, 133], we study the Price of Anarchy of first-price auctions under full information or Bayesian settings. Our main concern is the development of tools that provide *tight* bounds for the Price of Anarchy of these auctions. Our results complement the current

knowledge about simultaneous first-price auctions.

The current best upper bounds for the Price of Anarchy in first-price auction are $e/(e - 1) \approx 1.58$ for XOS valuations due to Syrgkanis and Tardos [133], and 2 for subadditive valuations Feldman et al. [64] (proven by different techniques). We provide *matching* lower bounds to those upper bounds, showing that *even* for the case of full information and mixed Nash equilibria the PoA is at least $e/(e - 1)$ for submodular⁷ valuations (and therefore for XOS) and 2 for subadditive valuations⁸.

We present an alternative proof of the upper bound of $e/(e - 1)$ for first-price auctions with *fractionally subadditive* valuations. This bound was shown before in [133] by using a general smoothness framework. Our approach does not adhere to their framework. A nice thing with our approach, is that it reveals the worst-case price distribution, that we then use as a building block for the matching lower bound construction (Chapter 4). An immediate consequence of our results is that the Price of Anarchy of these auctions *stays the same*, for mixed, correlated, coarse-correlated, and Bayesian Nash equilibria. Only for pure Nash equilibria it is equal to 1. Our findings suggest that smoothness may provide tight results for certain classes of auctions, using as a *base class* the class of mixed Nash equilibria, and not that of pure equilibria. This is in contrast to what is known for routing games, where the respective base class was the class of pure equilibria.

Then we generalise our results to a class of item bidding auctions that we call *bid-dependent* auctions (Chapter 5). Intuitively, a single item auction is bid-dependent if the winner is always the highest bidder, and a bidder's payment depends only on whether she gets the item or not and on her own bid. Note that both winner and losers may have to pay. Apart from the first-price auction (where the losers pay 0), another notable item-bidding auction that falls into this class is the simultaneous all-pay (first-price) auction [133], where all bidders (even the losers) are charged their bids. However, the second-price auction is not a bid-independent auction, since the bidder's payments may depend also on other bidder's bids, e.g. the highest bidder pays the second highest bid.

For submodular and XOS valuations we show upper bounds in the range be-

⁷In fact our lower bound holds even for the class of OXS valuations that is a strict subclass of submodular valuations. We refer the reader to Chapter 3 for a definition of OXS valuations and for their relation to other valuation classes.

⁸Independently, and after a preliminary version of [41], Roughgarden [124] showed a general method to provide lower bounds for the Price of Anarchy of auctions. We discuss it and compare it to our work in Section 2.2.

tween $e/(e - 1)$ and 2 that depend on a parameter of the auction. We further show that the lower bound of $e/(e - 1)$ of the first-price auction can be generalised for all bid-dependent auctions. For subadditive valuations, we show that the PoA of simultaneous bid-dependent auctions is exactly 2, by showing tight upper and lower bounds. We show that the upper bound technique due to Feldman et al. [64] for first-price auctions, can be applied to *all* mechanisms of this class. Interestingly, although one might expect that first-price auctions perform strictly better than all-pay auctions, our results suggest that all simultaneous bid-dependent auctions perform equally well. We note that our upper bound for subadditive valuations extends the previously known upper bound of 2 for all-pay auctions that was only known for XOS valuations [133].

Multi-Unit Auction. We further apply our techniques on *discriminatory price multi-unit auction* (Chapter 6). We complement the results of de Keijzer et al. [56] for the case of subadditive valuations, by providing a matching lower bound of 2, for the standard bidding format. For the case of submodular valuations, we were able to provide a lower bound of 1.099. We were also able to reproduce their upper bound of $e/(e - 1)$ for submodular bids, using our non-smooth approach. Note that the previous lower bound for such auctions was 1.0004 [56] for Bayesian Nash equilibria. Both our lower bounds hold for the case of mixed Nash equilibria and therefore for more general equilibrium concepts.

Bandwidth Allocation. Johari and Tsitsiklis [91] studied the efficiency of the proportional allocation mechanism. They noticed that it does not always result in an outcome that maximises the social welfare. On the other hand, they showed that this efficiency loss is bounded by a constant when agents' valuations are concave. More specifically, they proved that the proportional allocation mechanism admits a *unique pure* Nash equilibrium (a result due to Hajek and Gopalakrishnan [82]) with PoA of $4/3$ and there is no other mixed Nash equilibrium. Their result holds even when many resources are auctioned in simultaneous single-resource auctions, where players have combinatorial valuation functions over fractional allocations.

An essential assumption used by Johari and Tsitsiklis [91] is that agents have complete information of each other's valuations. However, in many realistic scenarios, the agents are only partially informed, which is expressed by using the Bayesian framework. A natural question is whether the efficiency loss is still bounded in the Bayesian setting. We give a negative answer to this question

by showing that the PoA over Bayesian equilibria is at least $\sqrt{m}/2$, where m is the number of resources (Chapter 7). This result complements the study by Caragiannis and Voudouris [30], where the PoA of single-resource proportional allocation games is shown to be at most 2 in the Bayesian setting.

Non-concave valuation functions were studied by Syrgkanis and Tardos [133] for both full information and Bayesian games. They showed that, when agents' valuations are lattice-submodular, the PoA for coarse correlated and Bayesian Nash equilibria is at most 3.73 by applying their general smoothness framework. We study subadditive valuations that is a superclass of lattice submodular valuations, but not of concave valuations [122]⁹, and we prove that the PoA over Bayesian Nash equilibria is at most 2 (Chapter 7). Moreover, we show optimality of the proportional allocation mechanism by proving that this bound is tight and cannot be improved by any *simple* mechanism, as defined in the framework of Roughgarden [124], or any *scale-free* mechanism¹⁰.

2.2 Literature Review

A long line of research aims to design simple auctions with good performance guarantee (see e.g. [85, 33]). The (in)efficiency of first-price price auctions has been observed in economics (cf. [100]) starting from the seminal work by Vickrey [138].

Simultaneous (item-bidding) Auctions. Bikhchandani [21] was the first who studied the simultaneous sealed bid auctions in full information settings and observed the inefficiency of their equilibria. Christodoulou, Kovács and Schapira [40] extended the concept of PoA to the Bayesian setting and applied it to simultaneous (item-bidding) auctions. Bikhchandani [21] and then Hassidim et al. [86] showed that, in case of general valuations, the pure Nash equilibria of first-price auctions are always efficient (whenever they exist), whereas regarding the second-price auctions, Fu, Kleinberg and Lavi [70] proved that the PoA is at most 2 under the no-overbidding assumption¹¹. For Bayesian Nash equilibria, Syrgkanis and Tardos [133] and Feldman et al. [64] showed improved

⁹In fact lattice-submodular and concave functions coincide in the single dimensional case.

¹⁰The scale free mechanisms are formally defined in Section 7.2.3. The basic property of a scale-free mechanism is that, if every bid is scaled by the same constant, the outcome remains unchanged.

¹¹The no-overbidding assumption is satisfied if for any bidder the sum of her bids over any set of items does not exceed her valuation of this set

Valuations	Lower Bound	Upper Bound
General, Pure	1	1 [21, 86]
General, M-B	$\Omega(\sqrt{m})$ [86]	$O(m)$ [86]
SA, M-B	2 [This thesis]	2 [64]
(XOS, SM, OXS) M-B	$e/(e-1)$ [This thesis]	$e/(e-1)$ [133]

Table 2.1: In the first column, the first argument refers to the valuation class and the second argument to the related equilibrium concept. SA and SM stand for subadditive and submodular valuations, respectively, and whenever ‘M-B’ appears the bounds hold for mixed, correlated, coarse correlated (defined in Chapter 3) or Bayesian Nash equilibria.

upper bounds on the PoA of first-price and second-price auctions. Syrgkanis and Tardos [133] came up with a general composability framework of smooth mechanisms, that proved to be quite useful, as it led to upper bounds for several settings, such as first price auctions, all-pay auctions and multi-unit auctions.

Only a few lower-bound results are known for the PoA of simultaneous auctions. For valuations that include *complementarities*, Hassidim et al. [86] presented an example with $\text{PoA} = \Omega(\sqrt{m})$ for the first-price auction; as mentioned in [64], similar lower bound can be derived for the second-price auction, as well. Under the non-overbidding assumption, Bhawalkar and Roughgarden [19] gave a lower bound of 2.013 for the second-price auction with subadditive bidders and $\Omega(n^{1/4})$ bidders’ valuations are drawn from correlated distributions. In [64], similar results are shown under the weak non-overbidding assumption. We summarise the PoA results for the first-price auction in table 2.1.

Independently, Roughgarden [124] presented a very elegant methodology to provide PoA lower bounds via a reduction from communication or computational complexity lower bounds for the underlying optimisation problem. One consequence of this reduction is a general lower bound of 2 and $e/(e-1)$ for the PoA of *any* simple¹² auction (including simultaneous auctions) with subadditive and fractionally subadditive bidders, respectively. Therefore, there is an overlap with our results for these two classes of valuations.

We emphasise that this approach is incomparable to ours in the following sense. On the one hand, the results in [124] hold for more general formats of combinatorial auctions than the ones we study here. On the other hand, our $e/(e-1)$ lower bound holds even for more special valuation functions where the results of [124] are either weaker ($2e/(2e-1)$ for submodular valuations) or

¹²In a simple mechanism, the players’ action space should be at most sub-doubly-exponential in the number of items.

not applicable (for OXS valuations). Feige and Vondrák [63] showed that, for submodular valuations, a strictly higher than $1 - 1/e$ amount of the optimum social welfare can be obtained with polynomial communication¹³, and Nisan and Segal [117] showed that, for gross substitute valuations (and therefore for its subclass, OXS valuations) exact efficiency can be obtained with polynomial communication. These two results show that, either there exists an auction better than the first-price auction, or the technique of Roughgarden [124] does not provide tight lower bounds for these classes of valuations in contrast to our results. We also note that the PoA lower bound obtained by the reduction used in [124] can only be applied to approximate Nash equilibria while our results are proved via an explicit construction that apply to exact Nash equilibria. Further, our PoA lower bound proof for subadditive valuations uses a simpler construction than the proof in [124] and it holds even for the case of only 2 bidders and identical items (multi-unit auction). Finally, it should be stressed that none of our lower bounds for multi-unit auctions can be derived from [124].

Multi-Unit Auction. Markakis and Telelis [108] studied uniform price *multi-unit* auctions. De Keijzer et al. [56] bounded the PoA under the Bayesian setting for several formats of multi-unit auctions with first or second price rules. Auctions employing greedy algorithms were studied by Lucier and Borodin [105]. A number of works [102, 29, 126] studied the PoA of generalised second-price auctions in the full information and Bayesian settings and even with correlated bidders [106]. Chawla and Hartline [32] proved that for the generalised first-price auctions with symmetric bidders, the pure Bayesian Nash equilibrium is unique and always efficient.

Bandwidth Allocation. The efficiency of the proportional allocation mechanism has been extensively studied in the literature of network resource allocation. Besides the work mentioned so far, Johari and Tsitsiklis [92] studied a more general class of scalar-parametrised mechanisms and proved that the proportional allocation mechanism achieves the best PoA when the mechanism only chooses a single price.

Syrkkanis and Tardos [133], Caragiannis and Voudouris [30] and Christodoulou et al. [46] studied the efficiency of the proportional allocation mechanism in the setting where agents are constrained by budgets that represent the maximum

¹³The information revealed by the bidders, e.g. their valuation for sets of items, needs polynomial time to be exchanged.

payment they can afford. They used as a benchmark the *effective welfare* capping the contribution of each player to the welfare by their budget. Syrgkanis and Tardos [133] proved an upper bound on the PoA of 3.73 for lattice-submodular valuations by considering many resources under both the full information and the Bayesian setting. Caragiannis and Voudouris [30] improved this bound to 2.78 for a single resource. At last, Christodoulou et al. [46] further improved both bounds to $1 + \phi \approx 2.618$, where ϕ is the golden ratio, for the full information case with many resources and even for the more general subadditive class of valuations. Zhang [142] and Feldman, Lai and Zhang [67] studied the efficiency and fairness of the proportional allocation mechanism when agents aim at maximising non quasi-linear utilities subject to budget constraints.

Nguyen and Tardos [112] and Christodoulou et al. [46] studied the proportional allocation mechanism in the polyhedral environment, where there exists a collection of resources and the goal is to associate each agent with a *single* value, representing their level of activity. Nguyen and Tardos [112] and Christodoulou et al. [46] proved, respectively, that the PoA for pure Nash equilibria are $4/3$ for concave valuations and 2 for subadditive valuations. Correa, Schulz and Stier-Moses [55] showed a relationship in the efficiency loss between the proportional allocation mechanism and non-atomic selfish routing for not necessarily concave valuation functions.

CHAPTER 3

Preliminaries

3.1 Mechanisms and Valuations

In a combinatorial auction with n players (or bidders) and m items, every player $i \in [n]$ has a valuation for each subset of items, given by a valuation set function $v_i : 2^{[m]} \rightarrow \mathbb{R}_+$, where $[n]$ denotes the set $\{1, 2, \dots, n\}$. The multi-unit auction is a special case where all items are identical, or alternatively there are m units of the same item, and the valuation of each player can be represented more concisely by $v_i : [m] \rightarrow \mathbb{R}_+$. In bandwidth allocation (divisible resources) there are n agents who compete for m divisible resources with unit supply. Every agent $i \in [n]$ has a valuation function, $v_i : [0, 1]^m \rightarrow \mathbb{R}_+$.

A mechanism takes the agents' bids/strategies $\mathbf{b} = (b_1, \dots, b_n)$ as input and outputs a tuple (\mathbf{X}, \mathbf{p}) for combinatorial auctions, $(\boldsymbol{\xi}, \mathbf{p})$ for multi-unit auctions and (\mathbf{x}, \mathbf{p}) for divisible resources, where \mathbf{X} , $\boldsymbol{\xi}$ and \mathbf{x} denote the allocation of the resources to the players and \mathbf{p} is their payments. The payment vector $\mathbf{p} = \mathbf{p}(\mathbf{b}) = (p_1, \dots, p_n)$ specifies the agents' payments, where p_i is a real value that denotes the payment of player i .

Combinatorial auction. The allocation $\mathbf{X} = \mathbf{X}(\mathbf{b}) = (X_1, \dots, X_n)$ is a partition of the items, where X_i is the set allocated to player i (allowing empty sets X_i), so that each item is assigned to exactly one player. The valuation functions are monotone and normalised, that is, for every $S \subseteq T \Rightarrow v_i(S) \leq v_i(T)$, and $v_i(\emptyset) = 0$. We use the short notation $v_i(j) = v_i(\{j\})$.

Multi-unit auction. The allocation $\boldsymbol{\xi} = (\xi_1, \xi_2, \dots, \xi_n)$ is a vector of nonnegative integers, where ξ_i denotes the number of units allocated to player i , and

we require that $\sum_i \xi_i \leq m$. The valuation functions are non-decreasing, i.e. $v_i(s) \leq v_i(s+1)$ for all $s \in [m-1]$, and normalised, i.e. $v_i(0) = 0$.

Bandwidth Allocation. The vector $\mathbf{x} = \mathbf{x}(\mathbf{b}) = (x_1, \dots, x_n)$ specifies the allocation of resources, such that the total allocation of each resource doesn't exceed the unit supply. For every i , $x_i = (x_{i1}, \dots, x_{im}) = (x_{ij})_j$ denotes the allocation to agent i , where x_{ij} is the quantity she receives from recourse j . We require that for every j , $\sum_i x_{ij} \leq 1$. The valuations are monotonically non-decreasing, that is, for every two allocations, $x_i, x'_i \in [0, 1]^m$, where $\forall j \in [m] x_{ij} \leq x'_{ij}$, we have $v_i(x_i) \leq v_i(x'_i)$. We further assume that the valuations are normalised as $v_i((0, \dots, 0)) = 0$.

We represent the valuations of all agents, respectively for the three auctions, by using the vectors $\mathbf{v} = \mathbf{v}(\mathbf{X}) = (v_1(X_1), \dots, v_n(X_n))$, $\mathbf{v} = \mathbf{v}(\boldsymbol{\xi}) = (v_1(\xi_1), \dots, v_n(\xi_n))$ and $\mathbf{v} = \mathbf{v}(\mathbf{x}) = (v_1(x_1), \dots, v_n(x_n))$. Occasionally, instead of $v_i(X_i)$, $v(\xi_i)$ and $v_i(x_i)$ we may use the notation $v_i(\mathbf{X})$, $v_i(\boldsymbol{\xi})$ and $v_i(\mathbf{x})$, which represents agent i 's valuation under the allocation \mathbf{X} , $\boldsymbol{\xi}$ and \mathbf{x} , respectively.

For each player i , there is some possible set of valuations V_i such that $v_i \in V_i$. A valuation profile for all players is $\mathbf{v} = (v_1, v_2, \dots, v_n) \in \times_i V_i$. In the *Bayesian* setting, the valuation of each player i is drawn from V_i according to some known distribution D_i . We assume that the D_i 's are independent (and possibly different) over the players and we denote by $\mathbf{D} = \times_i D_i$ their product distribution. In the *full information* setting the valuation v_i is fixed and known by all other players for all $i \in [n]$. Note that the latter is a special Bayesian auction, in which player i has valuation v_i with probability 1.

3.2 Item-bidding Auctions

In an *item bidding* auction every agent i submits a non-negative bid b_{ij} for each resource j representing their willingness to pay for that resource; $b_i = (b_{i1}, \dots, b_{im})$ is a vector of the bids for each resource. In the multi-unit auction, where all items are identical, we consider the *standard bidding format* [56, 100] in which the bids are in decreasing order¹⁴, i.e. $b_{ij} \geq b_{ij+1}$ and b_{ij} denotes player i 's

¹⁴This is a reasonable restriction when the valuations of an additional unit decreases with the number of units already obtained, which is the case in submodular valuations. However, we also use the same bidding format with subadditive valuations, which is also justifiable in discriminatory auctions where the players pay their bid. The reason is that in real life it is very rare that extra charges occur by increasing the number of obtained units; on the contrary, it is very common that buyers receive discounts for buying more units.

willingness to pay for receiving an extra item given that she has already paid for $j - 1$ items. We denote the strategies of all players as $\mathbf{b} = (b_1, \dots, b_n)$ and by $\mathbf{b}_{-i} = (b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ we denote the strategies of all agents except for i . For simplicity, in combinatorial auctions sometimes we use the notation $b_i(S) = \sum_{j \in S} b_{ij}$, and $b_i(j) = b_{ij}$.

We study specific mechanisms for the combinatorial auction, the multi-unit auction and the bandwidth allocation which are the first-price auction, the discriminatory auction and the proportional allocation mechanism, respectively. The items are then allocated as follows:

First Price Auction: For each $j \in [m]$, the bidder i with the highest bid b_{ij} receives the item. In a case of a tie we consider an arbitrary randomised tie-breaking rule. Note that with such a rule, for any fixed $\mathbf{b} = (b_1, \dots, b_n)$, the probabilities for the players to get a particular item are fixed. If no ties appear, the allocation to player i is $X_i = \{j \in [m] : b_{ij} > \max_{k \neq i} (b_{kj})\}$. A player pays his own bid (the highest bid) for every item he receives, i.e. $p_i = b_i(X_i)$.

Discriminatory Auction: The units are allocated to the m highest bids, i.e. ξ_i is the number of player i 's bids that are among the m highest bids of all players. In *discriminatory pricing*, every bidder i pays the sum of his *winning* bids, i.e. $p_i = \sum_{j \leq \xi_i} b_{ij}$.

Proportional Allocation Mechanism: The allocation of any resource $j \in [m]$ to player i is denoted by $x_{ij} = \frac{b_{ij}}{\sum_{k \in [n]} b_{kj}}$. Each player pays the sum of her own bids, i.e. $p_i = \sum_{j \in [m]} b_{ij}$. When all agents bid 0, the allocation can be defined arbitrarily, but consistently.

The *utility* u_i of agent i is defined as the difference between her valuation for the received allocation and her payment: $u_i(\mathbf{X}(\mathbf{b}), \mathbf{p}(\mathbf{b})) = u_i(\mathbf{b}) = v_i(X_i(\mathbf{b})) - p_i(\mathbf{b})$ for combinatorial auctions, $u_i(\boldsymbol{\xi}(\mathbf{b}), \mathbf{p}(\mathbf{b})) = u_i(\mathbf{b}) = v_i(\xi_i(\mathbf{b})) - p_i(\mathbf{b})$ for multi-unit auctions and $u_i(\mathbf{x}(\mathbf{b}), \mathbf{p}(\mathbf{b})) = u_i(\mathbf{b}) = v_i(x_i(\mathbf{b})) - p_i(\mathbf{b})$ for bandwidth allocation. In the Bayesian setting, we use the notation $u_i^{v_i}$ in order to specify that v_i is the valuation function of agent i . For simplicity, we use shorter notation for expectations, e.g., we use $\mathbb{E}_{\mathbf{v}}$ instead of $\mathbb{E}_{\mathbf{v} \sim \mathbf{D}}$, $\mathbb{E}[u_i(\mathbf{b})]$ instead of $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})]$ and $u(\mathbf{B})$ instead of $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u(\mathbf{b})]$ whenever \mathbf{D} and \mathbf{B} are clear from the context.

3.3 Bidding Strategies and Equilibria

We use \mathbf{b} to denote a *pure* strategy profile; in a more general context, we denote a strategy profile as $\mathbf{B} = (B_1, \dots, B_n)$, where B_i is a probability distribution over all possible pure strategies of agent i . We study five standard equilibrium concepts: pure Nash, mixed Nash, correlated, coarse correlated and Bayesian Nash equilibria. The first four of them are for the *full information* setting and the last one is defined in the *Bayesian* setting. In each one of the following lines, a strategy profile \mathbf{B} forms the equilibrium notion on the left, if for every agent i and all bids b'_i it satisfies the inequality on the right:

Pure Nash: \mathbf{B} chooses \mathbf{b} with probability 1, $u_i(\mathbf{b}) \geq u_i(b'_i, \mathbf{b}_{-i})$.

Mixed Nash: $\mathbf{B} = \times_i B_i$, $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] \geq \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(b'_i, \mathbf{b}_{-i})]$.

Correlated: $\mathbf{B} = (B_i)_i$, $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})|b_i] \geq \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(b'_i, \mathbf{b}_{-i})|b_i]$.

Coarse correlated: $\mathbf{B} = (B_i)_i$, $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] \geq \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_i(b'_i, \mathbf{b}_{-i})]$.

Bayesian Nash: $\mathbf{B}(\mathbf{v}) = \times_i B_i(v_i)$, $\mathbb{E}_{\mathbf{v}_{-i}, \mathbf{b}}[u_i^{v_i}(\mathbf{b})] \geq \mathbb{E}_{\mathbf{v}_{-i}, \mathbf{b}_{-i}}[u_i(b'_i, \mathbf{b}_{-i})]$, $\forall v_i \in V_i$.

The first four classes of equilibria are in increasing order of inclusion. Moreover, any pure (or mixed) Nash equilibrium is also a Bayesian pure (or mixed) Nash equilibrium.

3.4 Price of Anarchy

We give the following definitions only with respect to combinatorial auctions; the same notions can be defined accordingly for the multi-unit auctions and the bandwidth allocation.

The most common global objective in such settings is to maximise the sum of the valuations of the players for their received sets of items, i.e., to maximise the *social welfare* $\text{SW}(\mathbf{X})$ of the allocation, where $\text{SW}(\mathbf{X}) = \sum_{i \in [n]} v_i(X_i)$. Therefore, for an *optimal allocation* $\mathbf{O}(\mathbf{v}) = \mathbf{O}^{\mathbf{v}} = \mathbf{O} = (O_1, \dots, O_n)$ the value $\text{SW}(\mathbf{O})$ is maximum among all possible allocations. In bandwidth allocation we denote the optimum allocation as $\mathbf{o}(\mathbf{v}) = \mathbf{o}^{\mathbf{v}} = \mathbf{o} = (o_1, \dots, o_n)$ and by $o_i = (o_{i1}, \dots, o_{im})$ we denote the optimal allocation to agent i . In the multi-unit auction we use the same notation, $\mathbf{o}(\mathbf{v}) = \mathbf{o}^{\mathbf{v}} = \mathbf{o} = (o_1, \dots, o_n)$, for the optimum allocation, but now o_i denotes the number of units allocated to player i in \mathbf{o} . Whenever the allocation rule \mathbf{X} is clear from the context, we use $\text{SW}(\mathbf{b})$ and $v_i(\mathbf{b})$ instead of

$\text{SW}(\mathbf{X}(\mathbf{b}))$ and $v_i(X_i(\mathbf{b}))$.

The *Price of Anarchy (PoA)* is defined as the worst case ratio between the social welfare in the optimum allocation and the social welfare in any equilibrium. We define the PoA with respect to pure Nash equilibria, mixed Nash (or correlated or coarse correlated) equilibria and Bayesian Nash equilibria¹⁵, respectively, as

$$\text{PoA} = \max_{\mathbf{v} \in \times_i V_i} \max_{\mathbf{b} \in \mathcal{N}} \frac{\text{SW}(\mathbf{O}(\mathbf{v}))}{\text{SW}(\mathbf{b})}; \quad \text{PoA} = \max_{\mathbf{v} \in \times_i V_i} \max_{\mathbf{B} \in \mathcal{MN}} \frac{\text{SW}(\mathbf{O}(\mathbf{v}))}{\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[\text{SW}(\mathbf{b})]};$$

$$\text{PoA} = \max_{\mathbf{D}} \max_{\mathbf{B} \in \mathcal{BN}} \frac{\mathbb{E}_{\mathbf{v} \sim \mathbf{D}}[\text{SW}(\mathbf{O}(\mathbf{v}))]}{\mathbb{E}_{\mathbf{v} \sim \mathbf{D}, \mathbf{b} \sim \mathbf{B}(\mathbf{v})}[\text{SW}(\mathbf{b})]},$$

where \mathcal{N} , \mathcal{MN} and \mathcal{BN} are respectively the sets of pure Nash equilibria, mixed Nash (or correlated or coarse correlated) equilibria and Bayesian Nash equilibria.

3.5 Types of Valuations

3.5.1 Valuations for Combinatorial Auctions

Our results concern different classes of valuation functions, which we define next, in increasing order of inclusion. Let $v : 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$, be a valuation function. Then, for arbitrary item sets $S, T \subseteq [m]$, v is called

– *additive*, if

$$v(S) = \sum_{j \in S} v(j);$$

– *multi-unit-demand* or *OXS*¹⁶, if, for some k , there exist k unit demand valuations v^1, \dots, v^k (defined as $v^r(S) = \max_{j \in S} v^r(j)$, for any $S \subseteq [m]$), such that

$$v(S) = \max_{S = \dot{\cup}_{r \in [k]} S_r} \sum_{r \in [k]} v^r(S_r);^{17}$$

– *submodular*, if

$$v(S \cup T) + v(S \cap T) \leq v(S) + v(T);$$

¹⁵The equilibrium concept will be clear from the context.

¹⁶OXS stands for OR-of-XOR-of-Singletons, where SUM is denoted by OR and MAX by XOR. XORs of singleton valuations are the unit demand valuations. In OXS valuations, the valuation of a set is given by the best way to split the set among several unit demand valuations.

¹⁷ $\dot{\cup}$ stands for disjoint union.

– *fractionally subadditive* or *XOS*¹⁸, if v is determined by a finite set of *additive* valuations f_γ for $\gamma \in \Gamma$, so that

$$v(S) = \max_{\gamma \in \Gamma} f_\gamma(S);$$

– *subadditive*, if

$$v(S \cup T) \leq v(S) + v(T).$$

It is well-known that each one of the above classes is strictly contained in the next class, e.g., an additive set function is always submodular but not vice versa, a submodular is always XOS, etc. [62]. As an equivalent definition, submodular valuations are exactly the valuations with *decreasing marginal values*, meaning that, for any $S \subseteq T$, $v(\{j\} \cup T) - v(T) \leq v(\{j\} \cup S) - v(S)$ holds for any item j .

3.5.2 Valuations for Multi-Unit Auctions

In multi-unit auction we only consider submodular and subadditive valuations. The valuation v is called:

– *submodular*, if the items have decreasing marginal values, that is, for every $s \leq t \leq m$,

$$v_i(t+1) - v_i(t) \leq v_i(s+1) - v_i(s);$$

– *subadditive*, if, for every $s \leq t \leq m$,

$$v_i(s+t) \leq v_i(s) + v_i(t).$$

It is easy to check that every submodular valuation is also subadditive.

3.5.3 Valuations for Bandwidth Allocation

Our focus is on two valuation functions: the concave and subadditive. A function $v : [0, 1]^m \rightarrow \mathbb{R}_+$ is

– *concave*, if for all $x, y \in [0, 1]^m$ and any $\lambda \in (0, 1)$, it is

$$v((1-\lambda)x + \lambda y) \geq (1-\lambda)v(x) + \lambda v(y).$$

– *subadditive*, if for all $x, y \in [0, 1]^m$, such that $x + y \in [0, 1]^m$, where $x + y$ is the

¹⁸XOS stands for XOR-of-OR-of-Singletons. ORs of singleton valuations are exactly the additive valuations and we take the XOR meaning the maximum of them.

componentwise sum of x and y , it is

$$v(x + y) \leq v(x) + v(y).$$

Remark 3. Lattice submodular functions used in [133] are subadditive (see Section 7.2). In the case of a single variable (single resource), any concave function is subadditive; more precisely, concave functions are equivalent to lattice submodular functions in this case. However, concave functions of many variables *may not* be subadditive [122].

CHAPTER 4

First Price Auction

In this chapter we study the first-price item-bidding auction where each player submits a bid for each item and then each item is given to the highest bidder who pays her bid for that item. We provide the exact PoA when players have the following classes of valuations: OXS, submodular, fractionally subadditive and subadditive.

4.1 Fractionally Subadditive Valuations

In this section we present a lower bound of $\frac{e}{e-1}$ for the PoA of mixed Nash equilibria in simultaneous first price auctions with OXS and therefore, submodular and fractionally subadditive valuations. This is a matching lower bound to the results by Syrgkanis and Tardos [133].

In order to explain the key properties of the instance that provides the tight lower bound, we first discuss a new approach to obtain the same upper bound for the PoA of the first-price auction as in [133]. While the upper bound that we derive with the help of this idea, can also be obtained based on the very general *smoothness* framework [126, 133], the approach we introduce here does not adhere to this framework¹⁹. The strength of our approach consists in its potential to lead to better (in this case tight) *lower* bounds, as we demonstrate subsequently.

¹⁹Roughly, because the pure deviating bid a that we identify, depends on the other players' bids \mathbf{b}_{-i} in the Nash equilibrium, whereas in the smoothness technique, it doesn't.

4.1.1 Upper Bound

In order to keep the presentation pure, and to focus on the main ingredients, we first illustrate this approach on a single item auction with full information.

One Item

Since there is only one item, we drop the index referring to the items, e.g. b_i is not a vector anymore, instead it is the bid of player i for the item.

Theorem 4. The PoA of mixed Nash equilibria in first-price single-item auctions is at most $\frac{e}{e-1}$.

Proof. Let $\mathbf{v} = (v_1, \dots, v_n)$ be the valuations of the players, and suppose that player h has the highest valuation, i.e. $v_h \in \max_{k \in [n]} v_k$. We fix a mixed Nash equilibrium $\mathbf{B} = (B_1, B_2, \dots, B_n)$. Let t denote the highest bid in $\mathbf{b}_{-h} \sim \mathbf{B}_{-h}$, i.e. $t = \max_{i \neq h} b_i$; in other words t is the threshold such that only if player h bids above it, she gets the item. Further let $F(x)$ be the cumulative distribution function (CDF) of t , that is, $F(x) = \mathbb{P}_{\mathbf{b}_{-h} \sim \mathbf{B}_{-h}}[t \leq x]$. The following lemma prepares the ground for the selection of an appropriate bid which serves as a deviation from the equilibrium.

Lemma 5. For any pure strategy a of player h , $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_h(\mathbf{b})] \geq F(a)(v_h - a)$.

Proof. If F is continuous in a , then $F(a) = \mathbb{P}[t \leq a] = \mathbb{P}[t < a]$, tie-breaking in a does not matter, and $F(a)$ equals also the probability that bidder h gets the item if he bids a . Therefore, $F(a)(v_h - a) = \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[u_h(a, \mathbf{b}_{-i})] \leq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_h(\mathbf{b})]$, since \mathbf{B} is a Nash equilibrium. If F is not continuous in a ($\mathbb{P}[t = a] > 0$), then, as a CDF, it is at least right-continuous. By the previous argument $\mathbb{E}[u_h(\mathbf{b})] \geq F(x)(v_h - x)$ holds for every $x = a + \epsilon$ where F is continuous, and the lemma follows by taking $\epsilon \rightarrow 0$. \square

Since in a Nash equilibrium the expected utility of every (other) player is non-negative, by summing over all players,

$$\sum_{i=1}^n \mathbb{E}[u_i(\mathbf{b})] \geq F(a)(v_h - a). \quad (4.1)$$

On the other hand, for any fixed bidding profile \mathbf{b} we have $u_i(\mathbf{b}) = v_i(\mathbf{b}) - p_i(\mathbf{b})$, where $p_i(\mathbf{b}) = b_i$ whenever b_i is a winning bid, and $p_i(\mathbf{b}) = 0$ otherwise.

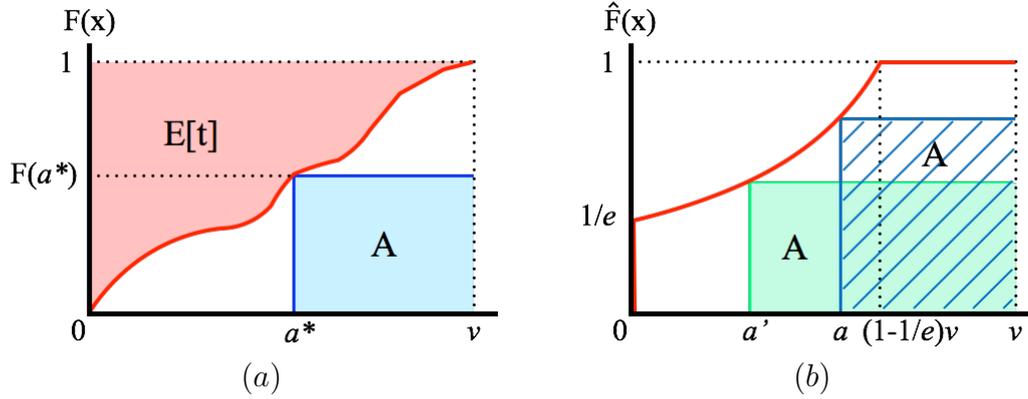


Figure 4.1: Figure (a) is a schematic illustration of the expression $F(a^*)(v - a^*) + \mathbb{E}[t]$, where $A = F(a^*)(v - a^*)$. Figure (b) shows the CDF $\hat{F}(x)$, which makes all the inequalities of Lemma 6 tight, i.e. for every $x \in [0, (1 - \frac{1}{e})v]$, $F(x)(v - x) = A = \frac{v}{e}$.

By taking expectations with respect to $\mathbf{b} \sim \mathbf{B}$, and summing over the players, $\mathbb{E}[\sum_i u_i(\mathbf{b})] = \mathbb{E}[\sum_i (v_i(\mathbf{b}) - p_i(\mathbf{b}))] = \mathbb{E}[\sum_i (v_i(\mathbf{b})) - \max_i b_i] \leq \mathbb{E}[SW(\mathbf{b})] - \mathbb{E}[t]$. By combining this with (4.1), we obtain,

$$\mathbb{E}[SW(\mathbf{b})] = \mathbb{E} \left[\sum_{i=1}^n u_i(\mathbf{b}) \right] + \mathbb{E}[t] \geq F(a)(v_h - a) + \mathbb{E}[t], \quad (4.2)$$

for any (deviating) bid a . (Analogues of this derivation are standard in the simultaneous auctions literature.) We choose the bid a^* that *maximises* the right hand side of (4.2), i.e. $a^* = \arg \max_a F(a)(v - a)$ (see Figure 4.1(a) for an illustration). Then, in order to upper bound the PoA, we look for the *maximum* value of λ , such that,

$$F(a^*)(v_h - a^*) + \mathbb{E}[t] \geq \lambda v_h. \quad (4.3)$$

In fact, we want to bound the sum of the two shaded regions of Figure 4.1(a). The following lemma settles the maximum value of such λ as $1 - \frac{1}{e}$ for mixed equilibria²⁰. This will complete the proof of the theorem, since by (4.2) and $SW(\mathbf{O}) = v_h$ we obtain $\mathbb{E}[SW(\mathbf{b})] \geq (1 - \frac{1}{e})SW(\mathbf{O})$.

Lemma 6. For any non-negative random variable t with CDF F , and any fixed

²⁰If \mathbf{B} is a pure equilibrium, then it is easy to verify that F is a step function, furthermore $a^* = t$, and inequality (4.3) boils down to $1 \cdot (v_h - a^*) + a^* = 1 \cdot v_h$

value v , it is true that

$$F(a^*)(v - a^*) + \mathbb{E}[t] \geq \left(1 - \frac{1}{e}\right)v.$$

Proof. Set $A = F(a^*)(v - a^*)$, for $a^* = \arg \max_a F(a)(v - a)$. We use the fact that the expectation of a non-negative random variable t with CDF F can be calculated as $\mathbb{E}[t] = \int_0^\infty (1 - F(x))dx$.

Thus,

$$\begin{aligned} F(a^*)(v - a^*) + \mathbb{E}[t] &\geq A + \int_0^{v-A} (1 - F(x))dx \\ &= A + (v - A) - \int_0^{v-A} F(x)dx \\ &\geq v - \int_0^{v-A} \frac{A}{v-x} dx = v + A \ln\left(\frac{A}{v}\right) \\ &\geq v + \frac{v}{e} \ln\left(\frac{1}{e}\right) = \left(1 - \frac{1}{e}\right)v, \end{aligned}$$

where the last inequality holds because $A \ln\left(\frac{A}{v}\right)$ is minimised for $A = \frac{v}{e}$. \square

This completes the proof of Theorem 4. \square

Worst-case price distribution. The CDF $F(x)$ that makes all the inequalities of (the proof of) Lemma 6 tight (see Figure 4.1(b)), is

$$\hat{F}(x) = \begin{cases} \frac{v}{e^{(v-x)}} & , \text{ for } x \leq \left(1 - \frac{1}{e}\right)v \\ 1 & , \text{ for } x > \left(1 - \frac{1}{e}\right)v \end{cases}$$

Observe that for $x \leq \left(1 - \frac{1}{e}\right)v$, $\hat{F}(x)(v - x) = \frac{v}{e}$ and for $x > \left(1 - \frac{1}{e}\right)v$, $\hat{F}(x)(v - x) = v - x < v - \left(1 - \frac{1}{e}\right)v = \frac{v}{e}$. So, the bid that maximises the quantity $\hat{F}(a)(v - a)$ is any value $a \in [0, \left(1 - \frac{1}{e}\right)v]$. The given distribution \hat{F} for t makes inequality (4.3) tight. Note that the inequality of Lemma 5 is also tight for *all* $a \in [0, \left(1 - \frac{1}{e}\right)v]$. In order to construct a (tight) lower bound for the PoA, we also need to tighten the inequality in (4.2). Intuitively, we need to construct a Nash equilibrium, where the CDF of t is equal to $\hat{F}(x)$ and b_h doesn't exceed t in expectation. We present a construction (with many items) in Section 4.1.2.

Remark 7. Here we discuss our technique and the smoothness technique that achieves the same upper bound [133]. In [133], a particular mixed bidding strategy A_i was defined for each player i , such that for any pure strategies \mathbf{b} , if

$t = \max_{k \neq i} b_k$ then $\mathbb{E}_{A_i}[u_i(A_i, t)] + t \geq v(1 - 1/e)$. If we denote $g(A, F) = \mathbb{E}_{A, F}[u_i(a, t) + t]$, it can be deduced that $\max_A \min_t g(A, t) \geq v(1 - 1/e)$. In Lemma 6 we show that $\min_F \max_a g(a, F) \geq v(1 - 1/e)$. Moreover, we prove that the inequality is tight by providing the *minimising* distribution \hat{F} , such that $\max_a g(a, \hat{F}) = v(1 - 1/e)$. By the Minimax Theorem, $\min_F \max_a g(a, F) = \max_A \min_t g(A, t) = v(1 - 1/e)$. One advantage of our approach is that it can be coupled with a worst-case distribution \hat{F} that serves as an optimality certificate of the method. Moreover, if one can convert \hat{F} to Nash Equilibrium strategy profile (see Section 4.1.2), a tight Nash equilibrium construction is obtained; this can be a challenging task, though.

Many Items

For completeness, we generalise the upper bound proof to many items and for more general informational and equilibrium concepts.

Coarse Correlated Equilibrium. We first prove that for all XOS (fractionally subadditive) valuation functions, the PoA is at most $\frac{e}{e-1}$ even in coarse correlated equilibria.

Theorem 8. The PoA for simultaneous first-price auctions with XOS valuations for coarse correlated equilibria is at most $\frac{e}{e-1} \approx 1.58$.

Proof. The proof highly relies on the definition of XOS valuations. We refer the reader to Section 3.5.1 for a reminder.

Lemma 9. Let S be any set of items, and f_i be a maximising additive function of S for player i with XOS valuation function v_i . Then for any strategy profile \mathbf{b} , where $b_{ij} = 0$ for $j \notin S$,

$$u_i(\mathbf{b}) \geq \sum_{j \in S} \mathbb{P}[j \in X_i(\mathbf{b})](f_i(j) - b_{ij}).$$

Proof. By the definition of XOS valuations, we have that $v_i(T) \geq f_i(T)$, for every $T \subseteq S$. Then,

$$\begin{aligned} u_i(\mathbf{b}) &\geq \sum_{T \subseteq S} \mathbb{P}[X_i(\mathbf{b}) = T](f_i(T) - b_i(T)) \\ &= \sum_{T \subseteq S} \sum_{j \in T} \mathbb{P}[X_i(\mathbf{b}) = T](f_i(j) - b_{ij}) \end{aligned}$$

$$\begin{aligned}
&= \sum_{j \in S} \sum_{T \subseteq S: j \in T} \mathbb{P}[X_i(\mathbf{b}) = T] (f_i(j) - b_{ij}) \\
&= \sum_{j \in S} \mathbb{P}[j \in X_i(\mathbf{b})] (f_i(j) - b_{ij}).
\end{aligned}$$

□

We fix a coarse correlated equilibrium $\mathbf{B} = (B_1, B_2, \dots, B_n)$. Let t_{ij} be the random variable indicating the highest bid on item j in $\mathbf{b}_{-i} \sim \mathbf{B}_{-i}$, i.e. $t_{ij} = \max_{k \neq i} b_{kj}$, and $F_{ij}(x) = \mathbb{P}[t_{ij} \leq x]$ be the CDF of t_{ij} .

Lemma 10. For any pure strategy b'_i of player i and any set of items S ,

$$\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] \geq \sum_{j \in S} F_{ij}(b'_{ij})(f_i(j) - b'_{ij}).$$

The proof is analogous to that of Lemma 5, by using also Lemma 9.

Let now f_i be a maximising additive valuation of player i for the optimal set O_i . By Lemma 6, for every item j , there exists a value a_{ij} such that $F_{ij}(a_{ij})(f_i(j) - a_{ij}) + \mathbb{E}[t_{ij}] \geq (1 - 1/e)f_i(j)$. Then, by summing over $j \in O_i$ and applying Lemma 10 for the pure strategy a_{ij} for $j \in O_i$ and 0 for the rest of the items, we get,

$$\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] + \sum_{j \in O_i} \mathbb{E}[t_{ij}] \geq \left(1 - \frac{1}{e}\right) \sum_{j \in O_i} f_i(j). \quad (4.4)$$

We further give a lower on the total payments to be used next:

$$\begin{aligned}
\sum_i \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[p_i(\mathbf{b})] &= \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} \left[\sum_j \max_k(b_{kj}) \right] = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} \left[\sum_i \sum_{j \in O_i} \max_k(b_{kj}) \right] \\
&\geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} \left[\sum_i \sum_{j \in O_i} \max_{k \neq i}(b_{kj}) \right] = \sum_i \sum_{j \in O_i} \mathbb{E}[t_{ij}].
\end{aligned} \quad (4.5)$$

By using (4.4) and (4.5) we are ready to complete the proof.

$$\begin{aligned}
\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[SW(\mathbf{b})] &= \sum_{i=1}^n \left(\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] + \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[p_i(\mathbf{b})] \right) \geq \sum_{i=1}^n \left(\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] + \sum_{j \in O_i} \mathbb{E}[t_{ij}] \right) \\
&\geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^n \sum_{j \in O_i} f_i(j) = \left(1 - \frac{1}{e}\right) \sum_{i=1}^n v_i(O_i)
\end{aligned}$$

$$= \left(1 - \frac{1}{e}\right) SW(\mathbf{O}).$$

□

Bayesian Nash Equilibrium. Consider any Bayesian bidding strategy \mathbf{B} and any valuation of player i , $v_i \sim D_i$, drawn independently from D_i . We denote by $\mathbf{C} = (C_1, C_2, \dots, C_n)$ the bidding distribution which includes both the randomness of both the bidding strategy $\mathbf{b} \sim \mathbf{B}$ and of the valuations $\mathbf{v} \sim \mathbf{D}$. It should be noted that \mathbf{C}_{-i} does *not* depend on some particular \mathbf{v}_{-i} , but merely on \mathbf{D}_{-i} and \mathbf{B}_{-i} . The utility of agent i with valuation v_i can be expressed by $u_i(\mathbf{B}_i(v_i), \mathbf{C}_{-i})$.

Let t_{ij} be the random variable indicating the maximum bid on item j of players other than i , i.e. $t_{ij} = \max_{k \neq i} b_{kj}$. However, now the randomisation of t_{ij} is derived by $\mathbf{b}_{-i} \sim \mathbf{C}_{-i}$. Let $F_{ij}(x)$ be the CDF of t_{ij} .

Given any valuation v_i , let $f_{v_i}^S$ be the maximising additive function for player i on set S . Similarly to Lemmas 5 and 10, we can prove the following.

Lemma 11. For any valuation v_i of player i , any pure strategy b'_i and any set of items S ,

$$\mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} [u_i^{v_i}(\mathbf{b})] = \mathbb{E}_{\mathbf{b} \sim (B_i(v_i), \mathbf{C}_{-i})} [u_i^{v_i}(\mathbf{b})] \geq \sum_{j \in S} F_{ij}(b'_{ij})(f_{v_i}^S(j) - b'_{ij}).$$

Theorem 12. The PoA of Bayesian Nash equilibria in simultaneous first price auctions is at most $\frac{e}{e-1}$.

Proof. Suppose B is any Bayesian Nash equilibrium and let t_{ij} be the random variable as defined above and F_{ij} be its CDF. We consider the following deviation for any player i . We fix the valuation of player i to be v_i and suppose that $\mathbf{w}_{-i} \sim \mathbf{D}_{-i}$. By Lemma 6, for any item j , there exists a value a_{ij} such that $F_{ij}(a_{ij})(f_{v_i}^{O_i(v_i, \mathbf{w}_{-i})}(j) - a_{ij}) + \mathbb{E}[t_{ij}] \geq (1 - 1/e)f_{v_i}^{O_i(v_i, \mathbf{w}_{-i})}(j)$. Then bid a_{ij} for every $j \in O_i(v_i, \mathbf{w}_{-i})$ and 0 for the rest of the items. Then, by applying Lemma 11 for $S = O_i(v_i, \mathbf{w}_{-i})$ and take the expectation over \mathbf{w}_{-i} ,

$$\mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} [u_i^{v_i}(\mathbf{b})] \geq \mathbb{E}_{\mathbf{w}_{-i}} \left[\sum_{j \in O_i(v_i, \mathbf{w}_{-i})} F_{ij}(a_{ij})(f_{v_i}^{O_i(v_i, \mathbf{w}_{-i})}(j) - a_{ij}) \right] \quad (4.6)$$

$$\geq \mathbb{E}_{\mathbf{w}_{-i}} \left[\sum_{j \in O_i(v_i, \mathbf{w}_{-i})} (1 - 1/e)f_{v_i}^{O_i(v_i, \mathbf{w}_{-i})}(j) - \mathbb{E}[t_{ij}] \right]. \quad (4.7)$$

By replacing \mathbf{w}_{-i} with \mathbf{v}_{-i} we get the following inequality,

$$\mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} [u_i^{v_i}(\mathbf{b})] + \mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} \left[\sum_{j \in O_i(\mathbf{v})} \mathbb{E}[t_{ij}] \right] \geq \left(1 - \frac{1}{e}\right) \mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} \left[\sum_{j \in O_i(\mathbf{v})} f_{v_i}(j) \right]. \quad (4.8)$$

We further give a lower bound on the total payments to be used next:

$$\begin{aligned} \sum_i \mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} [p_i(\mathbf{b})] &= \mathbb{E}_{\mathbf{b} \sim \mathbf{C}} \left[\sum_i p_i(\mathbf{b}) \right] = \mathbb{E}_{\mathbf{b} \sim \mathbf{C}} \left[\sum_j \max_k(b_{kj}) \right] \\ &= \mathbb{E}_{\substack{\mathbf{v} \sim \mathbf{D} \\ \mathbf{b} \sim \mathbf{C}}} \left[\sum_i \sum_{j \in O_i(\mathbf{v})} \max_k(b_{kj}) \right] \geq \sum_i \mathbb{E}_{\mathbf{v} \sim \mathbf{D}} \left[\sum_{j \in O_i(\mathbf{v})} \mathbb{E}[t_{ij}] \right]. \end{aligned} \quad (4.9)$$

By using (4.8) and (4.9) we are ready to complete the proof.

$$\begin{aligned} \mathbb{E}_{\mathbf{v} \sim \mathbf{B}(\mathbf{v})} [SW(\mathbf{b})] &= \sum_{i=1}^n \left(\mathbb{E}_{v_i} \left[\mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} [u_i^{v_i}(\mathbf{b})] \right] + \mathbb{E}_{\mathbf{v}} [p_i(\mathbf{b})] \right) \\ &\geq \sum_{i=1}^n \left(\mathbb{E}_{v_i} \left[\mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} [u_i^{v_i}(\mathbf{b})] \right] + \mathbb{E}_{\mathbf{v}} \left[\sum_{j \in O_i(\mathbf{v})} \mathbb{E}[t_{ij}] \right] \right) \\ &= \sum_{i=1}^n \left(\mathbb{E}_{v_i} \left[\mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} [u_i^{v_i}(\mathbf{b})] \right] + \mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} \left[\sum_{j \in O_i(\mathbf{v})} \mathbb{E}[t_{ij}] \right] \right) \\ &\geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^n \mathbb{E}_{v_i} \left[\mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} \left[\sum_{j \in O_i(\mathbf{v})} f_{v_i}(j) \right] \right] \\ &= \left(1 - \frac{1}{e}\right) \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} [v_i(O_i(\mathbf{v}))] \\ &= \left(1 - \frac{1}{e}\right) \mathbb{E}_{\mathbf{v}} [SW(\mathbf{O}(\mathbf{v}))]. \end{aligned}$$

□

4.1.2 Tight Lower Bound

Here we present a tight lower bound of $\frac{e}{e-1}$ for the PoA of mixed Nash equilibria in simultaneous first price auctions with OXS valuations. This implies a lower bound for submodular and fractionally subadditive (XOS) valuations. The following theorem (Theorem 13) is a special case of Theorem 22, however, in order

to make the presentation clearer we also give the proof of Theorem 13.

Theorem 13. The PoA of mixed Nash equilibria in simultaneous first price auctions with OXS valuations is at least $\frac{e}{e-1} \approx 1.58$.

Proof. We construct an instance with $n + 1$ players and n^n items. We define the set of items as $M = [n]^n$, that is, they correspond to all the different vectors $w = (w_1, w_2, \dots, w_n)$ with $w_i \in [n]$ (where $[n]$ denotes the set of integers $\{1, \dots, n\}$). Intuitively, they can be thought of as the nodes of an n dimensional grid, with coordinates in $[n]$ in each dimension.

We call player 0 the *dummy* player, and all other players $i \in [n]$ *real* players. We associate each *real* player i with one of the dimensions (directions) of the grid. In particular, for any fixed player i , her valuation for a subset of items $S \subseteq M$ is the size (number of elements) in the $n - 1$ -dimensional projection of S in direction i . Formally,

$$v_i(S) = |\{w_{-i} \mid \exists w_i \text{ s.t. } (w_i, w_{-i}) \in S\}|.$$

It is straightforward to check that v_i has decreasing marginal values, and is therefore submodular. It can be shown that those valuations are also OXS²¹. The *dummy* player 0 has valuation 0 for any subset of items.

Given these valuations, we describe a mixed Nash equilibrium $\mathbf{B} = (B_1, \dots, B_n)$ having a PoA arbitrarily close to $e/(e-1)$, for large enough n . The dummy player bids 0 for every item, and receives the item if all of the real players bid 0 for it. The utility and welfare of the dummy player is always 0. For real players the mixed strategy B_i is the following. Every player i picks a number $\ell \in [n]$ uniformly at random, and an x according to the distribution with CDF

$$G(x) = (n - 1) \left(\frac{1}{(1 - x)^{\frac{1}{n-1}}} - 1 \right),$$

where $x \in \left[0, 1 - \left(\frac{n-1}{n}\right)^{n-1}\right]$. Subsequently, she bids x for every item $w = (\ell, w_{-i})$, with $w_i = \ell$ as i^{th} coordinate, and bids 0 for the rest of the items, see Figure 4.2 for the cases of $n = 2$ and $n = 3$. That is, in any b_i in the support of B_i , the player bids a positive x only for an $n - 1$ dimensional slice of the items.

²¹In the definition of OXS valuations (Section 3.5.1), we set $k = n^{n-1}$ and for the unit-demand valuations corresponding to player i the following holds: if item j corresponds to $w = (w_1, w_2, \dots, w_n)$ then for each $r \in [k]$, $v_i^r(j) = 1$, if w_{-i} is the n -ary representation of r and $v_i^r(j) = 0$, otherwise.

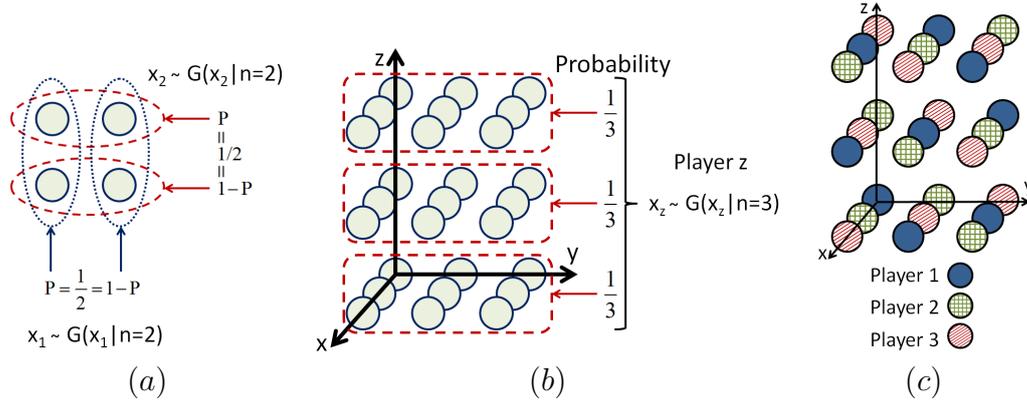


Figure 4.2: The figure illustrates the cases $n = 2$ and $n = 3$ ((a) and (b) respectively) for the lower bound example with OXS valuation functions. In (c) an optimal allocation for the case $n = 3$ is shown.

Observe that G has no mass points, so tie-breaking matters only in case of 0 bids for an item, in which case player 0 gets the item.

Let $F(x)$ denote the probability that bidder i gets a fixed item j , given that she bids $b_{ij} = x$ for this item, and the bids in \mathbf{b}_{-i} are drawn from \mathbf{B}_{-i} (due to symmetry, this probability is the same for all items $w = (\ell, w_{-i})$). For every other player k , the probability that she bids 0 for item j is $(n - 1)/n$, and the probability that j is in her selected slice but she bids lower than x is $G(x)/n$. Multiplying over the $n - 1$ other players, we obtain²²

$$F(x) = \left(\frac{G(x)}{n} + \frac{n - 1}{n} \right)^{n-1} = \frac{\left(\frac{n-1}{n} \right)^{n-1}}{1 - x}.$$

Notice that v_i is an additive valuation restricted to the slice of items that player i bids for in a particular b_i . Therefore, when player i bids x in b_i , her expected utility is $F(x)(1 - x)$ for one of these items, and comprising all items it is $\mathbb{E}[u_i(b_i)] = n^{n-1} F(x)(1 - x) = n^{n-1} \left(\frac{n-1}{n} \right)^{n-1} = (n - 1)^{n-1}$.

Next we show that \mathbf{B} is a Nash equilibrium. In particular, the bids b_i in the support of B_i maximise the expected utility of a fixed player i .

First, we fix an arbitrary w_{-i} , and focus on the set of items $C := \{(\ell, w_{-i}) \mid \ell \in [n]\}$, which we call a *column* for player i . Recall that i is interested in getting only one item within C , while her valuation is additive over items from different columns. Moreover, in a fixed \mathbf{b}_{-i} , every other player k submits the same bid for all items in C , because either the whole C is in the current slice of player k ,

²²Observe that $F(x)$, for large n , converges to $\hat{F}(x)$, which is the worst case distribution derived during the proof of the upper bound.

and she bids the same value x , or no item from the column is in the slice and she bids 0. Consider first a deviating bid, in which player i bids a positive value for more than one items in C , say (at least) the values $x \geq x' > 0$ where x is her highest bid in C . Then her expected utility for this column is strictly less than $F(x)(1-x)$, because her valuation is $F(x) \cdot 1$, but she might have to pay $x + x'$, in case she gets both items. Consequently, bidding x for only one item in C and 0 for the rest of C is more profitable.

Second, observe that restricted to a fixed column, submitting any bid $x \in \left[0, 1 - \left(\frac{n-1}{n}\right)^{n-1}\right]$ for one arbitrary item results in the constant expected utility of $\left(\frac{n-1}{n}\right)^{n-1}$, whereas a bid higher than $1 - \left(\frac{n-1}{n}\right)^{n-1}$ guarantees the item but pays more so the utility becomes strictly less than $\left(\frac{n-1}{n}\right)^{n-1}$ for this column. In summary, bidding for exactly one item from each column, an arbitrary (possibly different) bid $x \in \left[0, 1 - \left(\frac{n-1}{n}\right)^{n-1}\right]$ is a best response for i yielding the above expected utility, which concludes the proof that \mathbf{B} is a Nash equilibrium.

It remains to calculate the expected social welfare of \mathbf{B} , and the optimal social welfare. We define a random variable w.r.t. the distribution \mathbf{B} . Let $Z_j = 1$ if one of the real players $[n]$ gets item j , and $Z_j = 0$ if player 0 gets the item. Note that the social welfare is the random variable $\sum_{j \in M} Z_j$, and the expected social welfare is

$$\begin{aligned} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [SW(\mathbf{b})] &= \sum_j E[Z_j] \cdot 1 = n^n (1 - \mathbb{P}(\text{no real player bids for } j)) \\ &= n^n \left(1 - \left(\frac{n-1}{n}\right)^n\right). \end{aligned}$$

Finally, we show that the optimum social welfare is n^n . An optimal allocation can be constructed as follows: For each item (w_1, w_2, \dots, w_n) compute $r = (\sum_{i=1}^n w_i \bmod n)$. Allocate this item to the player $r + 1$. It is easy to see that this way the n items of *any* particular column $\{(\ell, w_{-k}) \mid \ell \in [n]\}$ (in any direction k) are given to the n different players, and that each player is allocated n^{n-1} items (Figure 4.2(c) shows the optimum allocation for $n = 3$). In other words, any two items allocated to the same player differ in at least two coordinates. In particular, they belong to different columns of *this player*, and all contribute 1 to the valuation of the player, which is therefore n^{n-1} . Since this valuation is maximum possible for every player, the obtained social welfare of n^n is optimal.

Thus, the PoA is $\frac{1}{(1 - (\frac{n-1}{n})^n)}$, and for large n it converges to $\frac{1}{(1 - \frac{1}{e})} \approx 1.58$. \square

4.2 Subadditive Valuations

We construct a lower bound of 2 on the PoA of mixed Nash equilibria when players have subadditive valuations. This lower bound matches the upper bound by Feldman et al. [64]. Theorem 14 is a special case of Theorem 27, however, in order to make the presentation clearer we also give the proof of Theorem 14.

Theorem 14. The PoA of mixed Nash equilibria in simultaneous first price auctions with subadditive valuations is at least 2.

Proof. Consider two players and m items with the following valuations: player 1 is a unit-demand player with valuation $v < 1$ (to be determined later) if she gets at least one item; player 2 has valuation 1 for getting at least one but less than m items, and 2 if she gets all the items.

Let \mathbf{B} be the following mixed bidding profile. Player 1 picks one of the m items uniformly at random, and bids x for this item and 0 for all other items. Player 2 bids y for each of the m items. The bids x and y are drawn from distributions with the following CDFs (inspired by [86]), respectively,

$$G(x) = \frac{(m-1)x}{1-x} \quad x \in [0, 1/m]; \quad F(y) = \frac{v-1/m}{v-y} \quad y \in [0, 1/m].$$

In the case of a tie, the item is always allocated to player 2. We are going to prove that \mathbf{B} is a mixed Nash equilibrium for every $v > 1/m$.

If player 1 bids any x in the range $(0, 1/m]$ for the one item, she gets the item with probability $F(x)$, since a tie appears with zero probability. Her expected utility for $x \in (0, 1/m]$ is $F(x)(v-x) = v-1/m$. Note that according to $G(x)$ she bids 0 with zero probability and thus her utility is still $v-1/m$ if she bids according to \mathbf{B} . Bidding something greater than $1/m$ results in a utility less than $v-1/m$. Regarding player 1, it remains to show that her utility while bidding for only one item is at least her utility while bidding for more items. Suppose player 1 bids x_j for each item j , $1 \leq j \leq m$ and w.l.o.g., assume that $x_j \geq x_{j+1}$. Player 1 gets at least one item if and only if $y < x_1$. So, with probability $F(x_1)$, she gets at least one item and she pays at least x_1 . Therefore, her expected utility is at most $F(x_1)(v-x_1) = v-1/m$, but it would be strictly less if she is charged nonzero payments for other items. This means that bidding only x_1 for one item and zero for the rest of them dominates the strategy we have assumed.

If player 2 bids a common bid y for all items, where $y \in [0, 1/m]$, she gets m items with probability $G(y)$ and $m - 1$ items with probability $1 - G(y)$. Her expected utility is therefore $G(y)(2 - my) + (1 - G(y))(1 - (m - 1)y) = G(y)(1 - y) + 1 - (m - 1)y = 1$. We show that player 2 cannot get a utility higher than 1 by using any deviating bids. Suppose now that player 2 bids y_j for each item j , for $1 \leq j \leq m$. Note first that player 2 has no incentive to bid higher than $1/m$ for any item, since bidding exactly $1/m$ would still guarantee the item and reduce the payment. Hence we assume that $y_j \leq 1/m$ and if player 1 chooses to bid for item j , player 2 receives item j with probability $G(y_j)$. Player 1 bids for item j (according to $G(x)$) with probability $1/m$. So, the expected utility of player 2 is

$$\begin{aligned}
& \frac{1}{m} \sum_{j=1}^m \left(G(y_j) \left(2 - \sum_{k=1}^m y_k \right) + (1 - G(y_j)) \left(1 - \sum_{\substack{k=1 \\ k \neq j}}^m y_k \right) \right) \\
&= \frac{1}{m} \sum_{j=1}^m \left(G(y_j)(1 - y_j) + 1 - \sum_{\substack{k=1 \\ k \neq j}}^m y_k \right) \\
&\leq \frac{1}{m} \sum_{j=1}^m \left(\frac{(m-1)y_j}{1-y_j} (1 - y_j) + 1 - \sum_{\substack{k=1 \\ k \neq j}}^m y_k \right) \\
&= \frac{1}{m} \sum_{j=1}^m \left(my_j + 1 - \sum_{k=1}^m y_k \right) \\
&= \frac{1}{m} \left(m \sum_{j=1}^m y_j + m - m \sum_{k=1}^m y_k \right) = 1.
\end{aligned}$$

Overall, we proved that \mathbf{B} is a mixed Nash equilibrium.

It is easy to see that the optimal allocation gives all items to player 2, and has social welfare of 2. In the Nash equilibrium \mathbf{B} , player 2 bids 0 with probability $1 - \frac{1}{mv}$, so, with at least this probability, player 1 gets one item. Therefore,

$$SW(\mathbf{B}) \leq \left(1 - \frac{1}{mv} \right) (v + 1) + \frac{1}{mv} 2 = 1 + v + \frac{1}{mv} - \frac{1}{m}$$

If we set $v = 1/\sqrt{m}$, then $SW(\mathbf{B}) \leq 1 + \frac{2}{\sqrt{m}} - \frac{1}{m}$. So, $\text{PoA} \geq \frac{2}{1 + \frac{2}{\sqrt{m}} - \frac{1}{m}}$ which, for

large m , converges to 2.

□

CHAPTER 5

Bid-Dependent Auctions

Here we generalise some results of the previous chapter to simultaneous bid-dependent auctions. Intuitively, a single item auction is *bid-dependent* if the winner is always the highest bidder, and a bidder's payment depends only on whether she gets the item or not, and on her *own* bid. For instance, the first-price auction and the all-pay auction are bid-dependent but the second-price auction is not.

We give upper bounds for the PoA in simultaneous bid-dependent auctions with fractionally subadditive bidder valuations. We parametrise those auctions based on their payment scheme. We use a single parameter θ , to be defined later, based on which we derive the upper bounds, which are between 1.58 and 2. For all those auctions we show a lower bound of 1.58, indicating that none of them can outperform the simultaneous first price auction.

Regarding subadditive valuations we prove a tight bound of 2 on the PoA of all simultaneous bid-dependent auctions.

5.1 Fractionally Subadditive Valuations

For a given simultaneous bid-dependent auction, we will denote by $q_j^w(x)$ and $q_j^l(x)$ a bidder's payment $p_{ij}(\mathbf{b})$ for item j when her bid for j is x , depending on whether she is the winner or a loser, respectively. Note that we assume $q_j^w(x)$ (resp. $q_j^l(x)$) to be the same for all bidders. Without this assumption the PoA is unbounded, as we show in the following example.

Example 15. Suppose there is a single item to be sold to two players with valuation $v_1 = 1$ and $v_2 = \epsilon$. The losing payment is 0 for both players but the winning

payments are different such that $q^w(x) = x$ for bidder 1 and $\bar{q}^w(x) = \epsilon \cdot x$ for bidder 2. If there is a tie, then the item is allocated to player 2. Now consider the bidding strategy $b_1 = b_2 = 1$. It is easy to see that it forms a Nash Equilibrium and has $PoA = 1/\epsilon$.

In order to guarantee the existence of reasonable Nash Equilibria, we also make the following natural assumptions about $q_j^w(x)$ and $q_j^l(x)$:²³

- $q_j^w(x)$ and $q_j^l(x)$ are non-decreasing, continuous functions of x and normalised, such that $q_j^l(0) = q_j^w(0) = 0$;
- $q_j^w(x) \geq q_j^l(x)$ for all $x \geq 0$;
- For any constant c , there exists some $x > 0$, such that $q_j^w(x) > c$ (to avoid the case that the payments are always less than the valuations, for which no Nash equilibria exist).

5.1.1 Upper Bounds

In this section we discuss the general upper bound for simultaneous bid-dependent auctions.

We define θ to be the worst case ratio between $q_j^l(x)$ and $q_j^w(x)$ over all j and x , i.e. $\theta = \max_{j \in [m]} \sup_{\{x: q_j^w(x) \neq 0\}} \{q_j^l(x)/q_j^w(x)\}$. If $q_j^l(x) = q_j^w(x) = 0$ for some x , we make the convention that $q_j^l(x)/q_j^w(x) = 0$.

Observe that $\theta \in [0, 1]$, due to the assumption that $q_j^l(x) \leq q_j^w(x)$. We will prove that (for $\theta \neq 1$) the PoA of coarse-correlated and Bayesian Nash equilibria of simultaneous bid-dependent auctions with fractionally subadditive bidders is at most $\frac{(\theta-1)^2}{\theta^2 - \theta + 1 - e^{\theta-1}}$. When we set $\theta = 0$ or $\theta \rightarrow 1$, we get back, respectively, the upper bounds of $e/(e-1)$ for first-price auctions, and 2 for all-pay auctions [133]. For simplicity, we define $1/\lambda(\theta) = \frac{(\theta-1)^2}{\theta^2 - \theta + 1 - e^{\theta-1}}$.

One Item

We start by proving a lemma for a single item, analogous to Lemma 6.

Lemma 16. Consider a single item bid-dependent auction with payment functions $q^w(x)$ and $q^l(x)$. Let \mathbf{B} be an arbitrary randomised bidding profile, and F_i denote the CDF of the random variable $t_i = \max_{k \neq i} b_k$, where $\mathbf{b} \sim \mathbf{B}$. Then for every bidder i , and non-negative value v , there exists a pure bidding strategy

²³Similar assumptions are also made in [20, 25, 104].

$a = a(v, \mathbf{B}_{-i})$ such that,

$$F_i(a) (v - q^w(a) + q^l(a)) - q^l(a) + \sum_{k \in [n]} p_k(\mathbf{B}) \geq \lambda(\theta) \cdot v,$$

where $p_k(\mathbf{B}) = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[p_k(\mathbf{b})]$ is the expected payment for player k .

Proof. Let $a = \arg \max_x \{F_i(x) (v - q^w(x) + q^l(x)) - q^l(x)\}$ and further let $A = F_i(a) (v - q^w(a) + q^l(a)) - q^l(a)$. In the following we use that F_i is the CDF of t_i , and since q^w is continuous and non-negative, $\mathbb{E}[q^w(t_i)] = \int_0^\infty (1 - F_i(x)) dq^w(x)$.

$$\begin{aligned} A + \sum_k p_k(\mathbf{B}) &\geq A + \mathbb{E}_{\mathbf{b}}[q^w(t_i)] \\ &= A + \int_0^\infty (1 - F_i(x)) dq^w(x) \\ &\geq A + \int_0^\infty \left(1 - \frac{A + q^l(x)}{v - q^w(x) + q^l(x)}\right) dq^w(x) \\ &\geq A + \int_0^\infty \left(\frac{v - A - q^w(x)}{v + (\theta - 1)q^w(x)}\right) dq^w(x) \\ &\geq A + \int_0^{v-A} \left(\frac{v - A - y}{v + (\theta - 1)y}\right) dy \end{aligned}$$

The second inequality follows from the definition of A and a and the third one is due to the fact that $q^l(x) \leq \theta \cdot q^w(x)$ for any x . For the last one, we use the fact that q^w is non-decreasing and show that there exists x_0 such that $q^w(x_0) \geq v - A$. That is true since for $x_0 \rightarrow \infty$, by definition of A , $A \geq F_i(x_0) (v - q^w(x_0) + q^l(x_0)) - q^l(x_0) = v - q^w(x_0)$ since $F_i(x_0) = 1$, meaning that $q^w(x_0) \geq v - A$. For completeness we also show that $v - A \geq 0$, by showing that $v \geq A \geq 0$: observe that $A = F_i(a)v - F_i(a)q^w(a) - (1 - F_i(a))q^l(a) \leq v$, since F_i is a CDF; moreover $A \geq F_i(0) (v - q^w(0) + q^l(0)) - q^l(0) = F_i(0)v \geq 0$.

In case $\theta < 1$,

$$A + \sum_k p_{kj}(\mathbf{B}) \geq A + \frac{(A + \theta(v - A))(\ln(A + \theta(v - A)) - \ln(v)) - (\theta - 1)(v - A)}{(\theta - 1)^2},$$

which is minimised for $A = \frac{v(\theta \cdot e^{1-\theta} - 1)}{(\theta - 1)e^{1-\theta}}$. The lemma follows by replacing A with this value.

In case $\theta = 1$, $A + \sum_k p_{kj}(\mathbf{B}) \geq A + \frac{(v-A)^2}{2v} \geq \frac{1}{2}v$, which coincides with the limit of $\lambda(\theta) = \frac{\theta^2 - \theta + 1 - e^{\theta-1}}{(\theta-1)^2}$ when $\theta \rightarrow 1$. \square

Many Items

In the following, let f_i^S be a maximising additive function of set S for player i with fractionally subadditive valuation function v_i . By the definition of fractionally subadditive valuations, we have that $v_i(T) \geq f_i^S(T)$, for every $T \subseteq S$ and $f_i^S(S) = v_i(S)$.

Lemma 17. For any set S of items, and any strategy profile \mathbf{b} , where $b_{ij} = 0$ for $j \notin S$,

$$u_i(\mathbf{b}) \geq \sum_{j \in S} \left(\mathbf{P}[j \in X_i(\mathbf{b})] (f_i^S(j) - q_j^w(b_{ij}) + q_j^l(b_{ij})) - q_j^l(b_{ij}) \right).$$

Proof.

$$\begin{aligned} u_i(\mathbf{b}) &\geq \sum_{T \subseteq S} \mathbf{P}[X_i(\mathbf{b}) = T] \left(f_i^S(T) - \sum_{j \in T} q_j^w(b_{ij}) - \sum_{j \in S \setminus T} q_j^l(b_{ij}) \right) \\ &= \sum_{T \subseteq S} \sum_{j \in T} \mathbf{P}[X_i(\mathbf{b}) = T] (f_i^S(j) - q_j^w(b_{ij})) - \sum_{T \subseteq S} \sum_{j \in S \setminus T} \mathbf{P}[X_i(\mathbf{b}) = T] q_j^l(b_{ij}) \\ &= \sum_{j \in S} \sum_{T \subseteq S: j \in T} \mathbf{P}[X_i(\mathbf{b}) = T] (f_i^S(j) - q_j^w(b_{ij})) \\ &\quad - \sum_{j \in S} \sum_{T \subseteq S: j \notin T} \mathbf{P}[X_i(\mathbf{b}) = T] q_j^l(b_{ij}) \\ &= \sum_{j \in S} \mathbf{P}[j \in X_i(\mathbf{b})] (f_i^S(j) - q_j^w(b_{ij})) - \sum_{j \in S} \mathbf{P}[j \notin X_i(\mathbf{b})] q_j^l(b_{ij}) \\ &= \sum_{j \in S} \left(\mathbf{P}[j \in X_i(\mathbf{b})] (f_i^S(j) - q_j^w(b_{ij}) + q_j^l(b_{ij})) - q_j^l(b_{ij}) \right). \end{aligned}$$

□

Coarse Correlated Equilibrium. We fix a coarse correlated equilibrium $\mathbf{B} = (B_1, B_2, \dots, B_n)$. For $\mathbf{b}_{-i} \sim \mathbf{B}_{-i}$, let $t_{ij} = \max_{k \neq i} b_{kj}$ be the random variable and $F_{ij}(x) = \mathbb{P}[t_{ij} \leq x]$ be the CDF of t_{ij} .

Lemma 18. For any set S of items, and any strategy profile b_i of player i , where $b_{ij} = 0$ for $j \notin S$,

$$u_i(\mathbf{B}) = \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [u_i(\mathbf{b})] \geq \sum_{j \in S} \left(F_{ij}(b'_{ij}) (f_i^S(j) - q_j^w(b'_{ij}) + q_j^l(b'_{ij})) - q_j^l(b'_{ij}) \right).$$

The proof is analogous to that of Lemma 5 by using Lemma 17.

Theorem 19. For bidders with fractionally subadditive valuations, the PoA of coarse correlated equilibria for bid-dependent auction is at most $1/\lambda(\theta)$.

Proof. Let \mathbf{B} be a coarse correlated equilibrium. For every player i , consider the maximising additive valuation, $f_i^{O_i}$ for his optimal set O_i . By Lemma 16, for every fixed player i and item j there exists a bid a_{ij} such that

$$F_{ij}(a_{ij}) (f_i^{O_i}(j) - q_j^w(a_{ij}) + q_j^l(a_{ij})) - q_j^l(a_{ij}) \geq \lambda(\theta) f_i^{O_i}(j) - \sum_k p_{kj}(\mathbf{B})$$

For player i , we consider the deviation where her bid is a_{ij} for every item in O_i and 0 for all other items, and apply Lemma 18. Combined with the above inequality (for all items in O_i), we obtain

$$u_i(\mathbf{B}) \geq \lambda(\theta) \sum_{j \in O_i} f_i^{O_i}(j) - \sum_{j \in O_i} \sum_k p_{kj}(\mathbf{B}) = \lambda(\theta) v_i(O_i^{\mathbf{v}}) - \sum_{j \in O_i} \sum_k p_{kj}(\mathbf{B}).$$

By summing over all players, we get

$$\sum_i u_i(\mathbf{B}) \geq \lambda(\theta) \sum_i v_i(O_i^{\mathbf{v}}) - \sum_{j \in [m]} \sum_k p_{kj}(\mathbf{B}) = \lambda(\theta) SW(\mathbf{O}) - \sum_k p_k(\mathbf{B}).$$

The theorem follows from $SW(\mathbf{B}) = \sum_i u_i(\mathbf{B}) + \sum_i p_i(\mathbf{B})$. \square

Bayesian Nash Equilibrium. Suppose that \mathbf{B} is a Bayesian Nash Equilibrium and the valuation of each player i is drawn according to $v_i \sim D_i$, where the D_i are independently distributed. We use the notation $\mathbf{C} = (C_1, C_2, \dots, C_n)$ to denote the bidding distribution in \mathbf{B} which involves the randomness of the valuations \mathbf{v} , and of the bidding strategy $\mathbf{B}(v)$, that is $b_i(v_i) \sim C_i$.

Let t_{ij} be the random variable indicating the maximum bid on item j of players other than i , i.e. $t_{ij} = \max_{k \neq i} b_{kj}$, when $\mathbf{b}_{-i} \sim \mathbf{C}_{-i}$. Let $F_{ij}(x)$ be the CDF of t_{ij} .

Similarly to Lemmas 5 and 18, we can prove the following.

Lemma 20. For any be an arbitrary set of items, S , and any player i with valuation v_i , let b'_i be a pure strategy such that $b'_{ij} = 0$ for $j \notin S$. Then,

$$\mathbb{E}_{\substack{\mathbf{v}_{-i} \\ \mathbf{b} \sim \mathbf{B}(\mathbf{v})}} [u_i^{v_i}(\mathbf{b})] = u_i^{v_i}(B_i(v_i), \mathbf{C}_{-i}) \geq \sum_{j \in S} (F_{ij}(b'_{ij})(f_{v_i}^S(j) - q_j^w(b'_{ij}) + q_j^l(b'_{ij})) - q_j^l(b'_{ij})).$$

Theorem 21. The PoA of Bayesian Nash equilibria for bid-dependent auction, when the bidders have fractionally subadditive, is at most $\lambda(\theta)$.

Proof. For any player i and any fractionally subadditive valuation $v_i \in V_i$, consider the following deviation: consider some $\mathbf{w}_{-i} \sim \mathbf{D}_{-i}$ and then for every $j \in O(v_i, \mathbf{w}_{-i})$ use the bid a_{ij} as defined in Lemma 16, such that

$$\begin{aligned} F_{ij}(a_{ij}) (f_{v_i}^{O(v_i, \mathbf{w}_{-i})}(j) - q_j^w(a_{ij}) + q_j^l(a_{ij})) - q_j^l(a_{ij}) + \sum_{k \in [n]} p_{kj}(B_i(v_i), \mathbf{C}_{-i}) \\ \geq \lambda(\theta) \cdot f_{v_i}^{O(v_i, \mathbf{w}_{-i})}(j). \end{aligned}$$

By applying Lemma 20 for $S = O_i(v_i, \mathbf{w}_{-i})$, taking expectation over v_i and \mathbf{w}_{-i} and summing over all players, we have that

$$\begin{aligned} & \sum_i \mathbb{E}_{\mathbf{v}}[u_i^{v_i}(\mathbf{B}(\mathbf{v}))] \\ &= \sum_i \mathbb{E}_{v_i} [u_i^{v_i}(B_i(v_i), \mathbf{C}_{-i})] \\ &\geq \sum_i \mathbb{E}_{v_i, \mathbf{w}_{-i}} \left[\sum_{j \in O_i(v_i, \mathbf{w}_{-i})} (F_{ij}(a_{ij}) (f_{v_i}^{O_i(v_i, \mathbf{w}_{-i})}(j) - q_j^w(a_{ij}) + q_j^l(a_{ij})) - q_j^l(a_{ij})) \right] \\ &= \sum_i \mathbb{E}_{\mathbf{v}} \left[\sum_{j \in O_i(\mathbf{v})} (F_{ij}(a_{ij}) (f_{v_i}^{O_i(\mathbf{v})}(j) - q_j^w(a_{ij}) + q_j^l(a_{ij})) - q_j^l(a_{ij})) \right] \\ &\geq \sum_i \mathbb{E}_{\mathbf{v}} \left[\sum_{j \in O_i(\mathbf{v})} \left(\lambda(\theta) \cdot f_{v_i}^{O_i(\mathbf{v})}(j) - \sum_k p_{kj}(B_i(v_i), \mathbf{C}_{-i}) \right) \right] \\ &= \lambda(\theta) \cdot \sum_i \mathbb{E}_{\mathbf{v}} [v_i(O_i^{\mathbf{v}})] - \sum_i \sum_j p_{kj}(\mathbf{C}). \end{aligned}$$

Then, $\mathbb{E}_{\mathbf{v}}[SW(\mathbf{B}(\mathbf{v}))] = \sum_i \mathbb{E}_{\mathbf{v}}[u_i^{v_i}(\mathbf{B}(\mathbf{v}))] + \sum_i \sum_j p_{kj}(\mathbf{C}) \geq \lambda(\theta) \cdot \mathbb{E}_{\mathbf{v}}[SW(\mathbf{O}^{\mathbf{v}})]$. \square

5.1.2 Lower Bound

Here we present a lower bound of $\frac{e}{e-1}$ for the PoA of all simultaneous bid-dependent auctions with OXS valuations and for mixed equilibria. This implies a lower bound for submodular and fractionally subadditive valuations, as well as for more general classes of equilibria.

Theorem 22. The PoA of simultaneous bid-dependent auctions with full information and OXS valuations is at least $\frac{e}{e-1} \approx 1.58$.

Proof. The proof is very similar to the one for simultaneous first price auctions (Section 4.1.2). Therefore, we only point out the differences. The same construction applies here; the only difference appears in the Nash equilibrium strategy profile.

We consider the same set of players and items as in Section 4.1.2; the valuation functions of the players are also the same as in Section 4.1.2, and the same tie breaking rule applies. As for the mixed Nash equilibrium \mathbf{B} , the dummy player still bids 0 for every item and every real player still picks an $n - 1$ dimensional slice in the same random way. However the bid x_j that she bids for every item j of that slice is drawn according to a distribution with the following *item-specific* CDF (we will show later that G_j is a valid CDF):

$$G_j(x) = n \left(\frac{\left(\frac{n-1}{n}\right)^{n-1} + q_j^l(x)}{1 - q_j^w(x) + q_j^l(x)} \right)^{\frac{1}{n-1}} - n + 1, \quad x \in [0, T_j].$$

where T_j is the bid such that $q_j^w(T_j) = 1 - \left(\frac{n-1}{n}\right)^{n-1}$; notice that due to our assumptions on q_j^w , there always exists such a value T_j . Note that we can no longer require that the bids of a player on different items are equal, since the CDFs G_j are different. Instead, we require that for every real player the bids x_j for different items in her slice are correlated in the following way: she chooses ρ uniformly at random from the interval $[0, 1]$, and then sets $x_j = G_j^{-1}(\rho)$, for every j in her slice. Note that for any two items j_1, j_2 of the slice, it holds that $G_{j_1}(x_1) = G_{j_2}(x_2) = \rho$ and x_{j_1} is not necessarily equal to x_{j_2} . However, for each item j in the slice, the way that x_j is chosen is equivalent to sampling it according to the CDF $G_j(x_j)$ (but in a correlated way to the other bids). The fact that each player's bids are such that the CDF values become equal, will be sufficient for proving that \mathbf{B} is a mixed Nash equilibrium.

The probability $F_j(x)$ that a player gets item j if she bids x for it is:

$$F_j(x) = \left(\frac{G_j(x)}{n} + \frac{n-1}{n} \right)^{n-1} = \frac{\left(\frac{n-1}{n}\right)^{n-1} + q_j^l(x)}{1 - q_j^w(x) + q_j^l(x)}, \quad x \in [0, T_j].$$

Recall that the valuation of player i is additive, restricted to the slice of items that she bids for in a particular b_i . Therefore, the expected utility of i when she bids x for item j is $F_j(x)(1 - q_j^w(x)) - (1 - F_j(x))q_j^l(x) = F_j(x)(1 - q_j^w(x) + q_j^l(x)) - q_j^l(x) = \left(\frac{n-1}{n}\right)^{n-1}$. By comprising all items, her expected utility for bidding $b_i \sim B_i$ is $\mathbb{E}[u_i(b_i)] = n^{n-1} \left(\frac{n-1}{n}\right)^{n-1} = (n-1)^{n-1}$.

Claim 23. \mathbf{B} is a Nash equilibrium.

Proof. First, we fix an arbitrary $w_{-i} \in [n]^{n-1}$, and focus on the set of items $C := \{(\ell, w_{-i}) \mid \ell \in [n]\}$, which we call a *column* for player i . Recall that i is interested in getting only one item within C , on the other hand his valuation is additive over items from different columns. Moreover, in a fixed \mathbf{b}_{-i} , every other player k submits bids x_j resulting in equal values of $G_j(x_j)$ for all items in C , because either the whole C is in the current slice of k , and he bids correlated bids on them, or no item from the column is in the slice and he bids 0.

Consider first a deviating bid, in which i bids a positive value for more than one items in C , say (at least) the values $x_1, x_2 > 0$ for items j_1, j_2 , respectively, and w.l.o.g. assume that $G_{j_1}(x_1) \geq G_{j_2}(x_2)$. We prove that if she loses item j_1 she should lose item j_2 as well: if she loses j_1 , then there must be a bidder k with bid $x'_1 > x_1$ for item j_1 , which implies $G_{j_1}(x'_1) > G_{j_1}(x_1)$. However, since the bids of player k are correlated (and j_2 is in his slice as well), for his bid x'_2 on j_2 it holds that $G_{j_2}(x'_2) = G_{j_1}(x'_1) > G_{j_1}(x_1) \geq G_{j_2}(x_2)$. Therefore, $x'_2 > x_2$ and player i cannot win item j_2 either, so bidding for item j_2 cannot contribute to the valuation, whereas the bidder might pay for more items than j_1 . Consequently, bidding for only one item in C and 0 for the rest of C is more profitable.

Second, observe that restricted to a fixed column, submitting any bid $x \in [0, T_j]$ for one arbitrary item j results in the constant expected utility of $\left(\frac{n-1}{n}\right)^{n-1}$, whereas by bidding higher than T_j the utility would be at most $1 - q_j^w(T_j) = \left(\frac{n-1}{n}\right)^{n-1}$ for this column. In summary, bidding for exactly one item j from each column, an arbitrary bid $x \in [0, T_j]$ is a best response for i yielding the above expected utility, which concludes the proof that \mathbf{B} is a Nash equilibrium. \square

The rest of the argument is exactly the same as in the proof for first price auctions. It remains to prove that the G_j 's are valid cumulative distribution functions. To this end it is sufficient to show that $G_j(T_j) = 1$ and that $G_j(x)$ is non-decreasing in $[0, T_j]$. For simplicity we skip index j .

$$\begin{aligned} G(T) &= n \left(\frac{\left(\frac{n-1}{n}\right)^{n-1} + q^l(T)}{1 - q^w(T) + q^l(T)} \right)^{\frac{1}{n-1}} - n + 1 \\ &= n \left(\frac{\left(\frac{n-1}{n}\right)^{n-1} + q^l(T)}{1 - \left(1 - \left(\frac{n-1}{n}\right)^{n-1}\right) + q^l(T)} \right)^{\frac{1}{n-1}} - n + 1 = 1. \end{aligned}$$

Now let $x_1, x_2 \in [0, T]$, and $x_1 > x_2$. In order to prove $G(x_1) \geq G(x_2)$, it is sufficient to prove that $\frac{\left(\frac{n-1}{n}\right)^{n-1} + q^l(x_1)}{1 - q^w(x_1) + q^l(x_1)} \geq \frac{\left(\frac{n-1}{n}\right)^{n-1} + q^l(x_2)}{1 - q^w(x_2) + q^l(x_2)}$.

$$\begin{aligned}
& \frac{\left(\frac{n-1}{n}\right)^{n-1} + q^l(x_1)}{1 - q^w(x_1) + q^l(x_1)} - \frac{\left(\frac{n-1}{n}\right)^{n-1} + q^l(x_2)}{1 - q^w(x_2) + q^l(x_2)} \\
= & \frac{q^l(x_1)(1 - \left(\frac{n-1}{n}\right)^{n-1} - q^w(x_2)) - q^l(x_2)(1 - \left(\frac{n-1}{n}\right)^{n-1} - q^w(x_1))}{(1 - q^w(x_1) + q^l(x_1))(1 - q^w(x_2) + q^l(x_2))} \\
& + \frac{\left(\frac{n-1}{n}\right)^{n-1} (q^w(x_1) - q^w(x_2))}{(1 - q^w(x_1) + q^l(x_1))(1 - q^w(x_2) + q^l(x_2))} \\
\geq & \frac{q^l(x_1)(1 - \left(\frac{n-1}{n}\right)^{n-1} - q^w(x_2)) - q^l(x_2)(1 - \left(\frac{n-1}{n}\right)^{n-1} - q^w(x_1))}{(1 - q^w(x_1) + q^l(x_1))(1 - q^w(x_2) + q^l(x_2))} \\
\geq & \frac{(q^l(x_1) - q^l(x_2))(1 - \left(\frac{n-1}{n}\right)^{n-1} - q^w(x_1))}{(1 - q^w(x_1) + q^l(x_1))(1 - q^w(x_2) + q^l(x_2))} \geq 0
\end{aligned}$$

The inequalities follow from the monotonicity of q^w and q^l , and from the fact that $1 - \left(\frac{n-1}{n}\right)^{n-1} \geq q^w(x_1)$ holds by the definition of T .

This completes the proof of Theorem 22. \square

5.2 Subadditive Valuations

Our results for subadditive valuations hold even for a class of auctions more general than bid-dependent auctions that were described in Section 5.1: we allow the payment rule to depend on the *rank* of the bid, where the r^{th} highest bid (with an arbitrary tie-breaking rule) has rank r . We use $q_j(x, r)$ to denote the payment that the bidder should pay for item j when her bid is x and gets rank r . In particular, given a mixed bidding strategy \mathbf{B} , bidder i 's expected payment $p_{ij}(\mathbf{B})$ for item j equals to $\mathbb{E}_{\mathbf{b}}[q_j(b_{ij}, r_i(\mathbf{b}^j))]$ where $r_i(\mathbf{b}^j)$ denotes the rank of b_{ij} among $\{b_{1j}, \dots, b_{nj}\}$. That leads to player i 's payment, $p_i(\mathbf{B}) = \sum_{j \in [m]} \mathbb{E}_{\mathbf{b}}[q_j(b_{ij}, r_i(\mathbf{b}^j))]$. Note that $q_j^w(x)$ from the previous subsection is $q_j(x, 1)$ here, and $q_j^l(x)$ can be different for different ranks. Analogous assumptions to the ones made on $q_j^w(x)$ and $q_j^l(x)$ can be made on $q_j(x, r)$ as well.

We next show that the PoA for mixed Nash, correlated, coarse-correlated and Bayesian equilibria is exactly 2.

5.2.1 Upper Bound

Lemma 24. For any simultaneous bid-dependent auction, subadditive valuation profile \mathbf{v} and randomised bidding profile \mathbf{B} , there exists a randomised bid vector a_i for each player i , sampled from a distribution $A_i(\mathbf{v}, \mathbf{B}_{-i})$, such that for the total expected utility and expected payments of the bidders it holds that

$$\sum_i u_i(A_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) \geq \frac{1}{2} \sum_i v_i(O_i^{\mathbf{v}}) - \sum_i \sum_j p_{ij}(\mathbf{B}),$$

where $O_i^{\mathbf{v}}$ is the optimal set of player i .

Proof. Under the profile \mathbf{v} , $O_i^{\mathbf{v}}$ is the set of items allocated to player i in the optimum. We denote by $h_j(\mathbf{b}) = \arg \max_i b_{ij}$ the bidder with the highest bid for item j , regarding the pure bidding \mathbf{b} . Let t_{ij} be the maximum of bids for item j among players other than i , and t_i be the vector such that its j^{th} coordinate equals t_{ij} if $j \in O_i^{\mathbf{v}}$, and 0 otherwise. Note that $t_i \sim T_i$ is an induced random variable of \mathbf{B}_{-i} . We define the randomised bid a_i to follow the same distribution T_i , i.e. $A_i(\mathbf{v}, \mathbf{B}_{-i}) = T_i$ (inspired by [64]).

We use the notation $v_i(b_i, t_i)$ and $W_i(b_i, t_i)$ to denote the player i 's valuation and winning set when she bids b_i and the prices are t_i , i.e., $v_i(S)$ and $W_i(S)$, respectively, where $S = \{j | b_{ij} \geq t_{ij}\}$.

$$\begin{aligned} & u_i(A_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) \\ &= \mathbb{E}_{a_i \sim A_i} \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}} [u_i(a_i, \mathbf{b}_{-i})] \\ &\geq \mathbb{E}_{a_i \sim T_i} \mathbb{E}_{t_i \sim T_i} [v_i(a_i, t_i)] - \sum_{j \in O_i^{\mathbf{v}}} \mathbb{E}_{a_i \sim T_i} [q_j(a_{ij}, 1)] \quad (\text{since } q_j(x, 1) \geq q_j(x, r)) \\ &= \mathbb{E}_{t_i \sim T_i} \mathbb{E}_{a_i \sim T_i} [v_i(t_i, a_i)] - \sum_{j \in O_i^{\mathbf{v}}} \mathbb{E}_{t_i \sim T_i} [q_j(t_{ij}, 1)] \quad (\text{swap } t_i \text{ and } a_i) \\ &= \frac{1}{2} \mathbb{E}_{t_i \sim T_i} \mathbb{E}_{a_i \sim T_i} [v_i(t_i, a_i) + v_i(a_i, t_i)] - \sum_{j \in O_i^{\mathbf{v}}} \mathbb{E}_{t_i \sim T_i} [q_j(t_{ij}, 1)] \\ &\geq \frac{1}{2} v_i(O_i^{\mathbf{v}}) - \sum_{j \in O_i^{\mathbf{v}}} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [q_j(b_{h_j(\mathbf{b})}(j), 1)] \\ &\geq \frac{1}{2} v_i(O_i^{\mathbf{v}}) - \sum_{j \in O_i^{\mathbf{v}}} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [q_j(b_{h_j(\mathbf{b})}(j), 1)] \\ &- \sum_{j \in O_i^{\mathbf{v}}} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} \left[\sum_k q_j(b_k(j), r_k(\mathbf{b}^j)) - q_j(b_{h_j(\mathbf{b})}(j), r_{h_j(\mathbf{b})}(\mathbf{b}^j)) \right] \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{2}v_i(O_i^{\mathbf{v}}) - \sum_{j \in O_i^{\mathbf{v}}} \sum_k \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [q_j(b_{kj}, r_k(\mathbf{b}^j))] \\
&= \frac{1}{2}v_i(O_i^{\mathbf{v}}) - \sum_{j \in O_i^{\mathbf{v}}} \sum_k p_{kj}(\mathbf{B}).
\end{aligned}$$

In the second inequality $v_i(t_i, a_i) + v_i(a_i, t_i) \geq v_i(O_i^{\mathbf{v}})$ is due to the subadditivity²⁴ of v_i and also $q_j(t_{ij}, 1) \leq q_j(b_{h_j(\mathbf{b})}(j), 1)$ since $q_j(\cdot, 1)$ is non-decreasing and $t_{ij} \leq b_{h_j(\mathbf{b})}(j)$. For the last inequality, notice that $\sum_k q_j(b_k(j), r_k(\mathbf{b}^j)) - q_j(b_{h_j(\mathbf{b})}(j), r_{h_j(\mathbf{b})}(\mathbf{b}^j)) \geq 0$, since from the sum of all payments for item j we subtracted the payment of the winner. The lemma follows by summing over all players. \square

Theorem 25. For bidders with subadditive valuations, the PoA of coarse correlated equilibria for any bid-dependent auction is at most 2.

Proof. Suppose \mathbf{B} is a coarse correlated equilibrium (notice that \mathbf{v} is fixed). By Lemma 24 and the definition of coarse correlated equilibrium, we have

$$\sum_i u_i(\mathbf{B}) \geq \sum_i u_i(A_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) \geq \frac{1}{2} \sum_i v_i(O_i^{\mathbf{v}}) - \sum_i \sum_j p_{ij}(\mathbf{B}).$$

By rearranging the terms, $SW(\mathbf{B}) = \sum_i u_i(\mathbf{B}) + \sum_i \sum_j p_{ij}(\mathbf{B}) \geq 1/2 \cdot SW(\mathbf{O})$. \square

Theorem 26. For bidders with subadditive valuations, the PoA of Bayesian Nash equilibria for any bid-dependent auction is at most 2.

Proof. Suppose \mathbf{B} is a Bayesian Nash Equilibrium and the valuation of each player i is $v_i \sim D_i$, where the D_i are independently distributed. We use the notation $\mathbf{C} = (C_1, C_2, \dots, C_n)$ to denote the bidding distribution in \mathbf{B} which includes the randomness of the valuations \mathbf{v} , and of the bidding strategy \mathbf{b} (like in the proof of Theorem 21). For any player i and any subadditive valuation $v_i \in V_i$, consider the following deviation: sample $\mathbf{w}_{-i} \sim D_{-i}$ and bid according to $A_i((v_i, \mathbf{w}_{-i}), \mathbf{C}_{-i})$ as defined in Lemma 24. By the definition of Nash equilibrium, we have $\mathbb{E}_{\mathbf{v}_{-i}}[u_i^{v_i}(B_i(v_i), \mathbf{B}_{-i}(\mathbf{v}_{-i}))] \geq \mathbb{E}_{\mathbf{w}_{-i}}[u_i^{v_i}(A_i((v_i, \mathbf{w}_{-i}), \mathbf{C}_{-i}), \mathbf{C}_{-i})]$. By taking expectation over v_i and summing over all players,

²⁴The inequality holds if $t_i \neq a_i$. If $t_i = a_i$ happens with non-zero probability, only for those cases, we increase the deviating bid by an arbitrarily small value.

$$\begin{aligned}
\sum_i \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{B}(\mathbf{v}))] &\geq \sum_i \mathbb{E}_{v_i, \mathbf{w}_{-i}} [u_i^{v_i}(A_i((v_i, \mathbf{w}_{-i}), \mathbf{C}_{-i}), \mathbf{C}_{-i})] \\
&= \mathbb{E}_{\mathbf{v}} \left[\sum_i u_i^{v_i}(A_i(\mathbf{v}, \mathbf{C}_{-i}), \mathbf{C}_{-i}) \right] \text{ (by relabelling } \mathbf{w}_{-i} \text{ by } \mathbf{v}_{-i}) \\
&\geq \frac{1}{2} \cdot \sum_i \mathbb{E}_{\mathbf{v}}[v_i(O_i^y)] - \sum_i \sum_j \mathbb{E}_{\mathbf{v}}[p_{kj}(\mathbf{B}(\mathbf{v}))] \\
&= \frac{1}{2} \cdot \mathbb{E}_{\mathbf{v}}[SW(\mathbf{O}^y)] - \sum_i \sum_j \mathbb{E}_{\mathbf{v}}[p_{kj}(\mathbf{B}(\mathbf{v}))],
\end{aligned}$$

where the last inequality follows by Lemma 24. Finally, we obtain $\mathbb{E}_{\mathbf{v}}[SW(\mathbf{B}(\mathbf{v}))] = \sum_i \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{B})] + \sum_i \sum_j \mathbb{E}_{\mathbf{v}}[p_{kj}(\mathbf{B}(\mathbf{v}))] \geq 1/2 \cdot \mathbb{E}_{\mathbf{v}}[SW(\mathbf{O}^y)]$. \square

5.2.2 Lower Bound

The lower bound holds even for only two different types of payments, q_j^w in the case of winning item j and q_j^l in the case of losing item j .

Theorem 27. For bidders with subadditive valuations, the PoA of mixed Nash equilibria for all simultaneous bid-dependent auctions is at least 2.

Proof. We consider 2 players and m items. Let v be a positive real value to be defined later. Player 1 has value v for every non-empty subset of items; player 2 values with 1 any non-empty strict subset of the items and with 2 the whole set of items. Consider now the mixed strategy profile \mathbf{B} , where player 1 picks item l uniformly at random and bids x_l for it and 0 for the rest of the items, whereas, player 2 bids y_j for every item j . For $1 \leq j \leq m$, x_j and y_j are drawn from distributions with the following CDFs $G_j(x)$ and $F_j(y)$, respectively:

$$G_j(x) = \frac{(m-1)q_j^w(x) + q_j^l(x)}{1 - q_j^w(x) + q_j^l(x)}, \quad x \in [0, T_j];$$

$$F_j(y) = \frac{v - 1/m + q_j^l(y)}{v - q_j^w(y) + q_j^l(y)}, \quad y \in [0, T_j],$$

where T_j is the bid such that $q_j^w(T_j) = 1/m$; notice that due to the assumptions on q_j^w , there always exists such a value T_j . Furthermore, in \mathbf{B} , the y_j 's are correlated in the following way: player 2 chooses ρ uniformly at random from the interval $[0, 1]$ and if $\rho \in \left[0, \frac{v-1/m}{v}\right)$ then $y_j = 0$, otherwise $y_j = F_j^{-1}(\rho)$,

for every $1 \leq j \leq m$.²⁵ Note that for every two items j_1, j_2 , it holds that $F_{j_1}(y_{j_1}) = F_{j_2}(y_{j_2})$. In case of a tie, player 2 gets the item. Due to the continuity of q_j^w and q_j^l , $G_j(x)$ and $F_j(x)$ are continuous and therefore none of the CDF have a mass point in any $x \neq 0$.

We next show that \mathbf{B} is a Nash equilibrium, and each of the F_j and G_j are valid cumulative distributions. The PoA bound can be then derived as follows. Player 2 bids 0 with probability $1 - \frac{1}{mv}$ so, $\mathbb{E}[SW(\mathbf{B})] \leq (1 - \frac{1}{mv})(1+v) + \frac{1}{mv}2 = 1 + v + \frac{1}{mv} - \frac{1}{m}$. For $v = 1/\sqrt{m}$, $\text{PoA} \geq \frac{2}{1 + \frac{2}{\sqrt{m}} - \frac{1}{m}}$ which, for large m , converges to 2.

Claim 28. \mathbf{B} is a Nash equilibrium.

Proof. If player 1 bids any x in the range of $(0, T_j]$ for a single item j and zero for the rest, her utility is $F_j(x)(v - q_j^w(x)) + (1 - F_j(x))(-q_j^l(x)) = F_j(x)(v - q_j^w(x) + q_j^l(x)) - q_j^l(x) = v - 1/m$. Since $G(0) = 0$, her utility is also $v - 1/m$ if she bids according to G . Suppose player 1 bids $x = (x_1, \dots, x_m)$, ($x_j \in [0, T_j]$) with at least two positive bids. W.l.o.g., assume $F_1(x_1) = \max_i F_i(x_i)$. If $y_1 \geq x_1$, player 1 doesn't get any item, since for every j , $F_j(y_j) = F_1(y_1) \geq F_1(x_1) \geq F_j(x_j)$ and so $y_j \geq x_j$ (recall that in any tie player 2 gets the item). If $y_1 < x_1$, player 1 gets at least the first item and has valuation v , but she cannot pay less than $q_1^w(x_1)$. So, this strategy is dominated by the strategy of bidding x_1 for the first item and zero for the rest. Bidding $x_j > T_j$ for any item guarantees the item but results in a utility less than $v - q_j^w(x_j) \leq v - q_j^w(T_j) = v - 1/m$, so it is dominated by the strategy of bidding exactly T_j for this item.

If player 2 bids any (y_1, \dots, y_m) such that $y_j \in [0, T_j]$, then (since player 1 bids positive for any particular item j with probability $1/m$) her expected utility is

$$\begin{aligned} & \frac{1}{m} \sum_{j=1}^m \left(G_j(y_j) \left(2 - \sum_{k=1}^m q_k^w(y_k) \right) + (1 - G_j(y_j)) \left(1 - \sum_{\substack{k=1 \\ k \neq j}}^m q_k^w(y_k) - q_j^l(y_j) \right) \right) \\ &= \frac{1}{m} \sum_{j=1}^m \left(1 + G_j(y_j) (1 - q_j^w(y_j) + q_j^l(y_j)) - \sum_{\substack{k=1 \\ k \neq j}}^m q_k^w(y_k) - q_j^l(y_j) \right) \end{aligned}$$

²⁵For each item j , the way player 2 chooses y_j is equivalent to picking it according to the CDF $F_j(y)$.

$$\begin{aligned}
&= \frac{1}{m} \sum_{j=1}^m \left(1 + (m-1)q_j^w(y_j) + q_j^l(y_j) - \sum_{\substack{k=1 \\ k \neq j}}^m q_k^w(y_k) - q_j^l(y_j) \right) \\
&= \frac{1}{m} \left(m + m \sum_{j=1}^m q_j^w(y_j) - m \sum_{k=1}^m q_k^w(y_k) \right) = 1.
\end{aligned}$$

Bidding greater than T_j for any item is dominated by the strategy of bidding exactly T_j for this item. Overall, \mathbf{B} is Nash equilibrium. \square

Claim 29. G_j and F_j are valid cumulative distributions.

Proof. It is sufficient to show that for every j , $G_j(T_j) = F_j(T_j) = 1$ and $G_j(x)$ and $F_j(x)$ are non-decreasing in $[0, T_j]$. In the following we skip index j .

$$\begin{aligned}
G(T) &= \frac{(m-1)q^w(T) + q^l(T)}{1 - q^w(T) + q^l(T)} = \frac{(m-1)\frac{1}{m} + q^l(T)}{1 - \frac{1}{m} + q^l(T)} = 1, \\
F(T) &= \frac{v - 1/m + q^l(T)}{v - q^w(T) + q^l(T)} = \frac{v - 1/m + q^l(T)}{v - 1/m + q^l(T)} = 1.
\end{aligned}$$

Now let $x_1 > x_2$, $x_1, x_2 \in [0, T_j]$. Then,

$$\begin{aligned}
&G(x_1) - G(x_2) \\
&= \frac{(m-1)q^w(x_1) + q^l(x_1)}{1 - q^w(x_1) + q^l(x_1)} - \frac{(m-1)q^w(x_2) + q^l(x_2)}{1 - q^w(x_2) + q^l(x_2)} \\
&= \frac{(m-1)(q^w(x_1) - q^w(x_2)) + m(q^w(x_1)q^l(x_2) - q^l(x_1)q^w(x_2))}{(1 - q^w(x_1) + q^l(x_1))(1 - q^w(x_2) + q^l(x_2))} \\
&\quad + \frac{q^l(x_1) - q^l(x_2)}{(1 - q^w(x_1) + q^l(x_1))(1 - q^w(x_2) + q^l(x_2))} \\
&= \frac{(m-1 + mq^l(x_2))(q^w(x_1) - q^w(x_2)) + m(\frac{1}{m} - q^w(x_2))(q^l(x_1) - q^l(x_2))}{(1 - q^w(x_1) + q^l(x_1))(1 - q^w(x_2) + q^l(x_2))} \\
&\geq 0;
\end{aligned}$$

$$\begin{aligned}
&F_j(x_1) - F_j(x_2) \\
&= \frac{v - 1/m + q^l(x_1)}{v - q^w(x_1) + q^l(x_1)} - \frac{v - 1/m + q^l(x_2)}{v - q^w(x_2) + q^l(x_2)} \\
&= \frac{(v - \frac{1}{m})(q^w(x_1) - q^w(x_2)) + \frac{1}{m}(q^l(x_1) - q^l(x_2)) + q^w(x_1)q^l(x_2) - q^l(x_1)q^w(x_2)}{(v - q^w(x_1) + q^l(x_1))(v - q^w(x_2) + q^l(x_2))}
\end{aligned}$$

$$= \frac{(v - \frac{1}{m} + q^l(x_2))(q^w(x_1) - q^w(x_2)) + (\frac{1}{m} - q^w(x_2))(q^l(x_1) - q^l(x_2))}{(v - q^w(x_1) + q^l(x_1))(v - q^w(x_2) + q^l(x_2))} \geq 0.$$

For both inequalities we use the monotonicity of q^w and q^l and the fact that $q^w(x) \leq 1/m$ and $q^l(x) \leq v = 1/\sqrt{m}$ for all $x \in [0, T]$. □

This finishes the proof of Theorem 27. □

CHAPTER 6

Discriminatory auctions

In this chapter, we study a special case of combinatorial auction, where all items are identical or units of the same item, called *multi-unit* auction. The players have a valuation for each number of units and they express their preferences by a bid vector where the j^{th} entry represents their willingness to pay for an extra unit when they have already paid for $j - 1$ units. We focus on a specific multi-unit auction, the *discriminatory* auction, where the units are allocated to the m highest bids and each player pays her bids for the number of units she receives, i.e. she pays the aggregation of her first k bids when she receives k units.

Here, we complement the results by de Keijzer et al. [56] for the case of subadditive valuations, by providing a matching lower bound of 2 for the standard bidding format. For the case of submodular valuations, we provide a lower bound of 1.099. As a technical contribution, we reprove the upper bound of $e/(e-1)$ due to [56] for submodular bids, using a similar approach to the previous chapters. Due to the different nature of this auction, the proof is not identical with the one for the first-price auction. Therefore, we present the complete proof of this upper bound.

We next introduce some extra notation regarding multi-unit auctions, to be used in the upper bound proof.

6.1 Preliminaries

Consider a discriminatory auction with submodular valuations, with n players and m items. Recall that $v_i(j)$ denotes the valuation of player i for j copies of the item. For any player i , we define $v_{ij} = \frac{v_i(j)}{j}$. It is easy to see that for submodular

functions, $v_{ij} \geq v_{i(j+1)}$ for all $j \in [m-1]$. Recall that we consider the standard bidding format used in [56, 100], where the bids are in decreasing order, i.e. $b_{ij} \geq b_{i(j+1)}$. Let $\beta_j(\mathbf{b})$ be the j^{th} lowest bid among the *winning* bids under the strategy profile \mathbf{b} . Consider any randomised bidding profile $\mathbf{B} = (B_1, \dots, B_n)$. For this \mathbf{B} , $\beta_j(\mathbf{b})$ is a random variable depending on $\mathbf{b} \sim \mathbf{B}$. We define the following functions:

$$\begin{aligned} F_{ij}(x) &= \mathbf{P}[\beta_j(\mathbf{b}_{-i}) \leq x], & \text{for } 1 \leq j \leq m; \\ G_{ij}(x) &= \mathbf{P}[\beta_j(\mathbf{b}_{-i}) \leq x < \beta_{j+1}(\mathbf{b}_{-i})] = F_{ij}(x) - F_{i(j+1)}(x), & \text{for } 1 \leq j \leq m-1. \end{aligned}$$

We define separately $G_{im}(x) = \mathbf{P}[\beta_m(\mathbf{b}_{-i}) \leq x] = F_{im}(x)$. Notice that $F_{ij}(x)$ is the CDF of $\beta_j(\mathbf{b}_{-i})$; moreover it holds that

$$\begin{aligned} F_{ij}(x) &= \sum_{k=j}^m G_{ik}(x), \\ \sum_{j=1}^{m'} F_{ij}(x) &= \sum_{j=1}^{m'} j G_{ij}(x) + \sum_{j=m'+1}^m m' G_{ij}(x). \end{aligned} \quad (6.1)$$

We further define $F_i^{\text{av}}(x) = \frac{1}{o_i^{\mathbf{v}}} \sum_{j=1}^{o_i^{\mathbf{v}}} F_{ij}(x)$, and let β_i^{av} be a random variable with $F_i^{\text{av}}(x)$ as CDF. $F_i^{\text{av}}(x)$ is a cumulative distribution function defined on \mathbb{R}^+ , since $F_i^{\text{av}}(0) = 0$, $\lim_{x \rightarrow +\infty} (F_i^{\text{av}}(x)) = 1$ and $F_i^{\text{av}}(x)$ is the average of non-decreasing functions, so it is itself a non-decreasing function. Moreover,

$$\begin{aligned} \mathbb{E}[\beta_i^{\text{av}}] &= \int_0^\infty (1 - F_i^{\text{av}}(x)) dx = \int_0^\infty \left(1 - \frac{1}{o_i^{\mathbf{v}}} \sum_{j=1}^{o_i^{\mathbf{v}}} F_{ij}(x)\right) dx \\ &= \frac{1}{o_i^{\mathbf{v}}} \sum_{j=1}^{o_i^{\mathbf{v}}} \int_0^\infty (1 - F_{ij}(x)) dx = \frac{1}{o_i^{\mathbf{v}}} \sum_{j=1}^{o_i^{\mathbf{v}}} \beta_j(\mathbf{B}_{-i}), \end{aligned}$$

where $\beta_j(\mathbf{B}_{-i}) = \mathbb{E}_{\mathbf{b}_{-i} \sim \mathbf{B}_{-i}}[\beta_j(\mathbf{b}_{-i})]$. Note that the above functions depend on some randomised bidding profile \mathbf{B}_{-i} and on \mathbf{v} . These will be clear from context when we use these functions below.

6.2 Submodular Valuations

We first give an alternative proof of the upper bound of $e/(e-1)$ proved by [56] for submodular valuations and then give our lower bound of 1.099.

6.2.1 Upper Bound

We start with a technical lemma, which will be used in order to prove our main results.

Lemma 30. For any submodular valuation profile \mathbf{v} and any randomised bidding profile \mathbf{B} , there exists a pure bidding strategy $\mathbf{a}_i(\mathbf{v}, \mathbf{B}_{-i})$ for each player i , such that:

$$\sum_{i=1}^n u_i(\mathbf{a}_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) \geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^n v_i(o_i^{\mathbf{v}}) - \sum_{j=1}^m \beta_j(\mathbf{B}).$$

Proof. Recall that $v_{ij} = \frac{v_i(j)}{j}$. Let a_i be the value that maximizes $(v_{io_i^{\mathbf{v}}} - a_i)F_i^{\text{av}}(a_i)$. Let $\mathbf{a}_i(\mathbf{v}, \mathbf{B}_{-i}) = (\underbrace{a_i, \dots, a_i}_{o_i^{\mathbf{v}}}, \underbrace{0, \dots, 0}_{m-o_i^{\mathbf{v}}})$ be the selected strategy profile for player i .

Observe that by the definition of G_{ij} , $G_{ij}(a_i)$ is the probability of a_i being the j^{th} lowest bid among winning bids under \mathbf{B}_{-i} . Therefore, if player i bids according to $\mathbf{a}_i(\mathbf{v}, \mathbf{B}_{-i})$, $G_{ij}(a_i)$ is the probability of player i getting exactly j items, if $j \leq o_i^{\mathbf{v}}$, and $o_i^{\mathbf{v}}$ items, if $j > o_i^{\mathbf{v}}$, under the bidding profile $(\mathbf{a}_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i})$. Similarly to Lemma 5, we get

$$\begin{aligned} u_i(\mathbf{a}_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) &\geq \sum_{j=1}^{o_i^{\mathbf{v}}} G_{ij}(a_i)(v_i(j) - ja_i) + \sum_{j=o_i^{\mathbf{v}}+1}^m G_{ij}(a_i)(v_i(o_i^{\mathbf{v}}) - o_i^{\mathbf{v}}a_i) \\ &= \sum_{j=1}^{o_i^{\mathbf{v}}} jG_{ij}(a_i)(v_{ij} - a_i) + \sum_{j=o_i^{\mathbf{v}}+1}^m o_i^{\mathbf{v}}G_{ij}(a_i)(v_{io_i^{\mathbf{v}}} - a_i) \\ &\geq (v_{io_i^{\mathbf{v}}} - a_i) \left(\sum_{j=1}^{o_i^{\mathbf{v}}} jG_{ij}(a_i) + \sum_{j=o_i^{\mathbf{v}}+1}^m o_i^{\mathbf{v}}G_{ij}(a_i) \right) \\ &= (v_{io_i^{\mathbf{v}}} - a_i) \sum_{j=1}^{o_i^{\mathbf{v}}} F_{ij}(a_i) = o_i^{\mathbf{v}}(v_{io_i^{\mathbf{v}}} - a_i)F_i^{\text{av}}(a_i) \\ &\geq \left(1 - \frac{1}{e}\right) o_i^{\mathbf{v}}v_{io_i^{\mathbf{v}}} - o_i^{\mathbf{v}}\mathbb{E}[\beta_i^{\text{av}}] \end{aligned}$$

$$\begin{aligned}
&= \left(1 - \frac{1}{e}\right) v_i(o_i^{\mathbf{v}}) - o_i^{\mathbf{v}} \mathbb{E}[\beta_i^{\text{av}}] \\
&= \left(1 - \frac{1}{e}\right) v_i(o_i^{\mathbf{v}}) - \sum_{j=1}^{o_i^{\mathbf{v}}} \beta_j(\mathbf{B}_{-i})
\end{aligned}$$

For the second inequality, $v_{ij} \geq v_{io_i}$ for submodular valuations and for the following equality, we used (6.1) where m' is set to $o_i^{\mathbf{v}}$. For the last inequality we apply Lemma 6, since a_i maximises the expression $(v_{io_i^{\mathbf{v}}} - a_i)F_i^{\text{av}}(a_i)$.

For any pure strategy profile \mathbf{b} and any valuation profile \mathbf{v} it holds that

$$\sum_{j=1}^m \beta_j(\mathbf{b}) \geq \sum_{i=1}^n \sum_{j=1}^{o_i^{\mathbf{v}}} \beta_j(\mathbf{b}) \geq \sum_{i=1}^n \sum_{j=1}^{o_i^{\mathbf{v}}} \beta_j(\mathbf{b}_{-i}).$$

By summing up over all players:

$$\begin{aligned}
\sum_{i=1}^n u_i(\mathbf{a}_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) &\geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^n v_i(o_i^{\mathbf{v}}) - \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} \left[\sum_{i=1}^n \sum_{j=1}^{o_i^{\mathbf{v}}} \beta_j(\mathbf{b}_{-i}) \right] \\
&\geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^n v_i(o_i^{\mathbf{v}}) - \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} \left[\sum_{j=1}^m \beta_j(\mathbf{b}) \right] \\
&= \left(1 - \frac{1}{e}\right) \sum_{i=1}^n v_i(o_i^{\mathbf{v}}) - \sum_{j=1}^m \beta_j(\mathbf{B}).
\end{aligned}$$

□

Theorem 31. The PoA of coarse correlated equilibria for the discriminatory auction is at most $\frac{e}{e-1}$, when the players' valuations are submodular.

Proof. Suppose \mathbf{B} is a coarse correlated equilibrium (in this case \mathbf{v} is fixed). By Lemma 30 and the definition of coarse correlated equilibrium, we have that

$$\begin{aligned}
\sum_{i=1}^n u_i(\mathbf{B}) &\geq \sum_{i=1}^n u_i(\mathbf{a}_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) \\
&\geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^n v_i(o_i^{\mathbf{v}}) - \sum_{j=1}^m \beta_j(\mathbf{B})
\end{aligned}$$

After rearranging the terms $SW(\mathbf{B}) = \sum_i u_i(\mathbf{B}) + \sum_j \beta_j(\mathbf{B}) \geq \left(1 - \frac{1}{e}\right) SW(\mathbf{o})$.

□

Theorem 32. The PoA of Bayesian Nash equilibria for the discriminatory auction is at most $\frac{e}{e-1}$, when the players' valuations are submodular.

Proof. Suppose \mathbf{B} is a Bayesian Nash Equilibrium and the valuation of each player i is $v_i \sim D_i$, where the D_i 's are independent distributions. We denote by $\mathbf{C} = (C_1, C_2, \dots, C_n)$ the bidding distribution in \mathbf{B} which includes the randomness of both the bidding strategy \mathbf{b} and of the valuations \mathbf{v} . For any agent i and any submodular valuation $v_i \in V_i$, consider the following deviation: sample $\mathbf{w}_{-i} \sim \mathbf{D}_{-i}$ and bid $\mathbf{a}_i((v_i, \mathbf{w}_{-i}), \mathbf{C}_{-i})$ as defined in Lemma 30. By the definition of the Bayesian Nash equilibrium, we have

$$\mathbb{E}_{\mathbf{v}_{-i}} [u_i^{v_i}(\mathbf{B}_i(v_i), \mathbf{B}_{-i}(\mathbf{v}_{-i}))] \geq \mathbb{E}_{\mathbf{w}_{-i}} [u_i^{v_i}(\mathbf{a}_i((v_i, \mathbf{w}_{-i}), \mathbf{C}_{-i}), \mathbf{C}_{-i})]$$

By taking expectation over v_i and summing up over all agents,

$$\begin{aligned} & \sum_{i=1}^n \mathbb{E}_{\mathbf{v}} [u_i(\mathbf{B}(\mathbf{v}))] \\ & \geq \sum_{i=1}^n \mathbb{E}_{v_i, \mathbf{w}_{-i}} [u_i^{v_i}(\mathbf{a}_i((v_i, \mathbf{w}_{-i}), \mathbf{C}_{-i}), \mathbf{C}_{-i})] \\ & = \mathbb{E}_{\mathbf{v}} \left[\sum_{i=1}^n u_i^{v_i}(\mathbf{a}_i(\mathbf{v}, \mathbf{C}_{-i}), \mathbf{C}_{-i}) \right] \text{ (by relabelling } \mathbf{w}_{-i} \text{ by } \mathbf{v}_{-i}) \\ & \geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^n v_i(o_i^{\mathbf{v}}) - \sum_{j=1}^m \beta_j(\mathbf{C}) \\ & = \left(1 - \frac{1}{e}\right) \sum_{i=1}^n v_i(o_i^{\mathbf{v}}) - \sum_{j=1}^m \mathbb{E}_{\mathbf{v}} [\beta_j(\mathbf{B}(\mathbf{v}))] \end{aligned}$$

So, $\mathbb{E}_{\mathbf{v}}[SW(\mathbf{B}(\mathbf{v}))] = \sum_i \mathbb{E}_{\mathbf{v}}[u_i(\mathbf{B}(\mathbf{v}))] + \sum_j \mathbb{E}_{\mathbf{v}}[\beta_j(\mathbf{B}(\mathbf{v}))] \geq (1 - \frac{1}{e}) \mathbb{E}_{\mathbf{v}}[SW(\mathbf{o}^{\mathbf{v}})]$. \square

6.2.2 Lower Bound

Theorem 33. The PoA of mixed Nash equilibria for discriminatory auctions with submodular valuations is at least 1.099.

Proof. We design a game with two players and two identical items. Player 1 has valuation (v, v) , i.e., her valuation is v if she gets one or two items; whereas player 2 has valuation $(1, 2)$, i.e., he is additive with value 1 for each item. We use the following distribution functions defined by Hassidim et al. [86]:

$$G(x) = \frac{x}{1-x}, \quad x \in [0, 1/2]; \quad F(y) = \frac{v-1/2}{v-y}, \quad y \in [0, 1/2].$$

Consider the following mixed strategy profile. Player 1 bids $(x, 0)$ and player 2 bids (y, y) , where x and y are drawn from $G(x)$ and $F(y)$, respectively. Noting that player 2 bids 0 with probability $F(0) = 1 - 1/2v$, we need a tie-breaking rule for the case of bidding 0, in which player 2 always gets the item. We claim that this mixed strategy profile is a mixed Nash equilibrium.

First we prove that playing $(x, 0)$ for player 1 is a best response for every $x \in [0, 1/2]$. Notice that the bid (x, x') with $x' \leq x$, is dominated by $(x, 0)$, since if player 1 gets at least one item, she should pay at least x and getting both items doesn't add to her utility.

$$u_1(x, 0) = F(x) \cdot (v - x) = v - 1/2.$$

Clearly, bidding higher than $1/2$ guarantees the item but the payment is higher.

Now we need to show that (y, y) is a best response for player 2, for every $y \in [0, 1/2]$. Consider any strategy (y, y') with $y, y' \in [0, 1/2]$ and $y \geq y'$.

$$\begin{aligned} \mathbb{E}[u_2(y, y')] &= \mathbf{P}[x \leq y'](2 - y - y') + \mathbf{P}[x > y'](1 - y) \\ &= G(y')(2 - y - y') + (1 - G(y'))(1 - y) \\ &= G(y')(1 - y') + 1 - y = 1 + y' - y \leq 1. \end{aligned}$$

Note that $\mathbb{E}[u_2(y, y)] = 1$ is maximum possible utility. Bidding strictly higher than $1/2$ for one or both items is not profitable, since her utility would be less than 1.

Now we calculate the expected social welfare of this Nash equilibrium.

$$\begin{aligned} \mathbb{E}[SW] &= \mathbf{P}[y \geq x]2 + \mathbf{P}[x > y](1 + v) \\ &= 2 - (1 - v)\mathbf{P}[x > y] \\ &= 2 - (1 - v) \int_0^{1/2} F(x) dG(x) \end{aligned}$$

This expression is maximised for $v = 0.643$. For this value of v , $\mathbb{E}[SW] = 1.818$. Since $SW(\mathbf{o}) = 2$, we get $\text{PoA} = 1.099$. \square

6.3 Subadditive Valuations

For subadditive valuations, de Keijzer et al. [56] showed an upper bound of 2 on the PoA of discriminatory auction, under the standard bidding format. We provide a tight lower bound of 2 which is similar to the lower bound of Section 4.2, adjusted to discriminatory auctions.

Theorem 34. For discriminatory auctions the PoA of mixed Nash equilibria is at least 2 for bidders with subadditive valuations.

Proof. Consider two players and m items with the following valuations: player 1 is a unit-demand player with valuation $v < 1$ if she gets at least one item; player 2 has valuation 1 for getting less than m but at least one items, and 2 if she gets all the items. Inspired by [86], we use the following distribution functions:

$$G(x) = \frac{(m-1)x}{1-x}, \quad x \in [0, 1/m]; \quad F(y) = \frac{v-1/m}{v-y}, \quad y \in [0, 1/m].$$

Consider the following mixed strategy profile, \mathbf{B} : player 1 bids $b_1 = (x, 0, \dots, 0)$ and player 2 bids $b_2 = (y, \dots, y)$, where x and y are drawn from $G(x)$ and $F(y)$, respectively. In case of a tie, the item is always allocated to player 2. We are going to prove that \mathbf{B} is a mixed Nash equilibrium for every $v > 1/m$.

If player 1 bids any x in the range $(0, 1/m]$ for one item, she gets the item with probability $F(x)$, since a tie occurs with zero probability. Her expected utility is $F(x)(v-x) = v-1/m$. So, for every $x \in (0, 1/m]$ her utility is $v-1/m$. If player 1 picks x according to $G(x)$, her utility is still $v-1/m$, since she bids 0 with zero probability. Bidding something greater than $1/m$ results in a utility less than $v-1/m$. Regarding player 1, it remains to show that her utility when bidding for only one item is at least as high as her utility when bidding for more items. Suppose player 1 bids (x_1, \dots, x_m) , where $x_i \geq x_{i+1}$, for $1 \leq i \leq m-1$. Player 1 doesn't get any item if and only if $y \geq x_1$. So, with probability $F(x_1)$, she gets at least one item and she pays at least x_1 . Therefore, her expected utility is at most $F(x_1)(v-x_1) = v-1/m$, but it would be strictly less if she had nonzero payments for other items with positive probability. This means that bidding only x_1 for one item and zero for the rest of them dominates the strategy (x_1, \dots, x_m) .

If player 2 bids y for all items, where $y \in [0, 1/m]$, she gets m items with probability $G(y)$ and $m-1$ items with probability $1-G(y)$. Her expected utility is $G(y)(2-my) + (1-G(y))(1-(m-1)y) = G(y)(1-y) + 1 - (m-1)y = 1$. Bidding

something greater than $1/m$ is dominated by bidding exactly $1/m$. Suppose now that player 2 bids (y_1, \dots, y_m) , where $y_i \geq y_{i+1}$ for $1 \leq i \leq m-1$. If $x \leq y_m$, player 2 gets all the items; otherwise she gets $m-1$ items and she pays her $m-1$ highest bids. So, her utility is

$$\begin{aligned}
& G(y_m) \left(2 - \sum_{i=1}^m y_i \right) + (1 - G(y_m)) \left(1 - \sum_{i=1}^{m-1} y_i \right) \\
&= G(y_m)(1 - y_m) + 1 - \sum_{i=1}^{m-1} y_i \\
&= my_m + 1 - \sum_{i=1}^m y_i \\
&\leq my_m + 1 - \sum_{i=1}^m y_m = 1.
\end{aligned}$$

Overall, we proved that \mathbf{B} is a mixed Nash equilibrium. It is easy to see that the social welfare in the optimum allocation is 2. In this Nash equilibrium, player 2 bids 0 with probability $1 - \frac{1}{mv}$, so, with at least this probability, player 1 gets one item.

$$SW(\mathbf{B}) \leq \left(1 - \frac{1}{mv} \right) (v + 1) + \frac{1}{mv} 2 = 1 + v + \frac{1}{mv} - \frac{1}{m}$$

If we set $v = 1/\sqrt{m}$, then $SW(\mathbf{B}) \leq 1 + \frac{2}{\sqrt{m}} - \frac{1}{m}$. So, $PoA \geq \frac{2}{1 + \frac{2}{\sqrt{m}} - \frac{1}{m}}$ which, for large m , converges to 2. \square

CHAPTER 7

Proportional Allocation Mechanism

We study the PoA of auctioning divisible resources by the simultaneous proportional allocation mechanism when the players have concave and subadditive valuations.

For the full information setting with concave valuations, Johari and Tsitsiklis [91] showed that there exists a unique pure Nash equilibrium with PoA of $4/3$ and there is no other mixed Nash equilibrium. Regarding the Bayesian setting, Caragiannis and Voudouris [30] showed for a single resource an upper bound of 2. We complement those results by showing that for the Bayesian setting with concave valuations over many resources, the PoA is at least $\sqrt{m}/2$ (Section 7.1), where m is the number of resources.

Then, in Section 7.2.1, we study subadditive valuations and show that the PoA is exactly 2 for all studied concepts of equilibria: pure Nash, mixed Nash, correlated, coarse correlated and Bayesian Nash equilibria. At last in Sections 7.2.2 and 7.2.3 we show that the proportional allocation mechanism admits the best PoA among a range of mechanisms.

7.1 Concave Valuations

In this section, we show that for concave valuations on multiple resources, Bayesian equilibria can be arbitrarily inefficient. More precisely, we prove that the PoA is $\Omega(\sqrt{m})$ in contrast to the constant bound for pure equilibria [91]. Therefore, there is a big gap between full and incomplete information settings.

Theorem 35. When the valuations are concave, the PoA of the proportional allocation mechanism for Bayesian equilibria is at least $\frac{\sqrt{m}}{2}$.

Proof. We consider an instance with m resources and 2 players with the following concave valuations. $v_1(\mathbf{x}) = \min_j \{x_j\}$ and $v_2(\mathbf{x})$ is drawn from a distribution D_2 , such that some resource $j \in [m]$ is chosen uniformly at random and then $v_2(\mathbf{x}) = x_j/\sqrt{m}$.

Let $\delta = 1/(\sqrt{m} + 1)^2$. We claim that $\mathbf{b}(\mathbf{v}) = (b_1, b_2(v_2))$ is a pure Bayesian Nash equilibrium, where $b_{1j} = \sqrt{\delta/m} - \delta$, $\forall j \in [m]$, and, if $j^* \in [m]$ is the resource chosen by D_2 , $b_{2j^*}(v_2) = \delta$ and $b_{2j} = 0$, for all $j \neq j^*$.

Under this bidding profile, player 1 bids the same value for all resources, and player 2 only bids positive value for a single resource associated with her valuation. Suppose that player 2 has positive valuation for resource j^* , i.e., $v_2(\mathbf{x}) = x_{j^*}/\sqrt{m}$. Then the rest $m - 1$ resources are allocated to player 1 and players are competing for resource j^* . Bidder 2 has no reason to bid positively for any other resource. If she bids any value b'_{2j^*} for resource j^* , her utility would be

$$u_2(\mathbf{b}_1, b'_{2j}) = \frac{1}{\sqrt{m}} \frac{b'_{2j^*}}{b_{1j^*} + b'_{2j^*}} - b'_{2j^*},$$

which is maximised for

$$b'_{2j^*} = \sqrt{\frac{b_{1j^*}}{\sqrt{m}}} - b_{1j^*}.$$

For $b_{1j^*} = \sqrt{\delta/m} - \delta$, the utility of player 2 is maximised for $b'_{2j^*} = 1/(\sqrt{m} + 1)^2 = \delta$ by simple calculations.

Since $v_1(\mathbf{x})$ equals the minimum of \mathbf{x} 's components, player 1's valuation is completely determined by the allocation of resource j^* . So the expected utility of player 1 under \mathbf{b} is

$$\mathbb{E}_{v_2}[u_1(\mathbf{b})] = \frac{\sqrt{\delta/m} - \delta}{\sqrt{\delta/m} - \delta + \delta} - m(\sqrt{\delta/m} - \delta) = (1 - \sqrt{m}\delta)^2 = \frac{1}{(\sqrt{m} + 1)^2} = \delta.$$

Suppose now that player 1 deviates to $b'_1 = (b'_{11}, \dots, b'_{1m})$. Then, her expected utility would be

$$\mathbb{E}_{v_2}[u_1(b'_1, b_2)] = \frac{1}{m} \sum_j \frac{b'_{1j}}{b'_{1j} + \delta} - \sum_j b'_{1j} = \frac{1}{m} \sum_j \left(\frac{b'_{1j}}{b'_{1j} + \delta} - m \cdot b'_{1j} \right)$$

$$\begin{aligned}
&\leq \frac{1}{m} \sum_j \left(\frac{\sqrt{\delta/m} - \delta}{\sqrt{\delta/m}} - m \cdot (\sqrt{\delta/m} - \delta) \right) \\
&= \frac{1}{m} \sum_j \left(1 - 2\sqrt{m \cdot \delta} + m \cdot \delta \right) = \frac{1}{m} \sum_j \left(1 - \sqrt{m \cdot \delta} \right)^2 \\
&= \frac{1}{m} \sum_j \left(\frac{1}{\sqrt{m} + 1} \right)^2 = \delta = \mathbb{E}_{v_2}[u_1(\mathbf{b})].
\end{aligned}$$

The inequality comes from the fact that $\frac{b'_{1j}}{b'_{1j} + \delta} - m \cdot b'_{1j}$ is maximised for $b'_{1j} = \sqrt{\delta/m} - \delta$. So we conclude that \mathbf{b} is a Bayesian equilibrium.

Finally, we compute the PoA. The expected social welfare under \mathbf{b} is

$$\begin{aligned}
\mathbb{E}_{v_2}[\text{SW}(\mathbf{b})] &= \frac{\sqrt{\delta/m} - \delta}{\sqrt{\delta/m} - \delta + \delta} + \frac{1}{\sqrt{m}} \frac{\delta}{\sqrt{\delta/m} - \delta + \delta} \\
&= 1 - \sqrt{m\delta} + \sqrt{\delta} = \frac{2}{\sqrt{m} + 1} < \frac{2}{\sqrt{m}}.
\end{aligned}$$

The optimal social welfare though is 1 by allocating to player 1 all resources. Therefore, $\text{PoA} \geq \frac{\sqrt{m}}{2}$. \square

7.2 Subadditive Valuations

In this section, we focus on players with subadditive valuations. We prove that the proportional allocation mechanism is at least 50% efficient for coarse correlated equilibria and Bayesian Nash equilibria, i.e., $\text{PoA} \leq 2$. We further show that this bound is tight and cannot be improved by any simple or scale-free mechanism. The upper bound is an improvement upon the 3.73 bound of Syrgkanis and Tardos [133] for the class of lattice submodular valuations which is a subclass of subadditive valuations²⁶.

Before proving our PoA bounds, we show that the class of subadditive functions is a superclass of lattice submodular functions. Based on the definition of Syrgkanis and Tardos [133], a valuation function is lattice submodular if and only if it is submodular on the product lattice of outcomes, i.e. $\forall x, y \in [0, 1]^m : v(x \vee y) + v(x \wedge y) \leq v(x) + v(y)$, where \vee and \wedge stand for the supremum (join) and infimum (meet), respectively.

²⁶The class of concave valuations is not a subclass of the subadditive valuations [122] (as the result of Section 7.1 indicates), however in the case of a single variable (single resource), the classes of lattice submodular and concave valuations coincide and therefore concave valuations are also subadditive.

Proposition 36. Any lattice submodular function v defined on $[0, 1]^m$ is sub-additive.

Proof. It has been shown in [133] that for any lattice submodular function $v(\mathbf{x})$, $\frac{\partial^2 v(\mathbf{x})}{(\partial x_j)^2} \leq 0$ and $\frac{\partial^2 v(\mathbf{x})}{\partial x_j \partial x_{j'}} \leq 0$. So the function $\frac{\partial v}{\partial x_j}(\mathbf{x})$ is non-increasing monotone for each coordinate $x_{j'}$. It suffices to prove that for any $\mathbf{x}, \mathbf{y} \in [0, 1]^m$, $v(\mathbf{x} + \mathbf{y}) - v(\mathbf{y}) \leq v(\mathbf{x}) - v(\mathbf{0})$. Let \mathbf{z}^k be the vector that $z_j^k = y_j$ if $j \leq k$ and $x_j + y_j$ otherwise. Note that $\mathbf{z}^0 = \mathbf{x} + \mathbf{y}$ and $\mathbf{z}^m = \mathbf{y}$. Similarly, we define \mathbf{w}^k to be the vector that $w_j^k = 0$ if $j \leq k$ and x_j otherwise. It is easy to see that $\mathbf{z}^k \geq \mathbf{w}^k$ for all $k \in [m]$. So we have,

$$\begin{aligned} v(\mathbf{x} + \mathbf{y}) - v(\mathbf{y}) &= \sum_{j \in [m]} v(\mathbf{z}^{j-1}) - v(\mathbf{z}^j) = \sum_{j \in [m]} \int_{y_j}^{x_j + y_j} \frac{\partial v}{\partial t_j}(t_j; \mathbf{z}_{-j}^j) dt_j \\ &\leq \sum_{j \in [m]} \int_{y_j}^{x_j + y_j} \frac{\partial v}{\partial t_j}(t_j - y_j; \mathbf{z}_{-j}^j) dt_j \leq \sum_{j \in [m]} \int_0^{x_j} \frac{\partial v}{\partial s_j}(s_j; \mathbf{w}_{-j}^j) ds_j = v(\mathbf{x}) - v(\mathbf{0}). \end{aligned}$$

The second equality is due to the definition of partial derivative and the inequalities is due to the monotonicity of $\frac{\partial v}{\partial x_j}(x)$. \square

7.2.1 Upper Bound

Similarly to the previous chapters, in order to prove PoA upper bounds we define a deviation with proper utility bounds and then use the definition of Nash equilibrium to bound players' utilities at equilibrium.

Lemma 37. Let \mathbf{v} be any subadditive valuation profile and \mathbf{B} be some randomised bidding profile. For any player i , there exists a randomised bidding strategy $a_i(\mathbf{v}, \mathbf{B}_{-i})$ such that:

$$\sum_i u_i(a_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) \geq \frac{1}{2} SW(\mathbf{o}^v) - \sum_i \sum_j \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [b_{ij}].$$

Proof. Let q_{ij} be the sum of the bids of all players except i on resource j , i.e., $q_{ij} = \sum_{k \neq i} b_{kj}$. Note that q_{ij} is a random variable that depends on $\mathbf{b}_{-i} \sim \mathbf{B}_{-i}$. Let Q_i be the probability distribution of $q_i = (q_{ij})_j$. We consider the bidding strategy $a_i(\mathbf{v}, \mathbf{B}_{-i}) = (o_{ij}^v \cdot b'_{ij})_j$, where $b'_i \sim Q_i$. Then,

$$u_i(a_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i})$$

$$\begin{aligned}
&= \mathbb{E}_{b'_i \sim Q_i} \mathbb{E}_{q_i \sim Q_i} \left[v_i \left(\left(\frac{o_{ij}^{\mathbf{v}} b'_{ij}}{o_{ij}^{\mathbf{v}} b'_{ij} + q_{ij}} \right)_j \right) - o_i^{\mathbf{v}} \cdot b'_i \right] \\
&\geq \frac{1}{2} \cdot \mathbb{E}_{q_i \sim Q_i} \mathbb{E}_{b'_i \sim Q_i} \left[v_i \left(\left(\frac{o_{ij}^{\mathbf{v}} b'_{ij}}{o_{ij}^{\mathbf{v}} b'_{ij} + q_{ij}} + \frac{o_{ij}^{\mathbf{v}} q_{ij}}{o_{ij}^{\mathbf{v}} q_{ij} + b'_{ij}} \right)_j \right) \right] - \mathbb{E}_{q_i \sim Q_i} [o_i^{\mathbf{v}} \cdot q_i] \\
&\geq \frac{1}{2} \cdot \mathbb{E}_{q_i \sim Q_i} \mathbb{E}_{b'_i \sim Q_i} \left[v_i \left(\left(\frac{o_{ij}^{\mathbf{v}} (b'_{ij} + q_{ij})}{b'_{ij} + q_{ij}} \right)_j \right) \right] - \mathbb{E}_{q_i \sim Q_i} [o_i^{\mathbf{v}} \cdot q_i] \\
&= \frac{1}{2} \cdot v_i(o_i^{\mathbf{v}}) - \sum_j \sum_{k \neq i} \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [o_{ij}^{\mathbf{v}} \cdot b_{kj}] \\
&\geq \frac{1}{2} \cdot v_i(o_i^{\mathbf{v}}) - \sum_j \sum_k \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [o_{ij}^{\mathbf{v}} \cdot b_{kj}].
\end{aligned}$$

The first inequality follows by swapping q_{ij} and b'_{ij} and using the subadditivity of v_i . The second inequality comes from the fact that $o_{ij}^{\mathbf{v}} \leq 1$. By summing up over all players and by using the fact that $\sum_{i \in [n]} o_{ij}^{\mathbf{v}} = 1$,

$$\begin{aligned}
\sum_i u_i(a_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) &\geq \frac{1}{2} \sum_i v_i(o_i^{\mathbf{v}}) - \sum_i \sum_j \sum_k \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [o_{ij}^{\mathbf{v}} \cdot b_{kj}] \\
&= \frac{1}{2} SW(\mathbf{o}^{\mathbf{v}}) - \sum_k \sum_j \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} \left[\sum_i (o_{ij}^{\mathbf{v}}) \cdot b_{kj} \right] \\
&= \frac{1}{2} SW(\mathbf{o}^{\mathbf{v}}) - \sum_i \sum_j \mathbb{E}_{\mathbf{b} \sim \mathbf{B}} [b_{ij}].
\end{aligned}$$

□

Theorem 38. The PoA of coarse correlated equilibria for the proportional allocation mechanism with subadditive valuations is at most 2.

Proof. Let \mathbf{B} be any coarse correlated equilibrium (note that \mathbf{v} is fixed). By Lemma 37 and the definition of the coarse correlated equilibrium, we have

$$\sum_i u_i(\mathbf{B}) \geq \sum_i u_i(a_i(\mathbf{v}, \mathbf{B}_{-i}), \mathbf{B}_{-i}) \geq \frac{1}{2} SW(\mathbf{o}) - \sum_i \sum_j \mathbb{E}[b_{ij}]$$

By rearranging terms, $SW(\mathbf{B}) = \sum_i u_i(\mathbf{B}) + \sum_i \sum_j \mathbb{E}[b_{ij}] \geq \frac{1}{2} \cdot SW(\mathbf{o})$. □

Theorem 39. The Bayesian PoA of the proportional allocation mechanism with subadditive players is at most 2.

Proof. Let \mathbf{B} be any Bayesian Nash Equilibrium and let $v_i \sim D_i$ be the valuation of each player i drawn independently from D_i . We denote by $\mathbf{C} =$

(C_1, C_2, \dots, C_n) the bidding distribution which includes the randomness of both the bidding strategy $\mathbf{b} \sim \mathbf{B}$ and of the valuations \mathbf{v} . For any player i and any subadditive valuation $v_i \in V_i$, consider the deviation $a_i(v_i; \mathbf{w}_{-i}, \mathbf{C}_{-i})$ as defined in Lemma 37, where $\mathbf{w}_{-i} \sim \mathbf{D}_{-i}$. By the definition of the Bayesian Nash equilibrium, we obtain

$$\mathbb{E}_{\mathbf{v}_{-i}} [u_i^{v_i}(\mathbf{B}_i(v_i), \mathbf{B}_{-i}(\mathbf{v}_{-i}))] = u_i^{v_i}(\mathbf{B}_i(v_i), \mathbf{C}_{-i}) \geq \mathbb{E}_{\mathbf{w}_{-i}} [u_i^{v_i}(a_i(v_i; \mathbf{w}_{-i}, \mathbf{C}_{-i}), \mathbf{C}_{-i})].$$

By taking the expectation over v_i and summing up over all players,

$$\begin{aligned} \sum_i \mathbb{E}_{\mathbf{v}} [u_i(\mathbf{B}(\mathbf{v}))] &\geq \sum_i \mathbb{E}_{v_i, \mathbf{w}_{-i}} [u_i^{v_i}(a_i(v_i; \mathbf{w}_{-i}, \mathbf{C}_{-i}), \mathbf{C}_{-i})] \\ &= \mathbb{E}_{\mathbf{v}} \left[\sum_i u_i^{v_i}(a_i(\mathbf{v}, \mathbf{C}_{-i}), \mathbf{C}_{-i}) \right] \geq \frac{1}{2} \mathbb{E}_{\mathbf{v}} [SW(\mathbf{o}^{\mathbf{v}})] - \sum_i \sum_j \mathbb{E}[b_{ij}]. \end{aligned}$$

So, $\mathbb{E}_{\mathbf{v}} [SW(\mathbf{B}(\mathbf{v}))] = \sum_i \mathbb{E}_{\mathbf{v}} [u_i(\mathbf{B}(\mathbf{v}))] + \sum_i \sum_j \mathbb{E}[b_{ij}] \geq \frac{1}{2} \mathbb{E}_{\mathbf{v}} [SW(\mathbf{o}^{\mathbf{v}})]$. \square

As we show next, the previous upper bounds of Theorems 38 and 39 are tight even for the *pure* Nash equilibrium and a *single* resource.

Theorem 40. The PoA of pure Nash equilibria for the proportional allocation mechanism with subadditive valuations is at least 2.

Proof. We consider a game with only two players and a single resource. The valuation of the first player is $v_1(x) = 1 + \epsilon \cdot x$, for some $\epsilon < 1$, if $x < 1$ and $v_1(x) = 2$ if $x = 1$. The valuation of the second player is $v_2(x) = \epsilon \cdot x$. One can easily verify that these two functions are subadditive and the optimal social welfare is 2. Consider the bidding strategies $b_1 = b_2 = \frac{\epsilon}{4}$. The utility of player 1, when she bids x and player 2 bids $\frac{\epsilon}{4}$, is given by $1 + \epsilon \cdot \frac{x}{x + \epsilon/4} - x$ which is maximised for $x = \frac{\epsilon}{4}$. The utility of player 2, when she bids x and player 1 bids $\frac{\epsilon}{4}$, is $\epsilon \cdot \frac{x}{x + \epsilon/4} - x$ which is also maximised when $x = \frac{\epsilon}{4}$. So (b_1, b_2) is a pure Nash Equilibrium with social welfare $1 + \epsilon$. Therefore, the PoA converges to 2 when ϵ goes to 0. \square

7.2.2 Lower Bound for Simple Mechanisms

Now, we show a lower bound that applies to all simple mechanisms, where the bidding space has size (at most) sub-doubly-exponential in m . More specifically,

we apply the general framework of Roughgarden [124], for showing lower bounds on the PoA for *all* simple mechanisms, via communication complexity reductions with respect to the underlying optimisation problem. In our setting, the problem is to maximise the social welfare by allocating divisible resources to players with subadditive valuations. We proceed by proving a communication lower bound for this problem in the following lemma.

Lemma 41. For any constant $\varepsilon > 0$, any $(2-\varepsilon)$ -approximation (non-deterministic) algorithm for maximising social welfare in resource allocation problem with subadditive valuations requires an exponential amount of communication.

Proof. We prove this lemma by reducing the communication lower bound for combinatorial auctions with general valuations (Theorem 3 of [113]) to our setting (see also [59] for a reduction to combinatorial auctions with subadditive players).

Nisan [113] used an instance with n players and m items, with $n < m^{1/2-\varepsilon}$. Each player i is associated with a set, T_i , of *bundles*, with $|T_i| = t$ for some $t > 0$. At every instance of this problem, the players' valuations are determined by sets, I_i , of bundles, where $I_i \subseteq T_i$ for every i . Given I_i , player i 's valuation on some subset S of items is $v_i(S) = 1$, if there exists some $R \in I_i$ such that $R \subseteq S$, otherwise $v_i(S) = 0$. In [113], it was shown that distinguishing between instances with optimal social welfare of n and 1 , requires t bits of communication. By choosing t exponential in m , their theorem follows.

We prove the lemma by associating any valuation v of the above combinatorial auction problem, to some appropriate subadditive valuation v' for our setting. For any player i and any fractional allocation $\mathbf{x} = (x_1, \dots, x_m)$, let $A_{x_i} = \{j | x_{ij} > \frac{1}{2}\}$. We define $v'_i(x_i) = v_i(A_{x_i}) + 1$ if $x_i \neq \mathbf{0}$ and $v'_i(x_i) = 0$ otherwise. It is easy to verify that v'_i is subadditive. Notice that $v'_i(x) = 2$ only if there exists $R \in I_i$ such that player i is allocated a fraction higher than $1/2$ for every resource in R . The value $1/2$ is chosen such that no two players are assigned more than that fraction from the same resource. This corresponds to the constraint of an allocation in the combinatorial auction where no item is allocated to two players.

Therefore, in the divisible goods allocation problem, distinguishing between instances where the optimal social welfare is $2n$ and $n + 1$ is equivalent to distinguishing between instances where the optimal social welfare is n and 1 in the corresponding combinatorial auction and hence requires exponential, in m , number of communication bits. \square

The PoA lower bound follows the general reduction described in [124].

Theorem 42. The PoA of ϵ -mixed Nash equilibria²⁷ of every simple mechanism, when players have subadditive valuations, is at least 2.

Remark 43. This result holds only for ϵ -mixed Nash equilibria. Considering exact Nash equilibria, we show a lower bound for all *scale-free* mechanisms in the following section.

7.2.3 Lower Bound for Scale-free Mechanisms

Here we prove a tight lower bound for all scale-free mechanisms including the proportional allocation mechanism. A mechanism (\mathbf{x}, \mathbf{q}) is said to be scale-free if a) for every player i , resource j and constant $c > 0$, $x_i(c \cdot \mathbf{b}_j) = x_i(\mathbf{b}_j)$. Moreover, for a fixed \mathbf{b}_{-i} , x_i is non-decreasing and positive whenever b_{ij} is positive. b) The payment for player i depends only on her bids $b_i = (b_{ij})_j$ and equals to $\sum_{j \in [m]} q_j(b_{ij})$ where q_i is non-decreasing, continuous, normalised ($q_i(0) = 0$), and there always exists a bid b_{ij} such that $q_i(b_{ij}) > 0$.

Theorem 44. The mixed PoA of scale-free mechanisms when players have subadditive valuations, is at least 2.

Proof. Given a mechanism (\mathbf{x}, \mathbf{q}) , we construct an instance with 2 players and m resources. Let V be a positive value such that V/m is in the range of both q_1 and q_2 . This can always be done due to our assumptions on q_i . Let T_1 and T_2 be the values such that $q_1(T_1) = q_2(T_2) = V/m$. W.l.o.g. we assume that $T_1 \geq T_2$. By monotonicity of q_1 , $q_1(T_2) \leq V/m$. Pick an arbitrary value $a \in (0, 1)$, and let $\alpha_1 = x_1(a, a)$ and $\alpha_2 = x_2(a, a)$. By the assumption that $x_i(\mathbf{b}_j) > 0$ for $b_{ij} > 0$, we have $\alpha_1, \alpha_2 \in (0, 1)$. Let $v = V/\sqrt{m}$. We define the players' valuations as:

$$v_1(x) = \begin{cases} 0, & \text{if } \forall j \in [m], x_j = 0, \\ v, & \text{if } \forall j x_j < \alpha_1, \exists k x_k > 0 \\ 2v, & \text{otherwise} \end{cases}$$

$$v_2(x) = \begin{cases} 0, & \text{if } \forall j \in [m], x_j = 0 \\ V, & \text{if } \exists j x_j < \alpha_2, \exists k x_k > 0 \\ 2V, & \text{otherwise} \end{cases}$$

We claim that the following mixed strategy profile \mathbf{B} is a Nash equilibrium. Player 1 picks resource l uniformly at random and bids $b_{1l} = y$, and $b_{1k} = 0$,

²⁷A bidding profile $\mathbf{B} = \times_i B_i$ is called ϵ -mixed Nash equilibrium if, for every player i and any bid b'_i , $\mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(\mathbf{b})] \geq \mathbb{E}_{\mathbf{b} \sim \mathbf{B}}[u_i(b'_i, \mathbf{b}_{-i})] - \epsilon$.

for $k \neq l$, where y is a random variable drawn by the cumulative distribution $G(y) = \frac{mq_2(y)}{V}$, $y \in [0, T_2]$. Player 2 bids $b_{2j} = z$ for every item j , where z is a random variable drawn by the cumulative distribution $F(z)$, defined as $F(z) = \frac{v - q_1(T_2) + q_1(z)}{v}$, $z \in [0, T_2]$. Note that $G(T_2) = 1$, $F(T_2) = 1$ and since q_i is continuous, G and F are continuous in $(0, T_2)$ and therefore both functions have no mass point in any $(0, T_2)$. We assume that if both players bid 0 for some resource, player 2 takes the whole resource. We are ready to show that \mathbf{B} is a Nash equilibrium.

If player 1 bids any y in the range $(0, T_2]$ for a single resource j and zero for the rest, then she gets an allocation of at least α_1 for resource j (that she values for $2v$), only if $y \geq z$, which happens with probability $F(y)$. If $y < z$ her value is v . Therefore, her expected valuation is $v + F(y)v$. So, for every $y \in (0, T_2]$ her expected utility is $v + F(y)v - q_1(y) = 2v - q_1(T_2)$. If player 1 picks y according to $G(y)$, her utility is still $2v - q_1(T_2)$, since she bids 0 with zero probability. Suppose player 1 bids $\mathbf{y} = (y_1, \dots, y_m)$, $y_j \in [0, T_2]$ for every j , with at least two positive bids, and w.l.o.g., assume $y_1 = \max_j y_j$. If $z > y_1$, player 1 has value v for the allocation she receives. If $z \leq y_1$, player 1 has value $2v$, but she may pay more than $q_1(y_1)$. So, this strategy is dominated by the strategy of bidding y_1 for the first resource and zero for the rest. Bidding greater than T_2 for any resource is dominated by the strategy of bidding exactly T_2 for that resource.

If player 2 bids $z \in [0, T_2]$ for all resources, she gets an allocation of at least α_2 for all the m resources with probability $G(z)$. So, her expected utility is $V + G(z)V - mq_2(z) = V$. Bidding greater than T_2 for any resource is dominated by bidding exactly T_2 for this resource. Suppose that player 2 bids any $\mathbf{z} = (z_1, \dots, z_m)$, with $z_j \in [0, T_2]$ for every j , then, since player 1 bids positively for any item with probability $1/m$, player's 2 expected utility is

$$\begin{aligned} \mathbb{E}[u_2] &= \frac{1}{m} \sum_j \left(V + G(z_j)V - \sum_k q_2(z_k) \right) \\ &= \frac{1}{m} \sum_j \left(V + mq_2(z_j) - \sum_k q_2(z_k) \right) \\ &= \frac{1}{m} \left(mV + m \sum_j q_2(z_j) - m \sum_k q_2(z_k) \right) = V. \end{aligned}$$

Overall, \mathbf{B} is Nash equilibrium. Therefore, it is sufficient to bound the expected social welfare in \mathbf{B} . Player 1 bids 0 with zero probability. So, whenever

player 2 bids 0, she receives exactly $m - 1$ resources, which she values for V . Player 2 bids 0 with probability $F(0) = 1 - \frac{q_1(T_2)}{v} \geq 1 - \frac{V}{mv} = 1 - \frac{1}{\sqrt{m}}$. Hence, $\mathbb{E}[\text{SW}(\mathbf{B})] \leq 2V - F(0) \cdot V + 2v \leq 2V \left(1 + \frac{1}{\sqrt{m}}\right) - V \left(1 - \frac{1}{\sqrt{m}}\right) = V \left(1 + \frac{3}{\sqrt{m}}\right)$. On the other hand, the social welfare in the optimum allocation is $2(V + v) = 2V \left(1 + \frac{1}{\sqrt{m}}\right)$ (player 1 is allocated α_1 proportion from one resource and the rest is allocated to player 2). We conclude that $PoA \geq 2 \frac{\left(1 + \frac{1}{\sqrt{m}}\right)}{\left(1 + \frac{3}{\sqrt{m}}\right)}$ which, for large m , converges to 2. \square

Part II

Cost-Sharing Networks

CHAPTER 8

Overview

This part is based on joint works with George Christodoulou and Stefano Leonardi. Chapters 10 and 11 are based on the paper [43], co-authored with George Christodoulou, which was published in the Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms in 2016. Chapter 12 is based on the paper [42], co-authored with George Christodoulou and Stefano Leonardi, which appeared in the Proceedings of the 9th International Symposium on Algorithmic Game Theory in 2016 and was invited to Special Issue of Theory of Computing Systems in 2017.

Network Cost-Sharing Games. We study a *multicast game* in a rooted undirected graph $G = (V, E)$ with a nonnegative cost c_e on each edge $e \in E$. A set of k *terminal vertices* or *players* t_1, \dots, t_k need to establish connectivity with the root t . Each player selects a path P_i and the outcome produced is the graph $H = \cup_i P_i$. The global objective is to minimise the cost, $\sum_{e \in H} c_e$, of this graph, which is the *Minimum Steiner Tree*. The cost of H should be covered by its users and the way the cost is split among them is dictated by a *cost-sharing protocol*. The players prefer paths that charge them with small cost, and therefore the solution will be a Nash equilibrium.

Cost-Sharing Design. Different cost-sharing protocols result in different quality of equilibria. In this work, we are interested in the design of protocols that induce good equilibrium solutions in the *worst-case*, therefore we focus on protocols that guarantee low PoA. Chen, Roughgarden and Valiant [36] initiated the *design* aspect for network cost-sharing games. They gave a characterisation

of protocols that satisfy some natural axioms and they thoroughly studied the PoA for the following two classes of protocols, that use different informational assumptions from the perspective of the designer.

Non-uniform protocols. The designer has full knowledge of the instance, that is, she knows both the network topology given by G and the costs c_e , and in addition the set of players' requests t_1, \dots, t_k . They showed that a simple priority protocol (see Example 45) has a constant PoA; the NE induced by the protocol simulate Prim's algorithm for the Minimum Spanning Tree (MST) problem, and therefore achieve constant approximation.

Uniform protocols. The designer needs to decide how to split the edge cost among the users *without knowledge of the underlying graph*. They showed that the PoA is $\Theta(\log k)$; both upper and lower bound comes from the analysis of the Greedy Algorithm for the Online Steiner Tree problem.

Cost-Sharing Design under Uncertainty. Arguably, there are situations where the former assumption is too optimistic while the latter is too pessimistic. We propose a model that lies in the middle-ground as a framework to design network cost-sharing protocols with good equilibria, when the designer has *incomplete information*. We assume that the designer has prior knowledge of the underlying metric, (given by the graph G and the shortest path metric induced by the costs c_e), but is *uncertain* about the requested subset of players. We consider three different models, the *adversarial*, the *stochastic* and the *Bayesian* models.

Adversarial Model. The designer *knows nothing* about the number or the positions of the t_i 's and has as goal to process the graph and choose a single, *universal* cost-sharing protocol that has low PoA against *all possible* requested subsets. Here, no distributional assumptions are made about arrivals of players, instead the worst-case approach is used similarly to competitive analysis of online algorithms. Once the designer selects the protocol, then an adversary will choose the requested subset of players and their positions in the graph (the t_i 's), in a way that *maximises* the PoA of the induced game.

Stochastic and Bayesian Models. The players' terminals are chosen according to some probability distribution which is known to the designer. The ground difference is that in the stochastic model the players act by fully knowing the instance, including the terminals of the other players, whereas in the Bayesian

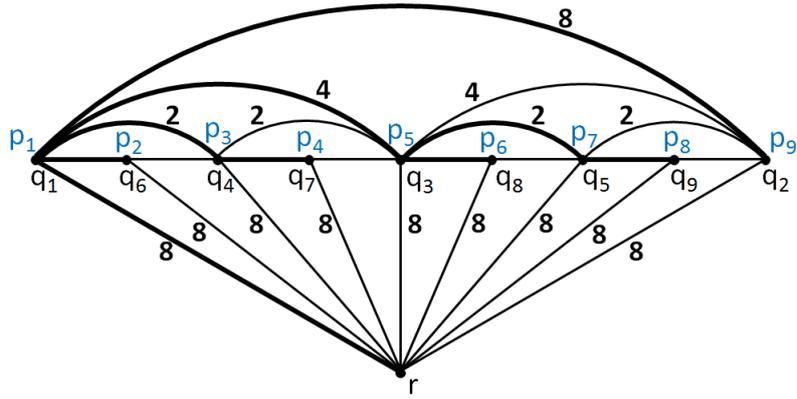


Figure 8.1: The figure is a variation of the worst-case graph for the Online Steiner Tree problem [87] that serves as the worst-case graph for the universal protocols of the network cost-sharing game. In edges with no written cost, we consider the unit cost; note that the graph can be easily generalised so that the number of vertices is arbitrarily big. We assume two orders on the vertices, denoted by q_i or p_i . The q -order is adversarially chosen and results in high PoA of $\Omega(\log k)$. The p -order results in constant PoA.

model, the players have also incomplete knowledge about the other terminals, the same way as the designer does and therefore, act in order to *minimise* their expected cost-share. The goal in both latter models is to choose a universal protocol where the expected worst-case cost in the Nash equilibrium is not far from the expected optimal cost.

Example 45. (Ordered protocols). An important special class with interesting properties is that of *ordered protocols*. The designer decides a total order of the users, and when a subset of players uses some edge, the full cost is covered by the player who comes first in the order. Any NE of the induced game corresponds to the solution produced by the Greedy Algorithm for the MST: each player is connected, via a shortest path, with the component of the players that come before him in the order. The analysis of the PoA in the uniform model boils down to the analysis of the Greedy Algorithm for the Online Steiner Tree problem, where the worst-case order is considered.

The following example demonstrates that even this special class of ordered protocols becomes very rich, once the designer has prior knowledge of the underlying metric space. Uniform protocols throw away this crucial component, the structure of the underlying metric, that universal protocols can use in their favour in order to come up with better PoA guarantees.

Uniform protocols. The designer chooses an order of the players $1, \dots, k$ with-

out prior knowledge of the graph. The adversary constructs a worst-case graph by simulating the adversary for the Greedy Algorithm of the Online Steiner Tree problem [87] and places the players accordingly. See for example Figure 8.1, the q labels. There is a Nash equilibrium that is formed by the bold edges, whereas the optimum solution is the path $(r, q_1, q_6, q_4, q_7, q_3, q_8, q_5, q_9, q_2)$. Therefore, the PoA of uniform ordered protocol is $\Omega(\log k)$ [36].

Universal protocols. The designer takes into account the graph; consider again the graph of Figures 8.1. Order the vertices according to the linear order dictated from the path p_1, \dots, p_9 (say from left to right). The adversary will choose k and the positions of the players (t_1, \dots, t_k) . Note that when a player chooses her path, all the players that are on her left will cover the cost of their chosen paths because they precede her in the order. Therefore, each player would connect with the previous player in the p -order via shortest path resulting in a cost that is exactly the same with that of the optimum solution²⁸. Overall, *no matter which subset of players the adversary chooses*, the PoA remains constant (in fact PoA= 1) as k grows.

Example 46. (Generalised weighted Shapley). In [36], it was shown that ordered protocols are essentially optimal among uniform protocols. In our model, the choice of the optimal method may depend on the underlying graph metric. Take the example in Figure 8.2. By using Shapley cost-sharing the adversary can choose to activate $\{v_1, v_2, v_3\}$ and it is a NE if v_1, v_3 connect directly to r and v_2 connects through v_1 . Regarding *any* ordered protocol, the square defined by the v_i 's contains a path of length 2 where the middle vertex comes *last* in the order. The adversary will select this triplet of players, say v_1, v_2, v_3 . In the NE, v_1 connects directly to r , v_3 and v_2 connect through v_1 . In both cases, (by ignoring ε) the cost of the NE is 5 and the minimum Steiner tree that connects those vertices with r has cost 4 and therefore, $\text{PoA} \geq 5/4$.

However, the following (generalised Shapley) protocol, has $\text{PoA} = 1$. Partition the players into two sets $S_1 = \{v_1, v_2\}, S_2 = \{v_3, v_4\}$. If players from both partitions appear on some edge, then the cost is charged only to players from S_1 . Players that belong to the same partition share the cost equally. One can

²⁸The optimum solution connects all players by using the unit-cost edges that connect the leftmost to the rightmost player and uses one edge of cost 8 in order to connect that component to r .

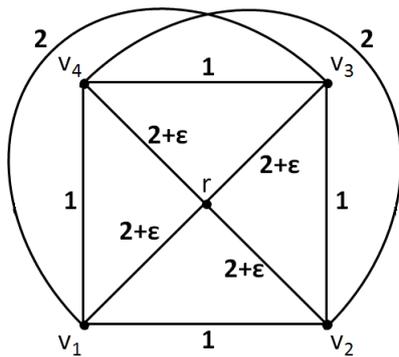


Figure 8.2: The figure shows an example where both the *best* ordered protocol and the Shapley protocol have $\text{PoA} \geq 5/4$, whereas there is an *intermediate* protocol with $\text{PoA}=1$; we set $\varepsilon > 0$ arbitrarily small.

verify that for all possible subsets of players this protocol produces only optimal equilibria.

8.1 Results

We propose a framework for the design of (universal) network cost-sharing protocols with good equilibria, in situations where the designer has incomplete information about the input. We consider three different models, the *adversarial*, the *stochastic* and the *Bayesian*. In all models, the designer has prior knowledge of the underlying metric but the requested subset of players is *not known* and is activated either in an adversarial manner (adversarial model) or is drawn from a known probability distribution (stochastic and Bayesian model). The central question we address is: *to what extent does prior knowledge of the metric help in good network design under uncertainty?*

For the adversarial model (Chapter 10), we first demonstrate that there exist classes of graph metrics where prior knowledge of the underlying metric can dramatically improve the performance of good network cost-sharing design. For *outerplanar* graph metrics, we provide a universal ordered cost-sharing protocol with constant PoA, against any choice of the adversary. This is in contrast to uniform protocols that ignore the graph and cannot achieve PoA better than $\Omega(\log k)$ in outerplanar metrics.

Our main technical result shows that there exist graph metrics, for which knowing the underlying metric does not help the designer, and *any universal*

protocol²⁹ has PoA of $\Omega(\log k)$. This matches the upper bound of $O(\log k)$ that can be achieved without prior knowledge of the metric [87, 36]. Our results for the adversarial model motivate the following question that is left open.

Open Question For which metric spaces can one design universal protocols with constant PoA? What sort of structural graph properties are needed to obtain good guarantees?

Then, we switch to the stochastic model (Chapter 11), where the players (terminal vertices) are activated according to some probability distribution that is known to the designer. We show that there exists a *randomised* ordered protocol that achieves constant PoA. If each player is activated independently with some probability, by using standard derandomisation techniques [141, 130], we produce a *deterministic* ordered protocol that achieves constant PoA. We remark, that the first result holds also for the *black-box* model, where the probability distribution is not known to the designer, but is allowed to draw independent (polynomially many) samples.

At last, for the Bayesian model (Chapter 12), we show that even for i.i.d. players, i.e. players with independent and identical prior distributions on the position of their terminal, there exists a lower bound of $\Omega(\sqrt{k})$ on the PoA of *any* (deterministic or randomised) cost-sharing protocol that satisfies certain natural axioms posed by [36]. One of the axioms that [36] required in their design space is that every cost-sharing protocol should satisfy *budget balance*, i.e. that the players' cost-shares cover *exactly* the cost of *any* solution. We relax this property by requiring budget-balance only in *all equilibria*. We provide a protocol that is *ex-post*³⁰ budget-balanced in equilibrium with constant PoA. We further present anonymous posted prices with the same upper bound on the PoA that are *ex-ante*³¹ budget-balanced in equilibrium; we discuss limitations of other concepts, such as budget-balance in equilibrium with “high” probability or bounded possible excess and deficit.

²⁹The result holds only for deterministic protocols. It is not known whether the same lower bound holds for randomised protocols. For the special case of an ordered protocol that orders the players uniformly at random, a logarithmic lower bound is known by Garg et al. [72].

³⁰In *ex-post* budget-balance we require budget-balance in every realisation of the game. If the *expected* excess and deficit are zero, the budget balance is called *ex-ante*.

³¹*Ex-post* budget-balanced cannot be obtained via anonymous prices.

8.2 Techniques

We prove our main lower bound for the adversarial model (Chapter 10) in two parts. In the first part (Section 10.2) we bound the PoA achieved by *any ordered* protocol. Our origin is a well-known “zig-zag” *ordered* structure that has been used to show a lower bound on the Greedy Algorithm of the Online Steiner Tree problem (see the labeled path $(q_1, q_6, q_4, \dots, q_2)$ in Figure 8.1). The challenge is to show that high dimensional hypercubes exhibit such a distance preserving structure *no matter how the vertices are ordered*. Section 10.2 is devoted to this task and we believe that this is of independent interest.

We show the existence proof by employing powerful tools from Extremal Combinatorics and in particular Ramsey Theory [78]. We are inspired by a Ramsey-type result due to Alon et al. [5], in which they show that for any given length $\ell \geq 5$, any r -edge colouring of a high dimensional hypercube contains a monochromatic cycle of length 2ℓ . Unfortunately, we cannot immediately use their results, but we show a similar Ramsey-type result for a different, carefully constructed structure; we assert that every 2-edge colouring of high dimensional hypercubes Q_n contains a monochromatic copy of that structure. Then, we prescribe a special 2-edge-colouring that depends on the ordering of Q_n , so that the special subgraph preserves some nice labelling properties. A suitable subset of the subgraph’s vertices can be 1-embedded into a hypercube of lower dimension. Recursively, we show existence of the desired *distance preserving* “zig-zag” structure.

In the second part (Section 10.3), we extend the lower bound to *all universal* cost-sharing protocols, by using the characterisation of [36]. At a high level, we use as basis the construction for the ordered protocol and create “multiple copies”³². The adversary will choose different subsets of players, depending on whether the designer chose protocols “closer” to Shapley or to ordered. In the latter case, we use arguments from Matching Theory to guarantee existence of ordered-like players in one of the hypercubes.

For the stochastic model (Chapter 11), we construct an approximate minimum Steiner tree over a subset of vertices which is drawn from the known probability distribution. This tree is used as a base to construct a spanning tree, which determines a total order over the vertices. We finally produce a deterministic

³²Note that the standard complexity measure, to analyse the inefficiency of equilibria, is the number of participants, k , and not the total number of vertices in the graph (see for example [6, 36]).

order by applying standard derandomisation techniques [141, 130].

Chapter 12 is devoted to the Bayesian setting. Regarding the cost-sharing protocols that are budget-balanced only in equilibrium (Section 12.2), we show an interesting connection between algorithms for *Oblivious Stochastic* optimisation problems and cost-sharing design with low PoA. We are able to enforce approximate solutions of the stochastic problem, as Bayesian Nash equilibria, with the same guarantees on the PoA. Although this connection is quite simple, it results in significant improvement on the PoA comparing to budget-balanced protocols.

8.3 Related Work

Following the work of [6, 7], a long line of research studies network cost-sharing games, mainly focusing on the PoS of the Shapley cost-sharing mechanism. [6] showed a tight $\Theta(\log k)$ bound for directed networks, while for undirected networks several variants have been studied [22, 23, 24, 35, 37, 58, 68, 103] but the exact value of PoS still remains a big open problem. For multicast games, an improved upper bound of $O(\log k / \log \log k)$ is known due to Li [103], while for broadcast games, a series of work [68, 101] lead finally to a constant due to Bilò et al. [24]. The PoA of some special equilibria has been also studied in [31, 34].

Chen, Roughgarden and Valiant [36] initiated the study of network cost-sharing design with respect to PoA and PoS. They characterised a class of protocols that satisfy certain desired properties (which was later extended by Gopalakrishnan, Marden and Wierman, in [76]), and they thoroughly studied PoA and PoS for several cases. Von Falkenhausen and Harks [139] studied singleton and matroid games with weighted players, while Gkatzelis, Kollias and Roughgarden [73] focus on weighted congestion games with polynomial cost functions. Gairing, Kollias and Kotsialou [71] studied the PoA and the PoS for cost-sharing methods on weighted congestion games for convex and polynomial cost functions, respectively.

Moulin and Shenker [111] studied cost-sharing games under mechanism design context; they characterised the budget-balanced and group strategyproof mechanisms and identify the one with minimum welfare loss. In similar context, other papers considered (group)strategy proof and efficient mechanisms and relaxed the budget-balanced constraint. Könemann et al. [97] showed, for the Steiner forest game, a 2-budget-balanced cost-sharing method, meaning that at

least half of the cost is covered by the players, and proved that this factor cannot be improved. Devanur, Mihail and Vazirani [57] and Immorlica, Mahdian and Mirrokni [89] studied other cost-sharing games, like the set cover and the facility location games showing positive and negative bounds on the fraction of the cost that is covered by the players.

Close in spirit to universal cost-sharing protocols is the notion of Coordination Mechanisms [39] that provides a way to improve the PoA in cases of incomplete information. The designer has to decide in advance local scheduling policies or increases in edge latencies, without knowing the exact input, and has been used for scheduling problems [1, 2, 10, 16, 27, 28, 39, 53, 88, 96] as well as for simple routing games [17, 48].

The underlying optimisation problem that we consider here is the minimum Steiner tree problem. It is a well-studied problem and known to be in NP-complete and the best known approximation is 1.39 [26]. As discussed in Example 45, the analysis of the equilibria induced by ordered protocols corresponds to the analysis of the Greedy Algorithm for the MST. In the uniform model, this corresponds to the analysis of the Greedy Algorithm [8, 87] for the (Generalised) Online Steiner Tree problem [4, 12, 136], which was shown to be $\Theta(\log k)$ -competitive by Imase and Waxman [87] ($O(\log^2 k)$ -competitive for the Generalised Online Steiner Tree problem by [8]). The universal model is closely related to universal network design problems [90], hence our choice for the term “universal”. In the universal TSP, given a metric space, the algorithm designer has to decide a *master* order so that tours that use this order have good approximation [13, 15, 44, 77, 83, 90, 120].

Much work has been done in stochastic models and we only mention the most related to our work. Karger and Minkoff [93] showed a constant approximation guarantee for the maybecast problem, where the designer needs to fix (before activation) some path for every vertex to the root. Garg et al. [72] gave bounds on the approximation of the stochastic online Steiner tree problem for several informational assumptions. A line of works [14, 77, 128, 130] has studied the *a priori* TSP. Shmoys and Talwar [130] assumed independent activations and demonstrated randomised and deterministic algorithms with constant approximations.

CHAPTER 9

Models and Preliminaries

A *multicast network cost-sharing game*, is specified by a connected undirected graph $G = (V, E)$, with a designated root r and nonnegative weight c_e for every edge e , a set of players $S = \{1, \dots, k\}$ and a cost-sharing protocol. Each player i is associated with a *terminal* t_i which she needs to connect with r ; t_i is called the *type* of player i . We denote by $\mathbf{t} = (t_1, \dots, t_k)$ the vector of players' terminals. We say that a vertex is *activated* if there exists some requested player associated with it. The cost-sharing protocol is defined by the *protocol designer* or just the designer.

9.1 Information Models

We consider the following information models from the perspective of the designer and the players:

Adversarial: The designer *knows nothing* about the set S of activated vertices.

The players though have *full knowledge* of the instance.

Stochastic: The players' types are drawn from some distribution D defined over V^k . The actual types are *unknown* to the designer, who is only aware of D . However, the players decide their strategies by *knowing* other players' types.

Bayesian: The players' types are drawn from some product distribution D defined over V^k . Both the designer and the players know only D . A natural assumption is that the players know their own type.

Regarding the stochastic and the Bayesian models, the type of player i is drawn

according to some known distribution, $D_i : V \rightarrow [0, 1]$ with $\sum_{v \in V} D_i(v) = 1$, possibly different for each player, that results in the joint distribution $D = \times_i D_i$.

9.2 Universal Cost-Sharing Protocols

The designer is not aware of the players' terminals and thus of the activated vertices, hence, she should define the cost-sharing protocol over a larger set N that includes all potential players; we will determine N later for each setting. Consider a set of players N , then a *cost-sharing method* $\xi_e : 2^N \rightarrow \mathbb{R}_+^{|N|}$ is a function of the set of players, $R \subseteq N$, using edge e and decides the cost-share for each such player $i \in R$. A natural rule is that the shares for players not included in R should always be 0. We use the notation $\xi_e(i, R)$ to denote the cost-share of player i under input R ; note that if $i \notin R$, then by default $\xi_e(i, R) = 0$. A *cost-sharing protocol* Ξ assigns, for every $e \in E$, some cost-sharing method ξ_e .

Following previous work [36, 139], we focus on cost-sharing protocols that satisfy the following natural properties:

- (1) *Budget-balance*: For every network game induced by the cost-sharing protocol Ξ , and every outcome of it, $\sum_{i \in R} \xi_e(i, R) = c_e$, for every edge e .
- (2) *Separability*: For every network game induced by the cost-sharing protocol Ξ , the cost-shares of each edge are completely determined by the set of players using it.
- (3) *Stability*: In every network game induced by the cost-sharing protocol Ξ , there exists at least one pure (Bayesian) Nash equilibrium, regardless of the graph structure.

We call a cost-sharing protocol Ξ *universal*, if it satisfies the above properties for any graph G , and it assigns the cost-sharing method $\xi_e : 2^N \rightarrow \mathbb{R}_+^{|N|}$ to any edge e based only on knowledge of G (without any knowledge of S) for the adversarial model, while in the stochastic and Bayesian models the method can in addition depend on D .

Due to the characterisation in [36], we restrict ourselves to the family of *generalised weighted Shapley protocols* in all, but the Bayesian, settings. Chen, Roughgarden and Valiant [36] characterise the linear protocols³³ satisfying the

³³A linear protocol, for every edge e of cost $c_e \geq 0$, assigns the method $c_e \cdot \xi$, where ξ is the method it assigns to any edge of unit cost.

three properties to be the generalised weighted Shapley protocols and they further showed that for any non-linear protocol, there exists a linear one with at most the same PoA. They assumed that no structural property of the network is known. We note that this is also our requirement, so that even if the network evolves over time, the protocol always guarantees the existence of a pure Nash equilibrium. We remark that, if additionally the existence of the equilibrium relies on the network structure, other protocols may also exist; however, this is out of the scope of this thesis and we refer the reader to [107] for an example in singleton games.

Determining N. For the adversarial setting, since the designer is not aware of \mathbf{t} , we will define the cost-sharing protocol for a set that includes all potential players. W.l.o.g. we can assume that each player is associated with a distinct vertex³⁴ and hence, we define the cost-sharing protocol with respect to V .

For the stochastic and Bayesian settings, it is not so clear any more if it is w.l.o.g. to assume that each player is associated with a distinct vertex; instead, we consider N to be the set of all possible pairs of player-type, i.e. $|N| = k|V|$, where k is the number of players.

Discussion. In a sense, the protocol is defined on the players' types and not on the players themselves, which is a ground difference between our model and the uniform protocols of [36]. Therefore, the protocol should decide the shares for any possible combination of players' types.

We fairly assume that when a player uses an edge, the protocol is aware of their origin and some global identification. For instance, if the players use the network in order to route some traffic, this would probably carry the address of origin and some id.

³⁴This is true, since on the one hand if the worst-case instance includes a vertex being the type of more than one player, the designer can use an arbitrary order among those players and charge only the first in the order; then the rest will follow the same path by paying zero. On the other hand, if the designer can benefit by considering any other charging method between players of the same vertex, the adversary can choose only one player per vertex.

9.3 Generalised Weighted Shapley Protocol

The *Generalised Weighted Shapley Protocol (GWSP)* (see also the definition of [76]) is defined by the players' weights (parameters) $\{w_1, \dots, w_n\}$ and an *ordered* partition of the players $\Sigma = (U_1, \dots, U_h)$. An interpretation of Σ is that for $i < j$, players from U_i “arrive” before players from U_j and they should cover the cost. More formally, for every edge e of cost c_e , every set of players R_e that uses e and for $s = \arg \min_j \{U_j | U_j \cap R_e \neq \emptyset\}$, the GWSP assigns the following method to e :

$$\xi_e(i, R_e) = \begin{cases} \frac{w_i}{\sum_{j \in U_s \cap R_e} w_j} c_e, & \text{if } i \in U_s \cap R_e \\ 0, & \text{otherwise} \end{cases}$$

In the special case that each U_i contains exactly one player, the protocol is called *ordered*. The order of U_i 's indicates a permutation of the players, denoted by π .

9.4 Pure Nash Equilibrium

We denote by $\mathcal{P}_i(t_i) = \mathcal{P}_i$ the strategy space of player i with terminal/type t_i , i.e. the set of all the paths connecting t_i to r . $\mathbf{P} = (P_1, \dots, P_k)$ denotes an *outcome* or a *strategy profile*, where $P_i \in \mathcal{P}_i$ for all $i \in S$. As usual, \mathbf{P}_{-i} denotes the strategies of all players but i . Let R_e be the set of players using edge $e \in E$ under \mathbf{P} . The cost-share of player i induced by ξ_e 's is equal to

$$c_i(\mathbf{P}) = \sum_{e \in P_i} \xi_e(i, R_e).$$

The players' objective is to minimise their cost-share $c_i(\mathbf{P})$.

A strategy profile $\mathbf{P} = (P_1, \dots, P_k)$ is a *Nash equilibrium (NE)* if for every player $i \in S$ and every strategy $P'_i \in \mathcal{P}_i$,

$$c_i(\mathbf{P}) \leq c_i(\mathbf{P}_{-i}, P'_i).$$

A strategy profile $\mathbf{P}(\mathbf{t}) = (P_1(t_1), \dots, P_k(t_k))$ is a *Bayesian Nash equilibrium*

(BNE) if for every player $i \in S$ with type t_i and every strategy $P'_i(t_i) \in \mathcal{P}_i(t_i)$,

$$\mathbb{E}_{\mathbf{t}_{-i} \sim D_{-i}} [c_i(\mathbf{P}(\mathbf{t}))] \leq \mathbb{E}_{\mathbf{t}_{-i} \sim D_{-i}} [c_i(\mathbf{P}_{-i}(\mathbf{t}_{-i}), P'_i(t_i))].$$

9.5 Price of Anarchy

The cost of an outcome $\mathbf{P} = (P_1, \dots, P_k)$ is defined as $c(\mathbf{P}) = \sum_{e \in \cup_i P_i} c_e$, while $\mathbf{O}(\mathbf{t}) = \mathbf{O} = (O_1, \dots, O_k) \in \arg \min_{\mathbf{P}} c(\mathbf{P})$ is an optimum solution. The *Price of Anarchy (PoA)* is defined as the worst-case ratio of the cost in a NE over the optimal cost in the game induced by S . In the adversarial model the *worst-case* S is chosen, while in the stochastic and Bayesian models S is drawn from a known distribution D . Formally, in the adversarial model we define the PoA of a protocol Ξ on G as

$$PoA(G, \Xi) = \max_{S \subseteq V \setminus \{r\}} \frac{\max_{\mathbf{P} \in \mathcal{N}} c(\mathbf{P})}{c(\mathbf{O})},$$

where \mathcal{N} is the set of all NE of the game induced by Ξ and S on G . In the stochastic and Bayesian models, the PoA of Ξ , given G and D is, respectively,

$$PoA(G, \Xi, D) = \frac{\mathbb{E}_{\mathbf{t} \sim D} [\max_{\mathbf{P} \in \mathcal{N}} c(\mathbf{P}(\mathbf{t}))]}{\mathbb{E}_{\mathbf{t} \sim D} [c(\mathbf{O}(\mathbf{t}))]}; \quad PoA(G, \Xi, D) = \max_{\mathbf{P} \in \mathcal{BN}} \frac{\mathbb{E}_{\mathbf{t} \sim D} [c(\mathbf{P}(\mathbf{t}))]}{\mathbb{E}_{\mathbf{t} \sim D} [c(\mathbf{O}(\mathbf{t}))]},$$

where \mathcal{N} is the set of all NE of the game induced by Ξ and \mathbf{t} on G , and \mathcal{BN} is the set of all BNE of the game induced by Ξ and D on G . In all models the objective of the designer is to come up with protocols that *minimise* the above ratios. Finally, the Price of Anarchy for a class of graph metrics \mathcal{G} , is defined for the adversarial model and for the stochastic (or Bayesian) models, respectively, as

$$PoA(\mathcal{G}) = \max_{G \in \mathcal{G}} \min_{\Xi(G)} PoA(G, \Xi); \quad PoA(\mathcal{G}) = \max_{G \in \mathcal{G}} \min_{\Xi(G, D)} \max_D PoA(G, \Xi, D).$$

9.6 Graph Notation

For every graph G , we denote by $V(G)$ and $E(G)$ the set of vertices and edges of G , respectively. For any $v, u \in V(G)$, (v, u) denotes an edge between v and u and $d_G(v, u)$ denotes the shortest distance between v and u in G ; if G is clear from the context, we simply write $d(v, u)$. A graph G is an *induced* subgraph of

H , if G is a subgraph of H and for every $v, u \in V(G)$, $(v, u) \in E(G)$ if and only if $(v, u) \in E(H)$. G is a *distance preserving (isometric)* subgraph of H , if G is a subgraph of H and for every $v, u \in V(G)$, $d_G(v, u) = d_H(v, u)$.

CHAPTER 10

Design Against the Adversary

The following sections are dedicated in answering the question of when prior knowledge of the underlying metric can help in the design against the adversary. In Section 10.1, we show that for the outerplanar graphs we can design a universal cost-sharing protocol that dramatically improves the PoA against uniform protocols that ignore the metric. In Section 10.2, we prove that for general graph metrics, no universal *ordered* protocol can outperform the best uniform protocols and then in Section 10.3, we generalise this result to include *all* universal protocols.

10.1 Outerplanar Graphs

In this section, we show that there exists a class of graph metrics, prior knowledge of which can dramatically improve the performance of good network cost-sharing design. For *outerplanar* graphs, we provide a universal cost-sharing protocol with constant PoA. In contrast, we stress that uniform protocols *cannot* achieve PoA better than $\Omega(\log k)$, because the lower bound for the greedy algorithm of the Online Steiner Tree problem can be embedded in an outerplanar graph (see Figure 10.1 for an illustration). An *outerplanar* graph is a planar graph where all the vertices belong to the outer face. For a *biconnected*³⁵ *outerplanar* graph the outer face forms a (unique) Hamiltonian cycle.

We next define an ordered cost-sharing protocol, Ξ_{tour} , and we show that it has constant PoA. We describe Ξ_{tour} only for metric spaces that are defined by *biconnected* outerplanar graphs. In order to define Ξ_{tour} for an outerplanar graph

³⁵A graph is biconnected if, after removing any vertex, the graph remains connected.

G that is not biconnected, we first turn it into an *equivalent*³⁶ biconnected graph G^* , by appropriately adding edges of sufficiently high cost h ; we can set h to be a value strictly greater than $\sum_{e \in E(G)} c_e$. Then, equivalence is obvious since we only add edges that cannot be used in either any NE or the minimum Steiner tree outcome. Hence, it is w.l.o.g. to consider only biconnected outerplanar graphs. It is known that every biconnected outerplanar graph admits a *unique* Hamiltonian cycle [134] that can be found in linear time [54].

Definition of Ξ_{tour} : Ξ_{tour} orders the vertices according to the cyclic order in which they appear in the Hamiltonian tour, starting from r and proceeding in a clockwise order π . In Figure 10.1, $\pi(r) < \pi(q_8) < \pi(q_4) < \pi(q_9) < \dots < \pi(q_{15})$.

In the following theorem we show that, for outerplanar graphs, the PoA of Ξ_{tour} is constant and more precisely is upper bounded by 2.

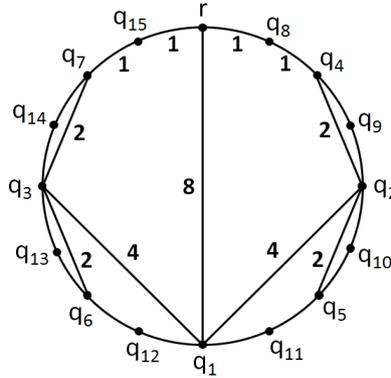


Figure 10.1: The figure shows an example of an outerplanar graph where the order $q_i < q_{i+1}$ results in PoA of $\Omega(\log k)$.

Theorem 47. The PoA of Ξ_{tour} in outerplanar graphs is at most 2.

Proof. Let $G = (V, E, r)$ be any biconnected outerplanar graph and S be the set of activated vertices. Let T^* be the minimum Steiner tree that connects $S \cup \{r\}$ and suppose that is rooted at r . We denote by $P_{T^*}(i, i')$ the unique path from i to i' in T^* and by T_v^* the subtree of T^* rooted at vertex v .

We first show the following claims that will be useful to complete the proof.

Claim 48. For any $i, i', j, j' \in V(T^*)$ such that $\pi(i) < \pi(j) < \pi(i') < \pi(j')$, the paths $P_{T^*}(i, i')$ and $P_{T^*}(j, j')$ are not vertex-disjoint, (i.e. $P_{T^*}(i, i')$ and $P_{T^*}(j, j')$ share a common vertex) .

³⁶We mean that any NE outcome and the minimum Steiner tree solution remain unchanged after the transformation and therefore the PoA of G is exactly the same with the one of G^* .

Proof. Consider the representation of G as a planar graph where the *unique* Hamiltonian tour of G is the outer face, meaning that all the edges of $P_{T^*}(i, i')$ and $P_{T^*}(j, j')$ are either edges of the Hamiltonian tour or chords of it. The Hamiltonian tour defines two paths between i and i' , one containing j and the other containing j' . Hence, the paths $P_{T^*}(i, i')$, $P_{T^*}(j, j')$ have either crossing edges or some common vertex. Due to the planarity of G the first case is excluded and the claim follows. \square

Claim 49. For any two vertex-disjoint subtrees T_v^* and T_u^* of T^* , rooted at vertices v and u , respectively, either all the vertices of T_v^* precedes all the vertices of T_u^* , or the opposite.

Proof. Assume on the contrary that w.l.o.g. there exist $i' \in V(T_v^*)$ and $j, j' \in V(T_u^*)$ such that $\pi(j) < \pi(i') < \pi(j')$. Since T_v^* and T_u^* are vertex-disjoint subtrees, first $r \notin V(T_v^*)$ and $r \notin V(T_u^*)$ and further the paths $P_{T^*}(r, i')$ and $P_{T^*}(j, j')$ should also be vertex-disjoint. Notice, though, that $\pi(r) < \pi(j) < \pi(i') < \pi(j')$ and so, by Claim 48 we end up with a contradiction. \square

Claim 50. For any two vertices $v, u \in V(T^*)$ where v is an ancestor of u , then v either precedes or follows all the vertices of the subtree T_u^* .

Proof. If v is the root then trivially v precedes all the vertices of T_u^* . Suppose now that $v \neq r$ and consider the case that v precedes u (the other case is similar). For the sake of contradiction, assume that there exists a vertex $u' \in V(T_u^*)$ such that u' precedes v and therefore, $\pi(r) < \pi(u') < \pi(v) < \pi(u)$. Note that $P_{T^*}(r, v)$ and $P_{T^*}(u, u')$ are vertex-disjoint and therefore, by Claim 48 we end up with a contradiction. \square

For convenience, we next refer to the set of the activated vertices as $S = \{1, 2, \dots, k\}$, based on their order π , from smaller label to larger, i.e. vertex i has the i^{th} smallest label among S . We further adopt the convention that $r = 0$.

Consider any NE, $\mathbf{P} = (P_i)_{i \in N}$. We bound from above the cost-share of each player at vertex $i \in [k]$ by the cost of the path in T^* that connects her with $i - 1$, i.e.,

$$c_i(\mathbf{P}) \leq c(P_{T^*}(i, i-1)) = \sum_{e \in P_{T^*}(i, i-1)} c_e.$$

Then, by summing over S ,

$$c(\mathbf{P}) = \sum_{i \in [k]} c_i(\mathbf{P}) \leq \sum_{i \in [k]} c(P_{T^*}(i, i-1)) = \sum_{i \in [k]} \sum_{e \in P_{T^*}(i, i-1)} c_e = \sum_{e \in E(T^*)} \sum_{i: e \in P_{T^*}(i, i-1)} c_e.$$

We argue next that by Claims 49 and 50 we can infer that, for each edge e of $E(T^*)$, there exist at most two paths $P_{T^*}(i, i - 1)$ containing e , leading to:

$$c(\mathbf{P}) \leq \sum_{e \in E(T^*)} 2c_e = 2c(T^*).$$

To explain the last argument, consider any edge $e = (v', v) \in E(T^*)$ and let v be the child of v' in T^* . For any vertex $u \notin V(T_v^*)$, either T_v^* and T_u^* are vertex-disjoint, or u is an ancestor of v in T^* . In either case, by Claims 49 and 50, u either precedes or follows *all* vertices of T_v^* . Let $\ell, h \in [k]$ be the vertices of $S \cap V(T_v^*)$ ³⁷ with the lowest and the highest labels, respectively (it is possible that $\ell = h$). It is easy to see that only the paths $P_{T^*}(\ell, \ell - 1)$ and $P_{T^*}(h + 1, h)$ use edge e (the second path exists only if $h < k$). \square

We next demonstrate that our analysis is tight.

Theorem 51. The PoA of Ξ_{tour} in outerplanar graphs is at least 2.

Proof. Consider a cycle graph $C = (V, E, r)$ with $2k$ vertices and unit-cost edges. Let the vertices $V = \{r = 0, 1, 2, \dots, 2k - 1\}$ be named based on their order π , from smaller to larger label. We consider the set of the k activated vertices to be $S = \{k, k + 1, \dots, 2k - 1\}$. It is a NE if the player on vertex k connects with r through the path $(0, 1, \dots, k)$ and each other player on vertex $k + i$, for $i \in \{1, \dots, k - 1\}$, connects with the vertex $k + i - 1$ and follow their path to the root. The cost of this NE is $2k - 1$.

The optimum solution would be to connect S with r through the path $(k, k + 1, \dots, 2k - 1, r)$ with cost k . Therefore, $\text{PoA} \geq \frac{2k-1}{2k}$, which for large k converges to 2. \square

10.2 Lower Bound of Ordered Protocols

The main result of this section is that, for general graph metrics, the PoA of *any* (deterministic) ordered protocol is $\Omega(\log k)$ which is tight. We formally define (Definition 53) the ‘zig-zag’ pattern of the lower bounds of the Greedy Algorithm of the Online Steiner Tree problem (see Example 45(b) and Figure 10.2). Then the main technical challenge is to show that *for any* ordering of the vertices of

³⁷There should exist at least one activated vertex in T_v^* , otherwise $e \notin E(T^*)$.

high dimensional hypercubes, there always exists a *distance preserving* path, such that the order of its vertices follows that zig-zag pattern. Finally, by connecting any two vertices of the hypercube with a direct edge of suitable cost, similar to the example in Figure 8.1, we get the final lower bound construction.

Definition 52 (Classes). For $s > 0$, and for a path graph $P = (v_0, \dots, v_{2^s})$ of $2^s + 1$ vertices, we define a partition of the vertices into $s + 1$ classes, D_0, D_1, \dots, D_s , as follows: Class 0 contains the endpoints of P , $D_0 = \{v_0, v_{2^s}\}$. For every $j \in \{1, \dots, s\}$, $D_j = \{v_i \mid (i \bmod 2^{s-j}) = 0 \text{ and } (i \bmod 2^{s-j+1}) \neq 0\}$.

For better interpretation of the notion of Classes, suppose that we construct P in $s + 1$ steps, where at each step $j \in \{0, \dots, s\}$ we introduce the vertices of class D_j as follows. At step 0 we connect the two endpoints v_0, v_{2^s} via an edge. At step 1 we replace the edge (v_0, v_{2^s}) by a two-length path $(v_0, v_{2^{s-1}}, v_{2^s})$, i.e. we place the vertex of D_1 between the existing vertices. Repeatedly, at each step j we replace each of the current edges by a two-length path, in the middle of which we place a vertex of D_j . As an example, consider the path $P = (v_0, v_1, \dots, v_8)$ of Figure 10.2, where $s = 3$. Then, $D_0 = \{v_0, v_8\}$, $D_1 = \{v_4\}$, $D_2 = \{v_2, v_6\}$ and $D_3 = \{v_1, v_3, v_5, v_7\}$. Note that always $|D_0| = 2$ and for $j \neq 0$, $|D_j| = 2^{j-1}$.

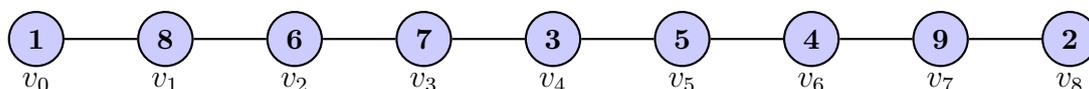


Figure 10.2: An example of a $P_3(\pi)$ path. The numbers correspond to the labels.

For $j > 0$ and $v_i \in D_j$, we define the *parents* of v_i as $\{w \mid d_P(v_i, w) = 2^{s-j}\}$, i.e. the closest vertices that belong to lower-labeled classes. In other words, if we consider the construction of P in $s + 1$ steps as above, the parents of v_i are its neighbours after step j . For example, for the path of Figure 10.2, the parents of v_4, v_5, v_6 are respectively the sets $\{v_0, v_8\}$, $\{v_4, v_6\}$, $\{v_4, v_8\}$. Remark that for all $v \notin \{v_0, v_{2^s}\}$, v has two parents belonging to lower-labeled classes than v and all vertices between v and any of its parents belong to higher-labeled classes than v . We are now ready to define the “zig-zag” pattern.

Definition 53 (Zig-zag pattern). We call a path graph, $P = (v_0, v_1, \dots, v_{2^s})$, with distinct integer labels π , *zig-zag*, and we denote it by $P_s(\pi)$, if for every $v \notin \{v_0, v_{2^s}\}$, v has greater label than both its parents w_1, w_2 , i.e. $\pi(v) > \pi(w_1)$ and $\pi(v) > \pi(w_2)$.

An example of such a path for $s = 3$ is shown in Figure 10.2. Our main result of this section is that there exist graphs, high dimensional hypercubes, such that for any order π , $P_s(\pi)$ always appears as a *distance preserving* subgraph. Our proof is existential and uses Ramsey theory.

Example 54. In order to give some intuition, we will first use a Ramsey-type result due to Alon et al. [5] to show that, for any π , $P_2(\pi)$ appears as a subgraph. Alon et al. [5] showed that for any given integer $\ell \geq 5$, any edge-colouring of a sufficiently high dimensional hypercube contains a monochromatic cycle of length 2ℓ . Let Q_n be the hypercube of [5] for $\ell = 5$, and notice that it is bipartite i.e. $Q_n = (A, B, E)$. For any ordering of vertices of Q_n we define a colouring as follows: for any edge (v, u) , with $v \in A$ and $u \in B$, if $\pi(v) < \pi(u)$, we paint the edge blue, otherwise we paint it red.

Suppose w.l.o.g. that the monochromatic cycle C_{10} of length 10 is blue (see also Figure 10.3 for an illustration). Then, for any $v \in A \cap V(C_{10})$ (continuous circles), its neighbours in C_{10} should have *higher* label (dashed circles). The vertices of $A \cap V(C_{10})$ can be 1-embedded into a cycle C_5 of length 5 (dotted cycle). We appropriately choose three consecutive vertices of C_5 , such that the middle one has higher ranking than the others in π ((5, 8, 1) in Figure 10.3). It is not hard to see that such a triplet is guaranteed because C_5 is a cycle. These three vertices with their intermediate ones in C_{10} form a path isomorphic to $P_2(\pi)$; that path in Figure 10.3 is the (5, 10, 8, 9, 1) .

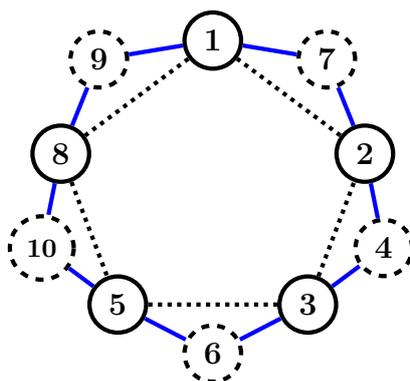


Figure 10.3: An example of $P_2(\pi)$ by using the result of [5].

There are two limitations in using the results of [5] in our proof. a) The induced monochromatic cycle of any length can only be used in order to prove the existence of a zig-zag pattern of length 4 and it doesn't help for paths of higher lengths as required for our lower bound. b) The induced zig-zag pattern is *not*

necessarily distance-preserving, because the monochromatic cycle derived by [5] is *not distance preserving*, which is a crucial property for our lower bound to hold. Therefore, in order to overcome those limits, we prove a similar Ramsey-type result, but for a different monochromatic subgraph with some special properties (to be described in Section 10.2.1).

Proof Overview The proof is by induction and in the inductive step our starting point is the n -th dimensional hypercube Q_n . Given an ordering/labelling π of the vertices of Q_n we first show that Q_n contains a subgraph W which is isomorphic to a ‘pseudo-hypercube’ Q_m^2 ($m < n$) where the labelling of its vertices satisfies a special property (to be described shortly). Q_m^2 is defined by replacing each edge of Q_m by a 2-edge path (of length two)³⁸.

Labelling property: For the subgraph W we require that all such newly formed 2-edge paths, are $P_1(\pi)$ paths, i.e. the label of the middle vertex is greater than the labels of the endpoints (Figure 10.4(a) shows such a labelling).

Next, we contract all such 2-edge paths of Q_m^2 into single edges, resulting in a graph isomorphic to Q_m ; this is the hypercube used for the next step. Note that each contracted edge still corresponds to a path in Q_n . Therefore, after s recursive steps, each edge corresponds to a 2^s path of Q_n . Further, note that such a path is a $P_s(\pi)$ path, due to the labelling property that we preserve at each step. We require that, at the end of the last inductive step, $Q_m = Q_1$ (a single edge), and (by unfolding it) we show that this edge corresponds to a distance preserving subgraph of the original graph/hypercube. At each step, we have $m < n$ and the relation between n and m is determined by a Ramsey-type argument.

We next describe the basic ingredients that we use to show existence of W . We apply a colouring scheme to the edges of Q_n that depends on the vertices’ order.

Colouring Scheme: Consider Q_n as a bipartite graph $Q_n = (A, B, E)$. For any edge (v, u) , with $v \in A$ and $u \in B$, if the v ’s label is smaller than u ’s, we paint the edge blue, otherwise we paint it red.

By a Ramsey-type argument we show that Q_n has a *monochromatic* subgraph isomorphic to a specially defined graph G_m ; G_m is carefully specified in such a way that it contains at least two subgraphs isomorphic to pseudo-hypercubes Q_m^2 . The special property of those two subgraphs is described next.

³⁸See Q_m^2 of Definition 57 and Figure 10.4(a) for an illustration

Let H_1 and H_2 be the two half cubes³⁹ of Q_n and let $V(H_1) = A$ and $V(H_2) = B$. Observe that if Q_m^2 is a subgraph of Q_n then the corresponding Q_m is an induced subgraph of either H_1 or H_2 . We carefully construct G_m such that it contains subgraphs W_1 and W_2 isomorphic to Q_m^2 , whose corresponding Q_m 's are induced subgraphs of H_1 and H_2 , respectively. The colour of G_m determines which of the W_1 and W_2 will serve as the desired W . In particular, if the colour is blue, then for every edge (v, u) , with $v \in V(H_1)$ and $u \in V(H_2)$, it should hold that v 's label is smaller than u 's and therefore the labelling property is satisfied for W_1 ; similarly, if the colour is red, W_2 serves as W .

Proof Roadmap The whole proof of the lower bound proceeds in several steps in the following sections. In Section 10.2.1 we give the formal definition of the subgraph G_m . Section 10.2.2 is devoted to show that every 2-edge colouring of a (suitably) high dimensional hypercube contains a monochromatic copy of G_m (Lemma 56), by using Ramsey theory. Then, in Section 10.2.3 we show that, for any ordering of the vertices of Q_n , we can define a special 2-edge-colouring, so that there exists a Q_m^2 subgraph of G_m that preserves the Labelling property (Lemma 58). At last, in Section 10.2.4, by a recursive application of the combination of the Ramsey-type result and the colouring, we prove the existence of the *zig-zag* path in high dimensional hypercubes (Theorem 59). We then show how to construct a graph that serves as lower bound for all ordered protocols (Theorem 61). This is done by connecting any two edges of the hypercube with a direct edge of appropriate cost, similar to the example in Figure 8.1.

Definitions and notation on Hypercubes We denote by $[a, b]$ (for $a \leq b$) the set of integers $\{a, a + 1, \dots, b - 1, b\}$, but when $a = 1$, we simply write $[b]$. We follow definitions and notation of [5]. Let Q_n be the graph of the n -dimensional hypercube whose vertex set is $\{0, 1\}^n$. We represent a vertex v of $V(Q_n)$ by an n -bit string $x = \langle x_1 \dots x_n \rangle$, where $x_i \in \{0, 1\}$. By $\langle xy \rangle$ or xy we denote the concatenation of an a -bit string x with an b -bit string y , i.e. $xy = \langle x_1 \dots x_a y_1 \dots y_b \rangle$. $x = \langle x_j \rangle_{j=1}^a$ is the concatenation of its a bits. An edge is defined between any two vertices that differ only in a single bit. We call this bit, *flip-bit*, and we denote it by ‘*’. For example, $x = \langle 11100 \rangle, y = \langle 11000 \rangle$ are two vertices of Q_5 and $(x, y) = \langle 11 * 00 \rangle$ is the edge that connects them. The distance between two vertices x, y is defined by their Hamming distance,

³⁹The two half-cubes of order n are formed from Q_n by connecting all pairs of vertices with distance exactly two and dropping all other edges.

$d(v, u) = |\{j : x_j \neq y_j\}|$. For a fixed subset of coordinates $R \subseteq [n]$, we extend the definition of the distance as follows,

$$d(x, y, R) = \begin{cases} d(x, y), & \text{if } \forall j \in R, x_j = y_j \\ \infty, & \text{otherwise.} \end{cases}$$

We define the *level* of a vertex x by the number of ‘ones’ it contains, $w(x) = \sum_{i=1}^n x_i$. We denote by L_i the set of vertices of level $i \in [0, n]$. We define the *prefix sum* of an edge $e = (x, y)$, where the flip-bit is in the j -th coordinate, by $p(e) = \sum_{i=1}^{j-1} x_i$. We represent any ordering π of $V(Q_n)$, by labelling the vertices with labels $1, \dots, 2^n$, where label i corresponds to ranking i in π .

10.2.1 Description of G_m

For a positive integer m , we define a graph $G_m = (V_m, E_m)$ that is an induced subgraph of Q_{4m} on $V_m = V_1 \cup V_2 \cup V_3 \subseteq V(Q_{4m})$. A vertex of V_1 is defined by $2m - 1$ concatenations of pairs $\langle 01 \rangle$ and $\langle 10 \rangle$ and a single pair $\langle 00 \rangle$ that appears in the *second* half of the string. A vertex of V_2 is defined by $2m$ concatenations of $\langle 01 \rangle$ and $\langle 10 \rangle$. A vertex of V_3 is defined by $2m - 2$ concatenations of $\langle 01 \rangle$ and $\langle 10 \rangle$, one pair $\langle 11 \rangle$ that appears on the first half of the string, and one pair $\langle 00 \rangle$ that appears on the second half. For example, for $m = 2$, $\langle 01 \ 10 \ 00 \ 10 \rangle \in V_1$, $\langle 01 \ 10 \ 10 \ 10 \rangle \in V_2$, $\langle 01 \ 11 \ 10 \ 00 \rangle \in V_3$. More formally, let $A = \{\langle 01 \rangle, \langle 10 \rangle\}$, then the subsets V_1, V_2, V_3 are defined as follows:

$$\begin{aligned} V_1 &:= V_1(m) = \{\langle a_j b_j \rangle_{j=1}^{2m} \mid \exists i \in [m+1, 2m] \text{ s.t.} \\ &\quad \langle a_i b_i \rangle = \langle 00 \rangle \text{ and } \forall j \neq i, \langle a_j b_j \rangle \in A\}, \\ V_2 &:= V_2(m) = \{\langle a_j b_j \rangle_{j=1}^{2m} \mid \forall j, \langle a_j b_j \rangle \in A\}, \\ V_3 &:= V_3(m) = \{\langle a_j b_j \rangle_{j=1}^{2m} \mid \exists i_1 \in [m], \exists i_2 \in [m+1, 2m] \text{ s.t.} \\ &\quad \langle a_{i_1} b_{i_1} \rangle = \langle 11 \rangle, \langle a_{i_2} b_{i_2} \rangle = \langle 00 \rangle \text{ and } \forall j \neq i_1, i_2, \langle a_j b_j \rangle \in A\}. \end{aligned}$$

Observe that G_m is bipartite with vertex partitions V_1 and $V_2 \cup V_3$, as vertices of V_1 belong to level $2m - 1$, while vertices of $V_2 \cup V_3$ to level $2m$.

Lemma 55. Every pair of vertices $x, x' \in V_1(m)$ with $d(x, x', [2m+1, 4m]) = 2$, have a unique common neighbour $y \in V_3(m)$. Also, every pair of vertices $x, x' \in V_2(m)$, with $d(x, x', [2m]) = 2$, have a unique common neighbour $y \in V_1(m)$.

Proof. Recall that (by definition) if $d(x, x', R) \neq \infty$ then x, x' should coincide in *all* R coordinates. For the first statement, observe that the premises of the

Lemma hold only if there exists $i \in [m]$ such that $x_{2i-1}x_{2i} = \langle 10 \rangle$ and $x'_{2i-1}x'_{2i} = \langle 01 \rangle$ (or the other way around), in which case the required vertex y from $V_3(m)$ has $y_{2i-1}y_{2i} = \langle 11 \rangle$; the rest of the bits are the same among x, x', y . For the second statement, the premises of the Lemma hold only if there exists an $i \in [m+1, 2m]$ such that $x_{2i-1}x_{2i} = \langle 10 \rangle$ and $x'_{2i-1}x'_{2i} = \langle 01 \rangle$ (or the other way around), in which case the required vertex y from $V_1(m)$ has $y_{2i-1}y_{2i} = \langle 00 \rangle$ and the rest of the bits are the same among x, x', y . \square

10.2.2 Ramsey-type Theorem

Lemma 56. For any positive integer m , and for sufficiently large $n \geq n_0 = g(m)$, any 2-edge colouring χ of Q_n , contains a monochromatic copy of G_m ⁴⁰.

Proof. The proof follows ideas of Alon et al. [5]. Consider a hypercube Q_n , with sufficiently large $n > 6m$ to be determined later, and some arbitrary 2-edge-colouring $\chi : E(Q_n) \rightarrow \{1, 2\}$. Let E^* be the set of edges between vertices of L_{4m-1} and L_{4m} (recall that $L_i = \{v | w(v) = i\}$).

Each edge $e \in E^*$ contains $4m - 1$ 1's, a flip-bit represented by $*$ and the rest of the coordinates are 0. Moreover, e is uniquely determined by its $4m$ non-zero coordinates $R_e \subseteq [n]$ and its prefix sum $p(e) \in [0, 4m - 1]$ (number of 1's before the flip-bit). Therefore, the colour $\chi(e)$ defines a colouring of the pair $(R_e, p(e))$, i.e. $\chi(e) = \chi(R_e, p(e))$. For each subset $R \subset [n]$ of $4m$ coordinates, we denote by $c(R) = (\chi(R, 0), \dots, \chi(R, 4m - 1))$ the colour induced by the edge colouring. The colouring of all subsets R defines a colouring of the complete $4m$ -uniform hypergraph of $[n]$ ⁴¹ using 2^{4m} colours.

By Ramsey's Theorem for hypergraphs [78], there exists $n_0 = g(m)$ such that for any $n \geq n_0$ there exists some subset $U \subset [n]$ of size $6m$ such that all $4m$ -subsets $R \subset U$ have the same colour $c(R) = c^*$. Therefore, for every $4m$ -subsets $R_1, R_2 \subset U$ and $p \in [0, 4m - 1]$, it is $\chi(R_1, p) = \chi(R_2, p) = c_p$. Since p takes $4m$ values and there are only two different colours, there must exist $2m$ indices $p_0, \dots, p_{2m-1} \in [0, 4m - 1]$ with the same colour $\chi(R, p_i) = \chi^*$, for all $R \subset U$, $|R| = 4m$ and $i \in [0, 2m - 1]$.

It remains to show that the graph formed by those edges contains a monochromatic copy of G_m . We will show this by placing the bits of each edges from E_m (the set of edges of G_m) to suitable coordinates of $[n]$ and filling the rest of the

⁴⁰The result could be extended to any (fixed) number of colours, but we need only two for our application.

⁴¹A k -uniform hypergraph is a hypergraph such that all its hyperedges have size k .

coordinates suitably by zeros and ones. More precisely, we insert blocks of 1's of suitable length among the bits of the edges of E_m , and all those bits are placed at the coordinates of U . The rest of the bits ($n - |U|$) are set to zero.

Let 1^r be a string of r 1's and define $\beta_i = 1^{p_i - p_{i-1} - 1}$ for $i \in [2m - 1]$, $\beta_0 = 1^{p_0}$ and $\beta_{2m} = 1^{4m - 1 - p_{2m-1}}$. For any edge $e = \langle a_j b_j \rangle_j \in E_m$, we insert β_0 at the beginning of the string, for $j \in [m]$ we insert β_j between a_j and b_j and for $j \in [m + 1, 2m]$ we insert the string β_j after b_j . The following illustrates these insertions:

$$\underbrace{1 \dots 1}_{p_0} a_1 \underbrace{1 \dots 1}_{p_1 - p_0 - 1} b_1 a_2 \dots a_m \underbrace{1 \dots 1}_{p_m - p_{m-1} - 1} b_m a_{m+1} b_{m+1} \underbrace{1 \dots 1}_{p_{m+1} - p_m - 1} a_{m+2} b_{m+2} \dots$$

$$\dots a_{2m} b_{2m} \underbrace{1 \dots 1}_{4m - 1 - p_{2m-1}}$$

Recall that each edge of E_m contains exactly $2m$ zero bits and $2m$ non-zero bits. Also notice that $\sum_j |\beta_j| = p_0 + \sum_{i=1}^{2m-1} (p_i - p_{i-1} - 1) + 4m - 1 - p_{2m-1} = -(2m - 1) + 4m - 1 = 2m$. Therefore, in total we have $6m$ bits (same as the size of U) and $4m$ non-zero bits (same as the size of R). We place these $6m$ bits precisely at the coordinates of U . The rest $n - 6m$ of the coordinates are filled with zeros.

It remains to show that for such edges the prefix of the flip-bit is *always* one of the p_0, \dots, p_{2m-1} . This would imply that all these edges are monochromatic. Furthermore, all but $4m$ coordinates are fixed and the $4m$ coordinates form exactly the sets $V_1(m), V_2(m), V_3(m)$; therefore, the monochromatic subgraph is isomorphic to G_m .

For any edge $e = \langle a_j b_j \rangle_j \in E_m$, let the flip-bit be at position:

- a_j for $j \in [m]$. Its prefix is $\sum_{i=0}^{j-1} |\beta_i| + (j - 1) = p_{j-1}$, where the term $j - 1$ corresponds to the number of pairs $\langle a_s b_s \rangle$ with $s < j$, each of which contributes to the prefix with a single 1.
- b_j for $j \in [m]$. Since $j \leq m$, $a_j = 1$. Then the prefix equals to $\sum_{i=0}^j |\beta_i| + (j - 1) + 1 = p_j$.
- a_j or b_j for $j \in [m + 1, 2m]$. For such j , $\langle a_j b_j \rangle \in \{ \langle 0* \rangle, \langle *0 \rangle \}$ and all other pairs belong to A . Therefore, the prefix is equal to $\sum_{i=0}^{j-1} |\beta_i| + (j - 1) = p_{j-1}$.

□

10.2.3 Colouring Based on the Labels

This part of the proof shows that for any ordering of the vertices of a hypercube Q_n , there is a 2-edge colouring with the following property: in the monochromatic G_m , either all the vertices of V_1 or all the vertices of V_2 have neighbours in G_m with only higher label. This implies a desired labelling property for a Q_m^2 subgraph of Q_n , the structure of which is defined next.

Definition 57. We define Q_n^r to be a subdivision of Q_n , by replacing each edge by a path of length s . Q_n^1 is simply Q_n . We denote by $Z(Q_n^r)$ the set of all pairs of vertices (x, x') , which correspond to edges of Q_n ; $P(x, x')$ is the corresponding path in Q_n^r . For every $(x, x') \in Z(Q_m^2)$, we denote by $\theta(x, x')$ the middle vertex of $P(x, x')$.

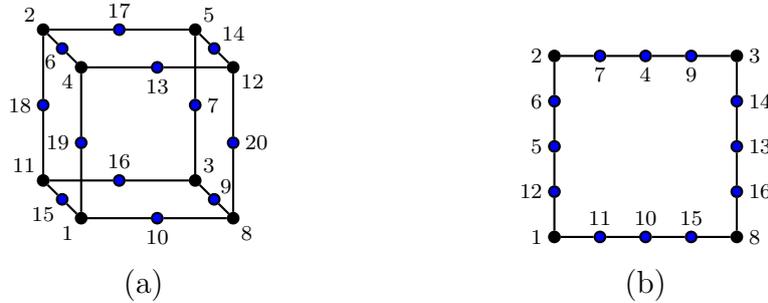


Figure 10.4: Examples of (a) Q_3^2 and (b) Q_2^4 . The labels on the nodes are examples of the labelling property, (a) after one inductive step, (b) after two inductive steps.

In the next lemma we show that for any ordering of the vertices of Q_n , there exists a subgraph isomorphic to Q_m^2 , such that the ‘middle’ vertices have higher label than their neighbours (Labelling Property).

Lemma 58. For any positive integer m , for all $n \geq n_0 = g(m)$ and for any ordering π of $V(Q_n)$, there exists a subgraph W of Q_n that is isomorphic to Q_m^2 , such that, for every $(x, x') \in Z(W)$, it is $\pi(\theta(x, x')) > \max\{\pi(x), \pi(x')\}$.

Proof. Choose a sufficiently large $n \geq n_0 = g(m)$ as in Lemma 56. Partition the vertices of Q_n into sets \mathcal{O}, \mathcal{E} of vertices of odd and even level, respectively. We colour the edges of Q_n as follows. For every edge $e = (z, z')$ with $z \in \mathcal{O}$ and $z' \in \mathcal{E}$, if $\pi(z) < \pi(z')$, then paint e blue. Otherwise paint it red. Therefore, for every blue edge, the endpoint in \mathcal{O} has smaller label than the endpoint in \mathcal{E} . The opposite holds for any red edge.

Lemma 56 implies that Q_n contains a monochromatic copy (blue or red) of G_m . Recall that this subgraph is bipartite between vertices of levels L_{4m-1} and L_{4m} and that $V_1 \subset L_{4m-1} \subset \mathcal{O}$ and $V_2 \cup V_3 \subset L_{4m} \subset \mathcal{E}$. Let $R \subset [n]$ be the subset of the $4m$ coordinates that correspond to vertices of G_m . Also let R_1 and R_2 be the subsets of the first $2m$ and the last $2m$ coordinates of R , respectively.

First suppose that the subgraph isomorphic to G_m is blue. An immediate implication of our colouring is that for every edge $(z, z') \in E_m$ with $z \in V_1$, $z' \in V_2 \cup V_3$ it must be $\pi(z) < \pi(z')$. Fix a $2m$ -bit string, str , that corresponds to a permissible bit assignment to the R_2 coordinates of some vertex in V_1 (see Section 10.2.1). Define W_{str} as the subset of vertices of V_1 where the R_2 coordinates are set to str . Recall that each of the first m pairs $\langle a_j b_j \rangle$, $j \in [m]$, of a vertex $z \in W_{str}$, may take any of the two bit assignments $\langle 01 \rangle$ and $\langle 10 \rangle$. Hence, $|W_{str}| = 2^m$.

Observe that we can embed W_{str} into Q_m with distortion 1 and scaling factor $1/2$, by mapping the first m pairs of bits into single bits; map $\langle 01 \rangle$ to 1 and $\langle 10 \rangle$ to 0. Every two vertices with distance d in Q_m , have distance $2d$ in Q_n . For every $x, x' \in W_{str} \subset V_1$ with $d(x, x') = 2$, it holds that $d(x, x', R_2) = 2$, since x, x' have the same R_2 coordinates. Lemma 55 implies that there exists $y = \theta(x, x') \in V_3$, such that $d(x, y) = d(x', y) = 1$, and therefore, $\pi(y) > \max\{\pi(x), \pi(x')\}$. Take the union $Y = \cup_y$ of all such vertices y , then $W_{str} \cup Y$ induces a subgraph W isomorphic to Q_m^2 , that fulfils the labelling requirements.

The case of G_m being red is similar. We focus only on the vertices of V_2 . Fix now a $2m$ -bit string, str , that corresponds to a permissible bit assignment of the R_1 coordinates of a vertex in V_2 . Define W_{str} as the subset of vertices of V_2 where the R_1 coordinates are set to str . Similarly, we can embed W_{str} into Q_m with distortion 1 and scaling factor $1/2$.

For every $x, x' \in W_s \subset V_2$ with $d(x, x') = 2$, where the R_1 coordinates are fixed to str , Lemma 55 implies that there exists $y = \theta(x, x') \in V_1$, such that $d(x, y) = d(x', y) = 1$, and therefore, $\pi(y) > \max\{\pi(x), \pi(x')\}$. Take the union $Y = \cup_y$ of all such vertices y , then $W_{str} \cup Y$ induces a subgraph W isomorphic to Q_m^2 , that fulfils the labelling requirements. \square

10.2.4 Lower Bound Construction

Now we are ready to prove the main theorem of this section.

Theorem 59. For every positive integer s , and for sufficiently large $n = n(s)$, there exists a graph Q_n such that, for *any* (deterministic) ordering π of its ver-

tices, it contains a zig-zag *distance preserving* path $P_s(\pi)$.

Proof. Let g be a function as in Lemma 56. We recursively define the sequence n_0, n_1, \dots, n_s , such that $n_s = 1$ and $n_{i-1} = g(n_i)$, for $i \in [s]$. We will show that Q_{n_0} ($n_0 = n(s)$) is the graph we are looking for.

Claim 60. For every $i \in [0, s]$, and for any vertex ordering π of Q_{n_0} , Q_{n_0} contains a subgraph isomorphic to $Q_{n_i}^{2^i}$, such that for every $(x, x') \in Z(Q_{n_i}^{2^i})$, $P(x, x')$ is a zig-zag path $P_i(\pi)$.

Proof. The proof is by induction on i . As a base case, $Q_{n_0}^{2^0} = Q_{n_0}$ is the graph itself. An edge is trivially a path $P_0(\pi)$, for any π . Suppose now that Q_{n_0} contains a subgraph isomorphic to $Q_{n_i}^{2^i}$, for some $i < r$, such that for every $q \in Z(Q_{n_i}^{2^i})$, $P(q)$ is a zig-zag path $P_i(\pi)$. It is sufficient to show that $Q_{n_i}^{2^i}$ contains a subgraph isomorphic to $Q_{n_{i+1}}^{2^{i+1}}$, such that for every $q \in Z(Q_{n_{i+1}}^{2^{i+1}})$, $P(q)$ is a zig-zag path $P_{i+1}(\pi)$.

For every $(x, x') \in Z(Q_{n_i}^{2^i})$, if we replace $P(x, x')$ with a direct edge $e = (x, x')$, the resulting graph is a copy of Q_{n_i} . Applying Lemma 58 on Q_{n_i} , guarantees the existence of a subgraph W isomorphic to $Q_{n_{i+1}}^2$ ($n_i = g(n_{i+1})$), where for every $(y, y') \in Z(W)$, $\pi(\theta(y, y')) > \max\{\pi(y), \pi(y')\}$. Each of the edges $(y, \theta(y, y'))$ and $(y', \theta(y, y'))$ of $Q_{n_{i+1}}^2$ are replaced by a path $P_i(\pi)$ in $Q_{n_i}^{2^i}$. Therefore, W is a copy of $Q_{n_{i+1}}^{2^{i+1}}$, with $P(y, y')$ being a zig-zag path $P_{i+1}(\pi)$. \square

We now argue that the resulting $P_s(\pi)$ is a distance preserving path. Our analysis indicate a sequence of hypercubes $Q_{n_0}, Q_{n_1}, \dots, Q_{n_s}$. Recall that in Lemma 58, in order to get $Q_{n_{i+1}}$ from Q_{n_i} we mapped $\langle 01 \rangle$ to 1 and $\langle 10 \rangle$ to 0 and the vertices of $Q_{n_{i+1}}$ did not differ in any other bit but the ones we mapped. Consider now the two vertices x, x' of $Q_{n_r} = Q_1$ with bit-strings $\langle 0 \rangle$ and $\langle 1 \rangle$, respectively. Their Hamming distance in their original bit representation (in Q_{n_0}) should be 2^s , the same with their distance in $P_s(\pi)$.

For instance, for $s = 4$, Table 10.1 shows the bit sequences in $Q_{n_3}, Q_{n_2}, Q_{n_1}$ and Q_{n_0} that correspond to the bits $\langle 0 \rangle$ and $\langle 1 \rangle$ of the vertices x, x' of $Q_{n_4} = Q_1$. In any Q_{n_i} , both bit sequences occupy exactly the same coordinates. The rest of

Table 10.1: Example of unfolding the bit mapping.

	Q_{n_4}	Q_{n_3}	Q_{n_2}	Q_{n_1}	Q_{n_0}
x	$\langle 0 \rangle$	$\langle 10 \rangle$	$\langle 0110 \rangle$	$\langle 10010110 \rangle$	$\langle 0110100110010110 \rangle$
x'	$\langle 1 \rangle$	$\langle 01 \rangle$	$\langle 1001 \rangle$	$\langle 01101001 \rangle$	$\langle 1001011001101001 \rangle$

the coordinates of x, x' are occupied by identical bits in all bit representations. Therefore, $d_{Q_{n_0}}(x, x') = 16 = 2^s$.

Moreover, if any two vertices of $P_s(\pi)$ are closer in Q_{n_0} than in $P_s(\pi)$, then this would contradict the fact that $d_{Q_{n_0}}(x, x') = 2^s$. \square

Finally we extend Q_n so that for any order π of its vertices, a path $P_s(\pi)$ exists along with the shortcuts similar to the example of Figure 8.1.

Theorem 61. Any (deterministic) ordered protocol on undirected graphs admits a PoA of $\Omega(\log k)$, where k is the number of activated vertices.

Proof. Let $k = 2^s + 1$ for some positive integer s . From Theorem 59, we know that for any vertex ordering π of $Q_{n(s)}$ there is a distance preserving path $P_s(\pi)$.

We use $Q_{n(s)}$ as a basis to construct the weighted graph $\tilde{Q}_{n(s)}$ with vertex set $V(\tilde{Q}_{n(s)}) = Q_{n(s)} \cup \{r\}$, where r is the designated root. We connect every pair of vertices x, y with a direct edge of cost $c_e = 2^s$, if t is one of its endpoints, otherwise its cost is $c_e = d_{Q_{n(s)}}(x, y)$ (similar to Figure 8.1).

The adversary selects to activate the vertices of $P_s(\pi)$, and the lower bound follows; in the NE the players choose their direct edges to connect with one of their parents (see at the beginning of Section 10.2 for the term “parent”). \square

10.3 Lower Bound for All Universal Protocols

In this section, we exhibit metric spaces for which no universal cost-sharing protocol admits a PoA better than $\Omega(\log k)$. Due to the characterisation of [36]⁴², we can restrict ourselves in generalised weighted Shapley protocols (GWSPs). We follow the notation of [36], and for the sake of self-containment we include here the most related definitions and lemmas.

⁴²Chen, Roughgarden and Valiant [36] characterise the linear protocols (i.e. for every edge e of cost $c_e \geq 0$, a linear protocol assigns the method $c_e \cdot \xi$, where ξ is the method it assigns to any edge of unit cost) satisfying the three properties of budget-balance, separability and stability to be the generalised weighted Shapley protocols. They further showed that for any non-linear protocol, there exists a linear one with at most the same PoA.

10.3.1 Cost-Sharing Preliminaries

A strictly positive function $f : 2^N \rightarrow \mathbb{R}^+$ is an *edge potential* on N , if it is strictly increasing, i.e. for every $R \subset S \subseteq N$, $f(R) < f(S)$, and for every $S \subseteq N$,

$$\sum_{i \in S} \frac{f(S) - f(S \setminus \{i\})}{f(\{i\})} = 1.$$

For simplicity, instead of $f(\{i\})$, we write $f(i)$. A cost-sharing protocol is called *potential-based*, if it is defined by assigning to each edge of cost c , the cost-sharing method ξ , where for every $S \subseteq N$ and $i \in S$,

$$\xi(i, S) = c \cdot \frac{f(S) - f(S \setminus \{i\})}{f(i)}.$$

Let Ξ_1 and Ξ_2 be two cost-sharing protocols for *disjoint* sets of vertices U_1 and U_2 , with methods ξ_1 and ξ_2 , respectively. The *concatenation* of Ξ_1 and Ξ_2 is the cost sharing protocol Ξ of the set $U_1 \cup U_2$, with method ξ defined as

$$\xi(i, S) = \begin{cases} \xi_1(i, S \cap U_1) & \text{if } i \in U_1, \\ \xi_2(i, S) & \text{if } S \subseteq U_2, \\ 0 & \text{otherwise.} \end{cases}$$

Note that the concatenation of two protocols for disjoint sets of vertices defines an order among these two sets. The GWSPs are concatenations of potential-based protocols.

Lemma 62. (Lemma 4.10 of [36]). Let f be an edge potential on N and ξ the induced (by f) cost-sharing method, for unit costs. For $k \geq 1$ and a constant α , with $1 \leq \alpha^{2k} \leq 1 + k^{-3}$, let $S \subseteq N$ be a subset of vertices with $f(i) \leq \alpha f(j)$, for every $i, j \in S$. If $|S| \leq k$, then for any $i, j \in S$,

$$\xi(i, S) \leq \alpha(\xi(j, S) + 2k^{-2}).$$

Lemma 63. (Lemma 4.11 of [36]). Let f be an edge potential on N , and ξ be the cost-sharing method induced by f , for unit cost. For any two vertices $i, j \in N$, such that $f(i) \geq \beta f(j)$, $\xi(i, \{i, j\}) \geq \beta/(\beta+1)$, and for every $S \supseteq \{i, j\}$, $\xi(j, S) \leq 1/(\beta+1)$.

10.3.2 Lower Bound

The following two technical lemmas will be used in our main theorem.

Lemma 64. Let X be a finite set of size msw^2 , for some $m, s, w > 0$, and X_1, \dots, X_m be a partition of X , with $|X_i| = sw^2$, for all $i \in [m]$. Then, for any colouring χ of X such that *no more* than w elements have the same colour, there exists a *rainbow* subset $S \subset X$ (i.e. $\chi(v) \neq \chi(u)$ for all $v, u \in S$), with $|S \cap X_i| = s$ for every $i \in [m]$.

Proof. Given the partition X_1, \dots, X_m of X and the colouring χ , we construct a bipartite graph $G = (A, B, E)$, where A is the set of colours used in χ . For every X_i we create a set B_i of size s ; then $B = \cup B_i$. If colour j is used in X_i , we add an edge (j, l) for all $l \in B_i$.

Each colour $j \in A$ appears in at most w distinct X_i sets, and since for each X_i there are s vertices (B_i), the degree of j is at most sw . On the other hand, each X_i has size sw^2 and hence, it has at least sw different colours. Therefore, the degree of each vertex of B is at least sw .

Consider any set $R \subseteq B$, and let $E(R)$ be the set of edges with at least one endpoint in R . If $N(R)$ denotes the set of neighbours of R , observe that $E(R) \subseteq E(N(R))$. By using the degree bound on vertices of B , $|E(R)| \geq sw|R|$ and by using the degree bound on vertices of A , $|E(N(R))| \leq sw|N(R)|$. Therefore, $|R| \leq |N(R)|$. By Hall's Theorem there exists a matching which covers every vertex in B . Each vertex in B_i is matched with a distinct colour and therefore in each X_i there exists a subset with at least s elements with distinct colours; let W_i be such a subset with exactly s elements. In addition the colours in different W_i subsets should be distinct by the matching. Then, $S = \cup W_i$. \square

Lemma 65. Let $X = (X_1, \dots, X_m)$ be a partition of $[m^2]$, with $|X_i| = m$, for all $i \in [m]$. Then, there exists a subset $S \subset [m^2]$ with exactly one element from each subset X_i , such no two distinct $x, y \in S$ are consecutive, i.e. for every $x, y \in S$, $|x - y| \geq 2$.

Proof. For every i , let $X_i = \{x_{i1}, \dots, x_{im}\}$. W.l.o.g we can assume that the x_{ij} 's are in increasing order with respect to j and in addition that X_i 's are sorted such that $x_{ii} < x_{ji}$, for all $j > i$ (otherwise rename the elements recursively to fulfil the requirement). Then, it is not hard to see that $S = \{x_{ii} | i \in [m]\}$ can serve as the required set. \square

Now we proceed with the main theorem of this section. We create a graph where every GWSP has high PoA. At a high level, we construct a high dimensional hypercube with sufficiently large number of potential players at each vertex (by adding many copies of each vertex connected via zero-cost edges). Moreover, we add shortcuts among the vertices of suitable costs and we connect each vertex with the root r via two parallel links with costs that differ by a large factor (see Figure 10.5). If the protocol induces a large enough set of potential players with *Shapley-like* values in some vertex, then it is a NE that all these players follow the most costly link to r . Otherwise, by using Lemmas 64 and 65 we show that there exists a set of potential players with *ordered-like* values, one at each vertex of the hypercube. Then, by using the results of Section 10.2, there exists a path where the vertices are *zig-zag*-ordered.

The separation into these two extreme cases was first used in [36]. The crucial difference, is that for their problem the protocol is specified independently of the underlying graph, and therefore the adversary knows the case distinction (ordered or Shapley) and bases the lower bound construction on that. However, our problem requires more work as the graph should be constructed *in advance*, and should work for both cases.

Theorem 66. There exist graph metrics, such that the PoA of any (deterministic) universal cost-sharing protocol is at least $\Omega(\log k)$, where k is the number of activated vertices.

Proof. Let $k = 2^{s-1} + 1$ be the number of activated vertices with $s \geq 4$, (so $k \geq 9$).

Graph Construction. We use as a base of our lower bound construction, a hypercube $Q := Q_n$, with edge costs equal to 1 and $n = n(s)$ as in Theorem 59. Based on Q , for $M = 16k^{12}2^{3n}$ we construct the following network with $N = 2^n M$ vertices, plus the designated root r . We add to Q direct edges/shortcuts as follows: for every two vertices v, u of distance 2^j , for $j \in [s]$, we add an edge/shortcut, (v, u) , with cost equal to $\hat{c}_j = 2^j \left(\frac{k-1}{k}\right)^j = \Omega(2^j)$. Moreover, for every vertex v_q of Q , we create $M - 1$ new vertices, each of which we connect with v_q via a zero-cost edge. Let V_q be the set of these vertices (including v_q). Finally, we add a root r , which we connect with every vertex v_q of Q , via two edges e_{q1} and e_{q2} , with costs $2k$ and $2k \cdot k/6$, respectively. We denote this new network by Q^* (see Figure 10.5).

We will show that any GWSP for Q^* has PoA of $\Omega(\log k)$. Any GWSP can be described by concatenations of potential-based cost-sharing protocols

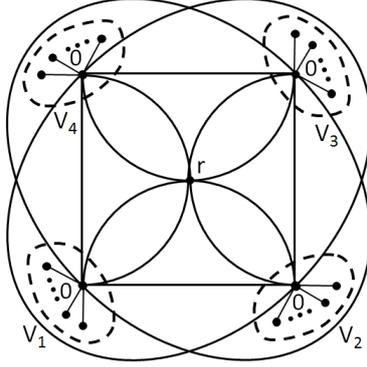


Figure 10.5: An example of Q^* for Q_2 as the base hypercube.

Ξ_1, \dots, Ξ_h for a partition of the $V(Q^*)$ into h subsets U_1, \dots, U_h , where Ξ_j is induced by some edge potential f_j . Following the analysis of Chen, Roughgarden and Valiant [36], we scale the f_j 's such that for every i, j , $f_j(i) \geq 1$. For nonnegative integers t and for $\alpha = (1 + k^{-3})^{\frac{1}{2k}}$, we form subgroups of vertices A_{jt} , for each U_j , as $A_{jt} = \{i \in U_j : f_j(i) \in [\alpha^t, \alpha^{t+1}]\}$ (note that some of A_{jt} 's may be empty).

The adversary proceeds in two cases, depending on the intersection of the A_{jt} 's with the V_q 's.

Shapley-like cost-sharing. Suppose first that there exist A_{jt} and V_q such that $|A_{jt} \cap V_q| \geq k$, and take a subset $R \subseteq A_{jt} \cap V_q$ with exactly k vertices. The adversary will request precisely the set R . We argue that there is a NE where all players follow the edge e_{q2} , with cost $2k \cdot k/6$.

Budget-balance implies that there exists some player $i^* \in R$ who is charged at most $1/k$ proportion of the cost. Moreover, Lemma 62 implies that, all $i \in R$ are charged at most $\alpha(1/k + 2k^{-2}) \leq 2 \cdot (3/k) = 6/k$ proportion of the cost. Therefore, no player's share is more than $2k$ and any alternative path would cost at least $2k$. However, the optimum solution is to use the parallel link e_{q1} of cost $2k$. Hence, the PoA is $\Omega(k)$ for this case.

Ordered-like cost-sharing. If there is no such R with at least k vertices, then $|A_{jt} \cap V_q| < k$ for all j, s and q , which means that each A_{jt} has size of at most $k2^n$. Recall that there are h potential-based protocols. For every $j \in [h]$, we group consecutive sets A_{jt} (starting from A_{j0}) into sets B_{jl} , such that each B_{jl} , (except perhaps from the last one), contains *exactly* $4k^5$ *nonempty* A_{jt} 's. The last B_{jl} contains at most $4k^5$ *nonempty* A_{jt} sets. Consider the lexicographic order among B_{jl} 's, i.e. $B_{jl} < B_{j'l'}$ if either $j < j'$ or $j = j'$ and $l < l'$. Rename these sets based on their total order as B_i 's. The size of each B_i is at most $4k^6 2^n$.

Now we apply Lemma 64 on the set N by considering the subsets V_q as the partition of N and as a colouring scheme, we colour all the vertices of each B_i with the same colour and use different colours among the sets B_i ⁴³. Lemma 64 guarantees that for each V_q there exists $V'_q \subset V_q$ of size 2^n , such that every $v \in \cup_q V'_q = V'$ belongs to a distinct B_i .

The order of B_i 's suggests an order of the vertices of V' . Since the V'_q 's form a partition of V' , Lemma 65 guarantees the existence of a subset $C \subset V'$, such that C contains exactly one vertex from each V'_q and there are no consecutive vertices in C . This means that C contains exactly one vertex from each set V_q and all these vertices belong to *different* and *non-consecutive* sets B_i .

To summarise, so far we know that:

- (i) for any pair of vertices $v, u \in C$, either v and u come from different U_j 's or their $f_j(v)$ and $f_j(u)$ values differ by a factor of at least $\alpha^{4k^5} \geq 8k + 1$ (since there exist at least $4k^5$ nonempty A_{jt} sets between the ones that v and u belong to).
- (ii) C is a copy of Q_n (by ignoring zero-cost edges).

Let π be the order of vertices of C (recall that they are ordered according to the B_i 's they belong to). Theorem 59 guarantees that there always exists at least one distance preserving path $P_s(\pi)$ (see Definition 53). Let S be the vertices of $P_s(\pi)$ excluding the last class D_s (see Definition 52). The adversary will activate precisely the set S ($|S| = k$). It remains to show that there exists a NE, the cost of which is a factor of $\Omega(\log k)$ away from optimum. We will refer to these vertices as $S = \{s_1, s_2, \dots, s_k\}$, based on their order π , from smaller label to larger, and let player i be associated with s_i .

Let \mathcal{P}^* be the class of strategy profiles $\mathbf{P} = (P_1, \dots, P_k)$ which are defined as follows:

- $P_1 = e_{11}$ and $P_2 = (s_1, s_2) \cup P_1$, where (s_1, s_2) is the shortcut edge between s_1 and s_2 .
- From $i = 3$ to k , let s_ℓ be any of s_i 's parents in the class hierarchy (we refer the reader to the beginning of Section 10.2); then $P_i = (s_i, s_\ell) \cup P_\ell$, where (s_i, s_ℓ) is the shortcut edge between s_i and s_ℓ .

⁴³For the Lemma 64 we set $w = 4k^6 2^n$ and $s = m = 2^n$. Recall that $|V_q| = M = sw^2$ as required and no more than w elements have the same colour, since the size of each B_i is at most $4k^6 2^n = w$.

We show in Claim 67 that there exists a strategy profile $\mathbf{P}^* \in \mathcal{P}^*$ which is a NE. \mathbf{P}^* has cost:

$$c(\mathbf{P}^*) = c(e_{11}) + \hat{c}_s + \sum_{j=1}^{s-1} |D_j| \cdot \hat{c}_{s-j} = \Omega(2^s) + \Omega(2^s) + \sum_{j=1}^{s-1} 2^{j-1} \cdot \Omega(2^{s-j}) = \Omega(s2^s).$$

However, there exists the solution $P_s(\pi) \cup e_{11}$, which has cost of $O(2^s)$. Therefore, the PoA is $\Omega(s) = \Omega(\log k)$.

Claim 67. There exists $\mathbf{P}^* \in \mathcal{P}^*$ which is a NE.

Proof. We prove the claim by using better-response dynamics. Note that any GWSP induces a potential game for which better-response dynamics always converge to a NE (see [36, 76]). We start with some $\mathbf{P}_1 \in \mathcal{P}^*$ and we prove that, after a sequence of players' *best*-responses, we end up in $\mathbf{P}_2 \in \mathcal{P}^*$. Proceeding in a similar way we eventually converge to \mathbf{P}^* , which is the required NE.

We next argue that for any $\mathbf{P} \in \mathcal{P}^*$, players 1 and 2, have no incentive to deviate from P_1 (argument (a)) and P_2 (arguments (b)), respectively. We further show that, given any strategy profile $\hat{\mathbf{P}}$, there exists some $\mathbf{P} \in \mathcal{P}^*$ such that: for every player $i \notin \{1, 2\}$, if $\mathbf{P}^i = (P_1, \dots, P_{i-1}, \hat{P}_{i+1}, \dots, \hat{P}_k)$ are the strategies of the other players, i prefers P_i to \hat{P}_i (arguments (c)-(e)). We define the desired \mathbf{P} recursively starting from $\hat{\mathbf{P}}$ as follows: $P_1 = e_{11}$, $P_2 = (s_1, s_2) \cup P_1$ and from $i = 3$ to k , $P_i \in A = \arg \min_{P'_i} \{c_i(\mathbf{P}^i, P'_i) \mid \exists (P'_{i+1}, \dots, P'_k) \text{ s.t. } (P_1, \dots, P_{i-1}, P'_i, \dots, P'_k) \in \mathcal{P}^*\}$. If $\hat{P}_i \in A$ then we set $P_i = \hat{P}_i$, otherwise we choose a path from A arbitrarily.

We first give some bounds on players' shares.

1. Let $R \subseteq S$ be any set of players that use some edge e of cost c_e and let i be the one with the smallest label. The total share of players $R \setminus \{i\}$ is upper bounded by $\sum_{i=1}^{|R|-1} \frac{1}{(8k+1)^{i+1}} \cdot c_e < \frac{c_e}{8k}$ (Lemma 63). Moreover, i 's share is at least $\frac{8k-1}{8k} c_e$.
2. The total cost of any P_i under \mathbf{P}^i , is at most $8k$. This is true because, for every player i' with $i' \leq i$, the first edge of $P_{i'}$ is a shortcut to reach one of $s_{i'}$'s parents, with cost at most 2^{s-j} , where D_j is the hierarchical class that $s_{i'}$ belongs to (we refer the reader to the beginning of Section 10.2 for the definition of classes). Therefore, the cost of P_i is at most $2k + \sum_{l=0}^{s-1} 2^{s-l} < 8k$.

3. By combining the above two arguments, under \mathbf{P}^i , the total share of player i for the edges of P_i at which she is not the first according to π , is at most $\frac{1}{8k} \cdot 8k \leq 1$.

Here, we give the arguments for players 1 and 2.

- (a) The share of player 1 under $\mathbf{P} \in \mathcal{P}^*$ is at most $2k$ and any other path would incur a cost strictly greater than $2k$.
- (b) The share of player 2 under $\mathbf{P} \in \mathcal{P}^*$ is at most $2^s + 1 = 2k - 1$ (argument 3), whereas if she doesn't connect through s_1 , her share would be at least $2k$. Moreover, if she connects to r through s_1 but by using any other path rather than the shortcut (s_1, s_2) , the total cost of that path is at least $2^s \left(\frac{k-1}{k}\right)^{s-1}$. Player 2 is first according to π at that path and by argument 1, her share is at least $2^s \frac{8k-1}{8k} \left(\frac{k-1}{k}\right)^{s-1} > \hat{c}_s$.

We next give the required arguments in order to show that P_i is a best response for player $i \neq \{1, 2\}$ under \mathbf{P}^i . In the following, let $s_i \in D_j$ and let s_ℓ be the parent of s_i such that $P_i = (s_i, s_\ell) \cup P_\ell$. Also let $s_{i'}$ be the predecessor of s_i , according to π , that is first met by following \hat{P}_i from s_i to r .

- (c) Suppose that $s_{i'} = s_\ell$.
- Assume that \hat{P}_i doesn't use the shortcut (s_i, s_ℓ) . The subpath of \hat{P}_i from s_i to s_ℓ contains edges at which i is first according to π of total cost at least $2^{s-j} \left(\frac{k-1}{k}\right)^{s-j-1}$. By argument 1, her share is at least $2^{s-j} \frac{8k-1}{8k} \left(\frac{k-1}{k}\right)^{s-j-1} > \hat{c}_{s-j}$.
 - Assume that \hat{P}_i doesn't use P_ℓ . The subpath of \hat{P}_i from s_ℓ to r contains edges at which i is first according to π of total cost at least 2 (the minimum distance between two activated vertices). By argument 1, her share is at least $2 \frac{8k-1}{8k} > 1$, where 1 is at most her share for P_ℓ (argument 3).

In both cases, $c_i(\mathbf{P}^i, P_i) < c_i(\mathbf{P}^i, \hat{P}_i)$.

- (d) Suppose that $s_{i'}$ is s_i 's other parent. If $\hat{P}_i \neq (s_i, s_{i'}) \cup P_{i'}$, the above arguments still hold and so $c_i(\mathbf{P}^i, P_i) < c_i(\mathbf{P}^i, \hat{P}_i)$. Otherwise, by the definition of P_i , either $P_i = \hat{P}_i$, or $c_i(\mathbf{P}^i, P_i) < c_i(\mathbf{P}^i, \hat{P}_i)$.

- (e) Suppose that $s_{i'}$ is not a parent of s_i . Player i 's share in P_i is at most \hat{c}_{s-j} for her first edge/shortcut and at most 1 for the rest of her path (argument 3). Note that any edge that is used by players that precedes i in π has cost at least \hat{c}_{s-j} . Therefore, in \hat{P}_i , player i is the first according to π for edges of total cost at least \hat{c}_{s-j+1} . This implies a cost-share of at least $\frac{8k-1}{8k}\hat{c}_{s-j+1}$ (argument 1). But for $k \geq 6$ and $j < s$, $\frac{8k-1}{8k}\hat{c}_{s-j+1} > \hat{c}_{s-j} + 1$.

We now describe a sequence of best-responses from some $\hat{\mathbf{P}} \in \mathcal{P}^*$ to \mathbf{P} (\mathbf{P} is constructed based on $\hat{\mathbf{P}}$ as described above). We follow the π order of the players and for each player we apply her best response. First note that players 1 and 2 have no better response, so $P_1 = \hat{P}_1$ and $P_2 = \hat{P}_2$. When we process any other player i , we have already processed all her predecessors in π and so, the strategies of the other players are \mathbf{P}^i . Therefore, P_i is the best response for i (it may be that $P_i = \hat{P}_i$, where no better response exists for i). The order that we process the vertices guarantees that $\mathbf{P} \in \mathcal{P}^*$. Best-response dynamics guarantee that eventually, no player could perform any best-response, resulting in the desired NE. \square

This completes the proof of Theorem 66. \square

CHAPTER 11

Stochastic Design

In this Chapter, we study the *stochastic* model, where the activated vertices, or alternatively the players' types, are no longer picked adversarially, but instead the players' types are drawn from some distribution D . We remind the reader that the designer is aware of merely D and not the actual types of the players; however, the players have full knowledge of the other players' types, as in Chapter 10. We design a randomised universal cost-sharing protocol with constant PoA.

Discussion on the Design. We note that in the whole chapter we only design ordered protocols. The total order is defined by combining an order, π_1 , among all the vertices of the graph and a global order, π_2 , among all players as following. For any $\mathbf{t} = \{t_1, \dots, t_k\}$, π_1 defines a partial order of players, i.e. it defines an order among players of different types, and then a total order is derived via π_2 that induces an order among players of the same type. Then, if more than one players have the same vertex as their type, the first player among them, based on π_2 , is only charged and the rest follow the same path with zero cost-shares. For the simplicity of the presentation, we only define an order among the activated vertices, i.e. the order π_1 . The global order π_2 can be defined arbitrarily based on players ids.

We suppose that the set of the activated vertices, S , is drawn from some probability distribution Π derived by D . In the stochastic model, D is chosen adversarially, so the same is considered for Π . We next design a randomised ordered protocol (Section 11.1) with constant PoA. In a slightly different model where each vertex is activated *independently* with some probability, by using standard

derandomisation techniques, we produce a deterministic ordered protocol (Section 11.2) that achieves constant PoA. We note that both the randomised and the deterministic protocols can be determined in polynomial time.

11.1 Randomised Protocol

We show that there exists a *randomised* ordered protocol that achieves constant PoA. This result holds even for the *black-box* model [130], meaning that the probability distribution is not known to the designer, however she is allowed to draw independent (polynomially many) samples.

The protocol’s design highly relies on approximation algorithms for the minimum Steiner tree problem and therefore, the resulting PoA upper bound (Corollary 69) depends on known approximation ratios for this problem. More precisely, given an α -approximate minimum Steiner tree, we show an upper bound of $2(\alpha + 2)$ (Theorem 68). The approximate tree is used in our algorithm as a base in order to construct a spanning tree, which finally determines an order of all vertices; the detailed algorithm is given in Algorithm 1. This algorithm and its slight variants have been used in different contexts: rent-or-buy problem [81], a priori TSP [130] and, stochastic Steiner tree problem [72].

Algorithm 1 Randomised order protocol Ξ_{rand}

Input: A rooted graph $G = (V, E, r)$ and an oracle for the probability distribution Π .

Output: Ξ_{rand} .

- Choose a random set of vertices R by drawing from distribution Π and construct an α -approximate minimum Steiner tree, $T_\alpha(R)$, over $R \cup \{r\}$.
 - Connect all other vertices $V \setminus V(T_\alpha(R))$ with their nearest neighbour in $V(T_\alpha(R))$ (by breaking ties arbitrarily).
 - Double the edges of that tree and traverse some Eulerian tour starting from r . Order the vertices based on their first appearance in the tour.
-

Theorem 68. Given an α -approximate solution of the minimum Steiner tree problem, Ξ_{rand} has PoA at most $2(\alpha + 2)$.

Proof. Let π be the order of V , defined by Ξ_{rand} , and S be the random set of activated vertices that require connectivity with r . For the rest of the proof we

denote by $MST(S)$ a minimum spanning tree over the vertices $S \cup \{r\}$ on the metric closure⁴⁴ of G .

Let s_1, \dots, s_r be the vertices of S as appeared in π and the strategy profile $\mathbf{P}_R(S) = (P_1, \dots, P_r)$ be a NE of set S . Under the convention that $s_0 = r$, $c_{s_i}(\mathbf{P}_R(S)) \leq d_G(s_i, s_{i-1})$ for all $s_i \in S$. We construct a tree $T_{R,S}$ from the $T_\alpha(R)$ of Algorithm 1, by connecting only all vertices of $S \setminus V(T_\alpha(R))$ with their nearest neighbour in $V(T_\alpha(R))$ (by breaking ties in accordance to Algorithm 1). Note that, by doubling the edges of $T_{R,S}$, there exists an Eulerian tour starting from r , where the order of the vertices S (based on their first appearance in the tour) is π restricted to the set S .⁴⁵ Therefore, $\sum_{s_i \in S} d_{T_{R,S}}(s_i, s_{i-1}) + d_{T_{R,S}}(s_0, s_r) = 2c(T_{R,S})$. By combining the above,

$$\begin{aligned} c(\mathbf{P}_R(S)) &= \sum_{s_i \in S} c_{s_i}(\mathbf{P}_R(S)) \leq \sum_{s_i \in S} d_G(s_i, s_{i-1}) \\ &\leq \sum_{s_i \in S} d_{T_{R,S}}(s_i, s_{i-1}) \leq 2c(T_{R,S}). \end{aligned} \quad (11.1)$$

Let $D_v(R)$ be the distance of v from its nearest neighbour in $(R \cup \{r\}) \setminus \{v\}$. In the special case that $v = r$, we define $D_v(R) = 0$. Then,

$$c(T_{R,S}) = c(T_\alpha(R)) + \sum_{v \in S \setminus R} D_v(R) \leq c(T_\alpha(R)) + \sum_{v \in S} D_v(R). \quad (11.2)$$

We use an indicator $I(v \in S)$ which is 1 when $v \in S$ and 0 otherwise; then $\sum_{v \in S} D_v(R) = \sum_v I(v \in S) D_v(R)$. By taking the expectation over R and S ,

$$\mathbb{E}_R[\mathbb{E}_S[c(T_{R,S})]] \leq \mathbb{E}_R[c(T_\alpha(R))] + \mathbb{E}_R[\mathbb{E}_S[\sum_{v \in V} I(v \in S) D_v(R)]].$$

Since S and R are independent samples we can bound the second term as:

$$\begin{aligned} \mathbb{E}_R[\mathbb{E}_S[\sum_{v \in V} I(v \in S) D_v(R)]] &= \sum_{v \in V} \mathbb{E}_S[I(v \in S)] \mathbb{E}_R[D_v(R)] \\ &= \sum_{v \in V} \mathbb{E}_S[I(v \in S)] \mathbb{E}_S[D_v(S)] \end{aligned}$$

⁴⁴The metric closure of an undirected graph G is the complete undirected graph on the vertex set $V(G)$, where the edge costs equal the shortest paths in G between the corresponding vertices.

⁴⁵This Eulerian tour matches the tour constructed by shortcutting the Eulerian tour of Algorithm 1 to contain only the vertices $R \cup S \cup \{r\}$.

$$\begin{aligned}
&= \mathbb{E}_S \left[\sum_{v \in V} I(v \in S) D_v(S) \right] \\
&= \mathbb{E}_S \left[\sum_{v \in S} D_v(S) \right] \leq \mathbb{E}_S [c(MST(S))]. \quad (11.3)
\end{aligned}$$

The third equality holds since $D_v(S)$ is the distance of v from its nearest neighbour in $(S \cup \{r\}) \setminus \{v\}$ and it is independent of the event $I(v \in S)$. For the inequality, note that $D_v(S)$ is upper bounded by the distance of v from its parent in the $MST(S)$.

Let T_S^* be the minimum Steiner tree over $S \cup \{r\}$, then it is well known that $c(MST(S)) \leq 2c(T_S^*)$. Overall,

$$\begin{aligned}
\mathbb{E}_R [\mathbb{E}_S [c(\mathbf{P}_R(S))]] &\leq 2 \mathbb{E}_R [\mathbb{E}_S [c(T_{R,S})]] \\
&\leq 2(\mathbb{E}_S [c(T_\alpha(S))] + \mathbb{E}_S [c(MST(S))]) \\
&\leq 2(\alpha + 2) \mathbb{E}_S [c(T_S^*)].
\end{aligned}$$

□

By applying the 1.39-approximation algorithm of [26] we get the following corollary.

Corollary 69. Ξ_{rand} has PoA at most 6.78.

11.2 Deterministic Protocol

We now consider a different model where each vertex v is activated independently with probability p_v , and w.l.o.g. set $p_r = 1$. The set of the activated vertices is sampled based on p_v 's, i.e., the probability that S is activated is $\Pi(S) = \prod_{v \in S} p_v \cdot \prod_{v \notin S} (1 - p_v)$. The p_v 's (and therefore Π), are chosen adversarially. We additionally assume that the probabilities p_v 's are known to the designer. We show that there exists a *deterministic* ordered protocol that achieves constant PoA.

Theorem 70. There exists a deterministic ordered protocol with PoA at most 16.

Proof. We use derandomisation techniques similar to [130, 141] and for completeness we give the full proof here. First we discuss how we can get a PoA of 6.78, if we drop the requirement of determining the protocol in polynomial time. Similar

to the proof of Theorem 68 we define the tree $T_{R,S}$ for the random activated set S as follows: we construct $T_{R,S}$ from the $T_\alpha(S)$ of Algorithm 1, by connecting only all vertices of $S \setminus V(T_\alpha(R))$ with their nearest neighbour in $V(T_\alpha(R))$ (by breaking ties in accordance to Algorithm 1). We apply the standard derandomisation approach of *conditional expectation method* on $T_{R,S}$. More precisely, we construct a deterministic set \hat{R} to replace the random set R in Algorithm 1, by deciding for each vertex of $V \setminus \{r\}$, one by one, whether it belongs to \hat{R} or not. The order we process the vertices is chosen arbitrarily. Assume that we have already processed the set $Q \subset V$ and we have decided that for its partition (Q_1, Q_2) , $Q_1 \subseteq \hat{R}$ and $Q_2 \cap \hat{R} = \emptyset$ (starting from $Q_1 = \{r\}$ and $Q_2 = \emptyset$). Let v be the next vertex to be processed. From the conditional expectations and the independent activations we know that

$$\begin{aligned} \mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset] &= \\ & \mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset, v \in R]p_v \\ & + \mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset, v \notin R](1 - p_v), \end{aligned}$$

meaning that

$$\begin{aligned} \text{either } & \mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset, v \in R] \\ & \leq \mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset], \\ \text{or } & \mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset, v \notin R] \\ & \leq \mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset]. \end{aligned}$$

In the first case we add v in Q_1 and in the second case we add v in Q_2 . Therefore, after processing all vertices, $Q_1 = \hat{R}$ and $E_S[c(T_{\hat{R},S})] \leq \mathbb{E}_{S,R}[c(T_{R,S})]$. If we replace the sampled R of Algorithm 1 with the deterministic set \hat{R} , we can get the same bound on the PoA with the randomised protocol of Theorem 68.

However, the value of $\mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset]$ seems difficult to be computed in polynomial time; the reason is that it involves the computation of $\mathbb{E}_R[c(T_\alpha(R))|Q_1 \subseteq R, Q_2 \cap R = \emptyset]$ which seems hard to be handled. To overcome this problem we use an estimator $EST(Q_1, Q_2)$ of $\mathbb{E}_{S,R}[c(T_{R,S})|Q_1 \subseteq R, Q_2 \cap R = \emptyset]$, which is constant away from the optimum $\mathbb{E}_S[c(T_S^*)]$, where T_S^* is the minimum Steiner tree over $S \cup \{r\}$.

Following [141, 130], we use the optimum solution of the relaxed Connected

Facility Location Problem on G in order to construct a feasible solution $\bar{\mathbf{y}}$ of the relaxed Steiner Tree Problem (STP) for a given set R . We show that the objective's value of the fractional STP for $\bar{\mathbf{y}}$ is constant away from $\mathbb{E}_S[c(T_S^*)]$ and that its (conditional) expectation over R can be efficiently computed. This quantity is used in order to construct the estimator $EST(Q_1, Q_2)$. We apply the method of conditional expectations on $EST(Q_1, Q_2)$ and after processing all vertices, by using the primal-dual algorithm [75], we compute a Steiner tree on Q_1 with cost no more than twice the cost of the fractional solution.

In the rooted Connected Facility Location Problem (CFLP), a rooted graph $G = (V, E, r)$ is given and the designer should select some facilities to open, including r , and connects them via some Steiner tree T . Every other vertex is assigned to some facility. The cost of the solution is M ($M > 1$) times the cost of T , plus the distance of every other vertex from its assigned facility. Our analysis requires to consider a slightly different cost of the solution, which is the cost of T , plus the distance of every other vertex v from its assigned facility multiplied by p_v . In the following LP relaxation of the CFLP, z_e and x_{ij} are 0-1 variables indicate, respectively, if $e \in E(T)$ and whether the vertex j is assigned to facility i . $\delta(U)$ denotes the set of edges with one endpoint in U and the other in $V \setminus U$, $d(i, j)$ denotes the minimum distance between vertices i and j in G and c_e is the cost of edge e .

LP1: CFLP		
	$\min \quad B + C$	
subject to	$\sum_{i \in V} x_{ij} = 1$	$\forall j \in V$
	$\sum_{e \in \delta(U)} z_e \geq \sum_{i \in U} x_{ij}$	$\forall j \in V, \forall U \subseteq V \setminus \{r\}$
	$B = \sum_{e \in E} c_e z_e$	
	$C = \sum_{j \in V} p_j \sum_{i \in V} d(i, j) x_{ij}$	
	$z_e, x_{ij} \geq 0$	$\forall i, j \in V \text{ and } \forall e \in E$

Let $(\mathbf{z}^* = (z_e^*)_e, \mathbf{x}^* = (x_{ij}^*)_{ij}, B^*, C^*)$ be the optimum solution of LP1.

Claim 71. $B^* + C^* \leq 3 \mathbb{E}_S[c(T_S^*)]$.

Proof. Given a set $S \subseteq V$, for every edge $e \in T_S^*$, let $z_e = 1$ and, for $e \notin T_S^*$, let $z_e = 0$. Moreover, for every $j \in V$, let $x_{ij} = 1$, if i is j 's nearest neighbour in $(S \cup \{r\}) \setminus \{j\}$. Set the rest of x_{ij} equal to 0. Note that this is a feasible solution of LP1 with objective value $B_S + C_S \leq c(T_S^*) + \sum_{v \in V} p_v D_v(S)$. By taking the

expectation over S ,

$$\begin{aligned}
B^* + C^* &\leq \mathbb{E}_S[B_S + C_S] \leq \mathbb{E}_S[c(T_S^*)] + \sum_{v \in V} \mathbb{E}_S[I(v \in S)] \mathbb{E}_S[D_v(S)] \\
&= \mathbb{E}_S[c(T_S^*)] + \mathbb{E}_S\left[\sum_{v \in S} D_v(S)\right] \leq \mathbb{E}_S[c(T_S^*)] + \mathbb{E}_S[c(MST(S))] \leq 3\mathbb{E}_S[c(T_S^*)].
\end{aligned}$$

□

By using the solution $(\mathbf{z}^* = (z_e^*)_e, \mathbf{x}^* = (x_{ij}^*)_{ij}, B^*, C^*)$, we construct a feasible solution for the following LP relaxation of the Steiner Tree Problem (STP) over some set $R \cup \{r\}$.

LP2: STP over $R \cup \{r\}$	
	$\min \sum_{e \in E} c_e y_e$
subject to	$\sum_{e \in \delta(U)} y_e \geq 1 \quad \forall U \subseteq V \setminus \{r\} : R \cap U \neq \emptyset$
	$y_e \geq 0 \quad \forall e \in E$

We define $a_{ij}(e) = 1$ if e lies in the shortest path between i and j and it is 0 otherwise. For every edge e we set $\bar{y}_e = z_e^* + \sum_{j \in R} \sum_{i \in V} a_{ij}(e) x_{ij}^*$.

Claim 72. $\bar{\mathbf{y}} = (\bar{y}_e)_e$ is a feasible solution for LP2.

Proof. The proof is identical with the one in [141] but we give it here for completeness. Consider any set $U \subseteq V \setminus \{r\}$ such that $R \cap U \neq \emptyset$ and let $\ell \in R \cap U$. It follows that

$$\begin{aligned}
\sum_{e \in \delta(U)} \bar{y}_e &\geq \sum_{e \in \delta(U)} z_e^* + \sum_{e \in \delta(U)} \sum_{j \in R} \sum_{i \in V} a_{ij}(e) x_{ij}^* \geq \sum_{i \in U} x_{i\ell}^* + \sum_{e \in \delta(U)} \sum_{i \in V} a_{i\ell}(e) x_{i\ell}^* \\
&\geq \sum_{i \in U} x_{i\ell}^* + \sum_{i \notin U} x_{i\ell}^* \sum_{e \in \delta(U)} a_{i\ell}(e) \geq \sum_{i \in U} x_{i\ell}^* + \sum_{i \notin U} x_{i\ell}^* = 1.
\end{aligned}$$

For the last inequality, note that $a_{i\ell}(e)$ should be 1 for at least one $e \in \delta(U)$ since $i \notin U$ and $\ell \in U$. □

Claim 73. Let $\bar{c}_{ST}(R)$ be the cost of the objective of LP2 induced by the solution $\bar{\mathbf{y}}$. Then $\mathbb{E}_R[\bar{c}_{ST}(R)] = B^* + C^*$.

Proof.

$$\mathbb{E}_R[\bar{c}_{ST}(R)] = \mathbb{E}_R\left[\sum_{e \in E} c_e (z_e^* + \sum_{j \in R} \sum_{i \in V} a_{ij}(e) x_{ij}^*)\right] = B^* + \mathbb{E}_R\left[\sum_{j \in R} \sum_{i \in V} \sum_{e \in E} c_e a_{ij}(e) x_{ij}^*\right]$$

$$= B^* + \mathbb{E}_R \left[\sum_{j \in R} \sum_{i \in V} d(i, j) x_{ij}^* \right] = B^* + \sum_{j \in V} p_j \sum_{i \in V} d(i, j) x_{ij}^* = B^* + C^*.$$

□

Observe that due to the expression of \bar{y} we can efficiently compute any conditional expectation

$$\mathbb{E}_R[\bar{c}_{ST}(R) | Q_1 \subseteq R, Q_2 \cap R = \emptyset];$$

this is because

$$\mathbb{E}_R \left[\sum_{j \in R} \sum_{i \in V} a_{ij}(e) x_{ij}^* | Q_1 \subseteq R, Q_2 \cap R = \emptyset \right] = \sum_{j \in Q_1} \sum_{i \in V} a_{ij}(e) x_{ij}^* + \sum_{j \notin Q_1 \cup Q_2} p_j \sum_{i \in V} a_{ij}(e) x_{ij}^*.$$

We further define $c_C(R) = \sum_{v \in V} p_v D_v(R)$. We can also efficiently compute any conditional expectation $\mathbb{E}[c_C(R) | Q_1 \subseteq R, Q_2 \cap R = \emptyset]$ (Claim 2.1 of [141]). We are ready to define our estimator:

$$EST(Q_1, Q_2) = 2 \mathbb{E}_R[\bar{c}_{ST}(R) | Q_1 \subseteq R, Q_2 \cap R = \emptyset] + \mathbb{E}_R[\bar{c}_C(R) | Q_1 \subseteq R, Q_2 \cap R = \emptyset].$$

Our goal is to define a deterministic set \hat{R} to replace the sampled R of Algorithm 1. We process the vertices one by one and we decide if they belong to \hat{R} by using the model conditional expectations on $EST(Q_1, Q_2)$. More specifically, assume that we have already processed the sets Q_1 and Q_2 (starting from $Q_1 = \{r\}$ and $Q_2 = \emptyset$) such that $Q_1 \subseteq \hat{R}$ and $Q_2 \cap \hat{R} = \emptyset$. Let v be the next vertex to be processed. From the conditional expectations and the independent activations we know that $EST(Q_1, Q_2) = p_v EST(Q_1 \cup \{v\}, Q_2) + (1 - p_v) EST(Q_1, Q_2 \cup \{v\})$. If $EST(Q_1 \cup \{v\}, Q_2) \leq EST(Q_1, Q_2)$ we add v to Q_1 , otherwise we add v to Q_2 . After processing all vertices and by using Claims 71 and 73,

$$\begin{aligned} EST(\hat{R}, V \setminus \hat{R}) &\leq EST(\{r\}, \emptyset) = 2 \mathbb{E}_R[\bar{c}_{ST}(R)] + \mathbb{E}_R[\bar{c}_C(R)] \\ &\leq 6 \mathbb{E}_S[c(T_S^*)] + \sum_{v \in V} p_v \mathbb{E}_R[D_v(R)] \\ &= 6 \mathbb{E}_S[c(T_S^*)] + \mathbb{E}_R \left[\sum_{v \in V} I(v \in R) D_v(R) \right] \\ &\leq 6 \mathbb{E}_S[c(T_S^*)] + \mathbb{E}_R[c(MST(R))] \leq 8 \mathbb{E}_S[c(T_S^*)]. \end{aligned}$$

Let $T_{PD}(\hat{R})$ be the Steiner tree over $\hat{R} \cup \{r\}$ computed by the primal-dual algo-

rithm [75]. Then,

$$EST(\hat{R}, V \setminus \hat{R}) = 2\bar{c}_{ST}(\hat{R}) + \sum_{v \in V} p_v D_v(\hat{R}) \geq c(T_{PD}(\hat{R})) + \mathbb{E}_S \left[\sum_{v \in S} D_v(\hat{R}) \right].$$

By combining inequalities (11.1) and (11.2) (after replacing R by \hat{R} and $T_\alpha(\hat{R})$ by $T_{PD}(\hat{R})$) with all the above, we have that

$$\mathbb{E}_S [c(\mathbf{P}_{\hat{R}}(S))] \leq 2 \left(c(T_{PD}(\hat{R})) + \mathbb{E}_S \left[\sum_{v \in S} D_v(\hat{R}) \right] \right) \leq 2EST(\hat{R}, V \setminus \hat{R}) \leq 16 \mathbb{E}_S [c(T_S^*)].$$

This finishes the proof of Theorem 70. □

CHAPTER 12

Bayesian Design

In this chapter, we discuss the Bayesian model. We first show that there exists a lower bound of $\Omega(\sqrt{k})$ on the PoA of *any* universal cost-sharing protocol (Section 12.1). The lower bound holds even when the players' types are i.i.d. meaning that they are drawn from independent and identical prior distributions, i.e. D is a product distribution of k identical distributions.

Budget-Balance in Equilibrium. One of the axioms of the universal cost-sharing protocols should satisfy, is *budget balance*. Although budget balance is a very natural requirement, apparently it puts considerable restrictions on the design space. However, since we expect that the players will end up in a Nash equilibrium, it is not clear why one should be interested to impose budget balance in non-equilibrium states; the players are going to deviate from such states anyway. We propose an alternative, relaxed requirement that we call *budget-balance in the equilibrium (BBiE)*. A BBiE cost-sharing protocol satisfies *ex-post* budget-balance in *all equilibria*; for any non-equilibrium profile we do not impose this requirement. This natural relaxation, enlarges the design space but maintains the desired property of balancing the cost in the equilibrium. More importantly, this amplification of the design space, allows us to design protocols that dramatically outperform the best possible PoA bounds obtained by budget-balanced protocols. In Section 12.2.1, we improve the PoA from $\Omega(\sqrt{n})$ to $O(1)$ by designing a BBiE cost-sharing protocol.

Posted Prices. We further examine the use of posted prices. It is a very common practice, especially in large markets and double auctions, for sellers

to use posted prices on their items. More closely to cost-sharing games is the model proposed by Kelly [94] regarding *bandwidth allocation*. Kelly’s mechanism processes players’ willingness to pay for the bandwidth and posts a price for the whole bandwidth. Then, each player pays a price proportional to the portion of bandwidth she wants to use. This can be seen as pricing an infinitesimal quantity of the bandwidth and the players, acting as price-takers, choose some number of quantities to buy. It turns out that it is in the best interest of the players to buy the whole bandwidth.

Posted prices have also been used for pricing in large markets. Kelso and Crawford [95] and Gul and Stacchetti [80] proved the existence of prices, for gross substitute valuations, that clear the market efficiently. Feldman, Gravin and Lucier [65] used anonymous posted prices in combinatorial auctions under the Bayesian setting and showed that when the valuations are fractionally sub-additive there exist prices that guarantee half of the expected optimum social welfare. Pricing bundles for combinatorial Walrasian equilibria was introduced by the same authors [66], who showed that half of the optimum social welfare can be achieved even for arbitrary valuation functions. Dynamic pricing schemes have been used by Cohen, Eden, Fiat and Jez [51] in several online settings to induce the same performance as the best online algorithm, and by Cohen-Addad, Eden, Feldman and Fiat [52] in matching markets in order to achieve the optimal social welfare, for any tie breaking rule.

The use of posted prices to serve as cost-sharing mechanism, is highly desirable, but not always possible to achieve; a price is posted for each resource and then the players behave as price takers, picking up the cheapest possible resources that satisfy their requirements. Such a mechanism is desirable because it is extremely easy to implement and also induces *dominant strategies*. The same constant bound achieved by a BBiE protocol can be further achieved by using *ex-ante* BBiE posted prices (Section 12.2.2). We stress that our results are implemented by *anonymous* posted prices. In Section 12.2.2, we further discuss the necessity of BBiE being satisfied *ex-ante*, i.e. in expectation, rather than ex-post like the BBiE protocols.

12.1 Fully Budget-Balance

In this section we show that the PoA of *any* universal cost-sharing protocol can be very high.

Theorem 74. The Bayesian PoA of any (deterministic or randomised) universal protocol, for the multicast game with $\Theta(n)$ vertices, is $\Omega(\sqrt{n})$, even for i.i.d. players.

Proof. We assume that the players' types are drawn independently and identically from some distribution D^* . We consider the graph of Figure 12.1 and we define D^* by drawing each of the $\{v_1, \dots, v_n\}$ with probability p , v with probability 0 and r with probability $1 - np$.

We set $p = 1 - \left(1 - \frac{1}{\sqrt{n}}\right)^{\frac{1}{n}}$, such that the probability that vertex v_i is drawn as the type of at least one player is $q_i = 1 - (1 - p)^n = \frac{1}{\sqrt{n}}$. We claim that, for any budget-balanced protocol, it is a Bayes-Nash equilibrium if any player with type v_i uses the direct edges (v_i, t) .

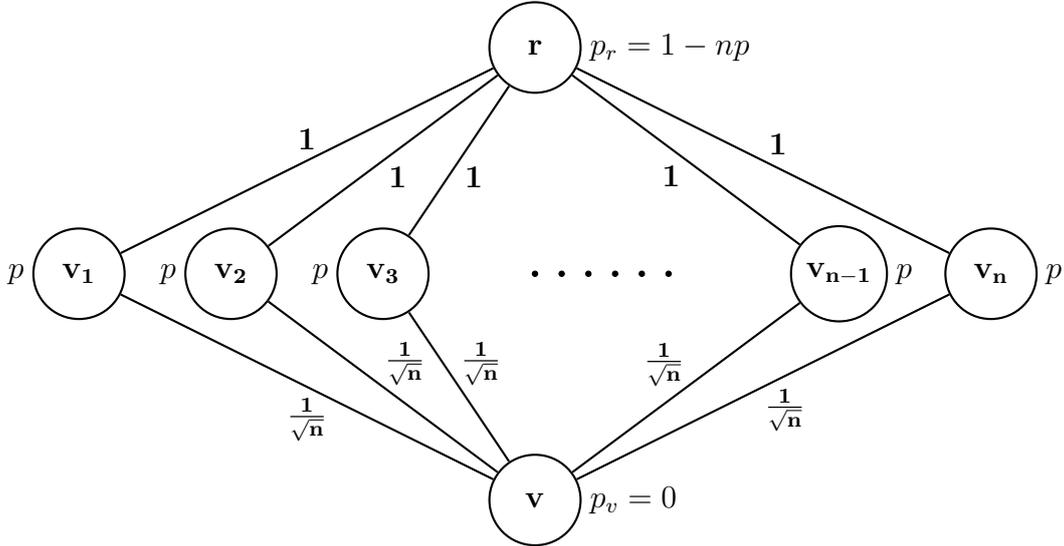


Figure 12.1: Lower bound on the PoA of any budget-balanced protocol.

Indeed, if player i uses any other path (v_i, v, v_j, t) her cost-share will be at least $\frac{2}{\sqrt{n}} + (1 - q_j) = 1 + \frac{1}{\sqrt{n}}$, which is greater than her current cost-share of at most 1. The expected social cost and optimum are: $\mathbb{E}[SC] = \sum_i q_i = \sqrt{n}$ and $\mathbb{E}[Opt] \leq \sum_i q_i \cdot \frac{1}{\sqrt{n}} + 1 = n^{\frac{1}{n}} + 1 = 2$. So, the Bayes PoA is at least $\frac{1}{2}\sqrt{n}$. \square

12.2 Budget-Balance in Equilibrium

In this section we drop the requirement of budget balance and instead we consider a more general class of cost-sharing protocols \mathcal{C} , where the requirement is to preserve the budget balance in the equilibrium (BBiE).

BBiE: In any pure (Bayesian) Nash equilibrium profile, the cost shares of the players choosing edge e should cover exactly the cost of e .

To show our results we will use an oblivious algorithm of the corresponding optimisation problem and we will enforce its solution by applying appropriate cost-sharing protocols and posted prices, e.g. choices, not consistent with this solution, are highly expensive.

The underlying optimisation problem of the multicast cost-sharing game is the minimum Steiner tree. The type of each player corresponds to an input component of the optimisation problem, i.e. some requested vertex, and the domain of her strategy space corresponds to the set of the paths connecting that requested vertex with the root r . An oblivious algorithm assigns an action for each input component, based on the prior distribution, and *independently of the realisation* of all other input components. In our case, an oblivious solution, maps each vertex to some path that connects it to r , and is used in *any realisation* of the input that contains this vertex.

12.2.1 BBiE Protocols

The following theorem associates the PoA of some cost-sharing game with the approximation ratio of an oblivious algorithm of the underlying optimisation problem. The theorem doesn't merely apply to the multicast cost-sharing game but it is more general and have been also used in the set cover cost-sharing game [42].

Theorem 75. Let G be any cost-sharing game and Π the underlying optimisation resource allocation problem. Suppose that the input of Π is chosen stochastically by the distribution $\pi_1 \times \pi_2 \times \dots \times \pi_k$, where π_i is the distribution from which the i^{th} input component is drawn from. Further suppose that the types of the players are drawn from $D = \times_i D_i$, where $D_i = \pi_i$ for all i . Then, given any oblivious algorithm of Π with approximation ratio ρ , there exists a cost-sharing protocol $\Xi \in \mathcal{C}$ for G with $\text{PoA} = O(\rho)$.

Proof. Suppose that E_i is the set of the resources allocated by the oblivious algorithm to the the input component that serves as the type of some player i . Even though it is not quite correct, we will say that R_i is the resources allocated to player i . Let S_e be the set of players that resource e is allocated to them.

We denote by h a very high value with respect to the parameters of the game. h should be larger than the total cost-share of any player by using any

budget-balanced protocol. Regarding the multicast game, it is safe to assume that $h > \sum_{e \in E} c_e$.

Then Ξ assigns the following cost-share to any player i for choosing any resource e , when the set of players choosing e is S ,

$$\xi_e(i, S) = \begin{cases} c_e/|S| & \text{if } i \in S_e \\ h & \text{if } i \notin S_e \\ 0 & \text{otherwise} \end{cases}$$

Note that Ξ assigns equal shares restricted to S_e and a high value h for other players. In fact, instead of equal shares we could use any budget-balanced protocol restricted to S_e , for instance any generalised weighted Shapley protocol.

Note that any player i using a resource $e \notin E_i$ should pay h . By the definition of h , this is *strictly* more than $\sum_{e' \in E_i} c_{e'}$, which is the maximum she may pay if she deviates to E_i . Therefore, the only Nash equilibria are for each player i to choose some subset of E_i , i.e follow the oblivious solution. This results to a PoA which is at most the same with the approximation ratio of the oblivious algorithm, so $\text{PoA} = O(\rho)$. Moreover, by the construction of Ξ , BBiE holds. \square

The spanning tree produced by Algorithm 1 provides an oblivious (randomised) algorithm with constant approximation. For completeness we restate the Algorithm here, adjusted in order to clearly extract an oblivious solution.

Algorithm 2 Randomised Oblivious Algorithm \mathcal{A}_{obl}

Input: A rooted graph $G = (V, E, r)$ and an oracle for D .

Output: Mapping of each vertex $v \in V$ to a path from v to r .

- Choose a random set of vertices R by drawing the types from distribution D and eliminating the duplicates. Construct an α -approximate minimum Steiner tree, $T_\alpha(R)$, over $R \cup \{r\}$.
 - Form a spanning tree T by connecting all other vertices $V \setminus V(T_\alpha(R))$ with their nearest neighbour in $V(T_\alpha(R))$ (by breaking ties arbitrarily).
 - Map each vertex $v \in V$ to the path connecting v with r in T .
-

By following the proof of Theorem 68, \mathcal{A}_{obl} results in a constant approximation of the expected minimum Steiner tree. Therefore, the following theorem holds, where $E[c(\mathcal{A}_{obl}(S))]$ is the expected cost induced by algorithm \mathcal{A}_{obl} under input S , where the expectation is over the randomisation of the algorithm, T_S^*

is the minimum Steiner tree on $S \cup r$ and α is an approximation ratio for the minimum Steiner tree problem.

Theorem 76. The approximation ratio of \mathcal{A}_{obl} when the input is drawn from D is $\alpha + 2$, i.e. $\mathbb{E}_{S \sim D}[E[c(\mathcal{A}_{obl}(S))]] \leq (\alpha + 2)E_{S \sim D}[c(T_S^*)]$.

By combining Theorems 75 and 76 and applying the best known approximation of the minimum Steiner tree [26], where $\alpha = 1.39$, we derive the following corollary:

Corollary 77. In the multicast game, there exists $\Xi \in \mathcal{C}$ (computed in polynomial time) with $\text{PoA} \leq 3.39$.

Garg et al. [72] showed a constant approximation on the online Steiner tree problem with stochastic input by using an oblivious algorithm. Their oblivious solution is the same with one derived by algorithm \mathcal{A}_{obl} . However, in their work they sampled the input set S from a distribution π^k , meaning that by using their result as it is, we could construct a BBiE protocol for i.i.d. players' types. Nevertheless, their algorithm holds for more general distributions $\pi_1 \times \pi_2 \times \dots \times \pi_k$ as we state in this thesis.

12.2.2 BBiE Posted Prices

In this section, we show how to set *anonymous* prices for the edges. Strict BBiE cannot be obtained by using anonymous posted prices, as we illustrate in Example 78. Instead, we require *ex-ante* budget-balance, meaning budget-balance on expectation. Similar bounds with the BBiE protocols can be shown here. But first, we argue that other natural variations of budget balance cannot be very promising: a) BBiE with “high” probability, b) bounded possible excess and deficit. Example 78 indicates that *any* anonymous posted prices may result in BBiE with probability at most $O(1/\sqrt{k})$. The same example serves to demonstrate that *no* posted prices can guarantee good bounds on possible excess and deficit, i.e. for *any* posted prices, there are cases where the total shares for some resource are either at least \sqrt{k} or at most $1/\sqrt{k}$ of the resource's cost. We stress that those restrictions holds even for i.i.d. types.

Example 78. Consider a graph with only three vertices v, u, r , where r is the designated root, and there are only two edges (v, t) and (u, t) , of unit cost. Further, consider k i.i.d. players whose type is the uniform distribution over the

two vertices v, u . The question that arises is how to set a price on an edge of unit cost, when each player may use it with probability $1/2$.

Let q be the price for edge e . If $1/q$ is not an integer in $[k]$, then budget-balance appears with zero probability. So, suppose that $1/q = a \in [k]$, then budget-balance appears only when a players use e and this happens with probability

$$\mathbb{P}[\# \text{ players} = a] = \binom{k}{a} \left(\frac{1}{2}\right)^a \left(1 - \frac{1}{2}\right)^{k-a} \leq \binom{k}{\lfloor k/2 \rfloor} \frac{1}{2^k} < \frac{1}{\sqrt{k}}.$$

Furthermore, for any price q for edge e , if $q \geq 1/\sqrt{k}$ then, in the case that all players use e , the total shares sum up to at least $k \cdot 1/\sqrt{k} = \sqrt{k}$. On the other hand, if $q < 1/\sqrt{k}$ then, in the case that only one player uses e , her share is at most $1/\sqrt{k}$. This means that we cannot guarantee good bounds on any possible excess and deficit.

Next we state our main theorem for this section, which is the existence of *anonymous* posted prices that are *ex-ante* BBiE.

Theorem 79. In the multicast game, there exist prices (computed in polynomial time) with PoA = 3.39.

Proof. We define, for any subset of types A , k_A to be the expected number of players having type in A and k_A^1 to be the expected number of players having type in A , given there exists at least one such player:

$$\begin{aligned} k_A &= \mathbb{E}_{\mathbf{t}}[|i : t_i \in A|] = k \sum_{i \in A} \pi_i; \\ k_A^1 &= \mathbb{E}_{\mathbf{t}}[|i : t_i \in A| \text{ given } |i : t_i \in A| \geq 1] = \frac{k \sum_{i \in A} \pi_i}{1 - \left(1 - \sum_{i \in A} \pi_i\right)^k}. \end{aligned}$$

Let T be the spanning tree constructed by \mathcal{A}_{obl} (Algorithm 2). We use T to set the posted prices. For each edge $e \in E(T)$, let $V(e)$ be the set of vertices that are disconnected from r in $T \setminus \{e\}$. We set the posted price p_e for each edge e as following:

$$\begin{aligned} \text{If } e \in E(T), & \quad p_e = c_e/k_{V(e)}^1, \\ \text{if } e \notin E(T), & \quad p_e = h > \sum_{e \in E} c_e. \end{aligned}$$

In the Bayes-Nash equilibrium each player will choose the unique path that connects her terminal with the root in T . The constant PoA follows by Theorem 76. Note that the expected sum of the prices for edge e is $k_{V(e)}^1 \cdot c_e / k_{V(e)}^1 = c_e$, whenever e is used and 0 otherwise; therefore the ex-ante BBiE is preserved. \square

Conclusion

In this thesis, we answered several questions in regards to mechanisms' efficiency in Auctions and Network Design and we further posed new open questions.

Auctions. The simultaneous first-price auction appears to be the best auction so far in terms of the PoA. We provided *tight* lower bounds which complement the current knowledge about the first-price auction for two important classes of valuation functions, namely fractionally subadditive and subadditive. For valuations without complementarities, i.e. subadditive valuations, the tight bounds hold for a more general class of auctions that includes all-pay auctions, and it is further extended to multi-unit auctions and divisible resources, like bandwidth.

Roughgarden [124] presented a very elegant methodology to provide PoA lower bounds for all *simple* auctions. One consequence is the indication that the simultaneous first-price auction is the most efficient (i.e. has the lowest PoA) simple auction for valuations without complementarities. However, a combination of [124] and this thesis' results indicates that the question of the most efficient auction remains, regarding the more specialised class of submodular valuations; either a different approach than the one in [124] is needed in order to prove optimality of the first-price auction or there is another auction that improves the PoA.

Network Design. Designing protocols for network cost-sharing games has drawn a lot of attention lately. In this thesis we posed and partially answered the following question: to what extent can prior knowledge of the underlying metric help in the design? For the general case of arbitrary metric spaces, we answered this question negatively. On the bright side, it seems that there are cases, such as the outerplanar graphs, where prior knowledge of the metric can dramatically improve the design. However, the design question still remains for

other significant metrics such as the Euclidean and planar graphs. Moreover, very few is known with respect to randomised protocols. It is possible that randomisation may help in the design and we have some serious indications that a randomised protocol may lead to a great improvement on the PoA in Euclidean metrics.

Of high interest, due to its simplicity and fairness, is the protocol where the cost of each edge is equally split among its users. Anshelevich et al. [6] showed that the quality of equilibria can be really poor, meaning that the PoA is very high. However, the Price of Stability of this game is not well-understood. It is an interesting open question to determine its exact value that is between constant and a sublogarithmic value [103].

Bibliography

- [1] Fidaa Abed, José R. Correa, and Chien-Chung Huang. Optimal coordination mechanisms for multi-job scheduling games. In *Proceedings of the 22nd Annual European Symposium on Algorithms, ESA 2014, Wroclaw, Poland, September 8-10, 2014*, pages 13–24, 2014.
- [2] Fidaa Abed and Chien-Chung Huang. Preemptive coordination mechanisms for unrelated machines. In *Proceedings of the 20th Annual European Symposium on Algorithms, ESA 2012, Ljubljana, Slovenia, September 10-12, 2012*, pages 12–23, 2012.
- [3] Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM J. Comput.*, 40(5):1211–1233, 2011.
- [4] Noga Alon and Yossi Azar. On-line steiner trees in the euclidean plane. *Discrete & Computational Geometry*, 10:113–121, 1993.
- [5] Noga Alon, Rados Radoicic, Benny Sudakov, and Jan Vondrák. A ramsey-type result for the hypercube. *Journal of Graph Theory*, 53(3):196–208, 2006.
- [6] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.
- [7] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.
- [8] Baruch Awerbuch, Yossi Azar, and Yair Bartal. On-line generalized steiner problem. *Theor. Comput. Sci.*, 324(2-3):313–324, 2004.

- [9] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. *SIAM J. Comput.*, 42(1):160–177, 2013.
- [10] Yossi Azar, Lisa Fleischer, Kamal Jain, Vahab S. Mirrokni, and Zoya Svitkina. Optimal coordination mechanisms for unrelated machine scheduling. *Operations Research*, 63(3):489–500, 2015.
- [11] Martin Beckmann, C. B. McGuire, and Christopher B. Winsten. *Studies in the economics of transportation*. Yale University Press, 1955.
- [12] Piotr Berman and Chris Coulston. On-line algorithms for steiner tree problems (extended abstract). In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing, STOC 1997, El Paso, Texas, USA, May 4-6, 1997*, pages 344–353, 1997.
- [13] Dimitris Bertsimas and Michelangelo Grigni. Worst-case examples for the spacefilling curve heuristic for the euclidean traveling salesman problem. *Operations Research Letters*, 8(5):241 – 244, 1989.
- [14] Dimitris J. Bertsimas, Patrick Jaillet, and Amedeo R. Odoni. A priori optimization. *Operations Research*, 38(6):1019–1033, 1990.
- [15] Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. Optimal lower bounds for universal and differentially private steiner trees and tsps. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011, Princeton, NJ, USA, August 17-19, 2011. Proceedings*, pages 75–86, 2011.
- [16] Sayan Bhattacharya, Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Coordination mechanisms from (almost) all scheduling policies. In *Innovations in Theoretical Computer Science, ITCS'14, Princeton, NJ, USA, January 12-14, 2014*, pages 121–134, 2014.
- [17] Sayan Bhattacharya, Janardhan Kulkarni, and Vahab S. Mirrokni. Coordination mechanisms for selfish routing over time on a tree. In *Proceedings (Part I) of the 41st International Colloquium on Automata, Languages and Programming, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014*, pages 186–197, 2014.

- [18] Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden. Weighted congestion games: The price of anarchy, universal worst-case examples, and tightness. *ACM Trans. Economics and Comput.*, 2(4):14:1–14:23, 2014.
- [19] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 700–709, 2011.
- [20] Kshipra Bhawalkar and Tim Roughgarden. Simultaneous single-item auctions. In *Proceedings of the 8th International Conference on Web and Internet Economics, WINE 2012, Liverpool, UK, December 10-12, 2012*, pages 337–349, 2012.
- [21] Sushil Bikhchandani. Auctions of heterogeneous objects. *Games and Economic Behavior*, 26(2):193 – 220, 1999.
- [22] Vittorio Bilò and Roberta Bove. Bounds on the price of stability of undirected network design games with three players. *Journal of Interconnection Networks*, 12(1-2):1–17, 2011.
- [23] Vittorio Bilò, Ioannis Caragiannis, Angelo Fanelli, and Gianpiero Monaco. Improved lower bounds on the price of stability of undirected network design games. *Theory Comput. Syst.*, 52(4):668–686, 2013.
- [24] Vittorio Bilò, Michele Flammini, and Luca Moscardelli. The price of stability for undirected broadcast network design with fair cost allocation is constant. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 638–647, 2013.
- [25] Michal Bresky. Pure Equilibrium Strategies in Multi-unit Auctions with Private Value Bidders. CERGE-EI Working Papers wp376, The Center for Economic Research and Graduate Education - Economic Institute, Prague, December 2008.
- [26] Jaroslav Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. An improved lp-based approximation for steiner tree. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 583–592, 2010.

- [27] Ioannis Caragiannis. Efficient coordination mechanisms for unrelated machine scheduling. *Algorithmica*, 66(3):512–540, 2013.
- [28] Ioannis Caragiannis and Angelo Fanelli. An almost ideal coordination mechanism for unrelated machine scheduling. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory, SAGT 2016, Liverpool, UK, September 19-21, 2016*, pages 315–326, 2016.
- [29] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM Conference on Electronic Commerce, EC 2011, San Jose, CA, USA, June 5-9, 2011*, pages 81–90, 2011.
- [30] Ioannis Caragiannis and Alexandros A. Voudouris. Welfare guarantees for proportional allocations. *Theory Comput. Syst.*, 59(4):581–599, 2016.
- [31] Moses Charikar, Howard J. Karloff, Claire Mathieu, Joseph Naor, and Michael E. Saks. Online multicast with egalitarian cost sharing. In *Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2008, Munich, Germany, June 14-16, 2008*, pages 70–76, 2008.
- [32] Shuchi Chawla and Jason D. Hartline. Auctions with unique equilibria. In *Proceedings of the 14th ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, pages 181–196, 2013.
- [33] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 311–320, 2010.
- [34] Chandra Chekuri, Julia Chuzhoy, Liane Lewin-Eytan, Joseph Naor, and Ariel Orda. Non-cooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications*, 25(6):1193–1206, 2007.
- [35] Ho-Lin Chen and Tim Roughgarden. Network design with weighted players. *Theory Comput. Syst.*, 45(2):302–324, 2009.

- [36] Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. Designing network protocols for good equilibria. *SIAM J. Comput.*, 39(5):1799–1832, 2010.
- [37] George Christodoulou, Christine Chung, Katrina Ligett, Evangelia Pyrga, and Rob van Stee. On the price of stability for undirected network design. In *Approximation and Online Algorithms, 7th International Workshop, WAOA 2009, Copenhagen, Denmark, September 10-11, 2009. Revised Papers*, pages 86–97, 2009.
- [38] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, STOC 2005, Baltimore, MD, USA, May 22-24, 2005*, pages 67–73, 2005.
- [39] George Christodoulou, Elias Koutsoupias, and Akash Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.
- [40] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. *J. ACM*, 63(2):11, 2016.
- [41] George Christodoulou, Annamária Kovács, Alkmini Sgouritsa, and Bo Tang. Tight bounds for the price of anarchy of simultaneous first-price auctions. *ACM Trans. Economics and Comput.*, 4(2):9, 2016.
- [42] George Christodoulou, Stefano Leonardi, and Alkmini Sgouritsa. Designing cost-sharing methods for bayesian games. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory, SAGT 2016, Liverpool, UK, September 19-21, 2016*, pages 327–339, 2016.
- [43] George Christodoulou and Alkmini Sgouritsa. Designing networks with good equilibria under uncertainty. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 72–89, 2016.
- [44] George Christodoulou and Alkmini Sgouritsa. An improved upper bound for the universal TSP on the grid. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, January 16-19, 2017*, pages 1006–1021, 2017.

- [45] George Christodoulou, Alkmini Sgouritsa, and Bo Tang. On the efficiency of all-pay mechanisms. In *Proceedings of the 23rd Annual European Symposium on Algorithms, ESA 2015, Patras, Greece, September 14-16, 2015*, pages 349–360, 2015.
- [46] George Christodoulou, Alkmini Sgouritsa, and Bo Tang. On the efficiency of the proportional allocation mechanism for divisible resources. In *Proceedings of the 8th International Symposium on Algorithmic Game Theory, SAGT 2015, Saarbrücken, Germany, September 28-30, 2015*, pages 165–177, 2015.
- [47] George Christodoulou, Alkmini Sgouritsa, and Bo Tang. On the efficiency of the proportional allocation mechanism for divisible resources. *Theory Comput. Syst.*, 59(4):600–618, 2016.
- [48] Giorgos Christodoulou, Kurt Mehlhorn, and Evangelia Pyrga. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica*, 69(3):619–640, 2014.
- [49] Brent N. Chun and David E. Culler. Market-based proportional resource sharing for clusters. Technical report, Berkeley, CA, USA, 2000.
- [50] Edward Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- [51] Ilan Reuven Cohen, Alon Eden, Amos Fiat, and Lukasz Jez. Pricing online decisions: Beyond auctions. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 73–91, 2015.
- [52] Vincent Cohen-Addad, Alon Eden, Michal Feldman, and Amos Fiat. The invisible hand of dynamic market pricing. In *Proceedings of the 17th ACM Conference on Economics and Computation, EC 2016, Maastricht, The Netherlands, July 24-28, 2016*, pages 383–400, 2016.
- [53] Richard Cole, José R. Correa, Vasilis Gkatzelis, Vahab S. Mirrokni, and Neil Olver. Inner product spaces for minsum coordination mechanisms. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 539–548, 2011.

- [54] Gérard Cornuéjols, Jean Fonlupt, and Denis Naddef. The traveling salesman problem on a graph and some related integer polyhedra. *Mathematical Programming*, 33(1):1–27, 1985.
- [55] José R. Correa, Andreas S. Schulz, and Nicolás E. Stier Moses. The price of anarchy of the proportional allocation mechanism revisited. In *Proceedings of the 9th International Conference on Web and Internet Economics, WINE 2013, Cambridge, MA, USA, December 11-14, 2013*, pages 109–120, 2013.
- [56] Bart de Keijzer, Evangelos Markakis, Guido Schäfer, and Orestis Telelis. Inefficiency of standard multi-unit auctions. In *Proceedings of the 21st Annual European Symposium on Algorithms, ESA 2013, Sophia Antipolis, France, September 2-4, 2013.*, pages 385–396, 2013.
- [57] Nikhil R. Devanur, Milena Mihail, and Vijay V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems*, 39(1):11–22, 2005.
- [58] Yann Disser, Andreas Emil Feldmann, Max Klimm, and Matús Mihalák. Improving the hk-bound on the price of stability in undirected shapley network design games. *Theor. Comput. Sci.*, 562:557–564, 2015.
- [59] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, 2010.
- [60] Dominic Dumrauf and Martin Gairing. Price of anarchy for polynomial wardrop games. In *Proceedings of the 2nd International Conference on Web and Internet Economics, WINE 2006, Patras, Greece, December 15-17, 2006*, pages 319–330, 2006.
- [61] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.
- [62] Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1):122–142, 2009.

- [63] Uriel Feige and Jan Vondrák. The submodular welfare problem with demand queries. *Theory of Computing*, 6(1):247–290, 2010.
- [64] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing Conference, STOC 2013, Palo Alto, CA, USA, June 1-4, 2013*, pages 201–210, 2013.
- [65] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 123–135, 2015.
- [66] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial walrasian equilibrium. *SIAM J. Comput.*, 45(1):29–48, 2016.
- [67] Michal Feldman, Kevin Lai, and Li Zhang. The proportional-share allocation market for computational resources. *IEEE Trans. Parallel Distrib. Syst.*, 20(8):1075–1088, 2009.
- [68] Amos Fiat, Haim Kaplan, Meital Levy, Svetlana Olonetsky, and Ronen Shabo. On the price of stability for designing undirected networks with fair cost allocations. In *Proceedings (Part I) of the 33rd International Colloquium on Automata, Languages and Programming, ICALP 2006, Venice, Italy, July 10-14, 2006*, pages 608–618, 2006.
- [69] Dimitris Fotakis, Spyros C. Kontogiannis, and Paul G. Spirakis. Selfish unsplittable flows. *Theor. Comput. Sci.*, 348(2-3):226–239, 2005.
- [70] Hu Fu, Robert Kleinberg, and Ron Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, Valencia, Spain, June 4-8, 2012*, page 586, 2012.
- [71] Martin Gairing, Konstantinos Kollias, and Grammateia Kotsialou. Tight bounds for cost-sharing in weighted congestion games. In *Proceedings (Part II) of the 42nd International Colloquium on Automata, Languages and Programming, ICALP 2015, Kyoto, Japan, July 6-10, 2015*, pages 626–637, 2015.

- [72] Naveen Garg, Anupam Gupta, Stefano Leonardi, and Piotr Sankowski. Stochastic analyses for online combinatorial optimization problems. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 942–951, 2008.
- [73] Vasilis Gkatzelis, Konstantinos Kollias, and Tim Roughgarden. Optimal cost-sharing in general resource selection games. *Operations Research*, 64(6):1230–1238, 2016.
- [74] Michel X. Goemans, Vahab S. Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2005, 23-25 October 2005, Pittsburgh, PA, USA*, pages 142–154, 2005.
- [75] Michel X. Goemans and David P. Williamson. A general approximation technique for constrained forest problems. *SIAM J. Comput.*, 24(2):296–317, 1995.
- [76] Ragavendran Gopalakrishnan, Jason R. Marden, and Adam Wierman. Potential games are *Necessary* to ensure pure nash equilibria in cost sharing games. *Math. Oper. Res.*, 39(4):1252–1296, 2014.
- [77] Igor Gorodezky, Robert D. Kleinberg, David B. Shmoys, and Gwen Spencer. Improved lower bounds for the universal and *a priori* TSP. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010, Barcelona, Spain, September 1-3, 2010. Proceedings*, pages 178–191, 2010.
- [78] Ronald L. Graham, Bruce L. Rothschild, and Joel H. Spencer. *Ramsey Theory, 2nd Edition*. Wiley Series in Discrete Mathematics and Optimization, 1990.
- [79] Theodore Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, July 1973.
- [80] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95 – 124, 1999.

- [81] Anupam Gupta, Amit Kumar, Martin Pál, and Tim Roughgarden. Approximation via cost sharing: Simpler and better approximation algorithms for network design. *J. ACM*, 54(3):11, 2007.
- [82] B. Hajek and G. Gopalakrishnan. Do Greedy Autonomous Systems Make for a Sensible Internet. Presented at the Conference on Stochastic Networks, Stanford University, 2002.
- [83] Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and Frank Thomson Leighton. Improved lower and upper bounds for universal TSP in planar metrics. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 649–658, 2006.
- [84] John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part ii. bayesian equilibrium points. *Management Science*, 14(5):320–334, 1968.
- [85] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce, EC 2009, Stanford, California, USA, July 6–10, 2009*, pages 225–234, 2009.
- [86] Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce, EC 2011, San Jose, CA, USA, June 5-9, 2011*, pages 295–296, 2011.
- [87] Makoto Imase and Bernard M. Waxman. Dynamic Steiner tree problem. *SIAM J. Discrete Math.*, 4(3):369–384, 1991.
- [88] Nicole Immorlica, Li (Erran) Li, Vahab S. Mirrokni, and Andreas S. Schulz. Coordination mechanisms for selfish scheduling. *Theor. Comput. Sci.*, 410(17):1589–1598, 2009.
- [89] Nicole Immorlica, Mohammad Mahdian, and Vahab S. Mirrokni. Limitations of cross-monotonic cost-sharing schemes. *ACM Trans. Algorithms*, 4(2), 2008.
- [90] Lujun Jia, Guolong Lin, Guevara Noubir, Rajmohan Rajaraman, and Ravi Sundaram. Universal approximations for tsp, steiner tree, and set cover. In

- Proceedings of the 37th Annual ACM Symposium on Theory of Computing, STOC 2005, Baltimore, MD, USA, May 22-24, 2005*, pages 386–395, 2005.
- [91] Ramesh Johari and John N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Math. Oper. Res.*, 29(3):407–435, 2004.
- [92] Ramesh Johari and John N. Tsitsiklis. Efficiency of scalar-parameterized mechanisms. *Operations Research*, 57(4):823–839, 2009.
- [93] David R. Karger and Maria Minkoff. Building steiner trees with incomplete global knowledge. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 613–623, 2000.
- [94] Frank Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8(1):33–37, 1997.
- [95] Alexander S. Kelso and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.
- [96] Konstantinos Kollias. Nonpreemptive coordination mechanisms for identical machines. *Theory Comput. Syst.*, 53(3):424–440, 2013.
- [97] Jochen Könemann, Stefano Leonardi, Guido Schäfer, and Stefan H. M. van Zwam. A group-strategyproof cost sharing mechanism for the steiner forest game. *SIAM J. Comput.*, 37(5):1319–1341, 2008.
- [98] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, STACS 1999, Trier, Germany, March 4-6, 1999*, STACS 99, pages 404–413, 1999.
- [99] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
- [100] Vijay Krishna. *Auction Theory*. Academic Press, 2002.
- [101] Euiwoong Lee and Katrina Ligett. Improved bounds on the price of stability in network cost sharing games. In *Proceedings of the 14th ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, pages 607–620, 2013.

- [102] Renato Paes Leme and Éva Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 735–744, 2010.
- [103] Jian Li. An $o(\log(n)/\log(\log(n)))$ upper bound on the price of stability for undirected shapley network design games. *Inf. Process. Lett.*, 109(15):876–878, 2009.
- [104] Alessandro Lizzeri and Nicola Persico. Uniqueness and existence of equilibrium in auctions with a reserve price. *Games and Economic Behavior*, 30(1):83–114, 2000.
- [105] Brendan Lucier and Allan Borodin. Price of anarchy for greedy auctions. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 537–553, 2010.
- [106] Brendan Lucier and Renato Paes Leme. GSP auctions with correlated types. In *Proceedings of the 12th ACM Conference on Electronic Commerce, EC 2011, San Jose, CA, USA, June 5-9, 2011*, pages 71–80, 2011.
- [107] Jason R. Marden and Adam Wierman. Distributed welfare games. *Operations Research*, 61(1):155–168, 2013.
- [108] Evangelos Markakis and Orestis Telelis. Uniform price auctions: Equilibria and efficiency. *Theory Comput. Syst.*, 57(3):549–575, 2015.
- [109] Paul Milgrom. Putting Auction Theory to Work: The Simultaneous Ascending Auction. *The Journal of Political Economy*, 108(2):245–272, 2000.
- [110] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124 – 143, 1996.
- [111] Hervé Moulin and Scott Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
- [112] Thành Nguyen and Éva Tardos. Approximately maximizing efficiency and revenue in polyhedral environments. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC 2007, San Diego, California, USA, June 11-15, 2007*, pages 11–19, 2007.

- [113] Noam Nisan. The communication complexity of approximate set packing and covering. In *Proceedings of the 29th International Colloquium on Automata, Languages and Programming, ICALP 2002, Malaga, Spain, July 8-13, 2002*, pages 868–875, 2002.
- [114] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing, STOC 1999, May 1-4, 1999, Atlanta, Georgia, USA*, pages 129–140, 1999.
- [115] Noam Nisan and Amir Ronen. Computationally feasible vcg mechanisms. In *Proceedings of the 2nd ACM Conference on Electronic Commerce, EC 2000*, pages 242–252, New York, NY, USA, 2000. ACM.
- [116] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [117] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *J. Economic Theory*, 129(1):192–224, 2006.
- [118] Christos H. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, STOC 2001, July 6-8, 2001, Heraklion, Crete, Greece*, pages 749–753, 2001.
- [119] Arthur C. Pigou. *The Economics of Welfare*. Macmillan and Co., 1920.
- [120] Loren K. Platzman and John J. Bartholdi III. Spacefilling curves and the planar travelling salesman problem. *J. ACM*, 36(4):719–737, 1989.
- [121] Robert C. Prim. Shortest connection networks and some generalizations. *Bell System Technology Journal*, 36:1389–1401, 1957.
- [122] R. A. Rosenbaum. Sub-additive functions. *Duke Mathematical Journal*, 17(3):227–247, 09 1950.
- [123] Robert W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.
- [124] Tim Roughgarden. Barriers to near-optimal equilibria. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science*,

- FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 71–80, 2014.
- [125] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32, 2015.
- [126] Tim Roughgarden. The price of anarchy in games of incomplete information. *ACM Trans. Economics and Comput.*, 3(1):6, 2015.
- [127] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
- [128] Frans Schalekamp and David B. Shmoys. Algorithms for the universal and a priori TSP. *Oper. Res. Lett.*, 36(1):1–3, 2008.
- [129] Lloyd Shapley and Martin Shubik. Trade using one commodity as a means of payment. *Journal of Political Economy*, 85(5):937–968, 1977.
- [130] David B. Shmoys and Kunal Talwar. A constant approximation algorithm for the a priori traveling salesman problem. In *Proceedings of the 13th International Conference on Integer Programming and Combinatorial Optimization, IPCO 2008, Bertinoro, Italy, May 26-28, 2008*, pages 331–343, 2008.
- [131] Ion Stoica, Hussein Abdel-wahab, Kevin Jeffay, Sanjoy K. Baruah, Johannes E. Gehrke, and C. Greg Plaxton. A proportional share resource allocation algorithm for real-time, time-shared systems. In *Proceedings of the 17th IEEE Real-Time Systems Symposium*, pages 288–299, 1996.
- [132] Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1):79–96, 2007.
- [133] Vasilis Syrgkanis and Éva Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing Conference, STOC 2013, Palo Alto, CA, USA, June 1-4, 2013*, pages 211–220, 2013.
- [134] Maciej M. Syslo. Characterizations of outerplanar graphs. *Discrete Mathematics*, 26(1):47 – 53, 1979.

- [135] Chih-Wei Tsai and Zsehong Tsai. Bid-proportional auction for resource allocation in capacity-constrained clouds. In *Proceedings of the 26th International Conference on Advanced Information Networking and Applications Workshops, WAINA 2012, Fukuoka, Japan, March 26-29, 2012*, pages 1178–1183, 2012.
- [136] Seeun Umboh. Online network design algorithms via hierarchical decompositions. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1373–1387, 2015.
- [137] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.
- [138] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.
- [139] Philipp von Falkenhausen and Tobias Harks. Optimal cost sharing for resource selection games. *Math. Oper. Res.*, 38(1):184–208, 2013.
- [140] John G. Wardrop. Road paper. some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1(3):325–362, 1952.
- [141] David P. Williamson and Anke van Zuylen. A simpler and better derandomization of an approximation algorithm for single source rent-or-buy. *Oper. Res. Lett.*, 35(6):707–712, 2007.
- [142] Li Zhang. The efficiency and fairness of a fixed budget resource allocation game. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming, ICALP 2005*, pages 485–496, Berlin, Heidelberg, 2005. Springer-Verlag.