Multi-Item Auctions and Fair Division

Vishnu V. Narayan

School of Computer Science

McGill University, Montreal

August, 2022

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Doctor of Philosophy

Abstract

The question of how to divide a collection of items amongst a set of agents is of central importance to society. There are two main directions from which this question is approached: a game-theoretic direction that studies the mechanisms – primarily auctions – that are used to divide items amongst agents, and a normative direction that studies the existence and computability of allocations that have desirable properties like fairness and high social welfare. In this thesis, we detail our contributions to both areas.

In Part I of this thesis, we analyze two prominent multi-item auctions, the sequential and simultaneous item-bidding auctions. We prove that the declining price anomaly is not guaranteed to hold in the equilibria of full-information sequential auctions with three or more buyers. We then analyze the risk-free profitability, i.e. the threshold payoff that a buyer can guarantee for itself, in sequential and simultaneous auctions, when the buyer's valuation function is in the subadditive set function class (and its subclasses).

In Part II, we discuss our contributions to the fair division problem, focusing on the envy-free allocation of indivisible items along with payments. We prove two conjectures of Halpern and Shah [*SAGT 2019*] and present additional upper bounds on the total quantity of subsidy sufficient to guarantee envy-freeness in any instance. We then study the tradeoffs between transfer payments, fairness, and welfare.

Abrégé

Savoir comment répartir un ensemble de biens à un groupe d'agent est central à la société. Il y a deux approches principales pour répondre à cette question : L'approche de la théorie des jeux où on étudie des mécanismes – principalement des enchères – qui sont utilisés pour répartir les biens, et l'approche normative qui étudie l'existence et la facilité de calculer des allocations ayant des propriétés désirables tel l'équitabilité et une valeur sociale élevée. Dans cette thèse, nous détaillons nos contributions dans les deux domaines.

Dans la première partie de cette thèse, nous analysons deux catégories majeures d'enchère avec plusieurs biens, les enchères en série et les enchères simultanées. Nous prouvons que les prix ne chutent pas nécessairement dans les équilibres des enchères en série avec au moins 3 acheteurs lorsque l'information est publique. Ensuite, nous analysons le profit sans risque, c'est-à-dire le profit de seuil qu'un acheteur peut se garantir, dans les enchères en série et simultanées, lorsque la fonction d'évaluation de l'acheteur est dans la classe des fonctions sous-additives (et ses sous-classes).

Dans la seconde partie, nous discutons nos contributions au problème des allocations équitables, nous nous concentrons sur les allocations d'objets sans envie d'objets indivisibles avec des paiements. Nous prouvons deux conjectures d'Halpern et Shah [*SAGT 2019*] et présentons des bornes supérieures supplémentaires sur la quantité totale de subventions suffisante pour garantir aucune envie dans chaque exemple. Nous étudions ensuite le compromis entre les paiements de transferts, l'équitabilité et le bien-être.

Acknowledgements

I am grateful to have had the opportunity to be supervised for my Ph.D. by Adrian Vetta, who is the best advisor one could hope to have. Adrian is very knowledgeable in several areas of theoretical computer science, is a highly skilled speaker, writer and educator, is tremendously patient and kind, and is endlessly generous in allocating his valuable time and unlimited energy to his students.

I have special gratitude for Joseph Cheriyan, who advised me through my M.Math. program and has continued to generously guide and support me in the following years.

I am lucky to have been surrounded by many excellent researchers and collaborators, including Adrian, Joseph, Fabrizio, Arindam, Enguerrand, Yaqiao, Ndiamé, Noah, Sarah, Ben, Siddharth, and Paritosh. I am especially grateful for the passion and knowledge of Denis Pankratov, who is delightful to work with, and of Bruce Reed, from whom I learned a great deal. I'm also very thankful for the invitation from Ariel Procaccia as a visiting Fellow at Harvard; the many wonderful members of Ariel's research group remain close colleagues and friends.

My time at McGill was enriched by a long list of lifelong friends with whom I have had many joyous and enriching conversations: among many others, this list includes Gautam, Johannes, Jack, Mashbat, Mete, Lianna, Yaqiao, Ndiamé, Richard, Vishal and Navin; thank you! Special thanks are due to Vadiraj for the many hundreds of weekend hours spent queuing vills.

I'm grateful for the financial support offered by McGill University, the Harold H. Helm Fellowship, the Murata Family Fellowship, the Trottier Institute for Science and Public

Policy, and the Government of Canada. I'm also grateful for the continuous assistance offered by the administrative staff at McGill (particularly to Ann and Kamini in the CS department for all of their help over the years.)

I have learned more than I can remember from the professors and course instructors that I have interacted with at RVCE, Waterloo and McGill. I'm thankful for the many friends and colleagues I've had at these places (too many to name). I'm especially grateful to Vijayaraghavan (Jama): his passion and desire to win at life will inspire me forever.

I am very grateful to my thesis examiners, Rupert Freeman and Sergey Norin, for the large amount of time and effort spent reviewing my thesis and for their feedback and suggestions for improvements. I'm also grateful to the members of my various Ph.D. committees, including Adrian Vetta, Yang Cai, Bruce Reed, Jérôme Waldispühl, Hamed Hatami, David Rolnick, Luc Devroye, Sergey Norin, Elliot Paquette, Robert Robere, Xiao-Wen Chang, and all of the administrative staff involved in my evaluations.

I am very thankful for the continuous support that I receive from my parents, sister, and brother-in-law.

Finally, I am grateful to my wife, whose unconditional love makes everything possible.

Contributions

The contents of this thesis are based on the following papers, which I coauthored. Authors are listed alphabetically and all authors contributed equally to each work.

- The Declining Price Anomaly is not Universal in Multi-Buyer Sequential Auctions (but almost is)
  With E. Prebet and A. Vetta. In *SAGT 2019* and *Theory of Computing Systems: SAGT 2019 Special Issue (2022)*.

- Risk-Free Bidding in Complement-Free Combinatorial Auctions
  With G. Rayaprolu and A. Vetta. In *SAGT 2019* and *Theory of Computing Systems: SAGT 2019 Special Issue (2022)*.

- One Dollar Each Eliminates Envy
  With J. Brustle, J. Dippel, M. Suzuki and A. Vetta. In *EC 2020*.

- Two Birds With One Stone: Fairness and Welfare via Transfers
  With M. Suzuki and A. Vetta. In *SAGT 2021*.

While completing my PhD, I coauthored other published research articles whose contents are not represented in this thesis, and a list of these follows.

- The Matching Augmentation Problem: A 7/4-Approximation Algorithm
  With J. Cheriyan, J. Dippel, F. Grandoni and A. Khan. In *Mathematical Programming (2020)*.

- Online Coloring and a New Type of Adversary for Online Graph Problems
  With Y. Li and D. Pankratov. In *WAOA 2020* and *Algorithmica (2022)*.

- The Speed and Threshold of the Biased Perfect Matching Game
  With N. Brustle, S. Clusiau, N. Ndiaye, B. Reed and B. Seamone. In *LAGOS 2021*.

- The Speed and Threshold of the Biased Hamilton Cycle Game
  With N. Brustle, S. Clusiau, N. Ndiaye, B. Reed and B. Seamone. In *LAGOS 2021*.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

The question of how to divide a collection of items amongst a set of agents has remained of central importance to society since antiquity. This is evidenced by an abundance of real world examples, ranging from classical problems like the division of inherited estates and land, border settlements, and the allocation of public resources and government spending, to modern considerations such as the distribution of computational resources, allocation of the electromagnetic spectrum, and the management of airport traffic. At a high level, there are two principal directions from which this question is approached. The first is a game-theoretic or mechanism design direction that studies the mechanisms – primarily, auctions – that are used to divide items amongst agents. The second is a normative direction that studies the existence and computability of allocations that have desirable properties like fairness and high social welfare. In this thesis, we present a brief overview of historical research efforts in these two directions, and an exposition of our contributions to both areas.

Typically, in the first approach, research efforts are directed towards the design of new mechanisms that achieve certain outcomes, such as high seller revenue or high social welfare. Additionally, many simple auction mechanisms see widespread deployment in the real world, and understanding how well these auctions perform is an important research goal. *Multi-item auctions*, in which a set of items is sold to a collection of bidders who typically place bids on individual (single) items, are some of the most natural and widely implemented auctions. In Part I of this thesis, we analyze two prominent multi-item auctions, the sequential and simultaneous item-bidding auctions, and derive some new results about their structure and equilibria.

In the second approach, the goal is to accurately define desirable properties that align with broad economic goals, and to study the conditions (typically on the agents' valuation functions) under which an allocation with these properties is present for every instance, and the conditions under which (and algorithms with which) such an allocation can be computed quickly. *Fairness* is perhaps the most widely studied property, and in Part II of this thesis, we discuss some of our recent contributions to the literature on the fair division problem.

Part I

Multi-Item Auctions

An auction is a process of buying, selling or exchanging goods or services through the use of a set of trading rules. Auctions have been used for several millennia for the sale of a variety of objects; Milgrom and Weber [60] mention many accounts that record the use of auctions dating back to antiquity. Today, the range and value of goods and services sold by auction has reached staggering heights, fuelled by the development of the internet and the advent of new applications such as electronic securities trading, the sale of electromagnetic spectrum licenses to telecommunications companies, online advertising, and online commerce.

Over the course of the last few decades, auctions have been studied in great detail by the academic community, and several books have been written about auction theory [49, 51]. In its most typical form, an auction is a competitive game in which a single seller, usually called the auctioneer, has a set of items (goods and/or services) for sale. The other participants, called bidders (or buyers), have some information about how much they value the items for sale. In a variety of settings, this information may be either public or private, and it may be either static or drawn from some distribution. The auctioneer chooses in advance some mechanism by which the items are to be sold, consisting of an allocation rule (that decides how items are to be allocated to bidders) and a price rule (that decides how much money the bidders are required to transfer to the auctioneer in exchange for the items). The bidders convey some information about their valuation to the auctioneer by placing bids on the items. The mechanism then chooses an allocation of the items to the bidders and a set of prices as a function of the bids, using the predetermined allocation and payment rules. In this thesis, we study auctions in the *independent private-values* model, where each agent has an independent and private valuation function over the set of items.

Single-item auctions are among the simplest types of auctions and are of central interest to economists. The canonical single-item auctions are the *first-price* and *second-price* auctions. As the name suggests, in the first-price auction each bidder independently submits a single bid, and the item is allocated to the highest bidder at a price equal to their bid. In the second-price auction each bidder submits a bid and the highest bidder is allocated (as before), but the price she pays is the second-highest bidder's bid, or the "second price". The equilibria of these auctions, their revenue guarantees, and the empirical behaviour of real-world bidders in these auctions are all well-understood (see, for example, the seminal works of Vickrey [77] and Myerson [62]).

In a multi-item auction, the buyers have (possibly non-additive) combinatorial valuation functions over a set of items. Among the two most prominent item-bidding auctions (in which the buyers make bids on individual items) are the *sequential* auction and the *simultaneous* auction. The sequential auction is perhaps the most natural method by which to sell multiple items. In these auctions, the items are ordered and sold one after another. Their simplicity arises from the fact that a standard single-item auction, such as a first-price or second-price auction, can be used to sell each item in the collection. Sequential auctions are therefore ubiquitous: they are not only typical of auction house and online sale environments, but are widely used to sell commercial real estate and are regularly deployed by governments worldwide to raise billions of dollars in the sale of spectrum licenses, cap-and-trade credits and other public resources.

However, while single-item auctions are very well understood from both a theoretical and practical perspective, the concatenation of these auctions, surprisingly, despite their widespread prevalence, is not. There are two main reasons for this. The first is that while a sequential auction itself is easy to understand and implement, the valuation functions that the agents have over the set of auctioned items typically aren't. For example, in an estate division involving complementary items such as a car and a garage, an agent may have very little value for the car without the garage, or vice-versa, but much greater value for the pair of items together. Similarly, for a pair of items that substitute one another, an agent's value for the pair is typically no greater than its value for a single item from the pair. The second reason is the sequential nature of item sales in these auctions. While this choice of mechanism might initially appear benign, it leads to the emergence of complicated strategic interactions among agents, as they now compete not only to buy a single item in each round, but also to enter the corresponding subgame for all future rounds. Similar complexities emerge in simultaneous item-bidding auctions where, as their name suggests, bidders place a vector of bids simultaneously (one for each item).

In the following chapter, we analyze sequential auctions: first, we present an intricate analysis of the structure of their equilibria, and give a characterization of a focal subgame-perfect equilibrium of these auctions. We then study sequential *multi-unit* auctions (i.e., auctions with a collection of identical items) and show that the declining price anomaly, a well-known and empirically widely-observed property of real-world sequential auctions, does not always occur in their equilibria when there are three or more buyers.

In the subsequent chapter, we analyze *risk-free* strategies in multi-item auctions, with a particular focus on subadditive valuations and on sequential auctions. Specifically, we find the exact threshold payoff that a buyer in a sequential auction can guarantee for itself in an adversarial setting.

Chapter 2

The Declining Price Anomaly

In a sequential multi-unit auction, identical copies of an item are sold one at a time, in sequence. In a private values model with *unit-demand*, risk neutral buyers, Milgrom and Weber [61, 78] showed that the sequence of prices forms a martingale. In particular, as one might intuitively anticipate, expected prices are constant over time. In contrast to this, on attending a wine auction, Ashenfelter [7] made the surprising observation that prices for identical lots declined over time: "The law of the one price was repealed and no one even seemed to notice!". This *declining price anomaly* was also noted in sequential auctions for the disparate examples of livestock (Buccola [21]), Picasso prints (Pesando and Shum [65]) and satellite transponder leases (Milgrom and Weber [61]). Interestingly, the possibility of decreasing prices in a sequential auction was raised by Sosnick [68] nearly sixty years ago. An assortment of reasons have been given to explain this anomaly. In the case of wine auctions, proposed causes include absentee buyers utilizing non-optimal bidding strategies (Ginsburgh [41]) and the *buyer's option rule* where the auctioneer may allow the buyer of the first lot to make additional purchases at the same price (Black and de Meza [20]). Minor non-homogeneities amongst the items can also lead to falling prices. For example, in the case of art prints the items may suffer slight imperfections or wear-and-tear; as a consequence, the auctioneer may sell the prints in decreasing order of quality (Pesando and Shum [65]). More generally, a decreasing price trajectory may arise due to risk-aversion, such as non-decreasing, absolute risk-aversion (McAfee and Vincent [58]) or aversion to price-risk (Mezzetti [59]; see also Hu and Zou [46]). Further potential economic and behavioural explanations have been provided in [41, 9, 74].

Of course, most of these explanations are context-specific. However, the declining price anomaly appears more universal. In fact, in practice the anomaly is ubiquitous: It has now been observed in sequential auctions for antiques (Ginsburgh and van Ours [42]), commercial real estate (Lusht [56]), condominiums (Ashenfelter and Genesove [8]), fish (Gallegati et al. [38]), flowers (van den Berg et al. [75]), fur (Lambson and Thurston [53]), lobsters (Salladarre et al. [67]), jewellery (Chanel et al. [28]), paintings (Beggs and Graddy [14]), stamps (Thiel and Petry [73]) and wool (Burns [23]). Despite this prevalence, and despite the fact that declining prices are common knowledge among participants at wine auctions, Ashenfelter [7] observed that even single-unit buyers often choose to buy at the higher price rather than save money by waiting for later rounds. To add to these empirical results, Milgrom and Weber (in unpublished work) showed that with single-unit demand bidders, expected prices should *increase* rather than decrease over time due to information release in earlier rounds.

Prompted by these inconsistencies, Gale and Stegeman [37] set out to find a more foundational explanation: is it possible that the equilibria of these auctions exhibit declining prices? In their groundbreaking work, they studied the equilibria of second-price sequential auctions with two multiunit-demand buyers, and showed that prices *always* weakly decrease over time at the *focal* equilibrium of this game, which is the unique subgame-perfect equilibrium that survives the iterated elimination of weakly dominated strategies. Moreover, this result applies regardless of the valuation functions of the buyers; the result also extends to the corresponding equilibrium in *first-price* sequential auctions. It is worth highlighting here two important aspects of the model studied by Gale and Stegeman [37]. First, Gale and Stegeman consider multiunit-demand buyers, whereas prior theoretical work focuses on the simpler setting of unit-demand buyers. As well as being of more practical relevance (see the many examples above), multiunit-demand buyers can implement more sophisticated bidding strategies. Consequently, it is not unreasonable that equilibria in the multiunit-demand setting may possess more interesting properties than equilibria in the unit-demand setting. Second, they study an auction with *full information*. The restriction to full information is extremely useful here as it separates away informational aspects, and allows us to focus on the strategic properties caused purely by the sequential sales of items and not by a lack of information.

The result of Gale and Stegeman [37] immediately motivates the question of whether the declining price anomaly is guaranteed to hold in general, that is, in the equilibria of sequential auctions with more than two buyers. In this chapter, we answer this question in the negative by exhibiting a sequential auction with three buyers and eight items where prices initially rise and then fall. In order to run our computations that find this counter-example (to the conjecture that prices are weakly decreasing for multi-buyer sequential auctions), we study in detail the structure of equilibria in sequential auctions. First, it is important to note that there is a fundamental distinction between sequential auctions with two buyers and sequential auctions with three or more buyers. In the former sequential auction, each subgame reduces to a standard *auction with independent valuations*. We explain this in Section 2.1.1, where we present the two-buyer full-information model of Gale and Stegeman [37]. In contrast, in a multi-buyer sequential auction each subgame reduces to an *auction with externalities*. Consequently, in order to study multi-buyer sequential auctions we must study the equilibria of auctions with externalities. A theory of such equilibria was more recently developed by Paes Leme et al. [64] via a correspondence with an ascending price mechanism. In particular, as we discuss in Section 2.1.3, this ascending price mechanism outputs a unique bid value, called the *dropout bid* β_i, for each buyer i. For first-price auctions it is known [64] that these dropout bids form a subgame perfect equilibrium and, moreover, the interval $[0, \beta_i]$ is the exact set of bids that survives *all* processes consisting of the iterated elimination of strategies that are weakly dominated. In contrast, we show that for second-price auctions it may be the case that no bids survive the iterated elimination of weakly dominated strategies; however, we prove that the interval $[0, \beta_i]$ is the exact set of bids for any losing buyer that survives *all* processes consisting of the iterated elimination of strategies that are weakly dominated *by a lower bid*.

Consequently, both first- and second-price sequential auctions have a focal subgame-perfect equilibrium that survives a standard iterated elimination process. Our goal then is to answer the question of whether this equilibrium exhibits declining prices. By extensive computational search, we find a counterexample for the case of three bidders and eight items, where the price first increases and then decreases along the equilibrium path. That is, we show that the declining price anomaly does not always hold in the equilibria of these auctions. In Section 2.2 we describe the counter-example, which applies to both the first-price and second-price sequential auction settings. We emphasize that there is nothing unusual about our example: the form of the buyers' valuation functions is standard, namely, weakly decreasing marginal valuations. Furthermore, the non-monotonic price trajectory does not arise because of the use of an artificial tie-breaking rule in allocating each item; the three most natural tie-breaking rules, see Section 2.1.4, all induce the same non-monotonic price trajectory. We present an even stronger result in Section 2.3: for *any* tie-breaking rule, there is a sequential auction on which it induces a non-monotonic price trajectory.

This lack of weakly decreasing prices provides an explanation for why multi-buyer sequential auctions have been hard to analyze quantitatively. We provide a second explanation in Section 2.3.3. There we present a three-buyer sequential auction that does satisfy weakly decreasing prices but which has subgames where some agent has a negative value from winning against one of the two other agents. Again, this contrasts with the two-buyer case where every agent always has a non-negative value from winning against the other agent in every subgame.

Finally in Section 2.4, we describe the results obtained via our large scale experimentation. These results show that whilst the declining price anomaly is not universal, exceptions are extremely rare. From a randomly generated dataset of over six million sequential auctions and a variety of tie-breaking rules only a 0.000183 proportion of the instances produced non-monotonic price trajectories. Consequently, these computations are consistent with the practical examples discussed in the introduction. Of course, it is perhaps unreasonable to assume that subgame perfect equilibria arise in practice; we remark, though, that the use of simple bidding algorithms by bidders may also lead to weakly decreasing prices in a multi-buyer sequential auction. For example, Rodriguez [66] presents a method called the *residual monopsonist procedure* inducing this property in restricted settings.

2.1 The Sequential Auction Model

Here we present the full-information sequential auction model. There are T identical items and n buyers. Exactly one item is sold in each time period over T time periods. Buyer i has a value $V_i(k)$ for winning exactly k items. Thus $V_i(k) = \sum_{\ell=1}^{k} v_i(\ell)$, where $v_i(\ell)$ is the marginal value buyer i has for obtaining an ℓth item. This induces an extensive form game. Gale and Stegeman [37], for the two-buyer case, and Paes Leme et al. [64], for the multiple-buyer *first-price* case, show that this sequential game has a focal subgame perfect equilibrium. In this section, we will show that an exact analogue of their results *does not* hold for the *second-price* case with multiple buyers. However, the main purpose of this section is to show that, in effect, the equilibrium of Paes Leme et al. [64] is also a focal equilibrium in the second-price setting, and their results can be extended to this setting with a small technical modification.

To analyze this game it is informative to begin by considering the two-buyer case, as studied by Gale and Stegeman [37], which we do in the following subsection.

2.1.1 The Two-Buyer Case

Since this is a multi-unit auction with identical items, during the auction, the relevant history is the number of items each buyer has currently won. We may compactly represent the extensive form ("tree") of the auction using a directed graph with a node (x_1, x_2) for any pair of non-negative integers that satisfies $x_1 + x_2 \leq T$. The node (x_1, x_2) induces a subgame with $T - x_1 - x_2$ items for sale and where each buyer i already possesses x_i items. Note there is a *source node*, $(0, 0)$, corresponding to the whole game, and *sink nodes* (x_1, x_2), where $x_1 + x_2 = T$. The values Buyer 1 and Buyer 2 have for a sink node (x_1, x_2) are $\Pi_1(x_1, x_2) = V_1(x_1)$ and $\Pi_2(x_1, x_2) = V_2(x_2)$, respectively. We want to evaluate the values (utilities) at the source node $(0, 0)$. We can do this recursively working from the sinks upwards. Take a node (x_1, x_2), where $x_1 + x_2 = T - 1$. This node corresponds to the final round of the auction, where the last item is sold, given that each buyer i has already won x_i items. The node (x_1, x_2) will have directed arcs to the sink nodes $(x_1 + 1, x_2)$ and $(x_1, x_2 + 1)$; these arcs correspond to Buyer 1 and Buyer 2 winning the final item, respectively. For the case of second-price auctions, it is then a weakly dominant strategy for Buyer 1 to bid its marginal value $v_1(x_1 + 1) = V_1(x_1 + 1) - V_1(x_1)$; similarly for Buyer 2. Of course, this marginal value is just $v_1(x_1 + 1) = \Pi_1(x_1 + 1, x_2) - \Pi_1(x_1, x_2 + 1)$, the difference in value between winning and losing the final item. If Buyer 1 is the highest bidder at (x_1, x_2), that is, $\Pi_1(x_1 + 1, x_2) - \Pi_1(x_1, x_2 + 1) \geq \Pi_2(x_1, x_2 + 1) - \Pi_2(x_1 + 1, x_2)$, then we have that

$$\begin{aligned}
\Pi_1(x_1, x_2) &= \Pi_1(x_1 + 1, x_2) - \big(\Pi_2(x_1, x_2 + 1) - \Pi_2(x_1 + 1, x_2)\big) \\
\Pi_2(x_1, x_2) &= \Pi_2(x_1 + 1, x_2)
\end{aligned}$$

That is, Buyer 1's value for the node (x_1, x_2), which is $\Pi_1(x_1, x_2)$, is its value for the node $(x_1 + 1, x_2)$ minus Buyer 2's bid for the next item. Symmetric formulas apply if Buyer 2 is the highest bidder at (x_1, x_2). Hence we may recursively define a value for each buyer for each node. The iterative elimination of weakly dominated strategies then leads to a subgame perfect equilibrium [37, 10].

`Example`: Consider a two-buyer sequential auction with two items, where the marginal valuations are $\{v_1(1), v_1(2)\} = \{10, 8\}$ and $\{v_2(1), v_2(2)\} = \{6, 3\}$. This game is illustrated in Figure 2.1. The base case with the values of the sink nodes is shown in Figure 2.1(a). The first row in each node refers to Buyer 1 and shows the number of items won (in plain text) and the corresponding value (in bold); the second row refers to Buyer 2. The outcome of the second-price sequential auction, solved recursively, is then shown in Figure 2.1(b). Arcs are labelled by the bid value; here arcs for Buyer 1 point left and arcs for Buyer 2 point right. Solid arcs represent winning bids and dotted arcs are losing bids. The equilibrium path is shown in bold.

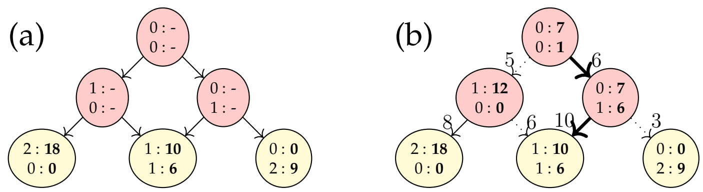

Figure 2.1: Second-Price Sequential Auction

Observe that the declining price anomaly is exhibited in this example: on the equilibrium path, Buyer 2 wins the first item for a price 5 and Buyer 1 wins the second item for a price 3. As stated, this example is not an exception. Gale and Stegeman [37] showed that weakly decreasing prices are a property of 2-buyer sequential auctions.

Theorem 2.1.1. *[37] In a 2-buyer second-price sequential auction there is a unique equilibrium that survives the iterated elimination of weakly dominated strategies. Moreover, at this equilibrium prices are weakly declining.*

We remark that the subgame perfect equilibrium that survives the iterated elimination of weakly dominated strategies is unique in terms of the values at the nodes. Moreover, given a fixed tie-breaking rule, the subgame perfect equilibrium also has a unique equilibrium path in each subgame.

In addition, Theorem 2.1.1 also applies to first-price sequential auctions. In this case, to ensure the existence of an equilibrium, we make the standard assumption that there is a fixed small bidding increment. That is, for any price p there is a unique maximum price smaller than p. Given this, for the example above, the subgame perfect equilibrium using a first-price sequential auction is as shown in Figure 2.2. Here we use the notation p^+ to denote a winning bid of value equal to p, and the notation p to denote a losing bid equal to maximum value smaller than p.

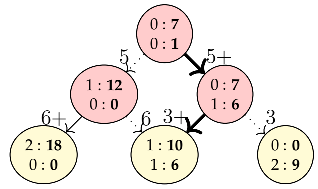

Figure 2.2: First-Price Sequential Auction

Observe that the resultant prices on the equilibrium path are more easily apparent in Figure 2.2 than in Figure 2.1. For this reason, the figures we present in this thesis will be for first-price auctions; equivalent figures can can be drawn for second-price auctions.

Following Gale and Stegeman [37], and prior to our work, the question of whether or not the declining price anomaly holds in the equilibria of sequential auctions with more than two buyers remained open. To resolve this question, we must first study equilibria in the full-information sequential auction model when there are more than two buyers.

2.1.2 The Multiple-Buyer Case

The underlying model of [37] extends in a straightforward manner to sequential auctions with $n \geq 3$ buyers. There is a node $(x_1, x_2, \ldots, x_n)$ for each set of non-negative integers satisfying $\sum_{i=1}^{n} x_i \leq T$. There is a directed arc from $(x_1, x_2, \ldots, x_n)$ to $(x_1, x_2, \ldots, x_{j-1}, x_j + 1, x_{j+1}, \ldots x_n)$ for each $1 \leq j \leq n$. Thus each non-sink node has n out-going arcs. This is problematic: whilst in the final time period each buyer has a value for winning and a value for losing, this is no longer the case recursively in earlier time periods. Specifically, buyer i has value for winning, but $n-1$ *(different) values* for losing depending upon the identity of the buyer $j \neq i$ who actually wins. Thus rather than each node corresponding to a standard auction, it now corresponds to an auction with *externalities*.

Formally, an auction with externalities is a single-item auction where each buyer i has a value $v_{i,i}$ for winning the item and, for each buyer $j \neq i$, buyer i has value $v_{i,j}$ if buyer j

wins the item. These auctions were first studied by Funk [36], Jehiel and Moldovanu [47] and Jehiel et al. [48]. Their motivations were applications where losing participants were not indifferent to the identity of the winning buyer; examples include firms seeking to purchase a patented innovation, take-over acquisitions of a smaller company in an oligopolistic market, and sports teams competing to sign a star athlete.

Thus in order to understand multi-buyer sequential auctions we must first understand equilibria in auctions with externalities. This is not a simple task: such an understanding was only recently provided by Paes Leme et al. [64].

2.1.3 Equilibria in Auctions with Externalities

An Ascending Price Mechanism

We can explain the result of [64] via an ascending price auction. Consider a two-buyer ascending price auction where the valuations of the buyers are v_1 and v_2, with $v_1 > v_2$. The requested price p starts at zero and continues to rise until the point where the second buyer drops out. Of course, this happens when the price reaches v_2, and so Buyer 1 wins for a payment $p^+ = v_2$. But this is exactly the outcome expected from a first-price auction: Buyer 2 loses with bid of p and Buyer 1 wins with a bid of p^+. To generalize this to multi-buyer settings we can view this process as follows. At a price p, buyer i remains in the auction as long as there is *at least one buyer j still in the auction* who buyer i is willing to pay a price p to beat; that is, $v_{i,i} - p > v_{i,j}$. The last buyer to drop out wins at the corresponding price. For example, in the two-buyer example above, Buyer 2 drops out at price $p = v_2$ as it would rather lose to Buyer 1 than win above that price. Therefore, at price p^+ there is no buyer still in the auction that Buyer 1 wishes to beat (because there are no other buyers remaining in the auction at all!). Thus Buyer 1 drops out at p^+ and, being the last buyer to drop out, wins at that price.

Observe that, even in the multi-buyer setting, this procedure produces a unique *dropout bid* β_i for each buyer i. To illustrate this, two auctions with externalities are shown in Figure 2.3. In these diagrams the label of an arc from buyer i to buyer j is $w_{i,j} = v_{i,i} - v_{i,j}$. That is, buyer i is willing to pay up to $w_{i,j}$ to win *if the alternative is that buyer j wins the item*. Now consider running our ascending price procedure for these auctions. In Figure 2.3(a), Buyer 1 drops out when the price reaches 18. Since Buyer 1 is no longer active in the auction, Buyer 4 drops out when the price reaches 23. At this point, Buyer 2 and Buyer 3 are left to compete for the item. Buyer 3 wins when Buyer 2 drops out at price 31. Thus the drop-out bid of Buyer 3 is 31^+. Observe that Buyer 2 loses despite having very high values for winning (against Buyer 1 and Buyer 4).

The example of Figure 2.3(b) with dropout bid vector $(\beta_1, \beta_2, \beta_3, \beta_4) = (24, 24, 24, 24^+)$ is more subtle. Here Buyer 2 drops out at price 24. But Buyer 3 only wanted to beat Buyer 2 at this price so it then immediately drops out at the same price. Now Buyer 1 only wanted to beat Buyer 2 and Buyer 3 at this price, so it then immediately drops out at the same price. This leaves Buyer 4 the winner at price 24^+.

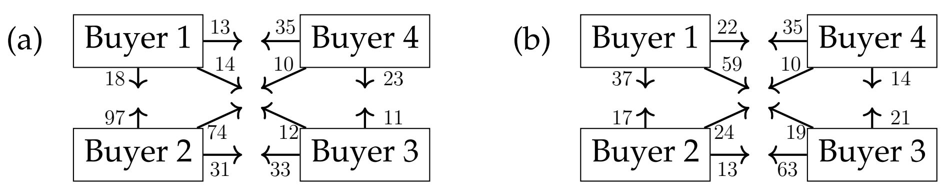

Figure 2.3: Drop-Out Bid Examples: *In these two examples the dropout bid vectors* $(\beta_1, \beta_2, \beta_3, \beta_4)$ *are* $(18, 31, 31^+, 23)$ *and* $(24, 24, 24, 24^+)$*, respectively*

Dropout Bids and Iterated Elimination of Weakly Dominated Strategies

As well as being solutions to the ascending price auction, the dropout bids have a much stronger property that makes them the natural and robust prediction for auctions with externalities. Specifically, Paes Leme et al. [64] proved that, for each buyer i, the interval $[0, \beta_i]$ is the set of strategies that survive *any sequence* consisting of the iterated elimination of weakly dominated strategies. This is formalized as follows. Take an n-buyer game with strategy sets $S_1, S_2, \ldots, S_n$ and utility functions $u_i : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$. Then $\{S_i^\tau\}_{i,\tau}$ is a *valid sequence* for the iterated elimination of weakly dominated strategies if for each τ there is a buyer i such that $S_j^\tau = S_j^{\tau-1}$ for each buyer $j \neq i$, and $S_i^\tau \subset S_i^{\tau-1}$ where for each strategy $s_i \in S_i^{\tau-1} \setminus S_i^\tau$ there is an $\hat{s}_i \in S_i^\tau$ such that $u_i(\hat{s}_i, s_{-i}) \geq u_i(s_i, s_{-i})$ for all $s_{-i} \in \prod_{j:j\neq i} S_j^\tau$, and with strict inequality for at least one s_{-i}. We say that a strategy s_i for buyer i *survives* the iterated elimination of weakly dominated strategies if for any valid sequence $\{S_i^\tau\}_{i,\tau}$ we have $s_i \in \bigcap_\tau S_i^\tau$.

Theorem 2.1.2. *[64] Given a first-price auction with externalities, for each buyer i, the set of bids that survive the iterated elimination of weakly dominated strategies is exactly $[0, \beta_i]$.*

An exact analogue of Theorem 2.1.2 does *not* hold for second-price auctions with externalities. Instead, it may be the case that no strategies survive the iterated elimination of weakly dominated strategies (we omit our three-buyer counterexample here).

Theorem 2.1.3. *There are second-price auctions with externalities where no strategies survive the iterated elimination of weakly dominated strategies.*

By considering examples that demonstrate this theorem, we can observe that in every example, the problem occurs when a strategy is deleted because it is weakly dominated by a *higher* bid. But this can never happen for a potentially winning bid in a first-price auction; thus Theorem 2.1.2 holds in first-price auctions when we restrict attention to sequences with the iterated elimination of strategies that are weakly dominated by a *lower* bid. Indeed, we prove that the corresponding theorem holds for second-price auctions.

Theorem 2.1.4. *Given a second-price auction with externalities, for each losing buyer i, the set of bids that survive the iterated elimination of strategies that are weakly dominated by a lower bid is exactly $[0, \beta_i]$.*

Proof. First we claim that for any losing buyer i and any price $p > \beta_i$ there is a sequence of iterative deletions of strategies that are weakly dominated by a lower bid that leads to the deletion of bid p from S_i^τ. Without loss of generality, we may order the buyers such that $\beta_1 \leq \beta_2 \leq \cdots \leq \beta_n$; in the case of a tie the buyers are placed in the order they were deleted by the tie-breaking rule. Initially $S_i^0 = [0, \infty)$, for each buyer i. We now define a valid sequence such that $S_i^i = [0, \beta_i]$. We proceed by induction on the label of the buyers. For the base case observe that for Buyer 1 we know $\beta_1 = \max_{j:j\neq i}(v_{i,i} - v_{i,j})$ is the highest price it wants to pay to beat anyone else. Suppose Buyer 1 bids $p > \beta_1$. Take any set of bids $b_{-1} \in \times_{j:j\geq 2} S_j^0$. We have three cases:

(i) Both bids p and β_1 are winning bids against b_{-1}. Then, as this is a second-price auction, Buyer 1 is indifferent between the two bids.

(ii) Both bids p and β_1 are losing bids against b_{-1}. Then Buyer 1 is indifferent between the two bids.

(iii) Bid p is a winning bid but β_1 is a losing bid against b_{-1}. Then since the winning price is at least β_1, Buyer 1 strictly prefers to lose rather than win. Moreover, since $S_j^0 = [0, \infty)$, there is a set of bids b_{-1} by the other buyers such that Buyer 1 strictly prefers to lose rather than win.

Thus the bid p is weakly dominated by the lower bid β_1. Since this applies to any $p > \beta_1$, in Step 1 we may delete every bid for Buyer 1 above β_1. Therefore $S_1^1 = [0, \beta_1]$ and $S_j^1 = [0, \infty]$ for each buyer $j \geq 2$.

For the induction hypothesis assume $S_j^{i-1} = [0, \beta_j]$, for all $j < i$ and $S_j^{i-1} = [0, \infty)$, for all $j \geq i$. Now take a losing buyer i and any set of bids $b_{-i} \in \times_{j:j\neq i} S_j^{i-1}$. Again, we have three cases:

(i) Both bids p and β_i are winning bids against b_{-i}. Then, as this is a second-price auction, buyer i is indifferent between the two bids.

(ii) Both bids p and β_1 are losing bids against b_{-i}. Then buyer i is indifferent between the two bids.

(iii) Bid p is a winning bid but β_i is a losing bid against b_{-i}. Then since β_i is a losing bid under the tie-breaking rule, it must be the case that the winning bid is from a buyer j where $j > i$. But, by definition of β_i, there is no buyer j, with $j > i$, that buyer i wishes to beat at price β_i.

So buyer i prefers the bid β_i to the bid p. Moreover, since any buyer $j : j > i$ has $S_j^{i-1} = [0, \infty)$, this preference is strict for some feasible choice of bids for the other buyers.

Thus, for buyer i, the bid p is weakly dominated by the lower bid β_i, and this applies to every $p > \beta_i$. Thus in Step i we may delete every bid for buyer i above β_i. Therefore $S^i_j = [0, \beta_i]$, for all $j < i+1$ and $S^{i-1}_j = [0, \infty)$, for all $j \geq i+1$. The claim then follows by induction. So, for any losing buyer i we have that no bid greater than β_i survives the iterated elimination of strategies that are weakly dominated by a lower bid.

Observe that the above arguments also apply for the winning buyer, that is, buyer n. Here, as there are no higher indexed buyers, it is not the case that β_n strictly dominates any bid $p > \beta_n$. Indeed, buyer n is indifferent between all bids in the range $[\beta_n, \gamma_n]$, where γ_n is the maximum value the buyer has for beating any buyer j with dropout bid $\beta_j = \beta_n$. Observe that γ_n does exist and is at least β_n by definition of the ascending price mechanism. Thus for the winning bidder no bid greater than γ_i survives the iterated elimination of strategies that are weakly dominated by a lower bid.

Second, we claim for any buyer i and any price $q < \beta_i$ there is no sequence of iterative deletions of strategies that are weakly dominated by a lower bid that leads to the deletion of bid q from the feasible strategy space of buyer i. If not, consider the first time τ that some buyer i has a value $q \in [0, \beta_i]$ deleted from S^τ_i. We may assume that q is deleted because it is was weakly dominated by a lower bid $p < q$. Now, by assumption, $[0, \beta_j] \subseteq S^{\tau-1}_j$, for each buyer j. Furthermore, by definition, there is some buyer k, with $k > i$ that buyer i wishes to beat at any price below β_i. In particular, Buyer i wishes to beat Buyer k at price p. But since $k > i$ we have $\beta_k \geq \beta_i$. Recall that $[0, \beta_k] \subseteq S^{\tau-1}_k$. It immediately follows that there is a set of feasible bids $b_k \in (p, q)$ and $b_j = 0$, for all $j \notin \{i, k\}$ such that Buyer i strictly prefers to win against these bids. Specifically, the bid q is not weakly dominated by the bid p, a contradiction. □

It follows that the dropout bids form the *focal* subgame perfect equilibrium for both first-price and second-price auctions with externalities.

We are now almost ready to be able to find equilibria in the sequential auction computations we will conduct. This, in turn, will allow us to present a sequential auction with non-monotonic prices. Before doing so, one final factor remains to be discussed regarding the transition from equilibria in auctions with externalities to equilibria in sequential auctions.

2.1.4 Tie-Breaking Rules

As stated, the dropout bid of each buyer is uniquely defined. However, our description of the ascending auction may leave some flexibility in the choice of winner. Specifically, it may be the case that, simultaneously, two or more buyers wish to drop out of the auction. If this happens at the end of the ascending price procedure, any of these buyers could be selected as the winner. An example of this is shown in Figure 2.4.

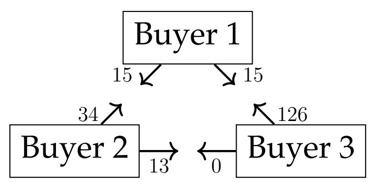

Figure 2.4: Tie-Breaking: *An example requiring tie-breaking to decide the winner; the drop-out bid vector is* $(15, 15, 15)$ *but there are two possible winners:* $(15, 15^{+}, 15)$ *or* $(15, 15, 15^{+})$

This observation implies that to fully define the ascending auction procedure we must incorporate a tie-breaking rule to order the buyers when more than one wish to drop out simultaneously. In a one-shot auction with externalities the tie-breaking rule only affects the choice of winner, but otherwise has no structural significance. However, in a sequential auction the choice of tie-breaking rule may have much more significant consequences. Specifically, because each node in the game tree corresponds to an auction with externalities, the choice of winner at one node may affect the valuations at nodes higher in the tree. In particular, the equilibrium path may vary with different tie-breaking rules, leading to different prices, winners, and utilities.

We show later in this section that there are a massive number of tie-breaking rules, even in small sequential auctions. We emphasize, however, that our main result (Section 2.3) holds regardless of the tie-breaking rule. That is, for *any* tie-breaking rule there is a sequential auction on which it induces a non-monotonic price trajectory. Before all of this, we first show that non-monotonic pricing occurs on the equilibrium path for perhaps the most natural choice of tie-breaking rule, namely `preferential-ordering`, where each buyer is given a distinct rank (index), and in the case of a tie the buyer with the highest index is eliminated. In our original paper, we analyze three of the most commonly used tie-breaking rules, namely the `preferential-ordering`, `first-in-first-out` and `last-in-first-out` rules, that correspond to the fundamental data structures of priority queues, queues, and stacks in computer science.

In order to understand the `preferential-ordering` rule it is useful to see how it applies on an example. Figure 2.5 presents a five-buyer example with the dropout vector $(\beta_1, \beta_2, \beta_3, \beta_4, \beta_5) = (40, 40, 40, 40, 40)$. On running the ascending price procedure, both Buyer 3 and Buyer 4 wish to drop out when the price reaches 40. Using our selected tie-breaking rule, the set of agents eligible for dropping out is $\{3, 4\}$ and we remove the highest index buyer, namely Buyer 4. With the removal of Buyer 4, neither Buyer 1 nor Buyer 5 have an incentive to continue bidding so they both decide to dropout. Thus the set of agents eligible for dropping out is now $\{1, 3, 5\}$ and `preferential-ordering` removes Buyer 5. Observe that with the removal of Buyer 5, that Buyer 2 no longer has an active participant it wishes to beat so the set of agents eligible for dropping out is updated to $\{1, 2, 3\}$. The `preferential-ordering` rule now first removes Buyer 3, then Buyer 2 and lastly Buyer 1. Thus Buyer 1 wins the item under the `preferential-ordering` rule. Observe that the `preferential-ordering` rule, `first-in-first-out` rule and `last-in-first-out` rule produce three distinct winners in this example!

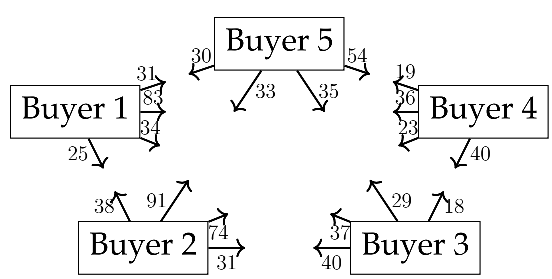

Figure 2.5: An Example to Illustrate the `preferential-ordering` Rule

We have now developed the tools required to implement our sequential auction computations. We describe these computations and their results in Section 2.4. Before doing so, we present in Section 2.2 one sequential auction obtained via these computations and verify that it leads to a non-monotonic price trajectory with the `preferential-ordering` rule (and, in fact, with the `first-in-first-out` and `last-in-first-out` rules, but we omit their analysis). We then explain in Section 2.3 how to generalize this conclusion to apply to every tie-breaking rule.

2.2 An Auction with Non-Monotonic Prices

Here we prove that the decreasing price anomaly is **not** guaranteed for sequential auctions with more than two buyers. Specifically, we exhibit a sequential auction that has three buyers and eight items, for which the equilibrium obtained by breaking ties with the `preferential-ordering` rule exhibits non-monotonic prices. We present the first-price version where at equilibrium the buyers bid their dropout values in each time period; as discussed previously, the same example extends to second-price auctions.

The valuations of the three buyers are defined as follows. Buyer 1 has marginal valuations $\{55, 55, 55, 55, 30, 20, 0, 0\}$, Buyer 2 has marginal valuations $\{32, 20, 0, 0, 0, 0, 0, 0\}$, and Buyer 3 has marginal valuations $\{44, 44, 44, 44, 0, 0, 0, 0\}$.

Let's now compute the extensive forms of the auction under the three tie-breaking rules. We begin with the `preferential-ordering` rule. To compute its extensive form, observe that Buyer 1 is guaranteed to win at least two items in the auction because Buyer 2 and Buyer 3 together have positive value for six items. Therefore, the feasible set of sink nodes in the extensive form representation are shown in Figure 2.6.

Given the valuations at the sink nodes we can work our way upwards recursively calculating the values at the other nodes in the extensive form representation. For example, consider the node $(x_1, x_2, x_3) = (4, 1, 2)$. This node has three children, namely $(5, 1, 2), (4, 2, 2)$ and $(4, 1, 3)$; see Figure 2.7(a). These induce a three-buyer auction as

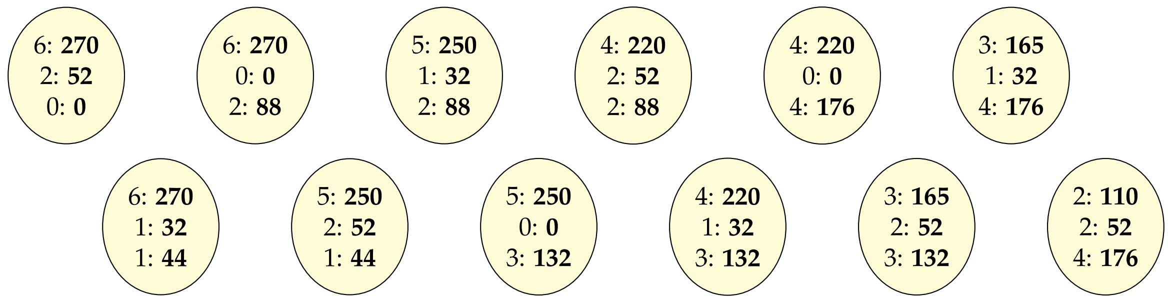

Figure 2.6: Sink Nodes of the Extensive Form Game

shown in Figure 2.7(b). This can be solved using the ascending price procedure to find the dropout bids for each buyer. Thus we obtain that the value for the node $(x_1, x_2, x_3) = (4, 1, 2)$ is as shown in Figure 2.7(c). Of course this node is particularly simple as, for the final round of the sequential auction, the corresponding auction with externalities is just a standard auction. That is, when the final item is sold, for any buyer i the value $v_{i,j}$ is independent of the buyer $j \neq i$.

Nodes higher up the game tree correspond to more complex auctions with externalities. For example, the case of the source node $(x_1, x_2, x_3) = (0, 0, 0)$ is shown in Figure 2.8. In this case, on applying the ascending price procedure, Buyer 1 is the first to dropout at price 15. At this point, both Buyer 2 and Buyer 3 no longer have a competitor that they wish to beat at this price, so they both want to dropout. With the `preferential-ordering` tie-breaking rule, Buyer 2 wins the item.

Using similar arguments at each node verifies the concise extensive form representation under the `preferential-ordering` tie-breaking rule shown in Figure 2.9. In this figure, the white nodes represent subgames where the sequential auction still has three active buyers; the pink nodes represent subgames with at most two active buyers; the yellow nodes are the sink nodes. Again, the equilibrium path with non-monotonic prices is shown in bold. Now consider this equilibrium path. Observe that Buyer 2 wins the first two items, Buyer 3 wins the next four items and Buyer 1 wins the final two items. The resultant price trajectory is $\{15, 17, 0, 0, 0, 0, 0, 0\}$. That is, the price rises and then falls to zero – a non-monotonic price trajectory.

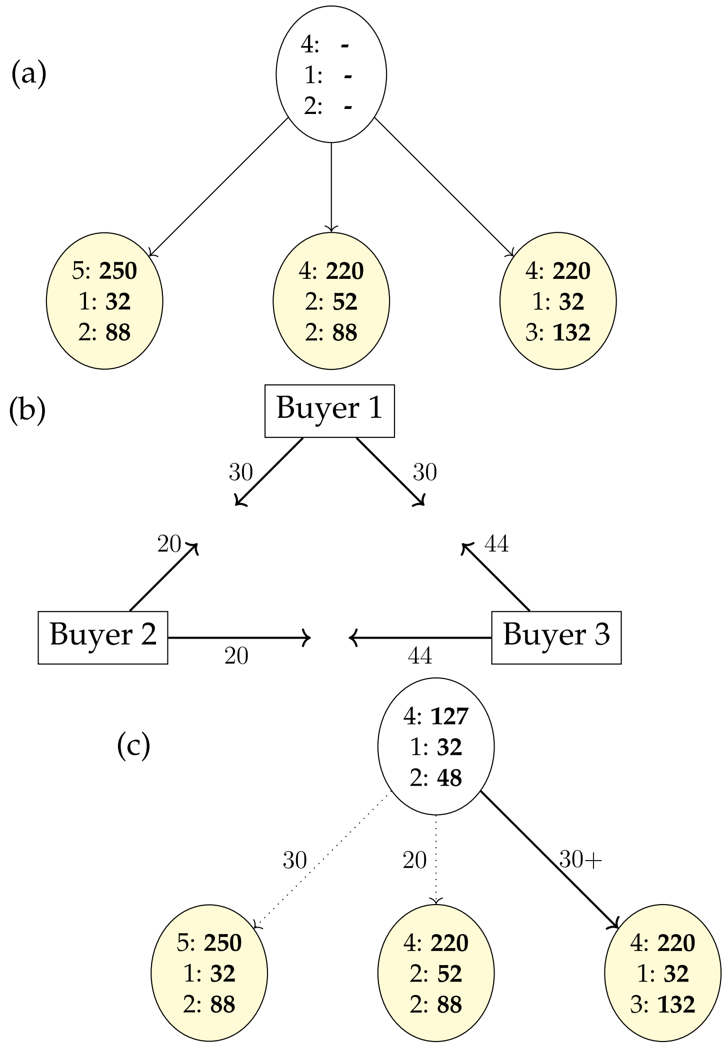

Figure 2.7: Solving a Subgame Above the Sinks

Exactly the same example works with the other two tie-breaking rules (figures omitted). The node values under `preferential-ordering` and `first-in-first-out` are exactly the same, despite the fact that these two rules do produce different winners at some nodes. By contrast, the `last-in-first-out` rule gives an extensive form where some nodes have different valuations than those produced by the other two rules. However, for all three rules the equilibrium path and price trajectory for the whole game, starting at the root, is the same. These observations will play a role when we prove that, for any tie-breaking rule, there is a sequential auction with non-monotonic prices.

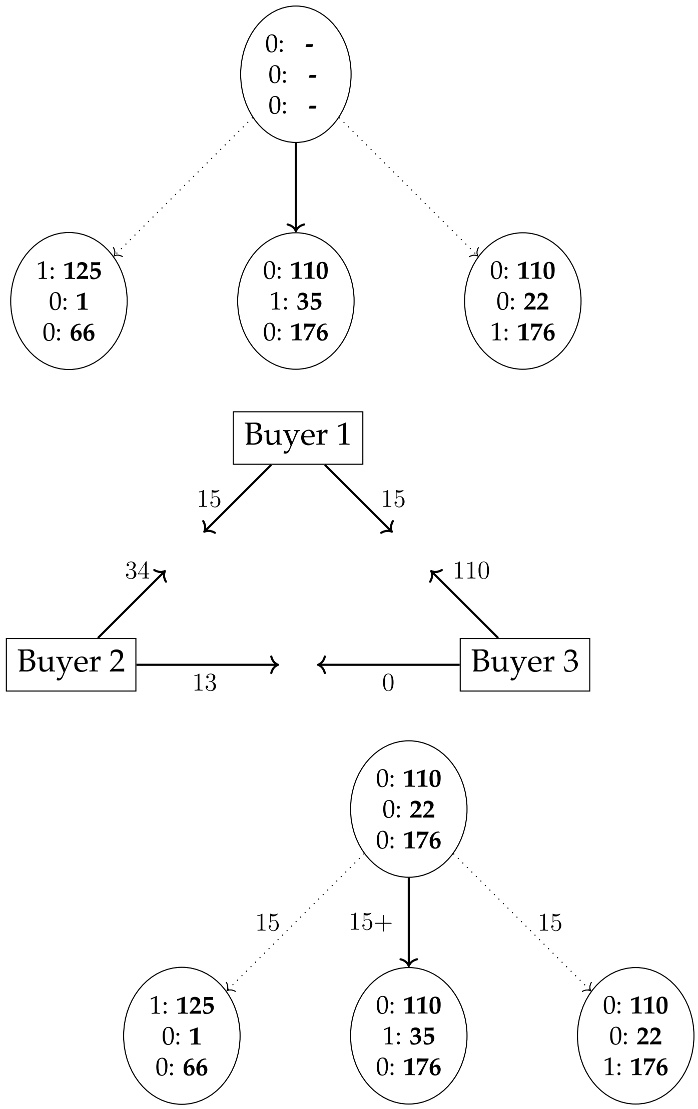

Figure 2.8: Solving the Subgame at the Root

Again, we emphasize that there is nothing inherently perverse about this example. The form of the valuation functions, namely decreasing marginal valuations, is standard. As explained, the equilibrium concept studied is the appropriate one for sequential auctions. Finally, the non-monotonic price trajectory is not the artifact of an aberrant tie-breaking rule. This result is rather surprising, since the result of Gale and Stegeman [37] implies that with two buyers, prices always decrease at equilibrium even when the valuation functions do not have decreasing marginals.

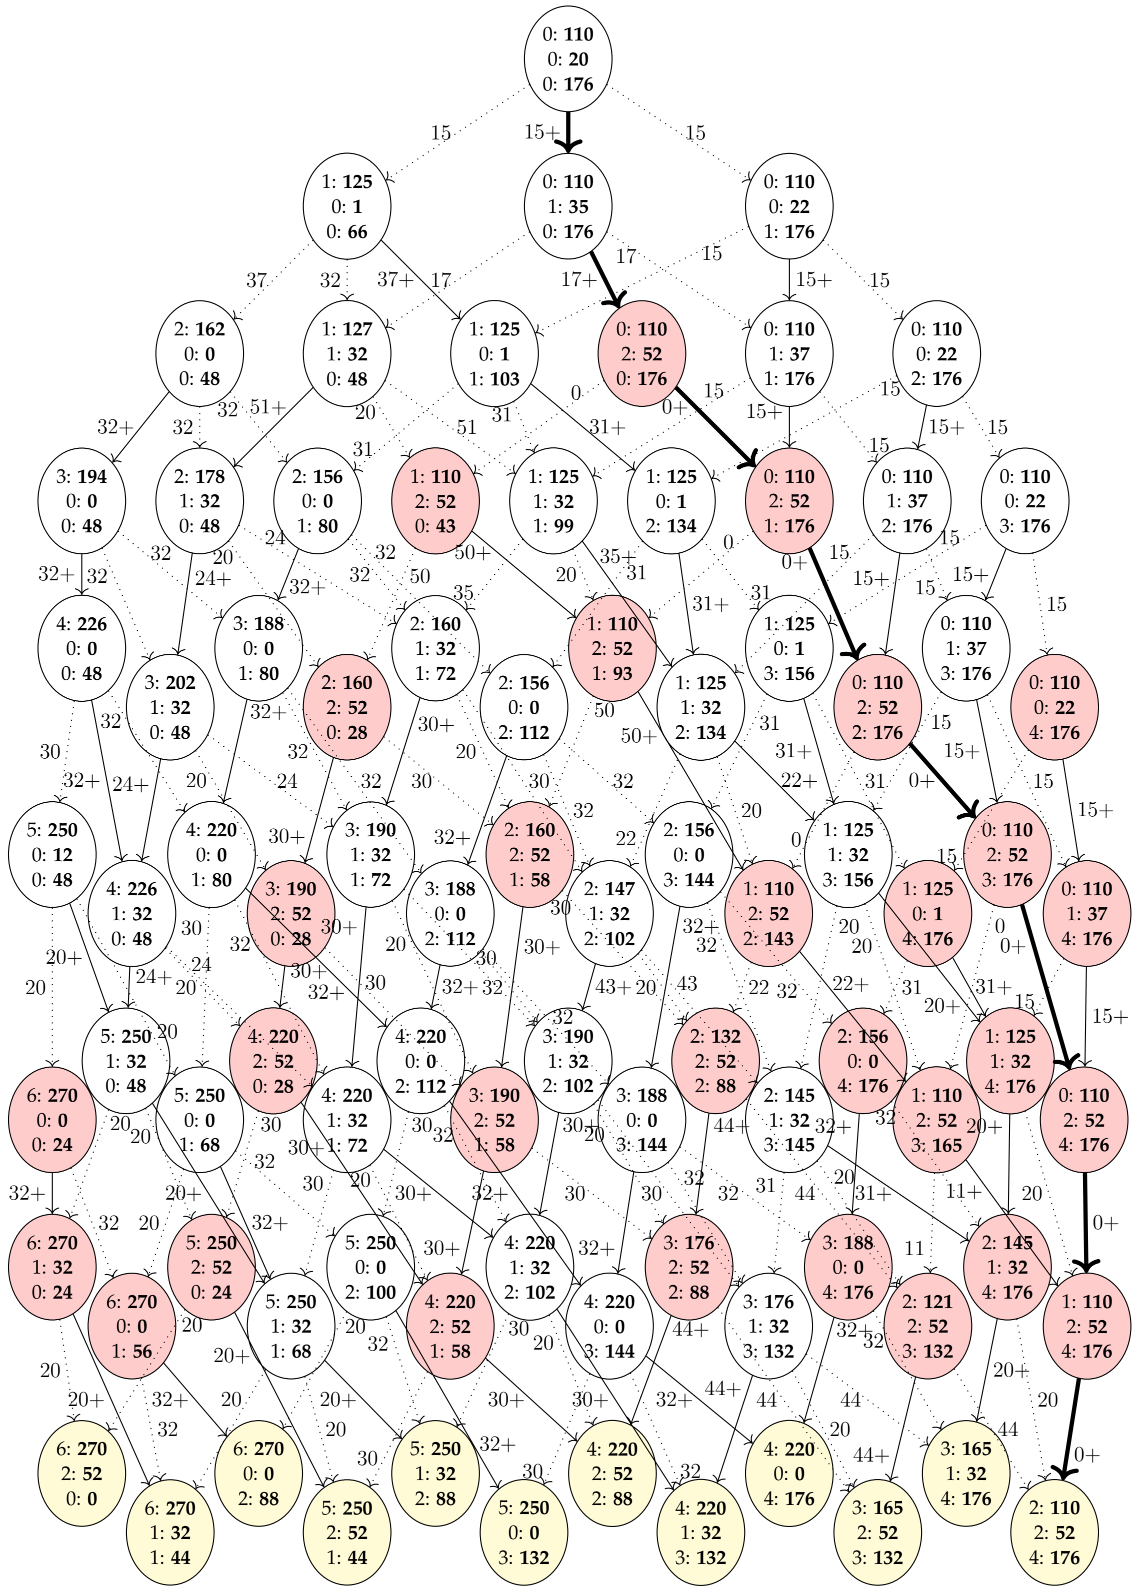

Figure 2.9: Non-Monotonic Prices with the `preferential-ordering` Rule

2.3 Non-Monotonic Prices under General Tie-Breaking Rules

We will now prove that non-monotonic prices are exhibited under any tie-breaking rule. In order to do this, we must first analyze the set of all tie-breaking rules.

2.3.1 Classifying the Set of Tie-Breaking Rules

Our definition of the set of tie-breaking rules will utilize the concept of an *overbidding graph*, introduced by Paes Leme et al. [64]. For any price p and any set of bidders S, the overbidding graph $G(S, p)$ contains a labelled vertex for each buyer in S and an arc (i, j) if and only if $v_{i,i} - p > v_{i,j}$. For example, recall the auction with externalities seen in Figure 2.5. This is reproduced in Figure 2.10 along with its overbidding graph $G(\{1, 2, 3, 4, 5\}, 40)$.

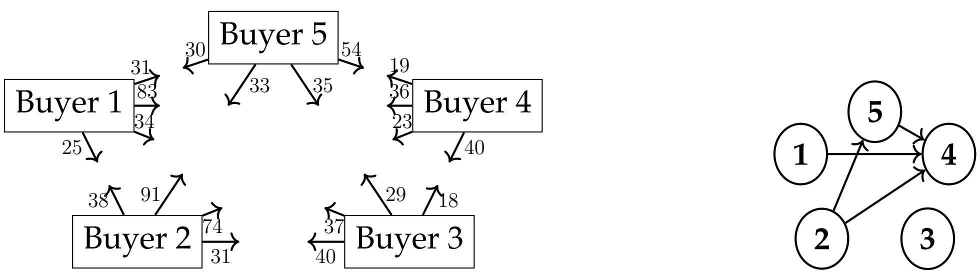

Figure 2.10: The Overbidding Graph $G(\{1, 2, 3, 4, 5\}, 40)$

First, recall that the drop-out bid β_i is unique for any buyer i, regardless of the tie-breaking rule. Consequently, whilst the tie-breaking rule will also be used to order buyers that are eliminated at prices below the final price p^*, such choices are irrelevant with regards to the final winner. Thus, the only relevant factor is how a decision rule selects a winner from amongst those buyers S^* whose drop-out bids are p^*. Second, recall that at the final price p^* the remaining buyers are eliminated one-by-one until there is a single winner. However, a buyer *cannot* be eliminated if there remains another buyer still in the auction that it wishes to beat at price p^*. That is, buyer i must be eliminated after buyer j if there is an arc (i, j) in the overbidding graph. Thus, the order of eliminations given by the tie-breaking rule must be consistent with the overbidding graph. In particular, the winner can only be selected from amongst the *source vertices*[1] in the overbidding graph $G(S^*, p^*)$. For example, in Figure 2.10 the source vertices are $\{1, 2, 3\}$. Note that this explains why the tie-breaking rules `preferential-ordering`, `first-in-first-out` and `last-in-first-out` chose Buyer 1, Buyer 2 and Buyer 3 as winners but none of them selected Buyer 4 or Buyer 5. Observe that the overbidding graph $G(S^*, p^*)$ is *acyclic*; if it contained a directed cycle then the price in the ascending auction would be forced to rise further. Because every directed acyclic graph contains at least one source vertex, any tie-breaking rule does have at least one choice for winner.

Thus a tie-breaking rule is simply a function $\tau : H \to \sigma(H)$. Here the domain of the function is the set of labelled, directed acyclic graphs and $\sigma(H)$ is the set of source nodes in H. Consequently, two tie-breaking rules are equivalent if they correspond to the same function τ. We are now ready to present our main result.

2.3.2 Non-Monotonic Prices for Any Tie-Breaking Rule

Theorem 2.3.1. *For any tie-breaking rule, there is a sequential auction with non-monotonic prices.*

Proof. We consider exactly the same example as before. That is, we have three buyers and eight items, and Buyer 1 has marginal values $\{55, 55, 55, 55, 30, 20, 0, 0\}$, Buyer 2 has marginals $\{32, 20, 0, 0, 0, 0, 0, 0\}$, and Buyer 3 has marginals $\{44, 44, 44, 44, 0, 0, 0, 0\}$.

First let's calculate how many tie-breaking rules there are for this auction. To count this we must consider all directed acyclic graphs with labels in $\{1, 2, 3\}$. Note that we must have at least two buyers with drop-out values equal to the final price p^* otherwise the auction would have terminated earlier. Thus it suffices to consider directed acyclic graphs with either two or three vertices. There are 8 such topologies that produce 34 labelled directed acyclic graphs and $12,288$ tie-breaking rules. This is illustrated in Table 2.1.

Luckily we do not need to examine all of these tie-breaking rules separately. The set of tie-breaking rules can be partitioned into exactly *ten* classes. Specifically, any tie-breaking

[1] A *source* is a vertex v with in-degree zero; that is, there no arcs pointing into v.

Directed Acylic Graph	# Labelled Graphs	# Sources
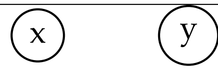	3	2
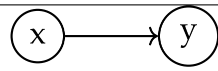	6	1
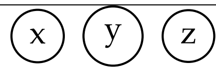	1	3
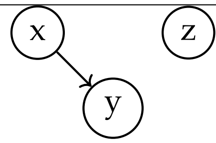	6	2
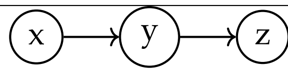	6	1
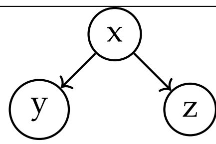	3	1
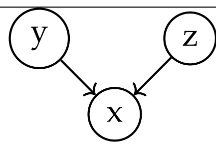	3	2
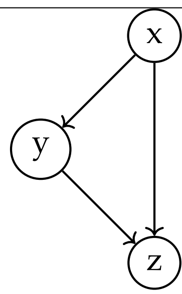	6	1
Total # Labelled DAGs	34	
Total # Tie-Breaking Rules	$1^{21} \cdot 2^{12} \cdot 3^{1} = 12,288$	

Table 2.1: Labelled Directed Acyclic Graphs

rule produces one of just ten possible (in terms of distinct node valuations) extensive forms for this sequential auction. Two of these we obtained via the three tie-breaking rules from before. We explain why there are only eight other feasible extensive forms. For any tie-breaking rule, as we work up from the sink nodes there are many nodes where tie-breaking is required. Despite this, the total number of distinct extensive forms does not blow-up multiplicatively. As previously alluded to, when we apply a tie-breaking rule there are two possibilities that arise. In the first possibility, the node valuations are the same regardless of which buyer is selected by the rule. For example, consider the node $(3, 0, 2)$ where Buyer 1 wins with `preferential-ordering` but Buyer 3 wins with `first-in-first-out`; in either case the node valuations are identical, namely $(188, 0, 112)$ as shown in Figures 2.9. Consequently, both rules lead to the same outcome. For our purpose, such nodes are of no importance.

In the second possibility, the node valuations do vary depending on which buyer is selected by the tie-breaking rule. However, of the 34 labelled directed acyclic overbidding graphs, only 4 affect the extensive form node valuations. These four critical overbidding graphs, which we call A, B, C and D, are shown in Figure 2.11.

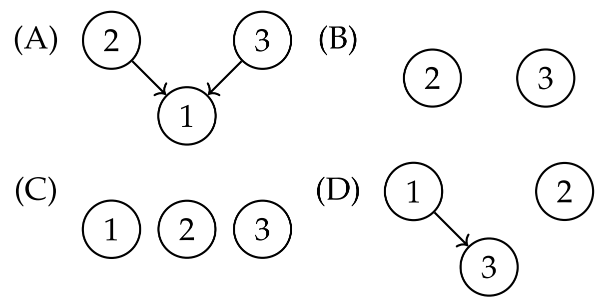

Figure 2.11: The Four Critical Overbidding Graphs

It is easy to verify that, working upwards from the sink nodes, the first such nodes where the choice of tie-breaking rule matters occur at depth 4, at the three nodes $(4, 0, 0)$, $(1, 0, 3)$ and $(0, 1, 3)$. The nodes $(1, 0, 3)$ and $(0, 1, 3)$ both correspond to the overbidding graph A whilst the node $(4, 0, 0)$ corresponds to the overbidding graph B. For the overbidding graph A the tie-breaking rule must select either the sink vertex 2 or the sink vertex 3 to win. Moreover, by definition, it must make the same choice at both $(1, 0, 3)$ and $(0, 1, 3)$. Furthermore, regardless of this choice, as we work up the extensive form the nodes $(1, 0, 2)$, $(0, 1, 2)$, $(0, 0, 3)$, $(0, 1, 1)$, $(0, 0, 2)$, $(0, 1, 0)$, $(0, 0, 1)$ and $(0, 0, 0)$ also all have the overbidding graph A and, thus, must also have the same winner.

The choice of winner at $(4, 0, 0)$ for overbidding graph B is also between Buyer 2 and Buyer 3, but in this case, the effect is more subtle. If Buyer 2 wins then the overbidding graph D is induced at node $(3, 0, 0)$, whereas if Buyer 3 wins then the overbidding graph

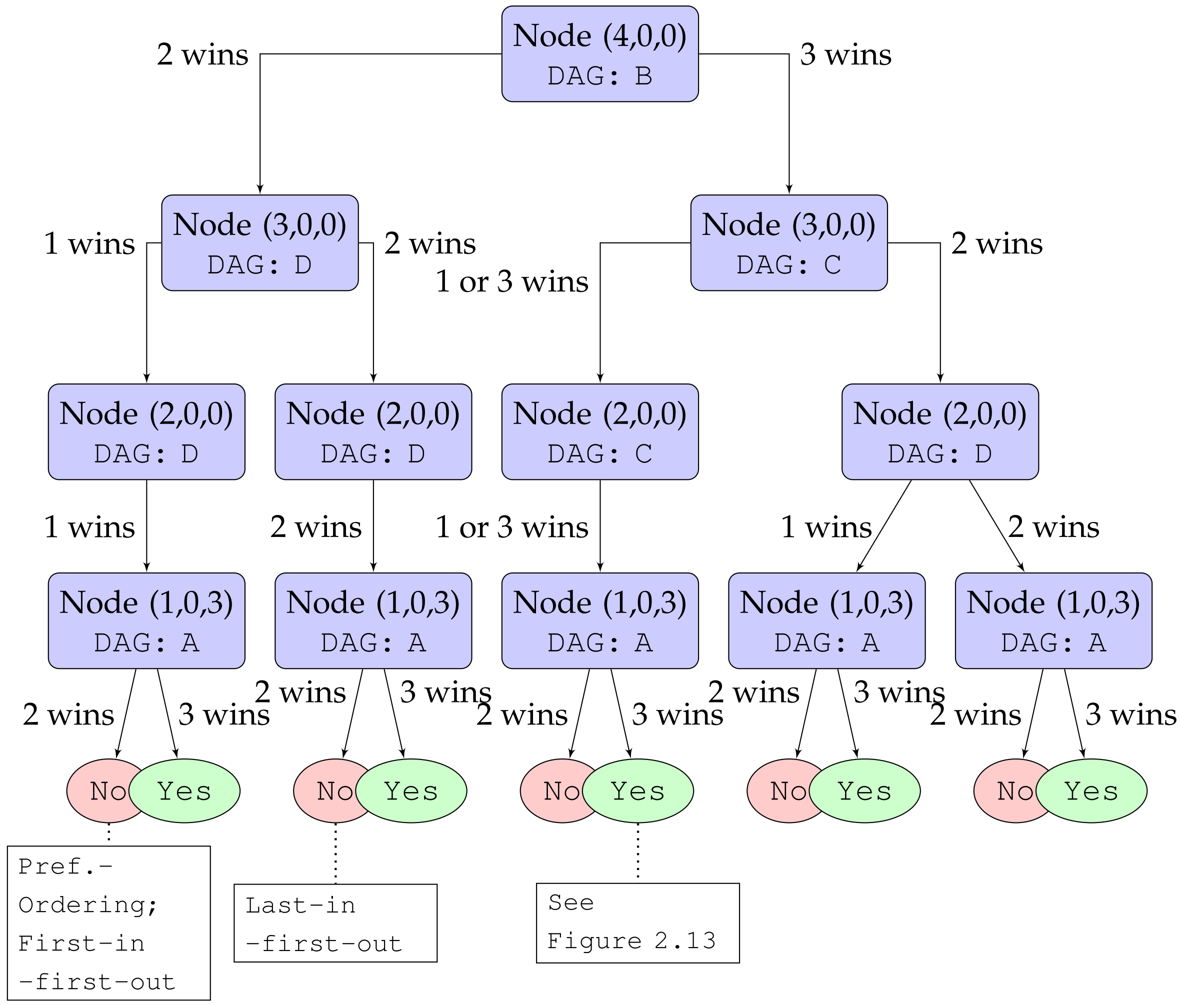

Figure 2.12: Monotonic Prices: Yes or No? *A Decision Tree Partitioning the Tie-Breaking Rules into Ten Classes*

C is induced at $(3, 0, 0)$. In the former case, the overbidding graph D arises at node $(2, 0, 0)$ regardless the choice of winner at $(3, 0, 0)$. In the latter case, there are three possible winners in the overbidding graph C at $(3, 0, 0)$. If Buyer 1 or Buyer 3 win these produce the same node valuations and give the overbidding graph C at $(2, 0, 0)$; if Buyer 2 wins this gives the overbidding graph D at $(2, 0, 0)$. A decision tree showing all the possible choices is shown in Figure 2.12. The reader may verify that these are the only decisions that affect the valuations at the nodes. Thus there are ten possible extensive forms, where `Yes/No` details whether or not a monotonic price trajectory is produced. Where the tie-breaking rules `preferential-ordering`, `first-in-first-out`, and `last-in-first-out` fit in this decision tree are highlighted in the figure.

Several observations are in order. First, not all of the classes of tie-breaking rule give non-monotonic price trajectories. An example of a tie-breaking rule with monotonic prices is shown in Figure 2.13. In fact, the choices made on the overbidding graphs B, C and D only affect valuations on nodes off the equilibrium path. The equilibrium path itself is determined uniquely by the choice made for the overbidding graph A. If the winner there is Buyer 2 then the prices are non-monotonic; if the winner there is Buyer 3 then the prices are monotonic.

We are now ready to complete the proof of the theorem. As we have just seen, any tie-breaking rule can be classified into one of ten classes depending upon its choices on this sequential auction. Five of the classes lead to non-monotonic prices on this instance. For the other five classes of tie-breaking rule we need to construct different examples on which they induces non-monotonic prices. But this is easy to do: take the same example but with the labels of Buyer 2 and Buyer 3 interchanged. The equilibrium paths for this sequential auction using any rule in the other five classes will then have non-monotonic price trajectories. □

2.3.3 Negative Utilities and Overbidding

We now discuss some interesting observations that arise from this specific sequential auction. First we recall another property of two-buyer sequential auctions: in each round of the auction each buyer has a non-negative value for winning the item over the other agent [37]. Interestingly, even this property fails to hold when there are more than two buyers.

Theorem 2.3.2. *There are multi-buyer sequential auctions with weakly decreasing marginals that have subgames in which one agent has a negative value for winning against another agent.*

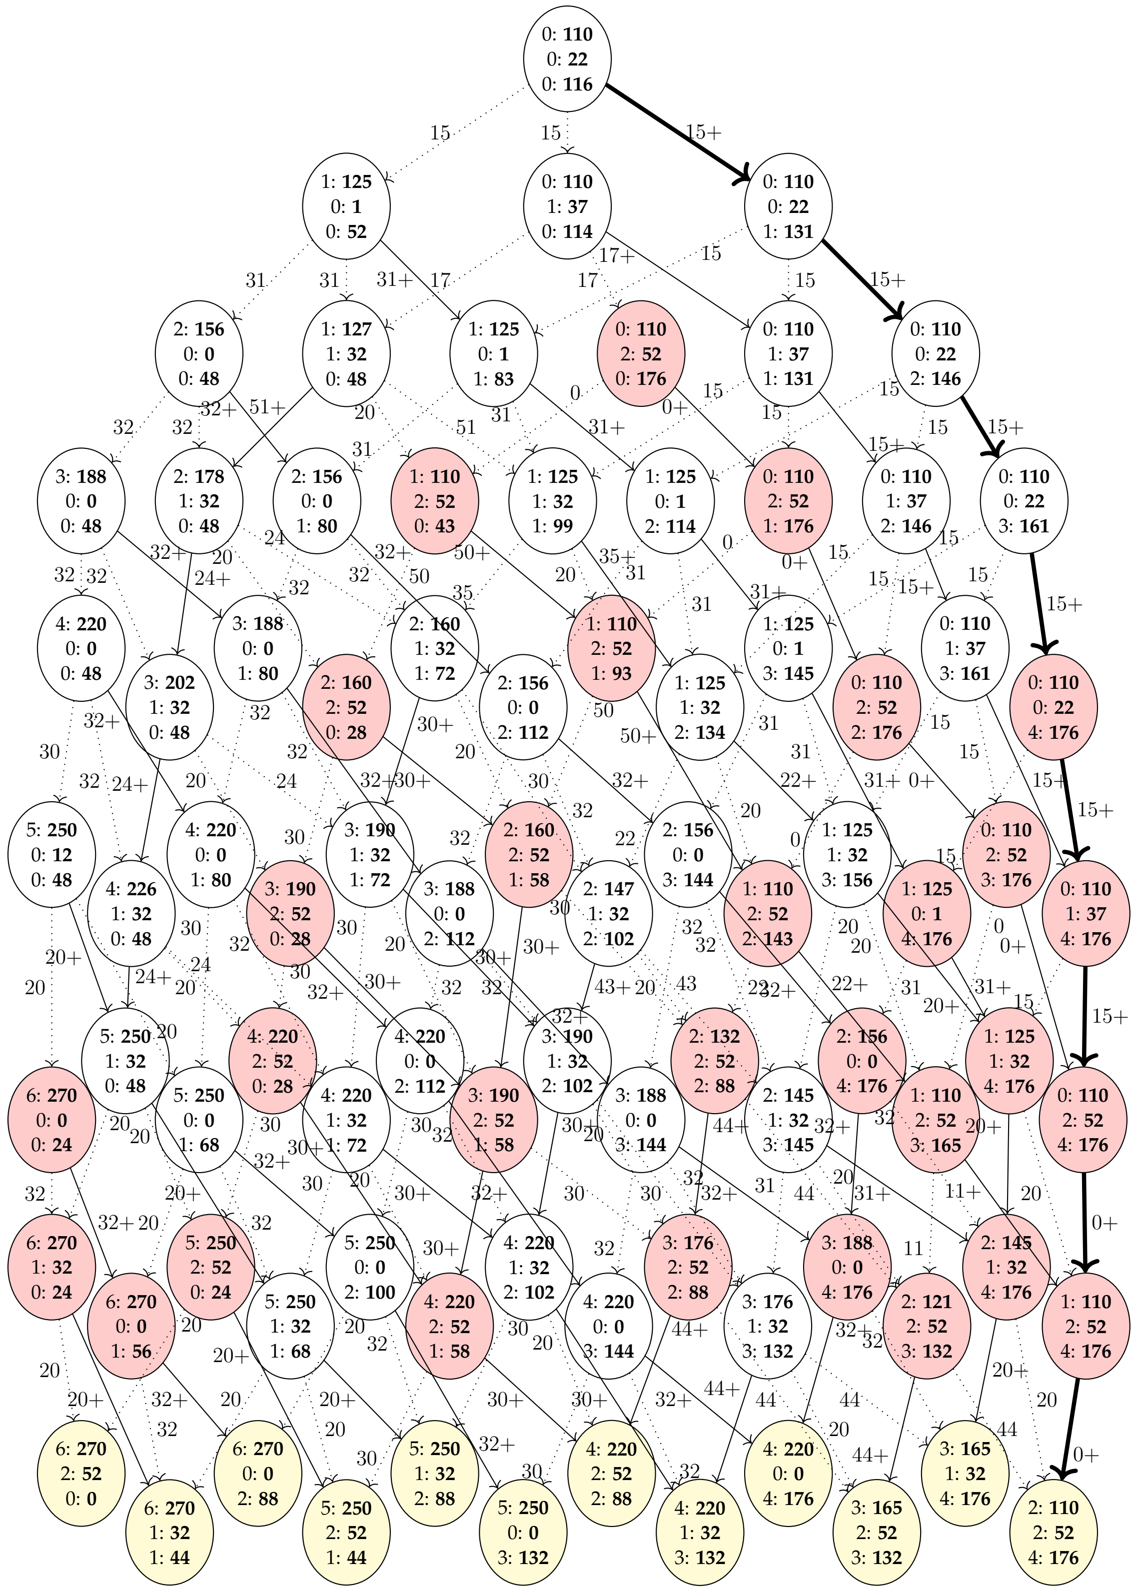

Figure 2.13: A Tie-Breaking Rule Resulting in Monotonic Prices

Proof. Consider again the sequential auction shown in Figure 2.13. Focus upon the auctions with externalities corresponding to the subgames rooted at the nodes $(0, 1, 0)$, $(0, 1, 1)$ and $(0, 1, 2)$. In all three cases, Buyer 3 has a negative value from winning over Buyer 2. For example, at node $(0, 1, 0)$ Buyer 3 has a utility of 131 for winning but a utility of 176 if Buyer 2 wins. (Note that Buyer 3 does have a positive value for defeating Buyer 1, specifically $131 - 48 = 83$.) Of course, this also implies there are sequential auctions with weakly decreasing marginal valuation functions where one agent has a negative value for winning the *first item* over one other agent. □

Second, observe in Figure 2.9 (see also Figure 2.8) that in the first round Buyer 3 has a value of $176 - 66 = 110$ for winning over Buyer 1. This far exceeds its marginal value of 44 for obtaining one item. Such "overbidding" also arises in two-buyer sequential auctions. The reader may wonder, however, whether this type of "overbidding" is responsible for the generation of non-monotonic price trajectories in multi-buyer auctions. This is not the case. To verify this we repeated all six million experiments described in Section 2.4 with the ascending price mechanism modified to exclude the possibility of a buyer bidding higher than their marginal value for their next unit of the good. The proportion of instances with non-monotonic price trajectories was similar (about 10% less). Moreover, there are instances where such "overbidding" does not arise but where the prices are non-monotonic.

2.4 Computational Results

Our computations were based on a dataset of over six million multi-buyer sequential auctions with non-increasing valuation functions randomly generated from different natural discrete probability distributions. Our goal was to observe the proportion of non-monotonic price trajectories in these sequential auctions and to see how this varied with (i) the number of buyers, (ii) the number of items, (iii) the distribution of valuation functions, and (iv) the tie-breaking rule. To do this, for each auction, we computed the subgame perfect equilibrium corresponding to the dropout bids and evaluated the prices along the equilibrium path to test for non-monotonicity.

We repeated this test for each of the three tie breaking rules described earlier, namely `preferential-ordering`, `first-in-first-out` and `last-in-first-out`. The results from our 6,240,000 randomly generated sequential auctions are shown in Figure 2.14. In these bar charts there is one bar for each combination of auction size and data structure (`preferential-ordering`, `first-in-firstout` and `last-in-first-out`). Each bar shows the number of auctions of that type that induced non-monotonic prices. For example, for sequential auctions with three buyers and five items that use the preferential-ordering tie-breaking rule, there were 7 auctions out of 240,000 that had non-monotonic prices. For four and five buyers there were 120,000 auctions of each type. We found no examples with less than 5 items that showed non-monotonicity, so the cases $T = 2, 3, 4$ are omitted. We omit the intricate details of the dataset generation and experimentation process (these are present in the full paper).

Observe that for a fixed number of buyers, there is a slight upward drift in the proportion of non-monotonic price trajectories as the number of items increases. Intuitively that seems unsurprising, as with longer price sequences there are more time periods at which deviations from monotonicity can arise. A very interesting question would be to study the limit of the proportion of non-monotonic price trajectories as the number of items gets very large. Unfortunately, due to the exponential explosion in the number of game tree nodes discussed in the previous sections, this question cannot be studied computationally. The main conclusion to be drawn from these computations is that non-monotonic prices are extremely rare. On the 6,240,000 auction instances, the `preferential-ordering` tie-breaking rule produced just 1,100 violations of the declining price anomaly. The `first-in-first-out` rule gave 986 violations and the `last-in-first-out` rule gave 1,334 violations. The overall observed rate of non-monotonicity across these over-18-million tests was 0.000183.

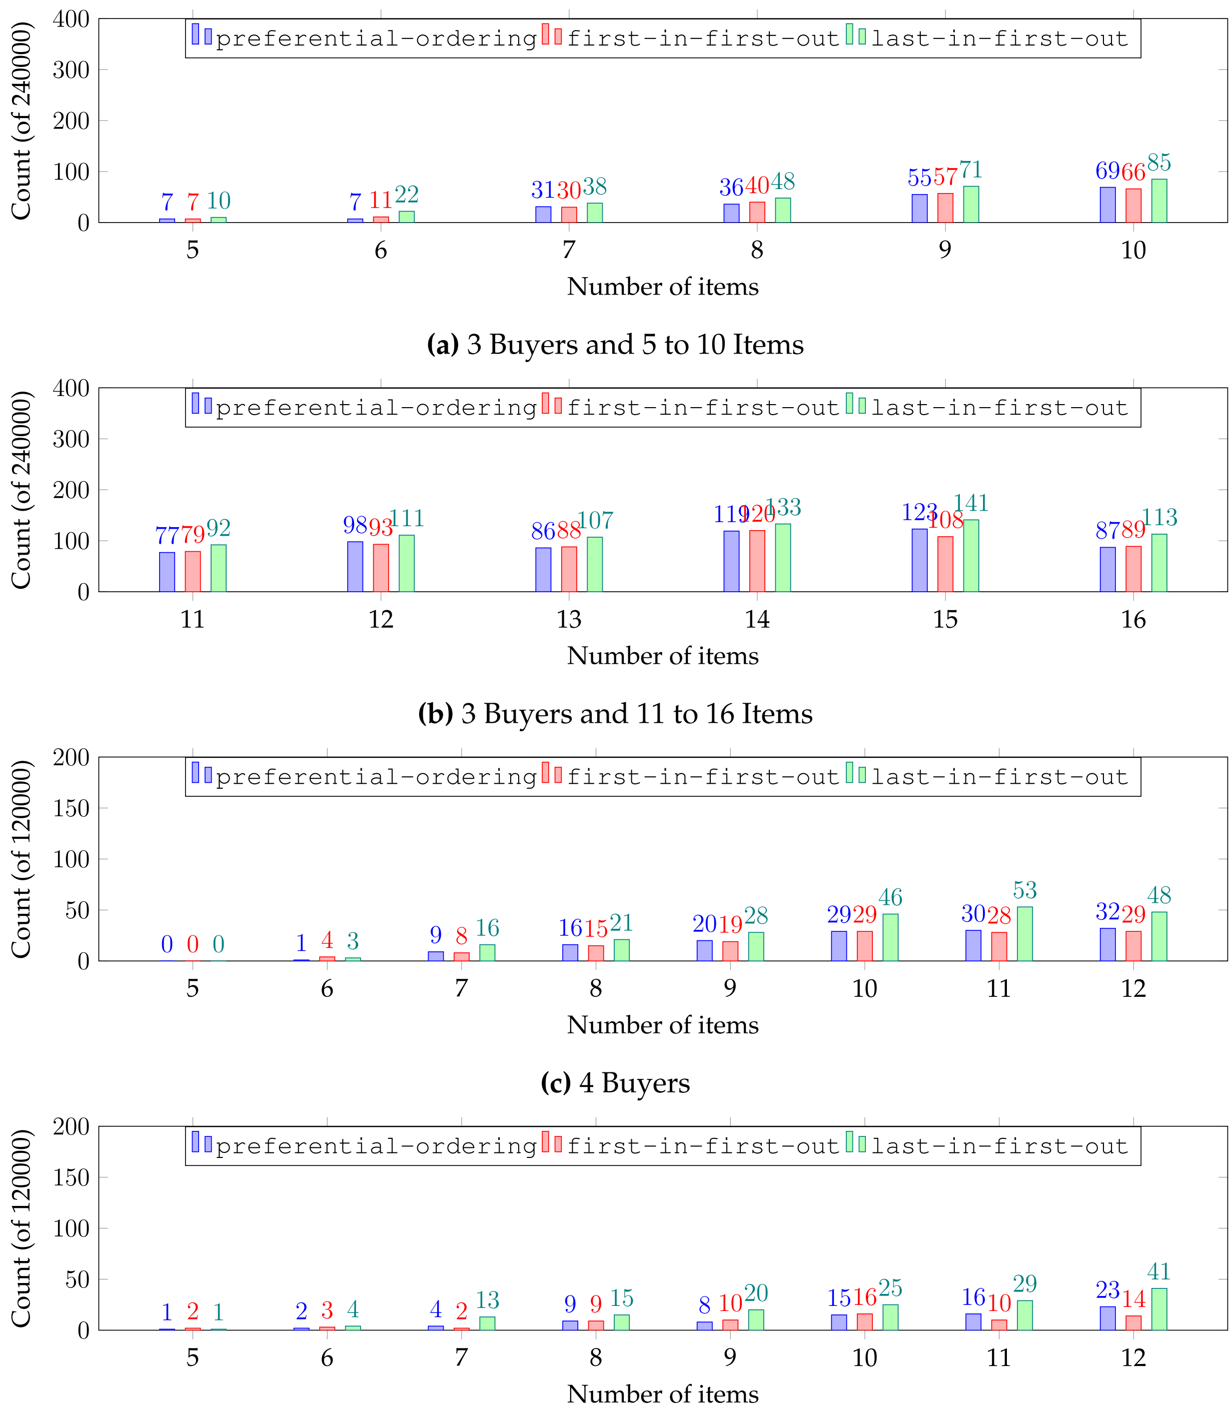

(a) 3 Buyers and 5 to 10 Items

(b) 3 Buyers and 11 to 16 Items

(c) 4 Buyers

(d) 5 Buyers

Figure 2.14: Bar Charts Showing the Frequency of Non-Monotonic Price Trajectories

Chapter 3

Risk-Free Bidding in Multi-Item Auctions

What strategy should a bidder use in a multi-item auction for a collection I of items? From a bidder's viewpoint, sequential auctions are perplexing for a variety of reasons. We observed in the previous chapter that the basic notion of equilibrium is a *subgame perfect equilibrium*, but that these equilibria are, in general, hard to compute. Intriguing structural properties can be derived for the equilibria of sequential auctions, but the recursive nature of this structure makes reasoning about equilibria complex. It follows that prescriptions (if such prescriptions can be computed) derived from the complete information setting, as studied in the aforementioned papers, are unlikely to extend reliably to more practical settings with incomplete information.

What strategy, then, as a participant in a sequential auction, should you use instead? Our goal in this section is to analyze one fundamental strategy that may be employed by a bidder in these auctions. We study this problem from the perspective of Bidder 1 in the following very general incomplete information setting. What is the maximum *risk-free profit* that Bidder 1 can guarantee for herself in a multi-item auction? Here, Bidder 1 knows her own entire valuation function but does not know the valuation function of the other agents. Clearly, if the other agents' bids are unrestricted then no guarantee is possible. Consequently, we impose a mild assumption on the other agents: Bidder i can spend at most some fixed budget B_i over the course of the multi-item auction. We will see that the critical case to analyze is when there are just two bidders (Bidder 1 and Bidder 2). We also assume that the only information Bidder 1 has on the other bidder is an estimate that his value for the entire collection of items is at most B_2; beyond this trivial upper bound, she has no specific information on the values the other bidder has for any subset of the items.

In the worst case, to maximize her guaranteed profit, we can model this problem as Bidder 1 competing in the auction against a single *adversary* who is incentivized to keep Bidder 1's utility low, and is willing to spend at most his budget B_2. This type of approach is analogous to that of a *safety strategy* in bimatrix games. This worst-case setting then corresponds to a special case of an auction with externalities, where Bidder 2 has no value for the items themselves, but is willing to bid on the items simply to prevent Bidder 1 from acquiring them. This situation often arises in practice, including in crucial matters of public safety and foreign policy. Jehiel et al. [48], in motivating their work on auctions with externalities, describe the situation in Ukraine after the breakup of the USSR: Ukraine inherited a huge nuclear arsenal, and Russia and the United States, while having no direct interest in acquiring Ukraine's weapons, were forced into action by the imminent danger of nuclear proliferation. In total, both countries paid billions of dollars in order for Ukraine to agree to dismantle its weapons and become a non-nuclear state.

In this work, we study the above auction setting and quantify the maximum risk-free profitability when the valuation function of Bidder 1 belongs to the class of subadditive (complement-free) functions and its subclasses. Interestingly, given the valuation class, tight bounds can be obtained that depend only on B_1 (the value Bidder 1 has for the entire set of items) and B_2. For example, the risk-free profitability of the class of fractionally subadditive (XOS) valuation functions is $(\sqrt{B_1} - \sqrt{B_2})^2$, for $B_2 \leq B_1$, and this bound is tight. Similarly, we present tight (but more complex) bounds for the class of subadditive valuation functions when the items are identical. We also analyze simultaneous auctions, for which we show that the risk-free profitability of the XOS class is at least $\frac{(B_1-B_2)^2}{2B_1}$ and $(B_1 - B_2)$ for first-price and second-price auctions, respectively.

As seen in the previous chapter, there is an extensive literature on sequential auctions. The study of incomplete information games was initiated by Milgrom and Weber [60, 78]. Theoretical studies on equilibria in complete information games include [37, 64, 63]. Given the abundance of sequential auctions in practice, there is also a very large empirical literature covering an assortment of applications ranging from antiques [42] to wine [7] and from fish [38] to jewellery [28].

Recently there has been a strong focus in the computer science community on the design of simple mechanisms. For multi-item auctions, *simultaneous auctions* are a notable example. These auctions are simple in that, as with a sequential auction, a standard single-item auction mechanism is used to sell each item. But in contrast with the sequential case, as the nomenclature suggests, these auctions are now held simultaneously rather than sequentially. Important streams of research in this area are concerned with the price of anarchy in simultaneous auctions and the hardness of computing an equilibrium (see, for example, [29, 19, 45, 33, 24]).

There has also been a range of papers examining the welfare of equilibria in sequential auctions. Paes Leme et al. [64] study the case of multi-bidder auctions. For sequential first-price auctions, they prove a factor 2 approximation guarantee for unit-demand bidders. In contrast, they show that equilibria can have arbitrarily poor welfare guarantees for bidders with submodular valuations. Feldman et al. [34] extend this result to the case where each bidder has either a unit-demand or additive valuation function.

Partly because of these negative results, a common assumption is that sequential auctions may not be a good mechanism by which to sell a collection of items. However, there are reasons to believe that, in practice, sequential auctions have the potential to proffer high welfare. For example, consider the influential paper of Lehmann et al. [54]. There, they present a simple greedy allocation mechanism with a factor 2 welfare guarantee for allocating items to agents with submodular valuation functions. One interesting implication of this result is that if the items are sold via a second-price sequential auction *and* every agent (assuming submodular valuations) truthfully bids their marginal value in each round then the outcome will have at least half the optimal social welfare.

In Section 3.1 we explain the sequential auction model and related definitions. We present our measure, the risk-free profitability of a bidder in incomplete information multi-bidder auctions, and explain how to quantify it via a two-bidder adversarial sequential auction. In Section 3.2 we present a simple sequential auction example (*uniform additive auctions*) to motivate the problem and to illustrate the difficulties that arise in designing risk-free bidding strategies, even in very small sequential auctions with at most three items.

Section 3.3 and Section 3.4 contain our main results. In Section 3.3 we begin by presenting tight upper and lower bounds on the risk-free profitability of a fractionally subadditive (XOS) bidder. For the lower bound, in Section 3.3.1 we exhibit a bidding strategy that guarantees Bidder 1 a profit of at least $(\sqrt{B_1} - \sqrt{B_2})^2$.

In Section 3.3.2 we describe a sequence of sequential auctions that provide an upper bound that is asymptotically equal to the aforementioned lower bound as the number of items increases. We prove these bounds for *first-price* sequential auctions, but nearly identical proofs show the bounds also apply for *second-price* sequential auctions. Next we prove that the risk-free profitability of an XOS bidder is lower in sequential auctions than in simultaneous auctions. Equivalently, a budgeted adversary is stronger in a sequential auction than in the corresponding simultaneous auction. Specifically, in Section 3.3.3, we prove that an XOS bidder has a risk-free profitability of at least $\frac{(B_1-B_2)^2}{2B_1}$ in a first-price simultaneous auction and of at least $B_1 - B_2$ in a second-price simultaneous auction. Several other interesting observations arise from these results. First, unlike for sequential auctions, the power of the adversary differs in a simultaneous auction depending on whether a first-price or second-price mechanism is used: the adversary is stronger in a first-price auction. Second, the risk-free strategies we present for simultaneous auctions require *no* information about the adversary at all. The performance of the strategy (its risk-free profitability) is a function of B_2, but the strategy itself does not require that Bidder 1 have knowledge of B_2 (nor an estimate of it). Third, for the case of first-price simultaneous auctions, it is necessary that Bidder 1 use randomization in her risk-free strategy. Finally, in Section 3.4 we study the risk-free profitability of a bidder with a subadditive valuation function. We give a possible explanation for why simple strategies fail to perform well in the general case. We then examine the special case where the items are identical. We derive tight lower and upper bounds for this setting.

3.1 The Model

3.1.1 Sequential Auctions and Valuation Functions

There are n bidders and a collection $I = \{a_1, \ldots, a_m\}$ of m items to be sold using a sequential auction. In the ℓth round of the auction item a_ℓ is sold via a first-price (or second-price) auction. We view the auction from the perspective of Bidder 1 who has a publicly-known valuation function $v_1 : 2^I \to \mathbb{R}_{\geq 0}$ that assigns a non-negative value to every subset of items. We denote v_1 by v where no confusion arises. This valuation function is assumed to satisfy $v(\emptyset) = 0$ and to be *monotone*, that is, $v(S) \leq v(T)$, for all $S \subseteq T$. When all the items have been auctioned, the *utility* or *profit* π_1 of Bidder 1 is her value for the set of items she was allocated minus the sum of prices of these items.

The sequential auction setting is captured by extensive form games. A *strategy* for player i is a function that assigns a bid b_i^t for the item a_t that depends on the previous bids $\{b_i^\tau\}_{i,\tau<t}$ of all players, and on the allocation of the first $t-1$ items. The profit of a strategy profile $\boldsymbol{b}$ for Bidder 1 is the profit Bidder 1 obtains when all bidders bid according to $\boldsymbol{b}$.

The question we then study is how much profit Bidder 1 can guarantee for herself. We examine the case where v is in the class of *subadditive* or *complement-free* valuation functions. Belonging to this class, of particular interest in this work are *additive* functions, *submodular* functions, and *fractionally subadditive* or *XOS* functions. These functions are defined as follows.

- `Subadditive (Complement-Free)`. A function v is subadditive if $v(S \cup T) \leq v(S) + v(T)$ for all $S, T \subseteq I$.
- `Additive (Linear)`. A function v is additive if $v(S) = \sum_{a \in S} v(a)$ for each $S \subseteq I$.
- `Submodular (Decreasing Marginal Valuations)`. A function v is submodular if $v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$ for all $S, T \subseteq I$.
- `Fractionally Subadditive (XOS)`. A function v is fractionally subadditive if there exists a nonempty collection of additive functions $\{\gamma_1, \ldots, \gamma_\ell\}$ on I such that for every $S \subseteq I$, $v(S) = \max_{j \in [\ell]} \gamma_j(S)$.[1]

Lehmann et al. [54] showed that these valuation classes form the following hierarchy:

ADDITIVE $\subseteq$ SUBMODULAR $\subseteq$ FRACTIONALLY SUBADDITIVE $\subseteq$ SUBADDITIVE

Other important classes in this hierarchy include unit-demand and gross substitutes valuation functions, but they will not be needed here.

3.1.2 Simultaneous Auctions

The simultaneous auction setting is similar to the sequential auction setting in that there are n bidders and a collection $I = \{a_1, \ldots, a_m\}$ of m items to be sold. Each bidder makes a bid on each item. Unlike the sequential case, there is now a single time period. Thus each bidder makes a vector of m bids – one for each item – simultaneously.

3.1.3 Bidding against an Adversary

To quantify the maximum profit that Bidder 1 can obtain, without loss of generality, we may *normalize* the valuation function (and corresponding auction) by scaling the values so that $v(I) = v_1(I) = 1$. Now the maximum guaranteed profit will depend on the strength of the other bidders. We quantify this by a parameter B: in the setting where each player

[1] This is the standard definition of XOS functions. Fractionally subadditive functions are defined in terms of fractional set covers; the equivalence between fractionally subadditive and XOS functions was shown by Feige [32].

$j \geq 2$ has valuation function v_j, B is the sum of the total values of the other bidders, i.e., $B = \sum_{j=2}^{n} v_j(I)$. This corresponds to an incomplete information auction where the only common knowledge are upper bounds on the value each agent has for the entire set of items. From the perspective of Bidder 1, it is apparent that the worst case arises when $n = 2$, and so $B = B_2 = v_2(I)$. Thus we may assume that $n = 2$, and we can view Bidder 1 as playing against an *adversary* with a budget B. To see this, observe that for a fixed $B = \sum_{j=2}^{n} v_j(I)$ if there are $n >= 3$ bidders then the worst case for Bidder 1 arises when the other bidders coordinate to act as a single adversary: however, if the budget is split between two or more other bidders then their ability to buy a single item of high value is potentially removed.

Thus, Bidder 1 seeks a strategy that works well against a rational bidder who, by monotonicity, has a value at most B for any subset of the items. We model the game that we will analyze as a zero-sum game with the following properties. Bidder 1's payoff in this game is simply her profit π_1 from the auctions, that is, her value for the set that she is allocated minus the sum of the prices of the items. Since Bidder 2 is viewed as an adversary, and the game is zero-sum, his payoff in this game, π_2, is equal to $-\pi_1$. Bidder 2 can evaluate his payoff during the game, i.e., Bidder 1's valuation function is common knowledge. Additionally, here the adversary's budget constraint is tight: for example, in the sequential case, in time step t, if Bidder 2 paid p_2^{t-1} for the items that have already sold, then his next bid b_2^t is at most $B - p_2^{t-1}$. We call this the *risk-free sequential auction game* $\mathcal{R}(v, B)$. The *guaranteed profit* for Bidder 1 is the minimum profit obtainable by playing a safety strategy in this game (i.e. the *value* of this game). For any normalized valuation v, we denote this profit by $\pi_1^*(v, B)$, or simply π_1^* where there is no ambiguity. For any class of set functions $\mathcal{C}$ and any budget $B \in (0, 1)$, we want to find the maximum profit Bidder 1 can guarantee in *all* m-item instances $\mathcal{R}(v, B)$ where $v \in \mathcal{C}$, which is precisely $\min_{v \in \mathcal{C}} \pi_1^*(v, B)$. We call this the *risk-free profitability* $\mathcal{P}(\mathcal{C}, B)$ of the class $\mathcal{C}$. For a budgeted adversary in a simultaneous auction, the analogue of this budget-constrained bidding is that the *sum* of the adversary's bids on the items is at most the budget B. We define risk-free profitability analogously for simultaneous auctions. The focus of this work is to quantify the risk-free profitability of the aforementioned classes of valuation functions for both sequential and simultaneous auction mechanisms. We note that for the sequential setting, in the course of our proofs we will show implicitly that the aforementioned zero sum game has a *subgame perfect equilibrium* in pure strategies. Consequently this is the solution concept that we will analyze for our main result.

Table 3.1 summarizes our obtained bounds by auction type, where the valuation functions of the agents are normalized as described previously.

Valuation Class	Lower Bound	Upper Bound
Sequential Auctions (First- and Second-Price)		
Additive	$(1-\sqrt{B})^2$	$(1-\sqrt{B})^2$
Submodular	$(1-\sqrt{B})^2$	$(1-\sqrt{B})^2$
XOS	$(1-\sqrt{B})^2$	$(1-\sqrt{B})^2$
Subadditive (Identical)	$t^*(B)$	$t^*(B)$, $B \in (0, \frac{1}{4})$ $(1-\sqrt{B})^2$, $B \in [\frac{1}{4}, 1)$
Simultaneous Auctions (First-Price)		
XOS	$1-B^2$, $B \in (0, 3-2\sqrt{2})$ $\frac{(1-B)^2}{2}$, $B \in [3-2\sqrt{2}, 1)$	$1-2B$, $B \in (0, \frac{1}{4})$ $\frac{2}{3}(1-B)$, $B \in [\frac{1}{4}, 1)$
Simultaneous Auctions (Second-Price)		
XOS	$(1-B)$	$(1-B)$

Table 3.1: Valuation Classes and their Risk-Free Profitability

3.2 Example: Uniform Additive Auctions

We now present a simple example of a sequential auction with an agent (Bidder 1) that strategizes against an adversary (Bidder 2), which will be helpful for two reasons. First, it illustrates some of the strategic issues facing the agent and, implicitly, the adversary in a sequential auction. Second, these examples form base cases in our proof in Section 3.3.2.

The auction is defined as follows. Bidder 1 has an additive valuation function where each item has exactly the same value. That is, for an auction with m items, we have that $v(a_t) = \frac{1}{m}$. The adversary Bidder 2 has a budget B. We call this the *uniform additive auction* on m items and denote it by $\mathcal{A}_m$. For our example, we are interested in uniform additive auctions where $m \leq 3$. We denote by b_i^j Bidder i's bid on item j.

`One Item.` First, consider the case $\mathcal{A}_1$. We have a single item a_1 with $v(\{a_1\}) = 1$ for Bidder 1. Let b_1 and b_2 be the bids placed on the item by Bidders 1 and 2, respectively. Clearly if $b_1 < B$ then the adversary's best response is to bid $b_2 = b_1^+$ and win the item. Then $\pi_1 = \pi_2 = 0$. On the other hand, if $b_1 \geq B$, then the adversary is constrained by his budget and cannot beat Bidder 1. Thus Bidder 1 wins and obtains a profit of $\pi_1 = 1 - b_1$. It follows that Bidder 1's risk-free strategy is to bid B and win the item at price B for a guaranteed profit of $\pi_1^* = 1 - B$. Clearly, if $B \geq 1$ then $\pi_1^* = 0$ since the adversary can prevent Bidder 1 from winning the item. Specifically, we have shown

$$\pi_1^* = \begin{cases} 1 - B & \text{if } 0 \leq B < 1 \\ 0 & \text{if } 1 \leq B \end{cases} . \tag{3.1}$$

`Two Items.` Now consider the case $\mathcal{A}_2$. So there are two items a_1 and a_2 and Bidder 1 has an additive valuation function with $v(\{a_1\}) = v(\{a_2\}) = \frac{1}{2}$ and $v(\{a_1, a_2\}) = 1$. We divide our analysis into three cases.

- $B < \frac{1}{4}$: If $B < \frac{1}{4}$, then Bidder 1 can bid B on each item and win both items at price B each, so her guaranteed profit is at least $1 - 2B > \frac{1}{2}$. If Bidder 1 bids less than B on either item, then Bidder 2 can win that item, ensuring that Bidder 1's profit is less than her value of the other item, that is $\frac{1}{2}$. Bidder 1's risk-free strategy is thus to bid B on both items for a profit $\pi_1^* = 1 - 2B$.

- $\frac{1}{4} \leq B < \frac{1}{2}$: If Bidder 1 bids $b_1^1 = d$ on a_1, with $0 \leq d \leq \frac{1}{2}$, then Bidder 2 can either win by bidding $b_2^1 > d$ or lose by bidding $b_2^1 < d$ (for now, we assume $d < B$). In the former case, the adversary's budget in the second auction is $B - b_2^1$, and there is only one item remaining. Then, reasoning as we did for the one item setting, it is easy to see that Bidder 1's profit from the second item is $\pi_1 = \frac{1}{2} - (B - b_2^1) = \frac{1}{2} - B + b_2^1$. This is minimized (with value $\frac{1}{2} - B + d$) when Bidder 2 bids an amount negligibly larger than d. In the latter case, the adversary loses the first item, so he has budget B in the second auction. Bidder 1's combined profit (on both items) is then $\pi_1 = (\frac{1}{2} - d) + (\frac{1}{2} - B) = 1 - B - d$. For $d = 0$ we have $\frac{1}{2} - B + d < 1 - B - d$ and for $d = B$ we have $\frac{1}{2} - B + d \geq 1 - B - d$, since $B \geq \frac{1}{4}$. But $\frac{1}{2} - B + d$ is increasing in d and $1 - B - d$ is decreasing in d. Therefore, assuming Bidder 2 plays a best response, we see that π_1 is maximized when the minimum of these values is maximized. That is $\pi_1^* = \max_{0 \leq d < B} \min\left[\frac{1}{2} - B + d, 1 - B - d\right]$. The optimal choice is $d = \frac{1}{4}$ giving $\pi_1^* = \frac{3}{4} - B$. Note that our assumption that $d < B$ is validated: if Bidder 1 bids an amount d that is greater than or equal to B on the first item the she will win both items for a total profit $(\frac{1}{2} - d) + (\frac{1}{2} - B) = 1 - d - B \leq 1 - 2B \leq \frac{3}{4} - B$.

- $\frac{1}{2} \leq B < 1$: Suppose Bidder 1 bids $b_1^1 = d$ on a_1, with $0 \leq d \leq \frac{1}{2}$, then Bidder 2 can either win by bidding $b_2^1 > d$ or lose by bidding $b_2^1 < d$. In the former case, the adversary's budget in the second auction is $B - b_2^1$, so Bidder 1's profit from the second item is $\pi_1 = \frac{1}{2} - (B - b_2^1) = \frac{1}{2} - B + b_2^1$. This is minimized (with value $\frac{1}{2} - B + d$) when Bidder 2 bids an amount negligibly larger than d. In the latter case, the adversary loses the first item, so he still has budget B for the second auction. Thus Bidder 1 loses the second item and makes no profit on it. Bidder 1's total profit is then $\frac{1}{2} - d$. Since $\frac{1}{2} - B + d$ is increasing in d and $\frac{1}{2} - d$ is decreasing in d, her profit π_1 is maximized at $d = \frac{B}{2}$. This gives a maximum guaranteed profit of $\pi_1^* = \frac{1}{2} - \frac{B}{2}$.

Putting this all together we have that

Budget	$0 \leq B < \frac{1}{4}$	$\frac{1}{4} \leq B < \frac{1}{2}$	$\frac{1}{2} \leq B < 1$	$1 \leq B$
Profit π_1^*	$1 - 2B$	$\frac{3}{4} - B$	$\frac{1}{2} - \frac{B}{2}$	0

(3.2)

Before proceeding to the case of the uniform additive auction with three items, we emphasize that even the very simple case $\mathcal{A}_2$ illustrates many of the strategic considerations that arise in more complex sequential auctions. To wit, in the first time period Bidder 1 faces the standard conundrum that bidding high increases her chances of winning but at the expense of receiving a smaller profit if she does win. More interestingly, in this adversarial setting, Bidder 1 has an additional incentive for bidding high. If she bids high *and* loses then she faces a weaker adversary in the subsequent time period. That is, by winning the first item at a high price the adversary's budget is significantly reduced in the auction for the second item. Counterintuitively, therefore, in adversarial sequential actions, Bidder 1 has an incentive to lose some of the items!

Interestingly, the adversary has perhaps even stronger incentives to lose than Bidder 1. Whilst winning the first item does hurt Bidder 1, by reducing his budget, this also reduces the strength of the adversary in the subsequent round. Thus, the optimal outcome for adversary is that he lose the first item at a high price; this keeps the profit of Bidder 1 low and increases the relative strength of the adversary in the second auction. This is in stark contrast with the simultaneous case, where both Bidder 1 and the adversary always have an incentive to win every item.

We remark that these basic motivations and incentives play a fundamental role in sequential auctions with many items. We now proceed to the case $\mathcal{A}_3$ of three items.

`Three Items`. Now there are three items a_1, a_2 and a_3 and Bidder 1 has an additive valuation function with $v(\{a_1\}) = v(\{a_2\}) = v(\{a_3\}) = \frac{1}{3}$. Applying a similar case analysis, her maximum guaranteed profits are then:

Budget	Profit π_1^*
$0 \leq B < \frac{1}{9}$	$1 - 3B$
$\frac{1}{9} \leq B < \frac{1}{6}$	$\frac{8}{9} - 2B$
$\frac{1}{6} \leq B < \frac{1}{3}$	$\frac{7}{9} - \frac{4B}{3}$
$\frac{1}{3} \leq B < \frac{5}{9}$	$\frac{7}{12} - \frac{3B}{4}$
$\frac{5}{9} \leq B < \frac{2}{3}$	$\frac{4}{9} - \frac{B}{2}$
$\frac{2}{3} \leq B < 1$	$\frac{1}{3} - \frac{B}{3}$
$1 \leq B$	0

(3.3)

We remark that this profit function is still piecewise linear and is so for $\mathcal{A}_m$ in general. However the complexity of the profit function grows rapidly as the number of items increases.

3.3 Tight Bounds for XOS Valuation Functions

In this section we prove tight bounds on the risk-free profitability of Bidder 1 with a fractionally subadditive (XOS) valuation function. In Sections 3.3.1 and 3.3.2 we study the sequential auction setting, and in Section 3.3.3 we consider the simultaneous case. Specifically, in Section 3.3.1, we show that the agent has a strategy in the normalized auction that gives a guaranteed profit of $(1-\sqrt{B})^2$ when the adversary has a budget of B. This is equivalent to a profit of $(\sqrt{B_1}-\sqrt{B_2})^2$ in the original (unnormalized) auction. Then, in Section 3.3.2, we prove that no strategy can guarantee a profit that is greater than this by an (asymptotically zero) additive quantity.

3.3.1 The XOS Lower Bound

It is quite straightforward to obtain a lower bound on the profitability of Bidder 1 when she has an XOS valuation function. She simply chooses the additive function that maximizes her valuation under the assumption that she wins every item. Then, for each item, she bids a fixed fraction of the value of this item under this additive function. For any XOS valuation, this strategy guarantees a profit of at least $(1-\sqrt{B})^2$ against any strategy utilized by an adversary with a budget of $B \in (0,1)$. Consequently, we have the following.

Theorem 3.3.1. $\mathcal{P}(XOS, B) \geq (1-\sqrt{B})^2$.

Proof. Let $I = \{a_1, \ldots, a_m\}$ be the set of auctioned items, and v be Bidder 1's XOS valuation function on I. Since v is XOS, there is a set $\{\gamma_1, \gamma_2, \ldots, \gamma_\ell\}$ of (normalized) additive set functions on I such that for any $S \subseteq I$ we have $v(S) = \max_{i\in[\ell]} \gamma_i(S)$. Let

$\gamma^* = \text{arg}\max_{i\in[\ell]} \gamma_i(I)$ be an additive function that induces the value of v on the entire set of items I. Thus $v(I) = \gamma^*(I)$. Moreover, by definition of v, we have that

$$v(S) \geq \gamma^*(S) \qquad\qquad \forall S \subseteq I. \tag{3.4}$$

Now consider the following strategy for Bidder 1. In time period t let Bidder 1 place a bid of $b_1^t = \sqrt{B} \cdot \gamma^*(a_t)$ on item $a_t \in I$, for all $t \in [m]$. Let $I_1 \subseteq I$ be the set of items won by Bidder 1 at the end of the auction (that is, by the end of time period m). Similarly let $I_2 \subseteq I$ be the set of items won by Bidder 2 at the end of the auction. Therefore $I_1 \cup I_2 = I$ and, by the additivity of γ^*, we have

$$\gamma^*(I_1) + \gamma^*(I_2) = 1. \tag{3.5}$$

It follows that during the sequential auction the adversary spent

$$\sum_{t\in I_2} b_2^t \;\geq\; \sum_{t\in I_2} b_1^t \;=\; \sum_{t\in I_2} \sqrt{B}\cdot\gamma^*(a_t) \;=\; \sqrt{B}\cdot\sum_{t\in I_2}\gamma^*(a_t) \;=\; \sqrt{B}\cdot\gamma^*(I_2). \tag{3.6}$$

Here the inequality arises because Bidder 2 won the items in I_2. The first equality follows by the definition of Bidder 1's safety strategy and the third equality follows by the additivity of γ^*. But the adversary has a total budget of B. Therefore, together this budget constraint on Bidder 2 and Inequality (3.6) imply that $\sqrt{B}\cdot\gamma^*(I_2) \leq B$. Hence, $\gamma^*(I_2) \leq \sqrt{B}$. From Equation (3.5) we then derive that

$$\gamma^*(I_1) \geq (1-\sqrt{B}). \tag{3.7}$$

Now define π_1 to be the total profit obtained by Bidder 1 using this safety strategy. Then

$$\begin{aligned}\pi_1 = v(I_1) - \sum_{t\in I_1} b_1^t &= v(I_1) - \sum_{t\in I_1}\sqrt{B}\cdot\gamma^*(a_t) \\ &= v(I_1) - \sqrt{B}\cdot\gamma^*(I_1) \geq (1-\sqrt{B})\cdot\gamma^*(I_1).\end{aligned} \tag{3.8}$$

Here the inequality follows from the property (3.4) applied to the subset I_1. Finally, combining (3.7) and (3.8) gives $\pi_1 \geq (1 - \sqrt{B})^2$, as required. $\square$

Next we will show this bound is tight by providing a construction where the adversary has a strategy limiting the profitability of Bidder 1 to this quantity. This is surprising because the bidding strategy described above is *non-adaptive* – it does not adapt to the history of the auction. Given the extra flexibility afforded by adaptive strategies, one would expect a priori the optimal risk-free strategy to be adaptive. In fact, the result that follows from this construction is doubly surprising, as it holds even if the adversary commits to his deterministic strategy in advance and Bidder 1 is allowed to randomize her strategy. Unlike the simultaneous case (which we discuss later), in this sequential setting the simple bidding strategy presented above is optimal for Bidder 1, and she can obtain no improvement with an adaptive or randomized strategy.

3.3.2 The XOS Upper Bound

In this section we present a matching upper bound, showing that the highest guaranteed profit of a risk-free strategy in a normalized two-player sequential auction with an XOS valuation is within a $\frac{1}{\sqrt{m}}$-additive factor of our lower bound. To do this, we present a sequential auction with an XOS valuation function where the game value is at most $(1 - \sqrt{B})^2 + \frac{1}{\sqrt{m}}$. In fact, rather surprisingly, this upper bound applies even for additive valuation functions. Specifically, we prove that for the uniform additive auction $\mathcal{A}_m$ the adversary has a strategy that ensures that the profit of Bidder 1 is at most $(1 - \sqrt{B})^2 + \frac{1}{\sqrt{m}}$. Consequently, the upper bound applies to every class of valuation functions that contains the additive functions. Together with the lower bound, this resolves the profitability of several well-studied classes, including the additive, submodular and gross substitutes valuation classes. We will see later on, in Section 3.4, that the situation is not as simple for subadditive valuations (that are not contained in XOS). We denote by XOS_m the class of

XOS functions on m items. The following theorem, together with the lower bound, gives our main result: $\mathcal{P}(XOS, B)$ is asymptotically equal to $(1-\sqrt{B})^2$ when $B \in (0,1)$.

Theorem 3.3.2. $\mathcal{P}(XOS_m, B) \leq (1-\sqrt{B})^2 + \frac{1}{\sqrt{m}}$.

Proof. We prove this result by induction on m. To do this, we start with a simple observation: after the first item has been sold in the uniform additive auction $\mathcal{A}_m$ then the sequential auction on items $\{a_2, \ldots, a_m\}$ is simply the auction $\mathcal{A}_{m-1}$ but with the additive values scaled by a multiplicative factor $\frac{m-1}{m}$; that is, the agent now has a value $\frac{1}{m}$ for each item rather than $\frac{1}{m-1}$ as in the unscaled $\mathcal{A}_{m-1}$. Consequently, by appropriately scaling the values *and* the budget of the adversary we will be able to analyze the auction $\mathcal{A}_m$ by studying the first round of that auction and then applying induction on the remaining rounds.

Formally, for any positive integer m let $f_m : \mathbb{R}_{\geq 0} \to [0,1]$ be a function giving the highest guaranteed profit $f_m(x)$ of a risk-free strategy given that the adversary has a budget $B = x$. Clearly, for all m, we have that $f_m(0) = 1$ and that $f_m(x) = 0$ for any $x \geq 1$. Set $f(x) = (1-\sqrt{x})^2$. Then we want to prove by induction that

$$f_m(x) \leq f(x) + \frac{1}{\sqrt{m}} \qquad \forall m \geq 1, \forall x \in (0,1). \tag{3.9}$$

`Base Cases`: For the base cases, consider $m \in \{1,2,3\}$. Note that we have already studied the auctions $\mathcal{A}_1, \mathcal{A}_2$ and $\mathcal{A}_3$ in Section 3.2. Specifically, we found that $f_1(x) = (1-x)$, and that $f_2(x)$ is given by (3.2) and $f_3(x)$ is given by (3.3). It can be easily verified (see Figure 3.1) that each of the above functions $f_m(x)$, $m \in \{1,2,3\}$, is at most $f(x) + \frac{1}{\sqrt{m}}$, for any $x \in [0,1]$. Consequently, the base cases hold.

`Induction Hypothesis`: Assume that $f_k(x) \leq f(x) + \frac{1}{\sqrt{k}}$ for all $k < m$.

`Induction Step`: Given the induction hypothesis we will now prove that $f_m(x) \leq f(x) + \frac{1}{\sqrt{m}}$. We will present a strategy for the adversary and prove that this strategy guarantees that Bidder 1 cannot make a profit greater than $f(x) + \frac{1}{\sqrt{m}}$ in the uniform

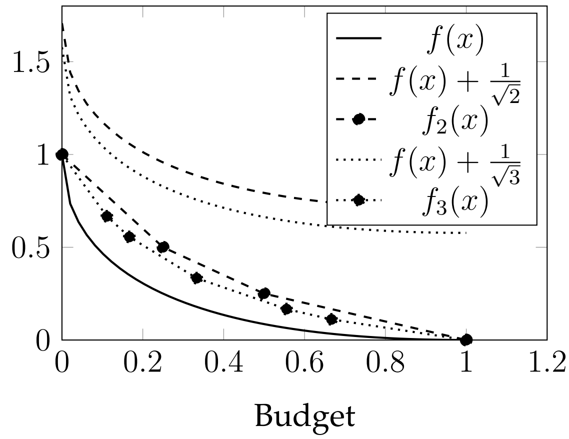

Figure 3.1: Plot of the Functions $f(x)$, $f(x)+\frac{1}{\sqrt{2}}$, $f_2(x)$, $f(x)+\frac{1}{\sqrt{3}}$ and $f_3(x)$

additive auction $\mathcal{A}_m$. Specifically, we consider the auction for the first item a_1 in $\mathcal{A}_m$, and we let $b_2^1 = \alpha \cdot \frac{1}{m}$ be the adversary's bid on this item. Since Bidder 1 has an additive value $\frac{1}{m}$ for this item, the adversary will never make a bid $b_2^1 > \frac{1}{m}$. Thus we may assume that the adversary makes a bid $b_2^1 = \alpha \cdot \frac{1}{m}$ for some $0 \leq \alpha \leq 1$. We then show that for some particular choice of α, even with an optimal response Bidder 1 does not make a profit greater than $f(x) + \frac{1}{\sqrt{m}}$. In determining her optimal response, Bidder 1 faces the dilemma of whether or not to outbid the adversary. Thus we have two possibilities:

• *Bidder 1 wins item a_1.*

In this case it is easy to see that Bidder 1 will bid $b_1^1 = b_2^{1+}$ (which is $b_2^1 + \epsilon$ for any negligibly small ϵ) as any higher bid will lead to a strictly smaller profit as this is a first-price auction. Thus, Bidder 1 makes an immediate profit of $\frac{1}{m} - \alpha \cdot \frac{1}{m} = \frac{1-\alpha}{m}$ on the first item. The rest of the sequential auction is a scaled version of $\mathcal{A}_{m-1}$. As discussed, the additive valuations of Bidder 1 are scaled by a multiplicative factor of $\frac{m-1}{m}$. Moreover, the budget of the adversary is also scaled. As the adversary lost the first item his budget remains x, which corresponds to a budget of $B = \frac{m}{m-1} \cdot x$ in the scaled auction $\mathcal{A}_{m-1}$. Therefore, given that the bidders play optimal strategies in the remaining rounds, the maximum profit Bidder 1 can make is:

$$g_m(x, \alpha) = \frac{1-\alpha}{m} + \frac{m-1}{m} \cdot f_{m-1}\left(\frac{mx}{m-1}\right). \tag{3.10}$$

• *Bidder 1 loses item a_1.*

If Bidder 1 loses the first item, then Bidder 1 makes no profit on a_1. Any bid $b_1^1 < b_2^1$ will lose the item. Since this is a first-price auction the adversary will pay b_2^1 if he wins regardless of the bid of Bidder 1. Thus Bidder 1 is indifferent between any bids less than b_2^1. Then, after the first round we again have a scaled version of $\mathcal{A}_{m-1}$ where the valuations of Bidder 1 are scaled by a factor of $\frac{m-1}{m}$. As the adversary won the first item for a price $b_2^1 = \alpha \cdot \frac{1}{m}$ his budget is now $x - \alpha \cdot \frac{1}{m}$, which corresponds to a budget of $B = \frac{m}{m-1} \cdot \left(x - \frac{\alpha}{m}\right) = \frac{mx-\alpha}{m-1}$ in the scaled auction $\mathcal{A}_{m-1}$. Therefore, given that the bidders play optimal strategies in the remaining rounds, the maximum profit Bidder 1 can make is:

$$h_m(x, \alpha) = \frac{m-1}{m} \cdot f_{m-1}\left(\frac{mx-\alpha}{m-1}\right). \tag{3.11}$$

Evidently, the best response of Bidder 1 to a bid $b_2^1 = \alpha \cdot \frac{1}{m}$ is given by the maximum of $g_m(x, \alpha)$ and $h_m(x, \alpha)$. Thus, the adversary should select α to minimize this maximum. Specifically,

$$\begin{aligned} f_m(x) &= \min_{0 \leq \alpha \leq 1} \max\Big(g_m(x, \alpha)\,,\, h_m(x, \alpha)\Big) \\ &= \min_{0 \leq \alpha \leq 1} \max\left(\frac{1-\alpha}{m} + \frac{m-1}{m} f_{m-1}\left(\frac{mx}{m-1}\right)\,,\, \frac{m-1}{m} f_{m-1}\left(\frac{mx-\alpha}{m-1}\right)\right). \end{aligned}$$

Thus, our goal is to prove that there exists a bid $b_2^1 = \tilde{\alpha} \cdot \frac{1}{m}$ by the adversary such that both $g_m(x, \tilde{\alpha})$ and $h_m(x, \tilde{\alpha})$ are at most $f(x) + \frac{1}{\sqrt{m}}$. This will ensure that the maximum guaranteed profit of Bidder 1 is $f_m(x) \leq f(x) + \frac{1}{\sqrt{m}}$ as required.

Our proof of this fact requires examination of three cases depending upon the magnitude of the budget of the adversary. In the first two cases, where the adversary has either a very low budget or a very high budget we can compute exactly the bids that Bidder 1 will make in her unique risk-free strategy. These two cases do not require the induction hypothesis (nor consideration of the functions $g_m(x, \tilde{\alpha})$ and $h_m(x, \tilde{\alpha})$) but constitute a part of our inductive step. The third case, where the adversary has an intermediate budget, is more difficult and represents one of the main technical contributions of this work.

`Low Budget Case:` $0 \leq x < \frac{1}{m^2}$.

When the adversary's budget x is less than $\frac{1}{m^2}$ then a risk-free strategy for Bidder 1 is to bid $b_1^t = x$ on item a_t, for every $t \in [m]$ (that is, Bidder 1 bids the entire budget of the adversary on each item). Bidder 1 will then win all of the items for a guaranteed profit of $1 - m \cdot x$. On the other hand if Bidder 1 bids an amount smaller than x on some item, then the adversary can win this item. Even if Bidder 1 wins all the remaining items her profit cannot exceed the total additive value of these remaining items which is $(m-1) \cdot \frac{1}{m} = 1 - \frac{1}{m}$. Because $x < \frac{1}{m^2}$, we have that $1 - \frac{1}{m} < 1 - m \cdot x$. As this is a first-price auction, Bidder 1 can never benefit by bidding strictly more than x on any item. If follows that the maximum profit the Bidder 1 can obtain is $f_m(x) = 1 - m \cdot x$.

It remains to show that $1 - mx \leq (1 - \sqrt{x})^2 + \frac{1}{\sqrt{m}}$ in this low budget case where $0 \leq x < \frac{1}{m^2}$. We prove this statement by partitioning the interval $[0, \frac{1}{m^2})$ at the two points $\frac{1}{1.4m^2}$ and $\frac{1}{1.1m^2}$ into a collection of three sub-intervals $\mathcal{I} = \{[0, \frac{1}{1.4m^2}), [\frac{1}{1.4m^2}, \frac{1}{1.1m^2}), [\frac{1}{1.1m^2}, \frac{1}{m^2})\}$. We can then verify separately in each sub-interval that when x falls inside this interval, we have $1 - mx \leq (1 - \sqrt{x})^2 + \frac{1}{\sqrt{m}}$.

For any sub-interval in this collection, let $c(m)$ and $d(m)$ be the endpoints of the sub-interval. When $c(m) \leq x < d(m)$, we have

$$\begin{aligned}
(1 - \sqrt{x})^2 + \frac{1}{\sqrt{m}} &= 1 + x - 2\sqrt{x} + \frac{1}{\sqrt{m}} \\
&\geq 1 + c(m) - 2\sqrt{d(m)} + \frac{1}{\sqrt{m}} \\
&\geq 1 \geq 1 - mx.
\end{aligned}$$

Here the first inequality arises because $x \geq c(m)$ and $x < d(m)$; the second inequality applies when $m = 3$ and $[c(m), d(m)) \in \mathcal{I}$. Now it is easy to verify that since $c(m) - 2\sqrt{d(m)} + \frac{1}{\sqrt{m}} = 0$ has no real roots greater than 3 for every $[c(m), d(m)) \in \mathcal{I}$, we have $1 + c(m) - 2\sqrt{d(m)} + \frac{1}{\sqrt{m}} > 1$ for all $m > 3$. Thus we have $1 - mx \leq (1 - \sqrt{x})^2 + \frac{1}{\sqrt{m}}$ for all values of x in the interval $[0, \frac{1}{m^2})$.

`High Budget Case:` $\frac{m-1}{m} < x \leq 1$.

Suppose that the budget x of the adversary is between $\frac{m-1}{m}$ and 1. Then a risk-free strategy for Bidder 1 is to bid $b_1^t = \frac{x}{m}$ on item a_t for every $t \in [m]$. The guaranteed profit of this strategy is $\frac{1-x}{m}$ because Bidder 1 will win exactly one item using this strategy. To see this, observe that if the adversary wins the first item at the price $\frac{x}{m}$ then his scaled budget in the subsequent subgame is exactly $\frac{m}{m-1} \cdot (x - \frac{x}{m}) = x$. If the adversary loses the first item then his scaled budget in the subsequent subgame is $\frac{m}{m-1} \cdot x > \frac{m}{m-1} \cdot \frac{m-1}{m} = 1$; so the adversary will win all the remaining items and hence Bidder 1 cannot win more than one item. But iterating this argument we see that if the adversary wins all of the first $m-1$ items then his (unscaled) budget for the final item is just $\frac{x}{m}$ and so Bidder 1 will win the final item for a profit of $\frac{1-x}{m}$ as required. It is easy to see that no bidding strategy for Bidder 1 guarantees a higher profit: with lower bids, Bidder 2 wins each item and ends up with a higher budget in the subsequent subgames; with higher bids, Bidder 1 wins a single item for smaller profit.

It remains to show that $f_m(x) = \frac{1-x}{m} \leq (1-\sqrt{x})^2 + \frac{1}{\sqrt{m}}$ when $\frac{m-1}{m} < x \leq 1$. We have

$$\begin{aligned}
(1-\sqrt{x})^2 + \frac{1}{\sqrt{m}} &= 1 + x - 2\sqrt{x} + \frac{1}{\sqrt{m}} \\
&\geq 1 + \frac{m-1}{m} - 2 + \frac{1}{\sqrt{m}} \\
&= \frac{1}{\sqrt{m}} - \frac{1}{m} \\
&\geq \frac{1}{m^2} \geq \frac{1-x}{m}.
\end{aligned}$$

Above, the first inequality holds because $\frac{m-1}{m} < x \leq 1$; the second inequality holds when $m \geq 3$; the third inequality holds since $x > \frac{m-1}{m}$. Thus we have $\frac{1-x}{m} \leq (1-\sqrt{x})^2 + \frac{1}{\sqrt{m}}$ when $\frac{m-1}{m} < x \leq 1$.

`Intermediate Budget Case:` $\frac{1}{m^2} \leq x \leq \frac{m-1}{m}$.

Recall that by the induction hypothesis $f_{m-1}(x) \leq f(x) + \frac{1}{\sqrt{m-1}}$. Our goal now is to prove that $f_m(x) = \min_{0 \leq \alpha \leq 1} \max\big(g_m(x,\alpha)\,,\, h_m(x,\alpha)\big) \leq f(x) + \frac{1}{\sqrt{m}}$ when $\frac{1}{m^2} \leq x \leq \frac{m-1}{m}$. Rather than calculate $f_m(x)$ exactly, our approach is to find a feasible choice $\tilde{\alpha}$ for the adversary that ensures that both $g_m(x, \tilde{\alpha})$ and $h_m(x, \tilde{\alpha})$ are at most $f(x) + \frac{1}{\sqrt{m}}$. To do this, we begin by investigating the properties of the functions $g_m(x, \alpha)$ and $h_m(x, \alpha)$. Using these properties, we find a candidate choice $\tilde{\alpha}$ which we first prove is feasible and second prove gives the desired upper bound.

Let's start by showing that $g_m(x, \alpha)$ and $h_m(x, \alpha)$ are both monotonic functions. To see this, observe that, for any fixed m, since the valuation function is additive and the space of available strategies for the adversary is constrained only by his budget, any strategy that is available to adversary with budget $\bar{x} < x$ is also available when his budget is x. Hence, the function f_m is non-increasing in x. Therefore $g_m(x, \alpha)$ is non-increasing in α and $h_m(x, \alpha)$ is non-decreasing in α, for any fixed x.

Now the minimum choice the adversary can make for α is zero. So suppose the adversary bids $b_2^1 = \alpha \frac{1}{m}$ for the item a_1 with $\alpha = 0$. Then clearly $g_m(x, 0) = \frac{1}{m} + \frac{m-1}{m} f_{m-1}\left(\frac{mx}{m-1}\right)$ and $h_m(x, 0) = \frac{m-1}{m} f_{m-1}\left(\frac{mx}{m-1}\right)$. Consequently, $g_m(x, 0) \geq h_m(x, 0)$.

On the other hand, consider the maximum choice the adversary can make for α. We denote this value by α_{max}. We have two cases.

- $x \geq \frac{1}{m}$

  Then the adversary may set $\alpha = 1$ and bid $\frac{1}{m}$ on the first item. In this case, both $g_m(x, 1)$ and $h_m(x, 1)$ are well defined, and we have $g_m(x, 1) = \frac{m-1}{m} f_{m-1}\left(\frac{mx}{m-1}\right)$ and $h_m(x, 1) = \frac{m-1}{m} f_{m-1}\left(\frac{mx-1}{m-1}\right)$. Because f_{m-1} is non-increasing, we have that $g_m(x, 1) \leq h_m(x, 1)$.

- $x < \frac{1}{m}$

  Now, by the budget constraint, the maximum possible value of α is mx. We want to show that $g_m(x, mx) \leq h_m(x, mx)$. To see this, suppose the adversary bids x on the first item (corresponding to the choice $\alpha = mx$) and loses. Bidder 1 then makes a profit of $\frac{1}{m} - x$ on the first item. The adversary can subsequently play the following strategy: bid x on every item until he wins an item. Of course, Bidder 1 will then win the remaining items for free after the adversary wins one item, because the budget of the adversary has then fallen to 0. Now, if Bidder 1's risk-free strategy is to win all of the items at price at least x, her (absolute) profit in this subgame is at most $\left(\frac{m-1}{m} - (m-1)x\right)$. If instead the adversary wins the k^{th} item, where $1 \leq k \leq m-1$, then Bidder 1's profit is at most $\frac{m-2}{m} - (k-1)x$, which is maximized at $k = 1$ with value $\frac{m-2}{m}$. In both cases, Bidder 1's total profit on all m items is either at most $1-mx$ or at most $\frac{m-1}{m} - x$. But these are both at most the profit Bidder 1 gets (namely, $\frac{m-1}{m}$) if she gives up the first item at price x and wins the remaining $m-1$ items for free. Thus $g_m(x, mx) \leq h_m(x, mx)$.

Set α_{max} to be the highest possible value of α for which both $g_m(x, \alpha)$ and $h_m(x, \alpha)$ are well-defined for all x. Therefore $\alpha_{max} = \min(1, mx)$. We have shown that $g_m(x, 0) \geq h_m(x, 0)$ and $g_m(x, \alpha_{max}) \leq h_m(x, \alpha_{max})$. Then, because $g_m(x, \alpha)$ is non-increasing in α and $h_m(x, \alpha)$ is non-decreasing in α for fixed x, our upper bound of $\max(g_m(x, \alpha), h_m(x, \alpha))$ is minimized at any bid $\bar{\alpha} \cdot \frac{1}{m}$ such that $0 \leq \bar{\alpha} \leq \alpha_{max}$ and $g_m(x, \bar{\alpha}) = h_m(x, \bar{\alpha})$. This is also precisely equal to a risk-free bid $\alpha^* \cdot \frac{1}{m}$ placed by Bidder 1 on the first item, since from her perspective, if the adversary plays a best response then she gets the *minimum* of $g_m(x, \alpha^*)$ and $h_m(x, \alpha^*)$, and this minimum is maximized when they are equal.

We now use the above observations to establish an upper bound on the highest guaranteed profit of a risk-free strategy. For an appropriately chosen bid $\tilde{\alpha}\frac{1}{m}$, we prove that both $g_m(x, \tilde{\alpha})$ and $h_m(x, \tilde{\alpha})$ are well-defined for all $x \in [\frac{1}{m^2}, \frac{m-1}{m}]$. We then prove that both these values are at most $f(x) + \frac{1}{\sqrt{m}}$. The facts rely on the four technical claims below.

The first two claims show that $\tilde{\alpha} = 1 - 2m(1-\sqrt{x}) + 2\sqrt{m(m-1)}(1-\sqrt{x})$. is a feasible choice for $\tilde{\alpha}$: specifically, $0 \leq \tilde{\alpha} \leq \alpha_{max}$. Let $\tilde{\alpha} = 1 - 2m(1-\sqrt{x}) + 2\sqrt{m(m-1)}(1-\sqrt{x})$. We show the following.

Claim 3.3.3. *For any $x \in [\frac{1}{m^2}, \frac{m-1}{m}]$, $0 \leq \tilde{\alpha}$.*

Proof. To prove that $\tilde{\alpha}$ is non-negative, for $x \in [\frac{1}{m^2}, \frac{m-1}{m}]$, we require that

$$1 - 2m(1-\sqrt{x}) + 2\sqrt{m(m-1)}(1-\sqrt{x}) \geq 0$$

Equivalently, we want to show that

$$2 \cdot (1 - \sqrt{x}) \cdot \left(m - \sqrt{m(m-1)}\right) \leq 1.$$

Since $x \geq \frac{1}{m^2}$, we have

$$\begin{aligned}
2 \cdot (1 - \sqrt{x}) \cdot \left(m - \sqrt{m(m-1)}\right) &\leq 2 \cdot \left(1 - \sqrt{\frac{1}{m^2}}\right) \cdot \left(m - \sqrt{m(m-1)}\right) \\
&= 2 \cdot \left(\frac{m-1}{m}\right) \cdot \left(m - \sqrt{m(m-1)}\right) \\
&= 2 \cdot \left((m-1) - (m-1)\sqrt{\frac{m-1}{m}}\right) \\
&= 2 \cdot (m-1) \cdot \left(1 - \sqrt{\frac{m-1}{m}}\right).
\end{aligned}$$

Thus we must show that $2(m-1)\left(1 - \sqrt{\frac{m-1}{m}}\right) \leq 1$. Equivalently we require that

$$\begin{aligned}
\frac{m-1}{m} &\geq \left(1 - \frac{1}{2 \cdot (m-1)}\right)^2 \\
&= 1 - \frac{1}{m-1} + \frac{1}{4 \cdot (m-1)^2} \\
&= \frac{m-2}{m-1} + \frac{1}{4 \cdot (m-1)^2}
\end{aligned}$$

To prove this, observe that $\frac{m-1}{m} - \frac{m-2}{m-1} = \frac{1}{m \cdot (m-1)}$, for all $m \geq 2$. Therefore

$$\begin{aligned}
\frac{m-1}{m} &= \frac{m-2}{m-1} + \frac{1}{m \cdot (m-1)} \\
&\geq \frac{m-2}{m-1} + \frac{1}{4 \cdot (m-1)^2}
\end{aligned}$$

where the inequality holds for all $m \geq 2$. This proves that $\tilde{\alpha} \geq 0$. □

Claim 3.3.4. *For any* $x \in [\frac{1}{m^2}, \frac{m-1}{m}]$, $\tilde{\alpha} \leq \alpha_{max}$.

Proof. We partition the proof into two cases.

(i) Assume that $x \geq \frac{1}{m}$. It follows that $\alpha_{max} = 1$ and so we must show

$$1 - 2m(1-\sqrt{x}) + 2\sqrt{m(m-1)}(1-\sqrt{x}) \leq 1$$

Equivalently, we require

$$2 \cdot (1-\sqrt{x}) \cdot \left(m - \sqrt{m(m-1)}\right) \geq 0$$

But this inequality holds because, by assumption, we have $x \in [\frac{1}{m}, \frac{m-1}{m}]$. In particular, for $m \geq 2$, we have both $(1-\sqrt{x}) > 0$ and $(2m - 2\sqrt{m(m-1)}) > 0$.

(ii) Assume that $x < \frac{1}{m}$. It now follows that $\alpha_{max} = m \cdot x$ and so we must show

$$1 - 2m(1-\sqrt{x}) + 2\sqrt{m(m-1)}(1-\sqrt{x}) \leq m \cdot x$$

Equivalently, we require

$$2 \cdot (1-\sqrt{x}) \cdot \left(m - \sqrt{m(m-1)}\right) + m \cdot x - 1 \geq 0.$$

To see this holds, observe that

$$\begin{aligned}
&2 \cdot (1-\sqrt{x}) \cdot \left(m - \sqrt{m(m-1)}\right) + m \cdot x - 1 \\
&\quad = \quad mx - 2\sqrt{x} \cdot \left(m - \sqrt{m(m-1)}\right) + \left(2 \cdot (m - \sqrt{m(m-1)}) - 1\right) \\
&\quad = \quad mx - 2\sqrt{mx} \cdot \left(\sqrt{m} - \sqrt{m-1}\right) + \left(2m - 1 - 2 \cdot \sqrt{m(m-1)}\right) \\
&\quad = \quad \left(\sqrt{mx} - (\sqrt{m} - \sqrt{m-1})\right)^2 \\
&\quad \geq \quad 0
\end{aligned}$$

The claim follows. $\square$

So we have a feasible choice for $\tilde{\alpha}$. To complete the proof that $f_m(x) \leq f(x) + \frac{1}{\sqrt{m}}$, we now show that both $g_m(x, \tilde{\alpha})$ and $h_m(x, \tilde{\alpha})$ are at most $f(x) + \frac{1}{\sqrt{m}}$.

Claim 3.3.5. $g_m(x, \tilde{\alpha}) \leq f(x) + \frac{1}{\sqrt{m}}$.

Proof. By our induction hypothesis,

$$\begin{aligned}
g_m(x, \tilde{\alpha}) &= \frac{1-\tilde{\alpha}}{m} + \frac{m-1}{m} \cdot f_{m-1}\left(\frac{mx}{m-1}\right) \\
&\leq \frac{1-\tilde{\alpha}}{m} + \frac{m-1}{m} \cdot \left(f\left(\frac{mx}{m-1}\right) + \frac{1}{\sqrt{m-1}}\right)
\end{aligned}$$

Thus it suffices to show

$$\frac{1-\tilde{\alpha}}{m} + \frac{m-1}{m} \cdot \left(f\left(\frac{mx}{m-1}\right) + \frac{1}{\sqrt{m-1}}\right) \leq f(x) + \frac{1}{\sqrt{m}}$$

Equivalently, we require

$$f(x) + \frac{1}{\sqrt{m}} - \frac{1-\tilde{\alpha}}{m} - \frac{m-1}{m} \cdot \left(f\left(\frac{mx}{m-1}\right) + \frac{1}{\sqrt{m-1}}\right) \geq 0$$

To prove this, observe that

$$\begin{aligned}
f(x) &+ \frac{1}{\sqrt{m}} - \frac{1-\tilde{\alpha}}{m} - \frac{m-1}{m}\left(f\left(\frac{mx}{m-1}\right) + \frac{1}{\sqrt{m-1}}\right) \\
&= \left(1 + x - 2\sqrt{x}\right) - \frac{1-\tilde{\alpha}}{m} - \frac{m-1}{m}\left(1 + \frac{mx}{m-1} - 2\sqrt{\frac{mx}{m-1}}\right) \\
&\qquad + \frac{1}{\sqrt{m}} - \frac{\sqrt{m-1}}{m} \\
&= 1 + x - 2\sqrt{x} - \frac{1}{m} + \frac{\tilde{\alpha}}{m} - 1 + \frac{1}{m} - x + 2\sqrt{\frac{(m-1)x}{m}} \\
&\qquad + \frac{1}{\sqrt{m}} - \frac{\sqrt{m-1}}{m} \\
&= -2\sqrt{x} + \frac{\tilde{\alpha}}{m} + 2\sqrt{\frac{(m-1)x}{m}} + \frac{1}{\sqrt{m}}\left(1 - \sqrt{\frac{m-1}{m}}\right)
\end{aligned}$$

By the definition of $\tilde{\alpha}$ we then have that

$$\begin{aligned}
&f(x) + \frac{1}{\sqrt{m}} - \frac{1-\tilde{\alpha}}{m} - \frac{m-1}{m}\left(f\left(\frac{mx}{m-1}\right) + \frac{1}{\sqrt{m-1}}\right) \\
&\quad = -2\sqrt{x} + \left(\frac{1}{m} - 2 + 2\sqrt{x} + \frac{2\sqrt{m-1}}{\sqrt{m}}\left(1-\sqrt{x}\right)\right) \\
&\qquad\qquad + 2\sqrt{\frac{(m-1)x}{m}} + \frac{1}{\sqrt{m}}\left(1 - \sqrt{\frac{m-1}{m}}\right) \\
&\quad = \frac{1}{m} - 2 + 2\sqrt{\frac{m-1}{m}}(1-\sqrt{x}) + 2\sqrt{\frac{m-1}{m}}\sqrt{x} + \frac{1}{\sqrt{m}}\left(1 - \sqrt{\frac{m-1}{m}}\right) \\
&\quad = \frac{1}{m} - 2\cdot\left(1 - \sqrt{\frac{m-1}{m}}\right) + \frac{1}{\sqrt{m}}\left(1 - \sqrt{\frac{m-1}{m}}\right) \\
&\quad = \frac{1}{m} + \left(\frac{1}{\sqrt{m}} - 2\right)\cdot\left(1 - \sqrt{\frac{m-1}{m}}\right)
\end{aligned}$$

Now set $q(m) = \frac{1}{m} + \left(\frac{1}{\sqrt{m}} - 2\right)\cdot\left(1 - \sqrt{\frac{m-1}{m}}\right)$. Clearly, to show that $q(m)$ is non-negative, it suffices to show that $(\sqrt{m} + \sqrt{m-1})\cdot m\cdot q(m)$ is non-negative. To do this, note that

$$\begin{aligned}
&(\sqrt{m} + \sqrt{m-1})\cdot m\cdot q(m) \\
&\quad = (\sqrt{m} + \sqrt{m-1})\cdot\left(1 + (1-2\sqrt{m})\cdot(\sqrt{m} - \sqrt{m-1})\right) \\
&\quad = (\sqrt{m} + \sqrt{m-1}) - (2\sqrt{m}-1)\cdot(\sqrt{m} - \sqrt{m-1})\cdot(\sqrt{m} + \sqrt{m-1}) \\
&\quad = (\sqrt{m} + \sqrt{m-1}) - (2\sqrt{m}-1)\cdot 1 \\
&\quad = 1 + \sqrt{m-1} - \sqrt{m} \\
&\quad \geq 0
\end{aligned}$$

Here the final inequality holds for $m \geq 2$. Thus $g_m(x, \tilde{\alpha}) \leq f(x) + \frac{1}{\sqrt{m}}$. □

Claim 3.3.6. $h_m(x, \tilde{\alpha}) \leq f(x) + \frac{1}{\sqrt{m}}$.

Proof. We now prove $h_m(x, \tilde{\alpha}) \leq f(x) + \frac{1}{\sqrt{m}}$. By our induction hypothesis,

$$\begin{aligned} h_m(x, \tilde{\alpha}) &= \frac{m-1}{m} \cdot f_{m-1}\left(\frac{mx - \tilde{\alpha}}{m-1}\right) \\ &\leq \frac{m-1}{m} \cdot \left(f\left(\frac{mx - \tilde{\alpha}}{m-1}\right) + \frac{1}{\sqrt{m-1}}\right) \end{aligned}$$

Hence, we want to show that

$$\frac{m-1}{m} \cdot \left(f\left(\frac{mx - \tilde{\alpha}}{m-1}\right) + \frac{1}{\sqrt{m-1}}\right) \leq f(x) + \frac{1}{\sqrt{m}}$$

Equivalently, we require

$$f(x) - \frac{m-1}{m} \cdot f\left(\frac{mx - \tilde{\alpha}}{m-1}\right) + \frac{1}{\sqrt{m}} - \frac{m-1}{m}\left(\frac{1}{\sqrt{m-1}}\right) \geq 0$$

To begin, let's show that the first two terms are equal; that is, $f(x) - \frac{m-1}{m} \cdot f\left(\frac{mx-\tilde{\alpha}}{m-1}\right) = 0$.

$$\begin{aligned} f(x) - \frac{m-1}{m} \cdot f\left(\frac{mx - \tilde{\alpha}}{m-1}\right) &= \left(1 + x - 2\sqrt{x}\right) - \frac{m-1}{m} \cdot f\left(\frac{mx - \tilde{\alpha}}{m-1}\right) \\ &= \left(1 + x - 2\sqrt{x}\right) - \frac{m-1}{m}\left(1 + \frac{mx - \tilde{\alpha}}{m-1} - 2\sqrt{\frac{mx - \tilde{\alpha}}{m-1}}\right) \\ &= \left(1 + x - 2\sqrt{x}\right) - \left(1 - \frac{1}{m} + x - \frac{\tilde{\alpha}}{m} - 2 \cdot \frac{m-1}{m} \cdot \sqrt{\frac{mx - \tilde{\alpha}}{m-1}}\right) \\ &= -2\sqrt{x} + \frac{1}{m} + \frac{\tilde{\alpha}}{m} + 2 \cdot \frac{m-1}{m} \cdot \sqrt{\frac{mx - \tilde{\alpha}}{m-1}} \end{aligned}$$

To prove the RHS is indeed 0 we must show that

$$2\sqrt{x} - \frac{1 + \tilde{\alpha}}{m} = 2 \cdot \frac{m-1}{m} \cdot \sqrt{\frac{mx - \tilde{\alpha}}{m-1}}$$

First observe that both sides are nonnegative, because $\sqrt{x} > \frac{1}{m}$ and $\tilde{\alpha} < \alpha_{max}$. Next multiply each side by m and take the square. This leads to

$$
\begin{aligned}
&\big(2m \cdot \sqrt{x} - (1+\tilde{\alpha})\big)^2 - 4 \cdot (m-1) \cdot (mx - \tilde{\alpha}) \\
&= \big((2m \cdot \sqrt{x} - 1) - \tilde{\alpha})\big)^2 - 4 \cdot (m-1) \cdot (mx - \tilde{\alpha}) \\
&= \big(\tilde{\alpha}^2 - (4m \cdot \sqrt{x} - 2)\tilde{\alpha} + (4m^2 x - 4m\sqrt{x} + 1)\big) \\
&\qquad + (4(m-1) \cdot \tilde{\alpha} - 4(m-1)mx) \\
&= \tilde{\alpha}^2 + (2 - 4m \cdot \sqrt{x} + 4(m-1)) \cdot \tilde{\alpha} + (4m^2 x - 4m\sqrt{x} + 1 - 4(m-1)mx) \\
&= \big(\tilde{\alpha}^2 - (4m \cdot \sqrt{x} - 2)\tilde{\alpha} + (4m^2 x - 4m\sqrt{x} + 1)\big) \\
&\qquad + (4(m-1) \cdot \tilde{\alpha} - 4(m-1)mx) \\
&= \tilde{\alpha}^2 + (4(m-1) - 4m \cdot \sqrt{x} + 2) \cdot \tilde{\alpha} + (4m^2 x - 4m\sqrt{x} + 1 - 4(m-1)mx) \\
&= \tilde{\alpha}^2 + (4m - 4m \cdot \sqrt{x} - 2) \cdot \tilde{\alpha} + (4mx - 4m\sqrt{x} + 1) \\
&= 0
\end{aligned}
$$

The final equality does follow as $z = \tilde{\alpha}$ is indeed a solution to the quadratic equation $z^2 + (4m - 4m\sqrt{x} - 2)z + (4mx - 4m\sqrt{x} + 1) = 0$. Putting this all together then gives

$$
\begin{aligned}
f(x) - \frac{m-1}{m} \cdot f\left(\frac{mx - \tilde{\alpha}}{m-1}\right) + \frac{1}{\sqrt{m}} - \frac{m-1}{m}\left(\frac{1}{\sqrt{m-1}}\right)& \\
&= 0 + \frac{1}{\sqrt{m}} - \frac{m-1}{m}\left(\frac{1}{\sqrt{m-1}}\right) \\
&= \frac{1}{\sqrt{m}} - \frac{\sqrt{m-1}}{m} \\
&= \frac{1}{\sqrt{m}}\left(1 - \sqrt{\frac{m-1}{m}}\right) \\
&\geq 0
\end{aligned}
$$

The claim follows. $\square$

So we have the following upper bounds on g_m and h_m: $g_m(x, \tilde{\alpha}) \leq f(x) + \frac{1}{\sqrt{m}}$ and $h_m(x, \tilde{\alpha}) \leq f(x) + \frac{1}{\sqrt{m}}$. Since we also have $f_m(x) \leq \max(g_m(x, \tilde{\alpha}), h_m(x, \tilde{\alpha}))$, we have $f_m(x) \leq f(x) + \frac{1}{\sqrt{m}}$ when $\frac{1}{m^2} \leq x \leq \frac{m-1}{m}$. With this third case (intermediate budget) completed so is the proof of Theorem 3.3.2. □

3.3.3 Risk-Free Bidding in Simultaneous Auctions

In this section we consider risk-free bidding in a simultaneous auction. For a budgeted adversary in a simultaneous auction, the analogue of budget-constrained bidding is that the *sum* of the adversary's bids on the items is at most the budget B. Intuitively, a budgeted adversary is weaker in a simultaneous auction than in a sequential auction. This is because, in a sequential auction, an adversary has the option to "overbid" on an item but suffers no consequence *if he loses the item*. The issue then is whether or not the resultant broader range of strategies available to an adversary in a sequential auction makes it provably more powerful than the corresponding adversary in a simultaneous auction. We show this in the following theorems. We begin by analyzing the second-price case.

Theorem 3.3.7. *The two-player simultaneous second-price auction with a normalized XOS valuation function and an adversary with normalized budget $B \in (0, 1)$ has a risk-free strategy for Bidder 1 that guarantees a profit of at least $(1 - B)$.*

Proof. We prove this theorem by using the following strategy for Bidder 1. Bidder 1 bids truthfully according to the additive function γ^* defined in Section 3.3.1 – that is, for each item $a_j \in I$, she bids $\gamma_j^* = \gamma^*(\{a_j\})$. We show that, if Bidder 1 plays according to this strategy, then for any feasible strategy of the adversary, Bidder 1 makes a profit of at least $(1 - B)$. In particular, we consider the adversary's best response to this strategy.

Suppose the adversary's best response is to make a sequence $b_1, \ldots, b_m$ of bids on the respective items. Clearly, in a best response the adversary will not bid more than $\gamma^*(\{a_j\})^+$. Let $I_1 \subseteq I$ and $I_2 \subseteq I$ be the set of items allocated to Bidder 1 and Bidder 2 respectively. Then Bidder 1's profit is given by

$$\begin{aligned}
\pi_1 = v(I_1) - \sum_{j:a_j \in I_1} b_j &\geq \sum_{j:a_j \in I_1} (\gamma_j^* - b_j) \\
&= \sum_{j:a_j \in I_1} (\gamma_j^* - b_j) + \sum_{j:a_j \in I_2} (\gamma_j^* - b_j) \\
&= \sum_{j:a_j \in I} \gamma_j^* - \sum_{j:a_j \in I} b_j \\
&= 1 - \sum_{j:a_j \in I} b_j \geq 1 - B
\end{aligned}$$

Here the first inequality follows by definition of an XOS function; the second equality arises because the adversary bids $b_j = \gamma_j^*$ on each item $j \in I_2$ that he wins; the fourth equality follows by definition of γ^* and the fact the auction is normalized; the second inequality follows from the budget constraint. □

Observe that $(1-\sqrt{B})^2 < 1 - B$, for all $B \in (0,1)$. Ergo, the risk-free profitability of Bidder 1 is strictly greater in a second-price simultaneous auction than in a second-price sequential auction. Conversely the adversary is strictly weaker in the second-price simultaneous auction.

Next let's consider the case of *first-price* simultaneous auctions. Note that the proof of Theorem 3.3.7 was via the use of a pure strategy for Bidder 1. For first-price simultaneous auctions it is not possible to rely on a pure strategy to beat the profit bound of $(1-\sqrt{B})^2$; to do so, the bidder must use a randomized strategy. To verify this, the following claim shows that in the *uniform additive* simultaneous auction, no deterministic strategy for Bidder 1 can guarantee a profit that is asymptotically greater than $(1-\sqrt{B})^2$.

Claim 3.3.8. *For any pure strategy of Bidder 1, there exists a strategy for the adversary that (asymptotically) restricts Bidder 1's profit to* $(1-\sqrt{B})^2$.

Proof. Let $b_1^1, \ldots, b_1^m$ be Bidder 1's bids on the m items. We may assume without loss of generality that $b_1^i \leq b_1^j$ whenever $i < j$. Bidder 2's strategy is to win k^* items, where

$k^* = \max\{k : \sum_{i=1}^{k} b_1^i < B\}$. Let p^* be the price of the lowest-indexed item that Bidder 1 wins, i.e., $p^* = b_1^{k^*+1}$. Since we maximize over all possible k, we know that the adversary cannot afford to win the entire set $\{a_1, \ldots, a_{k^*+1}\}$. Let P be the total price paid by the adversary. This implies that P is more than $B - p^*$, otherwise the adversary could have won another item. So we have $P > B - p^*$.

On the other hand, the total price paid by the adversary is at most

$$P = b_1^1 + \ldots + b_1^{k^*} \leq k^* \cdot b_1^{k^*+1} = k^* p^*.$$

Combining the inequalities, we have $k^* > \frac{B-p^*}{p^*}$. So the number of items that the adversary wins is at least $\frac{B-p^*}{p^*}$; thus, Bidder 1 wins at most $m - \frac{B-p^*}{p^*} = m - \frac{B}{p^*} + 1$ items. Since the items are ordered by Bidder 1's bids, her price for each of these items is at least p^*. Consequently, because her valuation function is uniform additive, her profit is at most

$$\pi_1 \leq (m - \frac{B}{p^*} + 1) \cdot (\frac{1}{m} - p^*)$$

It is easy to verify that this is maximized when $p^* = \sqrt{\frac{B}{m(m+1)}}$, and that the maximum value is $(1 - \sqrt{B})^2 + O(\frac{1}{\sqrt{m}})$. $\square$

In the other direction, it is also true that no deterministic strategy for the adversary can guarantee that Bidder 1 makes a profit that is less than $(1 - B)$.

Claim 3.3.9. *For any pure strategy of the adversary, there exists a strategy for Bidder 1 that guarantees a profit of $(1 - B)$ for any XOS bidder.*

Proof. By the bidding constraint, the sum of the adversary's bids is at most B. Let $b_2^1, \ldots, b_2^m$ be the bids made by the adversary on the items. Bidder 1 simply bids b_2^{i+} on each item a_i

as long as b_2^i is less than γ_i^* and wins this set of items. Bidder 1's profit is then

$$\begin{aligned}
v_1(I_1) - \sum_{i:b_2^i<\gamma_i^*} b_2^i &\geq \sum_{i:b_2^i<\gamma_i^*} (\gamma_i^* - b_2^i) \\
&\geq \sum_{i:b_2^i<\gamma_i^*} (\gamma_i^* - b_2^i) + \sum_{i:b_2^i\geq\gamma_i^*} (\gamma_i^* - b_2^i) \\
&= \sum_{i:a_i\in I} (\gamma_i^* - b_2^i) \\
&= 1 - \sum_{i:a_i\in I} b_2^i \\
&\geq 1 - B
\end{aligned}$$

Here the first inequality arises as Bidder 1 has an XOS valuation; the second equality follows by definition of γ^* and the fact the auction is normalized; the last inequality follows from the budget constraint. □

Due to this asymmetry in pure strategies, simultaneous first-price auctions against an adversary have no equilibrium in deterministic strategies, and we must introduce randomization to improve the lower bound. In fact, there is a randomized strategy for Bidder 1 that guarantees an expected profit of at least $\frac{1}{2}(1-B)^2$ when her valuation function is XOS. The function $\frac{1}{2}(1-B)^2$ is greater than $(1-\sqrt{B})^2$ for $B > 3-2\sqrt{2}$, which is approximately 0.17, so the upper bound from the first-price sequential auction case does not apply to first-price simultaneous auctions.

Theorem 3.3.10. *The two-player simultaneous first-price auction with a normalized XOS valuation function and an adversary with normalized budget $B \in (0,1)$ has a (randomized) risk-free strategy for Bidder 1 that guarantees a profit of at least $\frac{(1-B)^2}{2}$ in expectation.*

Proof. Bidder 1 selects m independent random variables X_i, each drawn from the uniform distribution $U(0,1)$, and bids $X_i \cdot \gamma_i^*$ on the item a_i. Our goal is to show that no strategy of the adversary can prevent Bidder 1 from making a profit of $\frac{1}{2}(1-B)^2$ in expectation. For the adversary, we may limit our attention to bids that are at most γ_i^*. So we may parameterize the bids of the adversary by a vector of ratios $\boldsymbol{b} = (b_1, \ldots, b_m) \in (0,1)^m$ such that $\sum_{i=1}^{m} b_i \cdot \gamma^*(i) \leq B$. Let $S = \{a_i \in I | X_i > b_i\}$ be the set of items that Bidder 1 wins and let $\pi(\boldsymbol{b})$ be the random variable representing Bidder 1's utility when the adversary bids $\boldsymbol{b}$. We then have the following.

$$\begin{aligned}
\mathbb{E}[\pi(\boldsymbol{b})] &= \mathbb{E}_{X_1 \ldots X_m}\left[v(S) - \sum_{i: a_i \in S} \gamma_i^* \cdot X_i\right] \\
&\geq \mathbb{E}_{X_1 \ldots X_m}\left[\sum_{i: a_i \in S} \gamma_i^* \cdot (1 - X_i)\right] \\
&= \mathbb{E}_{X_1 \ldots X_m}\left[\sum_{i: a_i \in I} \gamma_i^* \cdot (1 - X_i) \cdot \mathbb{1}_{[b_i < X_i]}\right] \\
&= \sum_{i: a_i \in I} \gamma_i^* \cdot \mathbb{E}_{X_i}\left[(1 - X_i) \cdot \mathbb{1}_{[b_i < X_i]}\right] \\
&= \sum_{i: a_i \in I} \gamma_i^* \cdot \frac{1}{2} \cdot (1 - b_i)^2
\end{aligned}$$

Again, here the first inequality follows from the definition of γ^*; the second equality is due to linearity of expectation; the final equality holds because X_i is uniformly distributed.

The adversary of course seeks to find a strategy to minimize $\mathbb{E}[\pi(\boldsymbol{b})]$ for any fixed valuation function v. For any valuation function v, we denote this minimum value by $\pi^*(v)$. The inequalities above imply that $\pi^*(v) \geq \pi^*(\gamma^*) = \min_{\boldsymbol{b}} \frac{1}{2} \sum_{i: a_i \in I} \gamma_i^* \cdot (1 - b_i)^2$. For the following analysis, we may assume without loss of generality that we only consider items a_i where $\gamma_i^* > 0$ (Bidder 1 will lose the remaining items at a bid of 0). Now, $\pi^*(v)$ is lower-bounded by the optimal value of the following quadratic program.

$$\texttt{(Adversarial QP)} \qquad \begin{array}{rrcll}
\min & \frac{1}{2} \cdot \sum_{i: a_i \in I} \gamma_i^* & \cdot & (1 - b_i)^2 & \\
\text{s. t.} & b_i & \leq & 1 & \forall i \in [m] \\
& -b_i & \leq & 0 & \forall i \in [m] \\
& \sum_{i=1}^{m} b_i \cdot \gamma^*(i) & \leq & B &
\end{array}$$

The Lagrangian of this problem is

$$\mathcal{L}(\vec{b}, \vec{\lambda}) = \frac{1}{2} \sum_{i \in I} \gamma_i^* \cdot (1 - b_i)^2 + \lambda_{2m+1} \cdot \left(\sum_{i=1}^{m} b_i \cdot \gamma_i^* - B \right)$$
$$+ \sum_{i=1}^{m} \lambda_i \cdot (b_i - 1) - \sum_{i=1}^{m} \lambda_{m+i} \cdot b_i$$

This is differentiable w.r.t. each of the b_i, so we can compute these partial derivatives to be

$$\frac{\partial \mathcal{L}}{\partial b_i} = \gamma_i^* \cdot (b_i - 1) + \lambda_{2m+1} \cdot \gamma_i^* + \lambda_i - \lambda_{m+i}.$$

The dual objective function $g(\vec{\lambda}) = \inf_{\vec{b}} \mathcal{L}(\vec{b}, \vec{\lambda})$ can be computed by setting $\frac{\partial \mathcal{L}}{\partial b_i} = 0$. So:

$$\begin{aligned}
g(\vec{\lambda}) =& \frac{1}{2} \sum_{i=1}^{m} \gamma_i^* \left(\lambda_{2m+1} + \left(\frac{\lambda_i - \lambda_{m+i}}{\gamma_i^*} \right) \right)^2 \\
&+ \lambda_{2m+1} \left(\sum_{i=1}^{m} \left(1 - \lambda_{2m+1} - \left(\frac{\lambda_i - \lambda_{m+i}}{\gamma_i^*} \right) \right) \gamma_i^* - B \right) \\
&- \sum_{i=1}^{m} \lambda_i \left(\lambda_{2m+1} + \left(\frac{\lambda_i - \lambda_{m+i}}{\gamma_i^*} \right) \right) \\
&- \sum_{i=1}^{m} \lambda_{m+i} \left(1 - \lambda_{2m+1} - \left(\frac{\lambda_i - \lambda_{m+i}}{\gamma_i^*} \right) \right)
\end{aligned}$$

The constraints on the dual are simply $\vec{\lambda} \geq 0$. So setting $\lambda_i = 0 \; \forall \; i \in \{1 \ldots 2m\}; \lambda_{2m+1} = (1 - B)$ is feasible. Let $\vec{\lambda'}$ denote this vector. The dual objective for this feasible input is:

$$\begin{aligned}
g(\vec{\lambda'}) &= \frac{1}{2} \sum_{i=1}^{m} \gamma_i^* \left(\lambda_{2m+1} \right)^2 \\
&\qquad + \lambda_{2m+1} \left(\sum_{i=1}^{m} \left(1 - \lambda_{2m+1} \right) \gamma_i^* - B \right) && (\lambda_i = 0) \\
&= \frac{1}{2} \left(\lambda_{2m+1} \right)^2 + \lambda_{2m+1} \left(\left(1 - \lambda_{2m+1} \right) - B \right) && \left(\text{as } \sum_i \gamma_i^* = 1\right) \\
&= \frac{1}{2} \left(1 - B \right)^2 && (\text{substituting for } \lambda_{2m+1})
\end{aligned}$$

This dual solution lower bounds the primal minimization program, and we have

$$\min_{\boldsymbol{b}} \mathbb{E}[\pi(\boldsymbol{b})] \geq \frac{1}{2}(1-B)^2$$

as desired.

$\square$

Finally, we show that an analogue of Theorem 3.3.7 does *not* hold for simultaneous first-price auctions. Specifically, we prove that there exists a *randomized* strategy for the adversary that gives an upper bound on the profitability that is strictly smaller than $(1 - B)$, showing that for simultaneous auctions, the adversary's power is greater in the first-price case than the second-price case.

Theorem 3.3.11. *In first-price simultaneous auctions with XOS valuations, the adversary has a (randomized) strategy that restricts the risk-free profit of Bidder 1 to strictly less than* $(1 - B)$ *in expectation.*

Proof. We prove this claim by considering the uniform additive simultaneous auction on m items, where m is even. The adversary chooses a subset $S \subseteq I$ of the items, with $|S| = \frac{m}{2}$, uniformly from the subsets of this size. He then bids $\frac{w_1(B)}{m}$ on each element in S, and $\frac{w_2(B)}{m}$ on each element not in S, where w_1 and w_2 are as follows.

$$w_1(B) = \begin{cases} 2B & \text{if } 0 < B < \frac{1}{4} \\ \frac{1}{3} + \frac{2B}{3} & \text{if } \frac{1}{4} \leq B < 1 \end{cases}$$

$$w_2(B) = \begin{cases} 0 & \text{if } 0 < B < \frac{1}{4} \\ \frac{4B}{3} - \frac{1}{3} & \text{if } \frac{1}{4} \leq B < 1 \end{cases}$$

It is easily shown that this strategy is feasible for Bidder 2, and that Bidder 1's best response is to bid $\frac{w_1(B)}{m}$ on every item and win all the items. Bidder 1's profit is then

$$\pi_1^* = \begin{cases} 1-2B & \text{if } 0 < B < \frac{1}{4} \\ \frac{2}{3}(1-B) & \text{if } \frac{1}{4} \leq B < 1 \end{cases}$$

which is strictly smaller than $(1-B)$ when $0 < B < 1$. $\square$

We remark that the strategies used in proving Theorems 3.3.7 and 3.3.10 require no knowledge of the adversary's budget. Bidder 1 can implement them based solely on her own valuation function so these profit guarantees are extremely robust.

Thus, indeed, the adversary is weaker in a simultaneous auction than in the corresponding sequential auction (for example, the bound of Theorem 3.3.10 is larger than that of Theorem 3.3.1 for $B > 0.18$). In addition, unlike for sequential auctions, the power of the adversary differs in a simultaneous auction depending on whether a first-price or second-price mechanism is used: the adversary is stronger in a first-price auction. Finally, unlike the sequential case, it is essential to introduce randomization to obtain non-trivial bounds in the first-price simultaneous setting.

3.4 Bounds for Subadditive Valuation Functions

In this section we make a return to sequential auctions. We study the risk-free profitability of Bidder 1 when her valuation function is subadditive. Since there exist subadditive functions that are not XOS, the simple strategy from Section 3.3.1 is no longer guaranteed to work. Indeed, we present in Section 3.4.2 a class of examples of subadditive valuations whose risk-free profitability is strictly less than $f(B) = (1-\sqrt{B})^2$ for an adversary with budget $B < \frac{1}{4}$. The relationship between the classes of XOS functions and subadditive functions was explored by Bhawalkar and Roughgarden [19], via the class of *β-fractionally subadditive* valuation functions.

Proposition 3.4.1. *[19] Every subadditive valuation is* $\ln m$*-fractionally subadditive.*

It follows from this proposition that there exists a bid vector r satisfying $\sum_{j:a_j\in I} r_j = v(I)$ and $\sum_{j:a_j\in S} r_j \leq \ln m \cdot v(S)$ for each subset $S \subseteq I$. In [19], the authors also provide an example showing that this is tight. Consequently, if Bidder 1 plays a strategy analogous to the strategy from Section 3.3.1 on this example, using the bid vector r in place of the additive function γ^*, then any strict subset S of I that Bidder 1 wins is only guaranteed to have value $O(\frac{1}{\ln m})$ and, potentially, this guarantees a profit of only $O(\frac{1}{\ln m})$. This relationship indicates an inherent difficulty in showing a non-trivial lower bound on the profitability of subadditive valuations. However, we make progress on an important special case, namely subadditive valuations on identical items. Here, every subset S of I such that $|S| = k$ where $0 \leq k \leq m$ has the same value that we denote by $v(k)$. The earlier assumptions still hold, so $v(0) = 0$, and v is monotone. In Section 3.4.1, we present a strategy for Bidder 1 that gives a new lower bound on the profitability. Then, in Section 3.4.2, we prove that this lower bound is tight when the budget B is in $(0, \frac{1}{4})$. Moreover, the lower bound is tight at every B of the form $(\frac{k}{k+1})^2$ for any positive integer k, and we conjecture that this tightness extends to all $B \in (0, 1)$.

3.4.1 The Subadditive Lower Bound with Identical Items

We obtain our lower bound on the profitability of Bidder 1 with a simple strategy: Bidder 1 chooses a constant price $\tilde{p}$ and a target allocation $\tilde{q}$ in advance, and bids $\tilde{p}$ on every item, stopping when she wins $\tilde{q}$ items. We will need the following claim.

Claim 3.4.2. *For any set* $S \subseteq I$*, where* $|S| = q$*,* $v(S) \geq \frac{v(I)}{\lceil \frac{m}{q} \rceil}$*.*

Proof. Let S be a subset of I of size q. We want to show that $v(S) = v(q) \geq \frac{v(I)}{\lceil \frac{m}{q} \rceil}$. Consider any partition of the set I into $\ell = \lceil \frac{m}{q} \rceil$ sets $S_1, \ldots, S_\ell$, where $S_1 = S$, and each of the first $\lfloor \frac{m}{q} \rfloor$ sets $S_1, \ldots, S_{\lfloor \frac{m}{q} \rfloor}$ has size q.

We have two cases. If m is a multiple of q, then $\lceil \frac{m}{q} \rceil = \lfloor \frac{m}{q} \rfloor = \ell$. By subadditivity, we have

$$v(I) \;=\; v(S_1 \cup \ldots \cup S_\ell) \;\leq\; v(S_1) + \ldots + v(S_\ell) \;=\; \ell \cdot v(q)$$

Thus we have $v(q) \geq \frac{v(I)}{\ell} = \frac{v(I)}{\lceil \frac{m}{q} \rceil}$.

Now if m is not a multiple of q, then $\lceil \frac{m}{q} \rceil - \lfloor \frac{m}{q} \rfloor = 1$. Then $|S_\ell| = m - q \cdot \lfloor \frac{m}{q} \rfloor < q$, and we have

$$\begin{aligned}
v(I) &= v(S_1 \cup \ldots \cup S_\ell) \\
&\leq v(S_1) + \ldots + v(S_\ell) \\
&= (\ell - 1) \cdot v(q) + v(m - q \cdot \lfloor \frac{m}{q} \rfloor) \\
&\leq (\ell - 1) \cdot v(q) + v(q) \\
&= \ell \cdot v(q).
\end{aligned}$$

Here the first in equality follows by subadditivity and the second by monotonicity. So we have $v(q) \geq \frac{v(I)}{\ell} = \frac{v(I)}{\lceil \frac{m}{q} \rceil}$ as required. □

Now, for an appropriate choice of $\tilde{p}$ and $\tilde{q}$, we show that Bidder 1 can guarantee a profit of at least $t^*(B) - O(\frac{1}{m})$, where

$$t^*(B) = \max_{k \in \mathbb{Z}_{\geq 1}} t_k(B).$$

Interestingly, $t_k(B) = \frac{1}{k+1} - \frac{B}{k}$ is the tangent to our earlier lower bound of $f(B) = (1 - \sqrt{B})^2$ at $B = (\frac{k}{k+1})^2$ (see Figure 3.2). Denote by SI_m the subadditive valuations on m identical items. We show the following lower bound on the profitability in the subadditive case with identical items.

Theorem 3.4.3. $\mathcal{P}(SI_m, B) \geq t^*(B) - O(\frac{1}{m})$.

Proof. Consider a normalized instance of the subadditive sequential auction on identical items where the adversary has normalized budget B. First, for any choice of target allo-

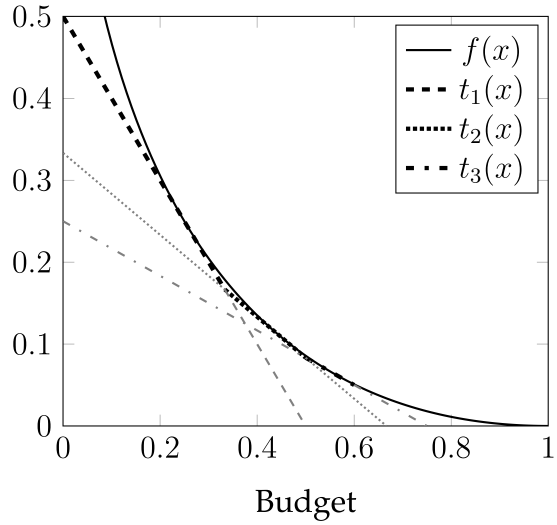

Figure 3.2: Plot of the Functions $f(x)$, $t_1(x)$, $t_2(x)$, and $t_3(x)$: $t^*(x)$ *is the piecewise linear function shown by the bold line segments*

cation q, we want to find the minimum price-per-item p_q so that the adversary's budget is insufficient to stop Bidder 1 from winning q items. In other words, we require the smallest possible price p_q so that the adversary cannot win more than $m - q$ items. That is, we need to let $p_q(m-q+1) = B+\delta$ for some negligibly small δ. We may ignore the negligible additive δ term and let $p_q(m - q + 1) = B$, from which we obtain $p_q = \frac{B}{m-q+1}$. So, for any q such that $0 \leq q \leq m$, in order to win exactly q items, Bidder 1 bids a price $p_q = \frac{B}{m-q+1}$ on every item and stops when she wins q items. By Claim 3.4.2, the value of this set in the normalized auction is at least $\frac{1}{\lceil \frac{m}{q} \rceil}$.

Finally, we want to choose a target allocation $\tilde{q}$ that maximizes Bidder 1's profit under this constant-price strategy. For the purpose of our analysis, for a fixed m, any optimal choice of q must be of the form $q = \lceil \frac{m}{k} \rceil$ for some positive integer $k \geq 2$, which is the smallest choice of q that has value at least $\frac{1}{k}$ according to our lower bound in Claim 3.4.2. Thus for any positive integer k, the profit obtained from winning exactly $q = \lceil \frac{m}{k} \rceil$ items with our constant-price strategy is at least

$$\begin{aligned}
v(q) - q \cdot p_q &\geq \frac{1}{\lceil \frac{m}{q} \rceil} - \frac{qB}{m - q + 1} \\
&= \frac{1}{\lceil \frac{m}{\lceil \frac{m}{k} \rceil} \rceil} - \frac{\lceil \frac{m}{k} \rceil B}{m - \lceil \frac{m}{k} \rceil + 1} \\
&\geq \frac{1}{\lceil \frac{m}{(\frac{m}{k})} \rceil} - \frac{(\frac{m}{k} + 1)B}{m - (\frac{m}{k} + 1) + 1}.
\end{aligned}$$

Therefore

$$\begin{aligned}
v(q) - q \cdot p_q &\geq \frac{1}{k} - B\left(\frac{m+k}{m(k-1)}\right) \\
&= \frac{1}{k} - \frac{B}{(k-1)} - \frac{1}{m} \cdot \left(\frac{Bk}{k-1}\right) \\
&= t_{k-1}(B) - \frac{1}{m} \cdot \left(\frac{Bk}{k-1}\right)
\end{aligned}$$

where $t_k(B) = \frac{1}{k+1} - \frac{B}{k}$ is the tangent to $f(B) = (1 - \sqrt{B})^2$ at $B = (\frac{k}{k+1})^2$. The second term in the above expression is an additive $O(\frac{1}{m})$ factor that we subtract from our lower bound, but this factor goes to 0 as the number of items increases. Bidder 1's strategy is then to choose $\tilde{q}$ to maximize this value, which is equivalent to maximizing over the sequence of tangents. This completes the proof of the lower bound. □

3.4.2 The Subadditive Upper Bound with Identical Items

In this section we present a matching upper bound for the range $0 < B < \frac{1}{4}$. For any B in this range, we have that $\max_{k \in \mathbb{Z}_{\geq 1}} t_k(B) = t_1(B)$, so our lower bound is just $t_1(B) = \frac{1}{2} - B$. We match this lower bound by constructing a sequence of valuation functions and corresponding strategies for the adversary that give us the following theorem.

Theorem 3.4.4. *$\mathcal{P}(SI_m, B) \leq t^*(B) + O(\frac{1}{\sqrt{m}})$ when $B \in (0, \frac{1}{4})$ and m is larger than some constant m_0 that depends only on B.*

Proof. For every budget $B = x$ with $0 < x < \frac{1}{4}$, let $\sigma(x) = \frac{8x}{1-4x}$. It is easy to see that $\sigma(x)$ is a (well-defined) positive real number for every x in $(0, \frac{1}{4})$. For every x in this interval and every positive integer m, we define the normalized identical-item auction on m items $\mathcal{S}_{x,m}$ in the following manner. We set $v(1) = \frac{1}{2+\sigma(x)}$, and for every $i \in \{2, \ldots, m-1\}$ we set $v(i) = \frac{1}{2+\sigma(x)} + \frac{(i-1)\sigma(x)}{d(2+\sigma(x))}$, where $d = m-2$. Finally, we set $v(m) = 1$. In other words, we first construct an unnormalized valuation function by setting the marginal value of obtaining a first item and an m^{th} item to 1, and the marginal value of getting an i^{th} item to $\frac{\sigma(x)}{d}$ for $2 \leq i \leq m-1$. Since the total value of the m items is $2+\sigma(x)$, we divide the value of every set by this factor to obtain the above normalized instance $\mathcal{S}_{x,m}$.

We have an important reason for choosing $\sigma(x) = \frac{8x}{1-4x}$: with this choice of $\sigma(x)$, the expression $\frac{\sigma(x)}{2+\sigma(x)}$, which shows up repeatedly in the following analysis, simplifies to the expression $4x$. Put differently, for any fixed x, we consider valuation functions where the marginal value of getting a first item or an m^{th} is $\frac{1-4x}{2}$, and the total marginal value of getting an additional $m-2$ items having already been allocated one item is $4x$. The reader may have observed that with this choice of $\sigma(x)$, this valuation function is not always subadditive. This is true: the function is subadditive if and only if $\sigma(x) \leq d$. When $\sigma(x) \leq d$, for any integers i, j with $i \geq 1$, $j \geq 1$, and $i + j \leq m$, it is easy to see that $v(i+j) \leq \frac{2}{2+\sigma(x)} + \frac{(i+j-2)\sigma(x)}{d(2+\sigma(x))}$, whereas $v(i) \geq \frac{1}{2+\sigma(x)} + \frac{(i-1)\sigma(x)}{d(2+\sigma(x))}$ and $v(j) \geq \frac{1}{2+\sigma(x)} + \frac{(j-1)\sigma(x)}{d(2+\sigma(x))}$, so we have $v(i+j) \leq v(i) + v(j)$ and the function is subadditive. For the other direction, if $\sigma(x) > d$ then $v(2) = \frac{1+\frac{\sigma(x)}{d}}{2+\sigma(x)} > \frac{2}{2+\sigma(x)} = 2v(1)$, so the function is not subadditive. Since $d = m-2$, we require $m \geq \sigma(x) + 2$ for subadditivity.

To find our upper bound, we will consider auction instances $\mathcal{S}_{x,m}$ where the adversary can play according to the following strategy. Initially, the adversary makes a bid of 0 on every item until Bidder 1 wins an item. Then, after Bidder 1 wins her first item, if the adversary has not yet won an item he makes a bid of $\frac{1+\sigma(x)}{d(2+\sigma(x))}$ on every item until Bidder 1 loses an item. At this point, each of the two bidders has won at least one item, so any allocation in the remaining subgame gives at least 1 item and at most $m-1$ items to Bidder 1. Consequently, the remaining subgame is simply a scaled instance of a uniform additive auction $\mathcal{A}_{m'}$ for some $m' \leq d$, and from this point on Bidder 2 simply plays an optimal strategy for the uniform additive auction. The adversary's bid of $\frac{1+\sigma(x)}{d(2+\sigma(x))}$ on each item induces a lower bound on the size of the auction instances that we consider: we require $x \geq \frac{1+\sigma(x)}{d(2+\sigma(x))}$ for this bid to be feasible. Substituting $d = m-2$ into the above inequality, we have $m \geq \frac{1+\sigma(x)}{x(2+\sigma(x))} + 2$. Combining this with the lower bound for subadditivity, for any fixed x we let $\mathcal{L}(x) = \max(\sigma(x)+2, \frac{1+\sigma(x)}{x(2+\sigma(x))} + 2)$, and we only consider auction instances $\mathcal{S}_{x,m}$ on $m \geq \mathcal{L}(x)$ items, since these are the only subadditive instances in which the above strategy is feasible. Observe that for any fixed $x \in (0, \frac{1}{4})$, $\mathcal{L}(x)$ is a constant. For any budget in this range, the following theorem provides an upper bound on the risk-free profit in a subadditive, identical-item sequential auction. This upper bound applies to the general subadditive case, showing that the profitability of subadditive functions is strictly less than that of additive, submodular or XOS functions.

Consider the auction instance $\mathcal{S}_{x,m}$ on m identical items. Our goal is to show that if Bidder 2 plays according to the above strategy in this auction, then for every strategy of Bidder 1, her risk free profit matches the lower bound. In particular, Bidder 1's best response to this strategy has a payoff of at most $t_1(x) + O(\frac{1}{\sqrt{m}})$. Suppose Bidder 1 plays a strategy such that j_1 is the index of the first item won by Bidder 1 and $j_2 > j_1$ is the index of the first item won by Bidder 2 after Bidder 1 wins an item (where necessary, we will consider separately the case where Bidder 2 does not win any item after Bidder 1 wins item a_{j_1}). We will show that for any feasible choice of j_1 and j_2, Bidder 1's payoff in any strategy that results in this choice is at most $t_1(x) + O(\frac{1}{\sqrt{m}}) = \frac{1}{2} - x + O(\frac{1}{\sqrt{m}})$. We have the following cases for j_1 and j_2.

`Case 1`: Bidder 2 wins the first item.

If Bidder 2 wins the first item, then $j_1 \geq 2$. Immediately after Bidder 1 wins item a_{j_1} (for a price of 0), the remaining subgame is the uniform additive auction $\mathcal{A}_{m'}$ where $m' = m - j_1$ and the remaining unscaled value is $\frac{\sigma(x)}{2+\sigma(x)} - \frac{(j_1-2)\sigma(x)}{d(2+\sigma(x))}$. Since Bidder 2 has only made bids equal to 0, the remaining unscaled budget is x. We first consider the case where Bidder 1 wins the second item, so $j_1 = 2$. There are $d = m-2$ items remaining, so the remaining value is simply $\frac{\sigma(x)}{2+\sigma(x)}$. Consequently the remaining subgame is an instance of the uniform additive auction $\mathcal{A}_d$ with scaled budget $\frac{2+\sigma(x)}{\sigma(x)} \cdot x$ when the value is scaled to 1. Since Bidder 1 makes a profit of $\frac{1}{2+\sigma(x)}$ on the item a_{j_1}, the risk-free profitability of any subgame with $j_1 = 2$ is

$$\begin{aligned}
\frac{1}{2+\sigma(x)} + \frac{\sigma(x)}{2+\sigma(x)} f_d\left(\frac{2+\sigma(x)}{\sigma(x)} \cdot x\right) &= \frac{1}{2+\sigma(x)} + 4x f_d\left(\frac{1}{4}\right) \\
&\leq \frac{1}{2+\sigma(x)} + 4x\left(f\left(\frac{1}{4}\right) + \frac{1}{\sqrt{d}}\right) \\
&\qquad \text{by Theorem 3.3.2} \\
&= \frac{1}{2+\sigma(x)} + 4x\left(\frac{1}{4} + \frac{1}{\sqrt{d}}\right) \\
&= \left(\frac{1}{2+\sigma(x)} + x\right) + \frac{4x}{\sqrt{d}} \\
&= \left(\frac{1}{2} - x\right) + \frac{4x}{\sqrt{d}} \\
&= t_1(x) + O\left(\frac{1}{\sqrt{m}}\right)
\end{aligned}$$

Observe that for any strategy where $j_1 > 2$, Bidder 1 is simply giving up items to the adversary at a price of 0, and the remaining subgame is a uniform additive auction where the unscaled budget of the adversary remains the same but the total number of items (and total value) decreases. Since $g_m(x, 0) \geq h_m(x, 0)$ (see Section 3.3.2), the risk-free profitability of any strategy where $j_1 > 2$ is upper bounded by the risk-free profitability of a strategy with $j_1 = 2$, and this profitability is at most $t^*(x) + O\left(\frac{1}{\sqrt{m}}\right)$ for an adversary with budget x.

`Case 2`: Bidder 1 wins the first $m-1$ items.

It remains to consider the cases where $j_1 = 1$, where Bidder 1 wins the first item. First, suppose Bidder 1 wins the first item (for a price of 0), and then wins the next $m-2$ items. After Bidder 1 wins the first item, since Bidder 2 has not won an item his strategy is to bid $\frac{1+\sigma(x)}{d(2+\sigma(x))}$ on each of the remaining items, so the total price paid by Bidder 1 for the next

$d = m - 2$ items is at least $\frac{1+\sigma(x)}{2+\sigma(x)}$. Hence Bidder 1's profit is at most $1 - \frac{1+\sigma(x)}{2+\sigma(x)} = \frac{1}{2+\sigma(x)} = \frac{1}{2} - 2x$, which is at most $t_1(x)$.

`Case 3:` $j_1 = 1$ and $j_2 \leq m - 1$.

The only remaining case is where Bidder 1 wins the first item for a price of 0, and Bidder 2 wins an item in $\{a_2, \ldots, a_{m-1}\}$. After Bidder 1 wins the first item, Bidder 2 bids $\frac{1+\sigma(x)}{d(2+\sigma(x))}$ on every item until he wins the item a_{j_2}. Bidder 1 pays a total of $(j_2 - 2)\frac{1+\sigma(x)}{d(2+\sigma(x))}$ for the first $j_2 - 1$ items, which have total value $\frac{1}{2+\sigma(x)} + (j_2 - 2)\frac{\sigma(x)}{d(2+\sigma(x))}$. Since the price of each of these items (which is $\frac{1+\sigma(x)}{d(2+\sigma(x))}$) is greater than the marginal value of each of these items (which is $\frac{\sigma(x)}{d(2+\sigma(x))}$), by choosing to increase j_2 Bidder 1 is winning items from the adversary at a price that is higher than their marginal value in the remaining uniform additive auction. Since $g_m(x, 1) \leq h_m(x, 1)$ (see Section 3.3.2), the risk-free profitability of any strategy where $j_2 > 2$ is lower than the risk-free profitability of a strategy where $j_2 = 2$. Consequently, we will fix $j_2 = 2$ and show that Bidder 1 makes a profit of at most $t^*(x) + O\left(\frac{1}{\sqrt{m}}\right)$ with this choice.

Now, after Bidder 2 wins item a_2, the remaining subgame is an instance of the uniform additive auction on $m - 2$ items where the total unscaled value is $\frac{\sigma(x)}{2+\sigma(x)}$ which is equal to $4x$, and the total unscaled budget is $x - \frac{1+\sigma(x)}{d(2+\sigma(x))}$, which simplifies to $\frac{dx-2x-\frac{1}{2}}{d}$. When the budget is scaled by $\frac{1}{4x}$, it becomes $\frac{dx-2x-\frac{1}{2}}{4dx}$. Finally, we also need to add the profit made by Bidder 1 from item a_1, which is $\frac{1}{2+\sigma(x)}$. Putting all this together, Bidder 1's profit is at most

$$\begin{aligned}
&\frac{1}{2+\sigma(x)} + 4x f_{m-2}\left(\frac{dx-2x-\frac{1}{2}}{4dx}\right) \\
&\qquad = \frac{1-4x}{2} + 4x\left[1 + \frac{dx-2x-\frac{1}{2}}{4dx} - 2\sqrt{\frac{dx-2x-\frac{1}{2}}{4dx}} + \frac{1}{\sqrt{m-2}}\right] \\
&\qquad\qquad\qquad\qquad \text{by Theorem 3.3.2} \\
&\qquad = \frac{1}{2} - 2x + 4x + \frac{dx-2x-\frac{1}{2}}{d} - 2\sqrt{\frac{4x(dx-2x-\frac{1}{2})}{d}} + \frac{4x}{\sqrt{m-2}} \\
&\qquad = \frac{1}{2} + 3x - \frac{(2x+\frac{1}{2})}{d} - 2\sqrt{4x^2 - \frac{4x(2x+\frac{1}{2})}{d}} + \frac{4x}{\sqrt{m-2}} \\
&\qquad = \frac{1}{2} + 3x - \frac{(2x+\frac{1}{2})}{d} - 2\sqrt{4x^2\left[1 - \frac{2x+\frac{1}{2}}{dx}\right]} + \frac{4x}{\sqrt{m-2}} \\
&\qquad = \frac{1}{2} + 3x - 4x\sqrt{1 - \frac{2x+\frac{1}{2}}{dx}} - \frac{(2x+\frac{1}{2})}{d} + \frac{4x}{\sqrt{m-2}}
\end{aligned}$$

Since $m \geq \mathcal{L}(x)$, $d \geq \frac{1+\sigma(x)}{x(2+\sigma(x))}$, so $dx \geq \frac{1+\sigma(x)}{2+\sigma(x)} = 2x + \frac{1}{2}$, so we have that $0 \leq 1 - \frac{2x+\frac{1}{2}}{dx} \leq 1$. Then

$$\begin{aligned}
&\frac{1}{2} + 3x - 4x\sqrt{1 - \frac{2x+\frac{1}{2}}{dx}} - \frac{(2x+\frac{1}{2})}{d} + \frac{4x}{\sqrt{m-2}} \\
&\qquad \leq \frac{1}{2} + 3x - 4x\left(1 - \frac{2x+\frac{1}{2}}{dx}\right) - \frac{(2x+\frac{1}{2})}{d} + \frac{4x}{\sqrt{m-2}} \\
&\qquad = \frac{1}{2} - x + \left[\frac{6x+\frac{3}{2}}{d} + \frac{4x}{\sqrt{m-2}}\right] \\
&\qquad = t_1(x) + O(\frac{1}{\sqrt{m}}).
\end{aligned}$$

This completes the proof of Theorem 3.4.4. □

An important consequence of the above result is that the lower bound for XOS valuations does not hold for subadditive valuations. This differentiates the class of subadditive valuations from the additive, submodular and XOS classes in that Bidder 1 can no longer guarantee a profit of $(1-\sqrt{B})^2$ when her valuation function is subadditive.

Completing the above proof concludes Part I of this thesis, where we analyze multi-item auctions. To summarize, in Chapter 2 we presented a refined analysis of equilibria of sequential auctions, and showed that the declining price anomaly does not always hold in their focal subgame-perfect equilibria. In Chapter 3, we analyzed the risk-free profitability of sequential and simultaneous auctions for some well-studied classes of valuation functions. Part II of this thesis focuses on a similar setting where agents have a combinatorial valuation function over a set of items, but studies the problem of fairly dividing the items among the agents.

Part II

Fair Division

We now shift our attention away from mechanisms that allocate items to agents, and towards the question of whether it is possible that every agent receives its due share. The predominant objectives of research in this area are to study the existence of allocations that achieve one of two broad goals: *fairness* or *efficiency*. At a high level, the fairness goal is to ensure that each agent receives its due share of the items, and the efficiency goal is to distribute the items in a way that satisfies some constraint on the aggregate utility achieved by all of the agents. Typically, the variety of problems that arise out of the question of how to define and achieve these goals – primarily, the first goal – are now collectively called the *fair division* problem.

But what does it mean for an allocation to be *fair*? This is often the first question that occurs to someone that come across this problem, and for good reason: we desire a definition of fairness that we can all agree mirrors the concept of fairness in the real-world, but one that also allows us to provide meaningful guarantees on the existence and computability of fair allocations.

The formal origin of fair division dates back to the 1940s, to the work of Banach, Knaster and Steinhaus [69]. Like much of the early literature on this problem, they focused on the *divisible* setting, where a single heterogeneous divisible item (conventionally, a cake) is to be fairly shared among a set of n agents with varying preferences over its pieces. With two agents, the folkloric *cut-and-choose* protocol ("I cut, you choose") achieves many intuitive fairness guarantees, but it is not easy to see how this process may be generalized. Banach, Knaster and Steinhaus [69] devised the *last diminisher* procedure to obtain a fair allocation. Their fairness objective was *proportionality*: an allocation is proportional if every agent is allocated a bundle (or piece of cake) of value at least $\frac{1}{n}$ of its total value for the grand bundle (entire cake). The subsequent pursuit of proportional cake divisions in a variety of settings led to the creation of popular algorithmic paradigms for cake-cutting such as the moving-knives procedures [31, 70].

The latter half of the 20^{th} century witnessed the creation of precise mathematical definitions for various fairness notions. In the following decades, *envy-freeness* (Gamow and

Stern [39], Foley [35]) emerged as the canonical fairness solution in economics. As the name might suggest, an allocation is envy-free if no agent prefers the piece of cake (or bundle of items) allocated to another agent to its own bundle (according to its own valuation function applied to both pieces). When the agents' valuation functions are of the forms most commonly considered in economic theory, such as *additive* valuations, envy-freeness *implies* proportionality and is thus a stronger fairness guarantee. This is because under these assumptions, for any agent and any partition of the cake into n pieces, some piece is worth at least $\frac{1}{n}$ of the value of the whole cake to that agent. Another classical fairness measure is *equitability*, where all agents should receive an allocation of the same value. When the items are divisible, Alon [4] showed for additive continuous valuation functions that allocations exist that simultaneously satisfy proportionality, equitability and envy-freeness.

In recent years, the research directions of the fair division community have undergone two major shifts. The first of these is towards the study of the *indivisible items* setting, where m items are to be integrally allocated amongst the n agents. Unfortunately, a simple example consisting of two agents and one indivisible item demonstrates that envy-freeness, proportionality and equitability are unachievable in this setting. Consequently, most research efforts have been directed at achieving approximate or relaxed fairness guarantees. The most prominent of these relaxations is the *EFk* guarantee. An allocation is *envy-free up to k items*, or EFk, if no agent envies another agent's bundle provided some k items are removed from that bundle. The EF1 guarantee is particularly notable, as EF1 allocations exist and can be computed in polynomial time if the valuation functions satisfy the mild assumption of monotonicity [55]. Another fairness guarantee that has attained recent popularity is the *maximin share guarantee*, introduced by Budish [22] and inspired by the cut-and-choose protocol. Suppose an agent partitions the items into n bundles and then receives its lowest-valued bundle. The corresponding value that the agent obtains by selecting its optimal partition is called its *maximin share*. The fairness objective then is to find an allocation where every agent receives a bundle of value at least some constant fraction of its maximin share. A large body of research produced over the last decade aims to achieve these guarantees or approximations thereof (see e.g. [52, 40]).

A parallel line of research attempts to circumvent the apparent unattainability of envy-freeness via the use of *payments*. In now classical work, Svensson [71], Maskin [57], and Tadenuma and Thomson [72] studied the indivisible item setting and asked if it is always possible to achieve an envy-free allocation simply by introducing a small quantity of a divisible item, akin to money, alongside the indivisible items. Their positive results were mirrored in follow-up work by Alkan et al. [3], Aragones [5], Klijn [50] and Haake et al. [43] which showed for various settings the existence of an envy-free *allocation with subsidy*. However, all of the above papers considered the restricted case where the number of items, m, is at most the number of agents n (or where the items were grouped into n fixed bundles). It was only recently that Halpern and Shah [44] extended these results to the general m-item setting. Specifically, they considered the setting in which n agents desire to partition a set of m items among themselves, and, without loss of generality, the value of each item is at most 1. Their main result is a proof of the existence of an *envy-freeable* allocation, i.e. one that can be made envy-free with payments, in every instance. They characterized the envy-freeable allocations in terms of the structure of the *envy graph*, whose nodes are the agents and whose arc weights represent the envies between pairs of agents. They then studied the problem of minimizing the amount of subsidy that is sufficient to guarantee envy-freeness. It is easy to see this *minimum subsidy* can be at least $n-1$ for all envy-freeable allocations. Indeed, consider the case of a single item which each agent values at exactly one dollar; evidently, every agent that does not receive the item must be compensated with a dollar. Based on computational analysis of over $100,000$ synthetic instances, Halpern and Shah [44] conjectured that with *additive* valuations, an envy-freeable allocation with a subsidy of at most $n-1$ dollars always exists. In addition, they conjectured that an allocation exists that is both envy-freeable *and* EF1 when the valuations are additive. In Chapter 4 we prove these conjectures.

The second major shift is an increased focus on economic efficiency. A popular notion of efficiency is *Pareto efficiency*, in which no agent's allocation can be improved without making some other agent worse off. A classical result of Varian [76] shows that in the divisible setting there always exists an allocation that is both envy-free and Pareto efficient. A large body of follow-up research aims to achieve Pareto efficiency alongside a variety of fairness guarantees. An alternate notion of efficiency arises when we maximize a *welfare function* that measures the aggregate utility of all agents. The most common welfare functions studied in the associated literature are the *utilitarian* social welfare (or simply the social welfare), which measures the sum of the agents' valuations, and the *Nash* social welfare, which measures the geometric mean of these valuations. The *price of fairness* of an instance, formally introduced by Caragiannis et al. [25], is the ratio of the welfare of an optimal allocation (without fairness constraints) to that of the best fair allocation. For the cake-cutting problem, Bei et al. [15] and Cohler et al. [30] study the problem of maximizing social welfare under proportionality and envy-freeness constraints. For the divisible setting, Bertsimas et al. [18] showed that the bounds of [25] are tight. Subsequent work on the price of fairness in the indivisible setting by Bei et al. [16] and Barman et al. [12] considers only the relaxed fairness guarantees (such as EF1 and $\frac{1}{2}$-MMS) that are always achievable in the indivisible setting. In Chapter 5, we study the problem of whether payments can be used to achieve envy-freeness and high welfare simultaneously.

Chapter 4

One Dollar Each

In classical work, Maskin [57] asked if envy-freeness can be achieved in the indivisible-items case via the addition of a single divisible good, namely *money*. He considered the case of a market with n agents and $m = n$ goods where each agent can be allocated *at most* one good, and has, without loss of generality, a value of at most one dollar for any specific good. Maskin [57] then showed that an envy-free allocation exists with the addition of $n-1$ dollars into the market. But what happens in the general setting where the number of agents and number of items may differ and where agents may be allocated more than one item? The purpose of this chapter is to understand this case. In this setting, Halpern and Shah [44] showed that an envy-freeable allocation always exists by characterizing the *envy-freeable* allocations in terms of the structure of the *envy graph* (see Section 4.2), whose nodes are the agents and whose arc weights represent the envies between pairs of agents. They then studied the problem of minimizing the amount of subsidy that is sufficient to guarantee envy-freeness. They proved that $m \cdot (n-1)$ dollars always suffice to support an envy-free allocation when the agents have *additive* valuation functions.

Based on the experimental analysis of over $100,000$ synthetic instances and over $3,000$ real-world instances of fair division, Halpern and Shah [44] conjecture that this upper bound can be improved to $n-1$ dollars. That is, for agents with additive valuations an envy-freeable allocation that requires a subsidy of at most $n-1$ dollars always exists. In addition, they conjecture that an allocation exists that is both envy-freeable (with perhaps a much larger subsidy) *and* EF1 for the fair division problem with additive valuations.

Conjecture 4.0.1. *[44] For additive valuations, there is an envy-freeable allocation that requires a total subsidy of at most $n - 1$ dollars.*

Conjecture 4.0.2. *[44] For additive valuations, there is an envy-freeable allocation that is EF1.*

In this chapter, we positively settle both Conjecture 4.0.1 and Conjecture 4.0.2: for additive valuation functions, precisely $n - 1$ dollars are sufficient to guarantee the existence of an envy-free allocation. In fact, our result is stronger in several ways. First, we present an algorithm that computes an allocation where the subsidy is not only at most $n - 1$ dollars in total, but also at most *one dollar* to each agent. Secondly, this allocation is also *envy-free up to one item (EF1)* – thus the same allocation settles the second conjecture from [44]. Thirdly, the allocation is *balanced*, that is, the cardinalities of the allocated bundles differ by at most one item. Furthermore, this envy-free allocation can be constructed in polynomial time. Formally, in Sections 4.3 and 4.4 we prove the following theorem.

Theorem 4.0.3. *For additive valuations there is an envy-freeable allocation where the subsidy to each agent is at most one dollar. (This allocation is also EF1, balanced, and can be computed in polynomial time.)*

It is easy to see that, when minimizing the total subsidy, at least one agent will not receive any payment. Thus Theorem 4.0.3 implies that the total subsidy required is indeed at most $n - 1$ dollars.

We also study the general case of monotone valuation functions. Requiring only the very mild assumption of monotonicity, we prove the perhaps surprising result that envy-free solutions still exist with a subsidy amount that is *independent* of the number of items m. Specifically, we prove that there is an envy-free allocation where each agent receives a subsidy of at most $2(n - 1)$ dollars, for a total subsidy of $O(n^2)$. Here the envy-free allocation can be constructed in polynomial time given a valuation oracle.

In Section 4.5 we consider the much more general setting where the agents have arbitrary monotone valuation functions. Analogously, without loss of generality, we may scale the valuations so that the marginal value of each item for any agent never exceeds one dollar. We prove the surprising result that envy-free solutions still exist with a subsidy of at most $2(n-1)$ dollars per agent – an amount that is again *independent* of the number of items m. Thus, the total subsidy required to ensure the existence of an envy-free allocation is at most $O(n^2)$. Note that the assumption of monotonicity is extremely mild and so the valuations the agents have for bundles of items may range from 0 to $\Omega(m)$ in quite an arbitrary manner. Consequently, it is somewhat remarkable that the total subsidy required to ensure the existence of an envy-free allocation is independent of m. In particular, when m is large the subsidy required is negligible in terms of m and thus, typically, also negligible in terms of the values of the allocated bundles. In this case, given a valuation oracle for each agent, the corresponding envy-free allocation and subsidies can be computed in polynomial time. In Section 4.5 we prove:

Theorem 4.0.4. *For monotonic valuations there is an envy-freeable allocation where the subsidy to each agent is at most $2(n-1)$ dollars. (Given a valuation oracle, this allocation can be computed in polynomial time.)*

In effect, our work implies that there is, in fact, a much stronger connection between the classical divisible goods (cake-cutting) setting and the indivisible goods setting than was previously known. While the classical guarantees (envy-freeness and proportionality) can be achieved with divisible goods, for the indivisible-goods setting much of the recent literature focuses on achieving weaker fairness properties. We show that by simply introducing a small subsidy that only depends on the number of agents, the much stronger classical guarantees *can* be achieved in the indivisible goods setting. Moreover, allocations that give these classical guarantees with a small bounded subsidy can be efficiently found.

4.1 The Fair Division with Subsidy Problem

There is a set $I = \{1, 2, \ldots, n\}$ of agents and a set $J = \{1, 2, \ldots, m\}$ of indivisible items. Each agent $i \in I$ has a valuation function v_i over the set of items. That is, for each bundle $S \subseteq J$ of items, agent i has value $v_i(S)$. The items are *goods*, so $v_i(S) \geq 0$ for any agent i and bundle $S \subseteq J$. We make the standard assumptions that the valuation functions are *monotonic*, that is, $v_i(S) \leq v_i(T)$ when $S \subseteq T$, and that $v_i(\emptyset) = 0$. An agent i and valuation function v_i are *additive* if, for each item $j \in J$, agent i has value $v_i(j) = v_i(\{j\})$, and for any collection $S \subseteq J$, agent i has value $v_i(S) = \sum_{j \in S} v_i(j)$. We denote the vector of valuation functions by $\boldsymbol{v} = (v_1, \ldots, v_n)$, and call $\boldsymbol{v}$ a *valuation profile*. Additionally, without loss of generality we scale each agent i's valuation function so that the maximum marginal value of any item j is at most 1 (that is, for each item j, $\max_{S \subseteq J \setminus \{j\}} v_i(S \cup \{j\}) - v_i(S) \leq 1$).[1] Specifically, for additive valuations, this implies $v_i(j) \leq 1$ for every agent i and item j.

An *allocation* is an ordered partition $\mathcal{A} = \{A_1, \ldots, A_n\}$ of the set of items into n bundles. Agent i receives the (possibly empty) bundle A_i in the allocation $\mathcal{A}$. The allocation $\mathcal{A}$ is *envy-free* if

$$v_i(A_i) \;\geq\; v_i(A_k) \qquad \forall i \in I, \forall k \in I.$$

That is, for any pair of agents i and k, agent i prefers its own bundle A_i over the bundle A_k. In the (envy-free) *fair division problem* the objective is to find an envy-free allocation of the items.

Unfortunately, this objective is generally impossible to satisfy. A natural relaxation of the objective arises by incorporating subsidies. Specifically, let $\boldsymbol{p} = (p_1, \ldots, p_n)$ be a non-negative subsidy vector, where agent i receives a payment $p_i \geq 0$. An *allocation with*

[1]Without scaling, our bound of one dollar to each agent becomes the maximum marginal value of an item.

payments $(\mathcal{A}, \boldsymbol{p})$ is then envy-free if

$$v_i(A_i) + p_i \geq v_i(A_k) + p_k \qquad \forall i \in I, \forall k \in I.$$

That is, each agent prefers its bundle plus payment over the bundle plus payment of every other agent. In the *fair division with subsidy problem* the objective is to find an envy-free allocation with payments whose total subsidy $\sum_{i \in I} p_i$ is minimized.

4.2 Envy-Freeability and the Envy Graph

For any fixed allocation $\mathcal{A}$, a payment vector $\boldsymbol{p}$ such that $\{\mathcal{A}, \boldsymbol{p}\}$ is envy-free does not always exist. To see this, consider an instance with a single item and agents $I = \{1, 2\}$ with values $v_1 < v_2$ for the item. Now take the fixed allocation where the item is given to agent 1. It follows that agent 2 must receive a payment of at least v_2 to eliminate its envy. But then, because $v_2 > v_1$, agent 1 is envious of the bundle plus payment allocated to agent 2. Thus, no payment vector can eliminate the envy of both agents for this allocation.

We call an allocation $\mathcal{A}$ *envy-freeable* if there exists a payment vector $\boldsymbol{p} = (p_1, \ldots, p_n)$ such that $\{\mathcal{A}, \boldsymbol{p}\}$ is envy-free. There is a nice graphical characterization for the envy-freeability of an allocation $\mathcal{A}$. The *envy graph*, denoted $G_{\mathcal{A}}$, for an allocation $\mathcal{A}$ is a complete directed graph with vertex set I. For any pair of agents $i, k \in I$ the weight of arc (i, k) in $G_{\mathcal{A}}$ is the envy agent i has for agent k under the allocation $\mathcal{A}$, that is, $w_{\mathcal{A}}(i, k) = v_i(A_k) - v_i(A_i)$.

An allocation is envy-freeable if and only if its envy graph does not contain a positive-weight directed cycle. More generally, Halpern and Shah [44] obtained the following theorem; we include their proof in order to familiarize the reader with the structure of envy-freeable allocations.[2]

Theorem 4.2.1. *[44] The following statements are equivalent.*

[2]Note that the construction used in [44] is similar to the construction from [5] for the n-item case.

(a) The allocation $\mathcal{A}$ is envy-freeable.

(b) The allocation $\mathcal{A}$ maximizes (utilitarian) welfare across all reassignments of its bundles to agents: for every permutation π of $I = [n]$, we have $\sum_{i \in I} v_i(A_i) \geq \sum_{i \in I} v_i(A_{\pi(i)})$.

(c) The envy graph $G_{\mathcal{A}}$ contains no positive-weight directed cycles.

Proof.

$(a) \Rightarrow (b)$: Let $\mathcal{A} = \{A_1, \ldots, A_n\}$ be envy-freeable. Then, by definition, there exists a payment vector $\boldsymbol{p}$ such that $v_i(A_i) + p_i \geq v_i(A_k) + p_k$, for any pair of agents i and k. Rearranging, we have $v_i(A_k) - v_i(A_i) \leq p_i - p_k$. Then, for any permutation π of $I = [n]$

$$\sum_{i \in I} \big(v_i(A_{\pi(i)}) - v_i(A_i)\big) \;\leq\; \sum_{i \in I} \big(p_i - p_{\pi(i)}\big) \;=\; \sum_{i \in I} p_i - \sum_{i \in I} p_{\pi(i)} \;=\; 0.$$

Thus the allocation $\mathcal{A}$ maximizes welfare over all reassignments of its bundles.

$(b) \Rightarrow (c)$: Assume $\mathcal{A}$ maximizes welfare over all reassignments of its bundles and take a directed cycle C in the envy graph $G_{\mathcal{A}}$. Without loss of generality $C = \{1, 2, \ldots, r\}$ for some $r \geq 2$. Now define a permutation π_C of I according to the following rules: (i) $\pi_C(i) = i + 1$ for each $i \leq r - 1$, (ii) $\pi_C(r) = 1$, and (iii) $\pi_C(i) = i$ otherwise. Then the weight of the cycle C in the envy graph satisfies

$$\begin{aligned}
w_{\mathcal{A}}(C) &= \sum_{(i,k) \in C} w_{\mathcal{A}}(i, k) \\
&= \sum_{i=1}^{r-1} (v_i(A_{i+1}) - v_i(A_i)) + (v_r(A_1) - v_r(A_r)) \\
&= \sum_{i=1}^{r-1} (v_i(A_{i+1}) - v_i(A_i)) + (v_r(A_1) - v_r(A_r)) + \sum_{i=r+1}^{n} (v_i(A_i) - v_i(A_i)) \\
&= \sum_{i \in I} v_i(A_{\pi(i)}) - v_i(A_i) \\
&\leq 0.
\end{aligned}$$

The inequality holds as $\mathcal{A}$ maximizes welfare over all bundle reassignments. Thus C has non-positive weight.

$(c) \Rightarrow (a)$: Assume the envy graph $G_{\mathcal{A}}$ contains no positive-weight directed cycles. Let $\ell_{G_{\mathcal{A}}}(i)$ be the maximum weight of any path (including the empty path) that starts at vertex i in $G_{\mathcal{A}}$. For each agent $i \in I$, set its payment $p_i = \ell_{G_{\mathcal{A}}}(i)$. Observe that $p_i \geq 0$ as the empty path has weight zero. The corresponding pair $(\mathcal{A}, \boldsymbol{p})$ is then envy-free. To see this, recall that there are no positive-weight cycles. Therefore, for any pair of agents i and k, we have

$$p_i \;=\; \ell_{G_{\mathcal{A}}}(i) \;\geq\; w_{\mathcal{A}}(i,k) + \ell_{G_{\mathcal{A}}}(k) \;=\; (v_i(A_k) - v_i(A_i)) + p_k.$$

Thus $v_i(A_i) + p_i \geq v_i(A_k) + p_k$ and the allocation $\mathcal{A}$ is envy-freeable. □

Theorem 4.2.1 is important for two reasons. The first is that whilst an allocation $\mathcal{A} = \{A_1, A_2, \ldots, A_n\}$ need not be envy-freeable, Condition (b) tells us that there is some permutation π of the bundles in $\mathcal{A}$ such that the resultant allocation, $\mathcal{A}^\pi = \{A_{\pi(1)}, \ldots, A_{\pi(n)}\}$, is envy-freeable. For example, consider again the simple one-item, two-agent instance above. If the item is allocated to agent 1 then the weight on the arc $(1,2)$ is $-v_1$ and the weight on the arc $(2,1)$ is v_2. Because $v_1 < v_2$, the envy graph has a positive-weight directed cycle $\{1,2\}$ and so, by Theorem 4.2.1, this allocation is not envy-freeable. However, suppose we fix the bundles and find a utility-maximizing reallocation of these fixed bundles. This reallocation assigns the item to agent 2 and now there is no positive-weight directed cycle in the resultant envy-free graph; consequently this allocation is envy-freeable by providing a subsidy in the range $[v_1, v_2]$ to agent 1.

Second, to calculate the subsidy vector $\boldsymbol{p}$ associated with an envy-freeable allocation, such as $\mathcal{A}^\pi$, it suffices to calculate the maximum-weight paths beginning at each vertex in its envy graph. (In fact, it is straightforward to prove that the heaviest-path weights *lower bound* the payment to each agent in any envy-free payment vector of an envy-freeable allocation [44].) Note that given any payment vector that eliminates envy, we may increase or decrease the payments to all agents equally while maintaining envy-freeness. As a consequence, in the payment vector that minimizes the total subsidy, there is at least one agent that receives a payment of 0. Together these arguments give the following very useful observation.

Observation 4.2.2. *For any envy-freeable allocation $\mathcal{A}$, the minimum total subsidy required is at most $(n-1) \cdot \ell_{G_{\mathcal{A}}}^{\max}$, where $\ell_{G_{\mathcal{A}}}^{\max}$ is the maximum weight of a directed path in the envy graph $G_{\mathcal{A}}$.*

Halpern and Shah [44] then prove:

Theorem 4.2.3. *[44] For any envy-freeable allocation $\mathcal{A}$, the minimum total subsidy required is at most $(n-1) \cdot m$.*

Proof. In a minimum subsidy vector, at least one agent requires no subsidy. Thus it suffices to show that the subsidy to any agent i is at most m. By Observation 4.2.2, it suffices to show that the heaviest path weight starting at any vertex is at most m. Without loss of generality, let the heaviest path be $P = \{1, 2, \ldots, r\}$. The subsidy made to agent 1 can then be upper bounded by

$$\ell_{G_{\mathcal{A}}}(1) \;=\; \sum_{(i,k)\in P} w_{\mathcal{A}}(i,k) \;=\; \sum_{i=1}^{r-1} (v_i(A_{i+1}) - v_i(A_i)) \;\le\; \sum_{i=1}^{r-1} v_i(A_{i+1}) \;\le\; \sum_{i=1}^{r-1} |A_{i+1}| \;\le\; |J| \;=\; m.$$

Here the second inequality holds because each agent has value at most one for any item. The third inequality is due to the fact that for the allocation $\mathcal{A}$ the bundles $\{A_1, A_2, \ldots, A_n\}$ are disjoint. Consequently $p_i \le m$ for each agent, as required. $\square$

For an arbitrary envy-freeable allocation $\mathcal{A}$ the bound in Theorem 4.2.3 is tight. To see this, consider the example where every agent has value 1 for each item, and the grand bundle (containing all items) is given to agent 1. This allocation is envy-freeable, and here each of the other $n-1$ agents requires a subsidy of m for envy-freeness. Ergo, to provide an improved bound on the total subsidy, we cannot consider any generic envy-freeable allocation. Instead, our task is to find a specific envy-freeable allocation where the heaviest paths in the associated envy graph have much smaller weight. In particular, for the case of additive agents, we want that these path weights are at most 1 rather than at most m. This is our goal in the subsequent sections of the chapter.

Before doing this, let us briefly discuss some computational aspects. Theorem 4.2.1 provides efficient methods to test if a given allocation is envy-freeable. For example, this can be achieved via a maximum-weight bipartite matching algorithm to verify Condition (b). Alternatively, Condition (c) can be tested in polynomial time using the Floyd-Warshall algorithm.[3] Finally, given an arbitrary non-envy-freeable allocation $\mathcal{A}$, one can efficiently find a corresponding envy-freeable allocation $\mathcal{A}^\pi$ by fixing the n bundles of the given allocation and computing a maximum-weight bipartite matching between the agents and the bundles.

4.3 An Allocation Algorithm for Additive Agents

In this section we present an allocation algorithm for the case of additive agents. Recall our task is to construct an envy-freeable allocation $\mathcal{A}$ with maximum path weight 1 in the envy graph $G_\mathcal{A}$. We do this via an allocation algorithm defined on the valuation graph for the instance. Given an instance with a set I of n agents and a set J of m items, the valuation graph H is the complete bipartite graph on vertex sets I and J, where edge (i, j) has weight $v_i(j)$. We denote by $H[\hat{I}, \hat{J}]$ the subgraph of H induced by $\hat{I} \subseteq I$ and $\hat{J} \subseteq J$. The allocation algorithm then proceeds in rounds where each agent is matched to exactly one item in each round. For the first round, we set $J_1 = J$. In round t, we then find a maximum-weight matching M_t in $H[I, J_t]$. If agent i is matched to item $j = \mu_i^t$ then we allocate item μ_i^t to that agent. We then recurse on the remaining items $J_{t+1} = J_t \setminus \cup_{i \in I}\{\mu_i^t\}$. The process ends when every item has been allocated. This procedure is formalized via pseudocode in Algorithm 1.

Suppose the algorithm terminates in T rounds. We assume that every agent receives an item in each round. For rounds 1 to $T - 1$ this is evident because agent i can be assigned a item for which it has zero value. For round T, we assume there are exactly n items remaining, possibly by adding dummy items of no value to any agent.

[3] In fact, a simple reduction converts the problem of finding minimum payments for a fixed allocation into a shortest-paths problem and any efficient shortest-paths algorithm can be applied.

Algorithm 1: Iterated Matching Algorithm

$A_i \leftarrow \emptyset$ for all $i \in I$;
$t \leftarrow 1$; $J_1 \leftarrow J$;
while $J_t \neq \emptyset$ **do**
  Compute a *maximum-weight matching* $M^t = \{(i, \mu_i^t)\}_{i \in I}$ in $H[I, J_t]$;
  Set $A_i \leftarrow A_i \cup \{\mu_i^t\}$ for all $i \in I$;
  Set $J_{t+1} \leftarrow J_t \setminus \cup_{i \in I}\{\mu_i^t\}$;
  $t \leftarrow t + 1$;
end

This algorithm has many interesting properties. In this section we prove that it outputs an envy-freeable allocation $\mathcal{A}$. Furthermore, the allocation $\mathcal{A}$ is EF1, thus settling Conjecture 4.0.2. The allocation is also balanced in that (discarding any additional dummy items) the bundles that the agents receive differ in size by at most one item; in particular, each agent receives a bundle of size either $\lfloor \frac{m}{n} \rfloor$ or $\lceil \frac{m}{n} \rceil$. The allocation algorithm also clearly runs in polynomial time.

We also show in this section that any allocation $\mathcal{A}$ that is both envy-freeable and EF1 has a heaviest path weight in the envy graph of weight at most $n-1$. Thus, By Observation 4.2.2, the algorithm outputs an allocation that requires a subsidy of at most $(n-1)^2$. As claimed though, the heaviest path weight in $G_{\mathcal{A}}$ is in fact at most one and so the total subsidy needed is at most $n-1$. We defer the proof of this fact, our main result, to Section 4.4.

4.3.1 The Allocation Is Envy-freeable

Let's first see that the allocation $\mathcal{A}$ output by the algorithm after the final round T is envy-freeable.

Lemma 4.3.1. *The output allocation $\mathcal{A}$ is envy-freeable.*

Proof. Let M^t be the maximum matching found in round t and $\boldsymbol{\mu}^t = \{\mu_1^t, \mu_2^t, \ldots, \mu_n^t\}$ the corresponding items allocated in that round. By Theorem 4.2.1 it suffices to show that no directed cycle in the envy graph corresponding to the final allocation $\mathcal{A}$ has positive weight. Take any directed cycle C in the envy graph $G_{\mathcal{A}}$. Again, we may assume without loss of generality that $C = \{1, 2, \ldots, r\}$ for some $r \geq 2$. We have

$$\begin{aligned}
w_{\mathcal{A}}(C) &= \sum_{(i,k)\in C} w_{\mathcal{A}}(i,k) \\
&= \sum_{(i,k)\in C} [v_i(A_k) - v_i(A_i)] \\
&= \sum_{(i,k)\in C} \sum_{t=1}^{T} \left[v_i(\mu_k^t) - v_i(\mu_i^t)\right] \\
&= \sum_{(i,k)\in C} \sum_{t=1}^{T} w_{\boldsymbol{\mu}^t}(i,k) \\
&= \sum_{t=1}^{T} \sum_{(i,k)\in C} w_{\boldsymbol{\mu}^t}(i,k).
\end{aligned}$$

Let π_C be the permutation of I under which $\pi_C(i) = i + 1$ for each $i \leq r - 1$, $\pi_C(r) = 1$, and $\pi_C(i) = i$ otherwise. In each round t, since M_t is a maximum-weight matching, $\sum_{(i,k)\in C} w_{\boldsymbol{\mu}^t}(i,k)$ is non-positive: otherwise, the matching $\hat{M}^t$ obtained by allocating to each agent i the item $\mu^t_{\pi_C(i)}$ has greater weight than M_t, a contradiction. Thus $w_{\mathcal{A}}(C)$ is also non-positive. Consequently, by Theorem 4.2.1 the allocation produced by the algorithm is envy-freeable. □

4.3.2 The Allocation Is EF1

We say that an allocation $\mathcal{A}$ satisfies the *envy bounded by a single good* property, and is *EF1*, if for each pair i, k of agents, either $A_k = \emptyset$ or there exists an item $j \in A_k$ such that

$$v_i(A_i) \geq v_i(A_k \setminus \{j\}).$$

Next, let's prove the output allocation $\mathcal{A}$ is EF1.

Lemma 4.3.2. *The output allocation $\mathcal{A}$ is EF1.*

Proof. Let $\mathcal{A} = \{A_1, \ldots, A_n\}$. Recall, in any round t, the algorithm computes a maximum-weight matching M^t in $H[I, J_t]$ and allocates item μ_i^t to agent i. Thus $A_i = \{\mu_i^1, \ldots, \mu_i^T\}$ is the set of items allocated to agent i. Observe that $v_i(\mu_i^t) \geq v_i(j)$ for any item $j \in J_{t+1}$, the collection of items unallocated at the start of round $t+1$. Otherwise, we can replace the edge (i, μ_i^t) with (i, j) in M_t, to obtain a matching of greater weight in $H[I, J_t]$. Therefore, for any pair of agents i and k, we have

$$\begin{aligned}
v_i(A_i) &= v_i(\{\mu_i^1, \ldots, \mu_i^T\}) \\
&= v_i(\mu_i^1) + \cdots + v_i(\mu_i^{T-1}) + v_i(\mu_i^T) \\
&\geq v_i(\mu_i^1) + \cdots + v_i(\mu_i^{T-1}) \\
&\geq v_i(\mu_k^2) + \cdots + v_i(\mu_k^T) \\
&= v_i(A_k \setminus \{\mu_k^1\}).
\end{aligned}$$

Ergo, the output allocation $\mathcal{A}$ is EF1. □

Claim 4.3.3. *[44] Let $\mathcal{A}$ be both envy-freeable and EF1. Then the minimum total subsidy required is at most $(n-1)^2$.*

Proof. Since there is an agent that requires no subsidy, it suffices to prove that the maximum path weight in the envy graph $G_{\mathcal{A}}$ is at most $n-1$. But $\mathcal{A}$ is EF1. So agent i envies agent k by at most one, the maximum value of a single item. Thus every arc (i, k) has weight at most one, that is, $w_{\mathcal{A}}(i, k) \leq 1$. The result follows as any path contains at most $n-1$ arcs. □

Since we have shown that the output allocation $\mathcal{A}$ is both envy-freeable and EF1, it immediately follows by Claim 4.3.3 that it requires a total subsidy of at most $(n-1)^2$.

4.4 The Subsidy Required Is at Most One per Agent

In this section we complete our analysis of the additive setting. By the EF1 property of the output allocation $G_{\mathcal{A}}$ we have an upper bound of 1 on the weight of any arc in the envy graph $G_{\mathcal{A}}$. But this is insufficient to accomplish our goal of proving that the envy graph has maximum path weight 1. How can we do this?

4.4.1 Lower Bounding Path Weights

As a thought experiment, imagine that, rather than an upper bound of 1 on each arc weight, we have a lower bound of -1 on each arc weight. The subsequent lemma proves this would be a sufficient condition!

Lemma 4.4.1. *Let $\mathcal{A}$ be an envy-freeable allocation. If $w_{\mathcal{A}}(i,k) \geq -1$ for every arc (i,k) in the envy graph then the maximum subsidy required is at most one per agent.*

Proof. By Theorem 4.2.1, as $\mathcal{A}$ is an envy-freeable the envy graph $G_{\mathcal{A}}$ contains no positive-weight cycles. Let P be the maximum-weight path in $G_{\mathcal{A}}$. Without loss of generality, $P = \{1, 2, \ldots, i\}$ with weight $p_1 = \ell_{G_{\mathcal{A}}}(1)$. Now take the directed cycle $C = P \cup (i, 1)$. Because C has non-positive weight and every arc weight is at least -1, we obtain

$$0 \;\geq\; w_{\mathcal{A}}(C) \;=\; \ell_{G_{\mathcal{A}}}(1) + w_{\mathcal{A}}(i,1) \;\geq\; \ell_{G_{\mathcal{A}}}(1) - 1.$$

Therefore $\ell_{G_{\mathcal{A}}}(1) \leq 1$ and the maximum subsidy is at most one. □

At first glance, Lemma 4.4.1 seems of little use. We already know every arc in the envy graph has weight at most 1. Suppose in addition that every arc weight was at least -1. That is, $1 \geq w_{\mathcal{A}}(i,k) \geq -1$ for each arc (i,k). Consequently, $v_i(A_i) \leq v_i(A_k) + 1$ and $v_i(A_k) \leq v_i(A_i) + 1$. In instances with a large number of valuable items this means that every agent is essentially indifferent over which bundle in $\mathcal{A}$ they receive. It is unlikely that an allocation with this property even exists for every instance, and certainly not the case that our algorithm outputs such an allocation.

The trick is to apply Lemma 4.4.1 to a modified fair division instance. In particular we construct, for each agent i, a modified valuation function $\bar{v}_i$ from v_i. We then prove that the allocation $\mathcal{A}^{\boldsymbol{v}}$ output for the original valuation profile $\boldsymbol{v}$ is envy-freeable even for the modified valuation profile $\bar{\boldsymbol{v}}$. Next we show that with this same allocation, every arc weight is at least -1 in the envy graph under the modified valuation profile $\bar{\boldsymbol{v}}$. By Lemma 4.4.1, this implies that the maximum subsidy required is at most one for the valuation profile $\bar{\boldsymbol{v}}$. To complete the proof we show that the maximum subsidy required by each agent for the original valuation profile $\boldsymbol{v}$ is at most the subsidy required for $\bar{\boldsymbol{v}}$.

4.4.2 A Modified Valuation Function

Let $\mathcal{A}^{\boldsymbol{v}} = \{A_1^{\boldsymbol{v}}, \ldots, A_n^{\boldsymbol{v}}\}$ be the allocation output by our algorithm under the original valuation profile $\boldsymbol{v}$. We now create the modified valuation profile $\bar{\boldsymbol{v}}$. For each agent i, define $\bar{v}_i$ according to the rule:

$$
\begin{aligned}
\bar{v}_i(\mu_i^t) &= v_i(\mu_i^t) && \forall t \leq T \\
\bar{v}_i(\mu_k^t) &= \max\left(v_i(\mu_k^t), v_i(\mu_i^{t+1})\right) && \forall k \in I \setminus \{i\},\ \forall t \leq T-1 \\
\bar{v}_i(\mu_k^T) &= v_i(\mu_k^T) && \forall k \in I \setminus \{i\}.
\end{aligned}
$$

That is, the value $\bar{v}_i(j)$ remains the same for any item $j \in A_i^{\boldsymbol{v}}$ that was allocated to agent i by the algorithm. For any other item j, the value $\bar{v}_i(j)$ is the maximum of the original value $v_i(j)$ and the value of the item allocated to i by the algorithm in the round that immediately follows the round where j was allocated to some agent.

The following two observations are trivial but will be useful.

Observation 4.4.2. *For any agent i and item $j \in A_i^{\boldsymbol{v}}$, we have $v_i(j) = \bar{v}_i(j)$.* □

Observation 4.4.3. *For any agent i and item $j \notin A_i^{\boldsymbol{v}}$, we have $v_i(j) \leq \bar{v}_i(j)$.* □

We will show the bound on the subsidy by a sequence of claims based on the proof plan outlined above. First we show that $\mathcal{A}^{\boldsymbol{v}}$ is envy-freeable even under the modified valuation profile.

Claim 4.4.4. *The allocation $\mathcal{A}^{\boldsymbol{v}}$ output under the original valuation profile $\boldsymbol{v}$ is an envy-freeable allocation under the modified valuation profile $\bar{\boldsymbol{v}}$.*

Proof. By Theorem 4.2.1, to show that the allocation $\mathcal{A}^{\boldsymbol{v}}$ is envy-freeable under the modified valuation profile $\bar{\boldsymbol{v}}$ we must show that there is no positive-weight cycle in the envy graph using the modified values. So suppose cycle C has positive modified weight. To obtain a contradiction, first observe that, in the allocation $\mathcal{A}^{\boldsymbol{v}}$, agent i receives the bundle $\mathcal{A}_i^{\boldsymbol{v}} = \{\mu_i^1, \mu_i^2, \ldots, \mu_i^T\}$. Thus with respect to $\bar{\boldsymbol{v}}$ the envy agent i has for agent k is

$$\bar{v}_i(\mathcal{A}_k^{\boldsymbol{v}}) - \bar{v}_i(\mathcal{A}_i^{\boldsymbol{v}}) \;=\; \sum_{t=1}^{T} \bar{v}_i(\mu_k^t) - \sum_{t=1}^{T} \bar{v}_i(\mu_i^t) \;=\; \sum_{t=1}^{T} \left(\bar{v}_i(\mu_k^t) - \bar{v}_i(\mu_i^t)\right). \tag{4.1}$$

As the envy graph contains a positive-weight cycle C we have, by (4.1), that

$$0 < \sum_{(i,k)\in C} \bar{v}_i(\mathcal{A}_k^{\boldsymbol{v}}) - \bar{v}_i(\mathcal{A}_i^{\boldsymbol{v}}) \;=\; \sum_{(i,k)\in C} \sum_{t=1}^{T} \left(\bar{v}_i(\mu_k^t) - \bar{v}_i(\mu_i^t)\right) \;=\; \sum_{t=1}^{T} \sum_{(i,k)\in C} \left(\bar{v}_i(\mu_k^t) - \bar{v}_i(\mu_i^t)\right).$$

This implies there exists a round t such that

$$\sum_{(i,k)\in C} \bar{v}_i(\mu_k^t) > \sum_{(i,k)\in C} \bar{v}_i(\mu_i^t). \tag{4.2}$$

Now M^t is a maximum-weight matching in $H[I, J_t]$ for the original valuation profile $\boldsymbol{v}$. Let $\hat{M}^t$ be the matching formed from M^t by permuting around the cycle C the bundles of the agents in C. But then, by (4.2), the matching $\hat{M}^t$ has greater weight in $H[I, J_t]$ than the matching M^t for the modified profile $\bar{\boldsymbol{v}}$. Consequently, we will obtain our contradiction if we can prove that M^t is a maximum-weight matching in $H[I, J_t]$ even with respect to $\bar{\boldsymbol{v}}$.

This is true in the final round matching; clearly M^T is a maximum-weight matching in $H[I, J_T]$ because, by definition, $\bar{\boldsymbol{v}}$ and $\boldsymbol{v}$ have the same value for items in J_T. Thus, it remains to prove the statement for each round $t \leq T-1$. Now $\boldsymbol{\mu}^t = \{\mu_1^t, \mu_2^t, \ldots, \mu_n^t\}$ is the allocation of the items in round t. Again, for a contradiction, assume that matching M^t is not maximum in $H[I, J_t]$ for the valuation profile $\bar{\boldsymbol{v}}$. Then, by Theorem 4.2.1, the envy graph $G_{\boldsymbol{\mu}^t}$ contains a positive-weight directed cycle C. Without loss of generality, let $C = \{1, \ldots, r\}$.

We divide our analysis into two cases, depending on whether the weights on the arcs of C change when the valuation profile is modified from $\boldsymbol{v}$ to $\bar{\boldsymbol{v}}$. Specifically, we color an arc $(i, i+1)$ of C *blue* if $\bar{v}_i(\mu_{i+1}^t) = v_i(\mu_{i+1}^t)$, that is, agent i's value for the item allocated to agent $i+1$ does not change when the valuation profile is modified. Otherwise, we color the arc $(i, i+1)$ *red*. Observe that if the arc $(i, i+1)$ of C is red, then $\bar{v}_i(\mu_{i+1}^t) = v_i(\mu_i^{t+1}) > v_i(\mu_{i+1}^t)$, so in the original valuation function agent i strictly prefers the item that it is allocated in round $t+1$ to the item that agent $i+1$ is allocated in round t. In turn, this implies that the weight on any red arc is necessarily negative. We have the following two cases to consider.

(i) *Every arc of C is blue.* Let π_C be the permutation of I under which $\pi_C(i) = i+1$ for each $i \leq r-1$, $\pi_C(r) = 1$, and $\pi_C(i) = i$ otherwise. The matching $\mathcal{M}^t$ obtained by allocating to each agent i the item $\mu_{\pi_C(i)}^t$ has greater weight than M^t with respect to the original valuation profile $\boldsymbol{v}$, contradicting the assumption that the algorithm selected a matching of maximum weight.

(ii) *C contains a red arc.* In this case, C can be decomposed into a sequence of d directed paths $P_1, \ldots, P_d$ such that each directed path consists of a (possibly empty) sequence of blue arcs followed by exactly one red arc. Figure 4.1 shows an example of such a decomposition. In the figure, blue arcs are represented by solid lines and red arcs by dashed lines.

Now, since C has positive total weight, there is a directed path $P \in \{P_1, \ldots, P_d\}$ of positive total weight. Without loss of generality, let $P = \{1, 2, \ldots, k+1\}$. Thus in

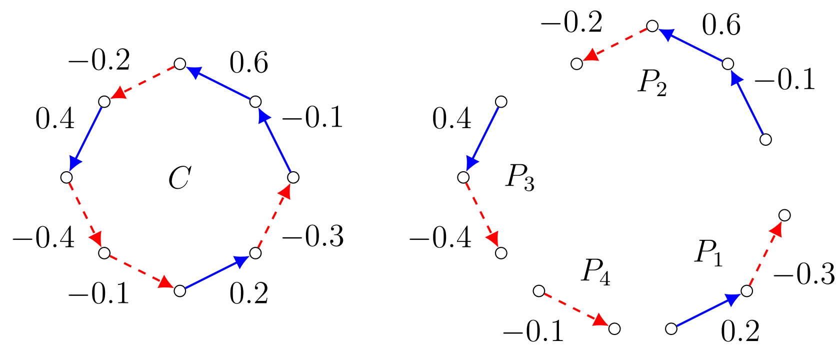

Figure 4.1: An example showing the decomposition of C into directed paths $P_1, \ldots, P_4$. In this example, P_2 has positive weight.

the envy graph $G_{\boldsymbol{\mu}^t}$ we have

$$w_{\boldsymbol{\mu}^t}(P) \;=\; \sum_{i=1}^{k} w_{\boldsymbol{\mu}^t}(i, i+1) \;>\; 0. \tag{4.3}$$

Construct a matching $\mathcal{M}^t = \{(i, \omega_i^t)\}_{i \in I}$ in the following manner. For each agent $i \geq k+1$, set $\omega_i^t = \mu_i^t$; that is, the end-vertex of the path P and all agents not on P are matched to the same item in $\mathcal{M}^t$ as in M^t. For each agent $i \leq k-1$, let $\omega_i^t = \mu_{i+1}^t$, that is in the allocation $\mathcal{M}^t$ agent i receives the item that agent $i+1$ receives in M^t. Finally, for agent k let $\omega_k^t = M_k^{t+1}$; that is, in $\mathcal{M}^t$ agent k receives the item it would have received in the next round in M^{t+1}.

Observe that every item allocated by $\mathcal{M}^t$ was available for allocation in round t and, thus, it was a feasible allocation to select in round t. Next let's compare the relative values of $\mathcal{M}^t$ and M^t under the original valuations $\boldsymbol{v}$. To do this, observe that by definition of $\mathcal{M}^t$ we have

$$\begin{aligned} v(\mathcal{M}^t) - v(M^t) &= \sum_{i=1}^{k} \left(v_i(\omega_i^t) - v_i(\mu_i^t) \right) \\ &= \sum_{i=1}^{k-1} \left(v_i(\omega_i^t) - v_i(\mu_i^t) \right) + \left(v_k(\omega_k^t) - v_k(\mu_k^t) \right) \\ &= \sum_{i=1}^{k-1} \left(v_i(\mu_{i+1}^t) - v_i(\mu_i^t) \right) + \left(v_k(\mu_k^{t+1}) - v_k(\mu_k^t) \right). \end{aligned} \tag{4.4}$$

But $(k, k+1)$ is a red arc in $G_{\boldsymbol{\mu}^t}$. Therefore, it is the case that $v_k(\mu_k^{t+1}) > v_k(\mu_{k+1}^t)$. Plugging this into (4.4) gives

$$\begin{aligned} v(\mathcal{M}^t) - v(M^t) \;&>\; \sum_{i=1}^{k-1} \left(v_i(\mu_{i+1}^t) - v_i(\mu_i^t) \right) + \left(v_k(\mu_{k+1}^t) - v_k(\mu_k^t) \right) \\ &= \sum_{i=1}^{k} \left(v_i(\mu_{i+1}^t) - v_i(\mu_i^t) \right). \end{aligned} \tag{4.5}$$

But, by definition, $w_{\boldsymbol{\mu}^t}(i, i+1) = v_i(\mu_{i+1}^t) - v_i(\mu_i^t)$. So, together (4.3) and (4.5) imply

$$v(\mathcal{M}^t) - v(M^t) \;>\; \sum_{i=1}^{k} w_{\boldsymbol{\mu}^t}(i, i+1) \;>\; 0. \tag{4.6}$$

Thus $\mathcal{M}^t$ has greater weight than M^t under the original valuations $\boldsymbol{v}$. This contradicts the optimality of M^t. □

Claim 4.4.4 shows that the allocation $\mathcal{A}^{\boldsymbol{v}}$ produced by the algorithm on the original instance is an envy-freeable allocation in the modified instance. We next show that for this modified valuation profile the subsidy required is at most 1 for each agent. In particular the total subsidy is at most $n-1$.

Claim 4.4.5. *For the envy-freeable allocation $\mathcal{A}^{\boldsymbol{v}}$ the subsidy to each agent is at most 1 for the modified valuation profile $\bar{\boldsymbol{v}}$.*

Proof. Take the valuation profile $\bar{\boldsymbol{v}}$ and the allocation $\mathcal{A}^{\boldsymbol{v}} = \{A_1^{\boldsymbol{v}}, \ldots, A_n^{\boldsymbol{v}}\}$. We claim that for any arc (i,k) its modified weight $\bar{w}_{\boldsymbol{A}^{\boldsymbol{v}}}(i,k)$ in the envy graph is *at least* -1. To prove this take any pair of agents i and k. Then

$$
\begin{aligned}
\bar{w}_{\boldsymbol{A}^{\boldsymbol{v}}}(i,k) &= \bar{v}_i(A_k^{\boldsymbol{v}}) - \bar{v}_i(A_i^{\boldsymbol{v}}) \\
&= \sum_{t=1}^{T} \bar{v}_i(\mu_k^t) - \sum_{t=1}^{T} \bar{v}_i(\mu_i^t) \\
&= \sum_{t=1}^{T} \bar{v}_i(\mu_k^t) - \sum_{t=1}^{T} v_i(\mu_i^t) \\
&= \sum_{t=1}^{T-1} \max(v_i(\mu_k^t), v_i(\mu_i^{t+1})) + v_i(\mu_i^T) - \sum_{t=1}^{T} v_i(\mu_i^t) \\
&\geq \sum_{t=1}^{T-1} v_i(\mu_i^{t+1}) - \sum_{t=1}^{T-1} v_i(\mu_i^t). \qquad (4.7)
\end{aligned}
$$

We can simplify (4.7) and lower bound it via a telescoping sum:

$$
\begin{aligned}
\bar{w}_{\boldsymbol{A}^{\boldsymbol{v}}}(i,k) &\geq \sum_{t=1}^{T-1} \left(v_i(\mu_i^{t+1}) - v_i(\mu_i^t)\right) \\
&= v_i(\mu_i^T) - v_i(\mu_i^1) \\
&\geq -v_i(\mu_i^1) \\
&\geq -1. \qquad (4.8)
\end{aligned}
$$

Now by Claim 4.4.4, the allocation $\mathcal{A}^{\boldsymbol{v}}$ is envy-freeable with respect to the valuations $\bar{\boldsymbol{v}}$. Applying Lemma 4.4.1, because the arc weights are lower bounded by -1 the subsidy required per agent is then at most one for the modified valuation profile $\bar{\boldsymbol{v}}$.

Finally, since there is an agent whose payment is 0, the total subsidy required is upper bounded by $n-1$. □

The following claim shows that, for any agent, the subsidy for the original valuation profile is at most the subsidy required for the modified valuation function.

Claim 4.4.6. *For the allocation $\mathcal{A}^{\boldsymbol{v}}$ the minimum subsidy required by an agent given valuation profile $\boldsymbol{v}$ is at most the minimum subsidy required given valuation profile $\bar{\boldsymbol{v}}$.*

Proof. By Observation 4.4.2, $v_i(j) = \bar{v}_i(j)$ for any $j \in A_i^{\boldsymbol{v}}$. Therefore, by additivity,

$$\bar{v}_i(A_i^{\boldsymbol{v}}) \;=\; \sum_{j \in A_i^{\boldsymbol{v}}} \bar{v}_i(j) \;=\; \sum_{j \in A_i^{\boldsymbol{v}}} v_i(j) \;=\; v_i(A_i^{\boldsymbol{v}}). \tag{4.9}$$

On the other hand, Observation 4.4.3 states that $v_i(j) \leq \bar{v}_i(j)$ for any $j \notin A_i^{\boldsymbol{v}}$. Thus, for any pair i and k of agents, we have

$$\bar{v}_i(A_k^{\boldsymbol{v}}) \;=\; \sum_{j \in A_k^{\boldsymbol{v}}} \bar{v}_i(j) \;\geq\; \sum_{j \in A_k^{\boldsymbol{v}}} v_i(j) \;=\; v_i(A_k^{\boldsymbol{v}}). \tag{4.10}$$

Combining (4.9) and (4.10) gives

$$\bar{w}_{\boldsymbol{A}^{\boldsymbol{v}}}(i,k) \;=\; \bar{v}_i(A_k^{\boldsymbol{v}}) - \bar{v}_i(A_i^{\boldsymbol{v}}) \;\geq\; v_i(A_k^{\boldsymbol{v}}) - v_i(A_i^{\boldsymbol{v}}) \;=\; w_{\mathcal{A}^{\boldsymbol{v}}}(i,k).$$

Consequently, the weight of any arc (i,k) in the envy graph with the modified valuation profile is at least its weight with the original valuation profile. Therefore the weight of any path in the envy graph is greater with the modified valuation profile than with the original valuation profile. The claim follows. □

Together Claims 4.4.5 and 4.4.6 give our main result.

Theorem 4.0.3. *For additive valuations there is an envy-freeable allocation where the subsidy to each agent is at most one dollar. (This allocation is also EF1, balanced, and can be computed in polynomial time.)* □

4.5 Bounding the Subsidy for Monotone Valuations

We now consider the much more general setting where the valuations of the agents are arbitrary monotone functions. That is, the only assumptions we impose are that $v_i(S) \leq$

$v_i(T)$ when $S \subseteq T$ and the basic assumption that $v_i(\emptyset) = 0$. Without loss of generality, we may scale the valuations so that the marginal value of each item for any agent never exceeds one dollar. Our goal in this section is to show that there is an envy-freeable allocation in which the total subsidy required for envy-freeness is at most $2(n-1)^2$. In particular, the total subsidy required is independent of the number of items m. When $m > 2(n-1)$ this bound beats the bound $(n-1) \cdot m$ of [44] for additive valuations described in Theorem 4.2.3 and, more importantly, it applies to the far more general class of arbitrary monotone valuations.

Our method to compute the desired envy-freeable allocation begins with finding an EF1 allocation. The well-known *envy-cycles* algorithm of Lipton et al. [55] finds such an allocation in polynomial time given oracle access to the valuations, under the same mild conditions on the valuations. For completeness, we briefly describe the envy-cycles algorithm. The algorithm proceeds in a sequence of m rounds, allocating one item in each round. At any point during the algorithm, we denote by G the envy graph corresponding to the current allocation, and by H the subgraph of G that consists of all the agents and only the arcs that have positive weight, that is, positive envy. We call H the *auxiliary graph* of G. The algorithm relies on the following lemma.

Lemma 4.5.1. *[55] For any partial allocation $\mathcal{A}$ with auxiliary graph H, there is another partial allocation $\mathcal{A}'$ with auxiliary graph H' such that*

(i) H' is acyclic.

(ii) For each agent i, the maximum weight of an outgoing arc from i is less in $\mathcal{A}'$ than in $\mathcal{A}$.

The basic idea of the algorithm then is to maintain the following two invariants: (i) at each step, the partial allocation is EF1, and (ii) at the start and end of each round, the auxiliary graph H is acyclic. Since the auxiliary graph is a directed acyclic graph at the start of each round, it has a source vertex. The algorithm simply chooses this vertex and allocates the next item to the corresponding agent. Because no other agent envies this agent before this item is allocated, the envies of the other agents are bounded by the value of this item (so the allocation of this item maintains the EF1 invariant). Next, the algorithm identifies a directed cycle (if one exists) in the auxiliary graph H and redistributes bundles by rotating them around this cycle. It is easy to see that the EF1 guarantee is maintained after this redistribution of the bundles, and that the number of arcs in H strictly decreases. All cycles in H are then eliminated in sequence until H is acyclic and the round ends. When all items have been allocated, the final allocation is EF1.

This immediately raises the question of whether the resulting allocation is envy-freeable. By Claim 4.3.3, we know that if an allocation is both envy-freeable and EF1, then the total subsidy required for envy-freeness is $(n-1)^2$, since the weight of any path is at most $n-1$. Unfortunately, it is possible that the allocation output by the envy-cycles algorithm is not envy-freeable. However, we show that an EF1 allocation can still be used to produce an envy-freeable allocation that requires only a small increase in the subsidy! Specifically, the following key lemma shows that if we begin by fixing the bundles of an EF1 allocation and then redistribute these bundles to produce an envy-freeable allocation, the weight of any path increases to at most $2(n-1)$. By Theorem 4.2.1, an envy-freeable allocation can be found by computing a maximum-weight matching.

Lemma 4.5.2. *Let $\mathcal{A}$ be an EF1 allocation and $\mathcal{B}$ be the envy-freeable allocation corresponding to a maximum-weight matching between the agents and the bundles of $\mathcal{A}$. Then $\mathcal{B}$ can be made envy-free with a subsidy of at most $2(n-1)$ to each agent.*

Proof. Let $\mathcal{A} = \{A_1, \ldots, A_n\}$ be an EF1 allocation. So, for any pair i and k of agents, $v_i(A_k) - v_i(A_i) \leq 1$. Let π be a permutation of the bundles that maximizes $\sum_i v_i(A_{\pi(i)})$. Then, by Theorem 4.2.1, the allocation $\mathcal{B} = \{B_1, \ldots, B_n\} = \{A_{\pi(1)}, \ldots, A_{\pi(n)}\}$ is envy-freeable. Next, let P be a directed path in the envy graph $G_{\mathcal{B}}$. Without loss of generality, $P = \{1, 2, \ldots, r\}$ for some $r \geq 2$. Our goal is to show that the weight of P in $G_{\mathcal{B}}$ is at most $2(n-1)$. Clearly, the weight of P in $G_{\mathcal{A}}$ is at most $n-1$. Consider an arc $(i, i+1)$ of P. Since $\mathcal{A}$ is EF1, for any agent k, we have $v_i(A_k) - v_i(A_i) \leq 1$. Now, agent $i+1$ receives the bundle of agent $\pi(i+1)$ in the redistributed allocation $\mathcal{B}$. We have $v_i(A_{\pi(i+1)}) - v_i(A_i) \leq 1$

and, thus, $v_i(B_{i+1}) - v_i(A_i) \leq 1$. It follows that:

$$\begin{aligned}
w_{\mathcal{B}}(P) &= \sum_{(i,k)\in P} w_{\mathcal{B}}(i,k) \\
&= \sum_{i=1}^{r-1} \left(v_i(B_{i+1}) - v_i(B_i)\right) \\
&= \sum_{i=1}^{r-1} \left(v_i(B_{i+1}) - v_i(A_i) + v_i(A_i) - v_i(B_i)\right) \\
&\leq \sum_{i=1}^{r-1} \left(1 + v_i(A_i) - v_i(B_i)\right) \\
&\leq (n-1) + \sum_{i=1}^{r-1} \left(v_i(A_i) - v_i(B_i)\right). \qquad (4.11)
\end{aligned}$$

To complete the proof, it remains to show that $\sum_{i=1}^{r-1} \left(v_i(A_i) - v_i(B_i)\right)$ is at most $n-1$. Together with (4.11), this implies that $w_{\mathcal{B}}(P) \leq 2(n-1)$.

Since π maximizes $\sum_i v_i(A_{\pi(i)})$, we have $\sum_i v_i(B_i) \geq \sum_i v_i(A_i)$. The key observation is that, while the sum of values of the bundles received by all agents increases when we redistribute the bundles from $\mathcal{A}$ to $\mathcal{B}$, the value of the bundle received by any single agent increases by *at most* one because $\mathcal{A}$ is EF1. This then constrains the amount by which the total value for any subset of agents can *decrease*. Specifically, let $R \subseteq I$ be the set of agents i that receive a bundle B_i of smaller value than A_i, that is, $R = \{i \in I : v_i(B_i) < v_i(A_i)\}$. Let $S = I \setminus R$, so $S = \{i \in I : v_i(B_i) \geq v_i(A_i)\}$.

Now, we have two cases to consider.

(i) $|R| = 0$.

Then $\sum_{i=1}^{r-1} \left(v_i(A_i) - v_i(B_i)\right) \leq 0$ and the result follows.

(ii) $|R| \geq 1$.

Then $|S| \le n-1$, and we have

$$\begin{aligned}
\sum_{i\in[r-1]} (v_i(A_i) - v_i(B_i)) &= \sum_{i\in[r-1]\cap R} (v_i(A_i) - v_i(B_i)) + \sum_{i\in[r-1]\cap S} (v_i(A_i) - v_i(B_i)) \\
&\le \sum_{i\in[r-1]\cap R} (v_i(A_i) - v_i(B_i)) \\
&\le \sum_{i\in R} (v_i(A_i) - v_i(B_i)) \\
&\le \sum_{i\in S} (v_i(B_i) - v_i(A_i)) \\
&\le n-1.
\end{aligned}$$

The second to last inequality says that the total decrease in value for agents in R is at most the total increase in value for agents in S (since B is an optimal redistribution of the bundles). The final inequality follows from the fact that $|S| \le n-1$ and for each $i \in S$, $v_i(B_i) - v_i(A_i) \le 1$ since $\mathcal{A}$ is EF1. □

Together, Lemmas 4.5.1 and 4.5.2 bound the total subsidy sufficient for envy-freeness when the valuation functions are monotone.

Theorem 4.0.4. *For monotonic valuations there is an envy-freeable allocation where the subsidy to each agent is at most $2(n-1)$ dollars. (Given a valuation oracle, this allocation can be computed in polynomial time.)* □

Chapter 5

Fairness and Welfare

As we saw earlier, as a consequence of the impossibility of envy-freeness without payments, the majority of the community's recent research efforts have been directed towards achieving *approximate* or *relaxed* fairness guarantees, such as EF1 or approximate MMS. In the previous chapter, we studied the use of payments in fair item allocation. We showed that an allocation can be made *envy-free* via the use of a small subsidy. In this chapter, we ask and answer a natural follow-up question: can this tool be made to do more? Can we use it to simultaneously guarantee full envy-freeness while also achieving high *welfare*, and if so, how much in total transfer payments do we need for this? These questions are the focus of this chapter.

One contribution of our work is to extend the literature on subsidies and their application. However, rather than subsidies, we analyze the related concept of *transfer payments* between the agents for two reasons. First, a subsidy is an external source of added utility which, in the context of welfare, would bias any subsequent comparisons with the welfare-maximizing allocation *without* subsidies. A transfer payment is neutral in this regard. Second, subsidies require an external agent willing to fund the mechanism – a typically unrealistic hope. In contrast, transfer payments require the consent only of the agents who are already willing participants in the mechanism. Provided the cost of the payments are outweighed by the benefits of participation then giving consent for this is reasonable. We remark that subsidies and transfers are in a sense interchangeable. Given an envy-free allocation with subsidies, subtracting the average subsidy from each agent's individual payment gives payments which sum to zero, that is, transfer payments. Conversely, given transfer payments, adding an appropriate fixed amount to each payment induces non-negative subsidy payments.[1]

A second contribution is to extend the research on the price of fairness. Specifically, we impose no *balancing constraint* on the valuation functions of the agents. To understand this, note that a common assumption in the price of fairness literature is that the valuation function of each agent is scaled so that the value of the grand bundle of items is *equal* for all agents. In the context of fairness, this scaling is benign because it has no affect on the most widely used measures of fairness. For example, it does not change the (relative) envy between any pair of agents. However, in the context of efficiency or welfare, this scaling can dramatically alter the welfare of any allocation by restricting attention to balanced instances, where agents are of essentially equal importance in generating welfare. This is important because it is the elimination of unbalanced instances that allows non-trivial bounds on the price of fairness to be obtainable ([16, 12]). Indeed, as will be seen later in the chapter, it is the unbalanced instances that are typically the most problematic in obtaining both fairness and high welfare.

We do, for simplicity, continue to make the standard assumption in the literature on subsidies and assume that the maximum marginal value for an item for any agent is always at most one dollar. We emphasize that this assumption is benign with respect to both fairness and welfare: it does not affect the relative envy between agents, and it does not affect the welfare of an allocation (as all valuations can be scaled down uniformly). Expressing the transfers in dollar amounts allows for a consistent comparison with earlier work on the topic, and equivalent bounds for the original instance can be recovered by multiplying these expressions by the maximum marginal value of an item for any agent.

[1]Of course, whilst the correspondence between subsidies and transfers is simple, the switch to transfer payments does have a technical drawback: because transfer payments do not provide an (unnatural) external boost to welfare, obtaining welfare guarantees for the case of transfers is generally harder than for the case of subsidies.

We now present the main results in the chapter. We study the trade-off between fairness and efficiency in the presence of transfer payments for the class of ρ-mean welfare functions, with particular focus on the two most important special cases, namely the *Nash social welfare* and *utilitarian social welfare* functions. An allocation is *envy-freeable* if it can be made envy-free with the addition of subsidies (or, equivalently, transfer payments). Our first observation is that to achieve both fairness and high welfare, it is not sufficient to simply find an envy-freeable allocation – making transfer payments is necessary. In fact, no non-zero welfare guarantee is achievable for all ρ without considering transfers in the computation of the welfare. Letting W^ρ denote ρ-mean welfare, we have:

Observation 5.0.1. *For any $\epsilon > 0$, there exist instances where the welfare of every envy-freeable allocation A satisfies $\frac{\mathrm{W}^\rho(A)}{\mathrm{W}^\rho(A^*)} \leq \epsilon$.*

Here A^* is the welfare-maximizing allocation. The observation applies even in the case of additive valuations with Nash social welfare functions. Consequently, the focus on allocations with transfers is justified. For ρ-mean welfare functions, we show that positive welfare guarantees are achievable with transfers.

Corollary 5.2.3. *For subadditive valuations, there exists an envy-free allocation with transfers (A, t) such that $\frac{\mathrm{W}^\rho(A,t)}{\mathrm{W}^\rho(A^*)} \geq \frac{1}{n}$ and with a total transfer $\sum_i |t_i|$ of at most $2n^2$. This allocation can be computed in polynomial time.*

Here n is the number of agents. Note that the total transfer is independent of the number m of items. This implies, as m grows, that the transfer payments are negligible in terms of the number of items (and of total welfare). In particular, our ultimate objective is to obtain both envy-freeness and high welfare using negligible transfers. Of course, the welfare guarantee of $\frac{1}{n}$ does not signify high welfare. So we investigate whether improved bounds can be obtained for the important special cases of $\rho = 0$ (Nash social welfare) and $\rho = 1$ (utilitarian social welfare). Strong guarantees on welfare can be obtained for the former. Specifically, there exists an envy-free allocation with transfers with a Nash social welfare that is at least an $e^{-1/e} \approx 0.6922$ fraction of the optimal welfare.

Theorem 5.3.1. *For general valuations, there exists an envy-free allocation with transfers (A, t) such that $\frac{\mathrm{NSW}(A,t)}{\mathrm{NSW}(A^*)} \geq e^{-1/e}$.*

Furthermore, for additive valuations, such constant factor welfare guarantees can be obtained with negligible transfer payments.

Theorem 5.3.3. *For additive valuations, given an α-approximate allocation to maximum Nash social welfare, there exists a polynomial time computable envy-free allocation with transfers (A, t) such that $\frac{\mathrm{NSW}(A,t)}{\mathrm{NSW}(A^*)} \geq \frac{1}{2}\alpha \cdot e^{-1/e}$ with a total transfer $\sum_i |t_i|$ of at most $2n^2$.*

In sharp contrast, for utilitarian social welfare, the factor $\frac{1}{n}$ welfare threshold is tight. To achieve any welfare guarantee greater than $\frac{1}{n}$ requires non-negligible transfer payments. Specifically, we show

Corollary 5.4.2. *For any $\alpha \in \left[\frac{1}{n}, 1\right]$, there exists an instance with additive valuations such that any envy-free allocation with transfers (A, t) satisfying $\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \geq \alpha$ requires a total transfer $\sum_{i \in N} |t_i|$ of at least $\frac{1}{4}\left(\alpha - \frac{1}{n}\right)^2 m$.*

In fact, there exist instances for which any EFk allocation with $k = o(m)$ has a welfare guarantee of at most $\frac{1}{n} + o(1)$ (Lemma 5.4.1). This implies that EFk allocations cannot provide higher welfare with moderate transfers.

On the positive side, we can design algorithms to produce envy-free allocations with welfare guarantee α whose total transfer payment is comparable to the minimum amount possible, quantified in terms of the maximum value $\max\limits_i v_i(A_i^*)$ any agent has in the welfare-maximizing allocation.

Theorem 5.4.4. *For additive valuations, for any $\alpha \in (0, 1]$, there is a polynomial time computable envy-free allocation with transfers (A, t) such that $\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \geq \alpha$ with total transfer $\sum_{i \in N} |t_i| \leq n(\alpha \max\limits_i v_i(A_i^*) + 2)$.*

Theorem 5.4.5. *For general valuations, for any $\alpha \in \left(0, \frac{1}{3}\right]$, there is an envy-free allocation with transfers (A, t) such that $\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \geq \alpha$ with total transfer $\sum_{i \in N} |t_i| \leq 2n^2 \left(3\alpha \max_i v_i \left(A_i^*\right) + 2\right)$.*

In Section 5.1, we present our model of the fair division problem with transfers. Section 5.2 contains an exposition of the prior results in the literature that will be useful, along with our preliminary results on the ρ-mean welfare of envy-free allocations with transfers. In Section 5.3, we present our results on Nash social welfare, and in Section 5.4 we present our results on utilitarian social welfare.

5.1 The Model and Preliminaries

Let $M = \{1, \cdots, m\}$ be a set of m indivisible items and let $N = \{1, \cdots, n\}$ be a set of agents. Let $v_i : 2^M \to \mathbb{R}$ be agent i's valuation function, where $v_i(\emptyset) = 0$. We make the standard assumption that each valuation function is *monotone*, satisfying $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$. Additionally, following previous work on subsidies, just like in the previous chapter we uniformly scale (w.l.o.g.) the valuation functions by the same factor for each agent so that the maximum marginal value of any item is at most 1. Besides general monotone valuations, we are also interested in well-known classes of valuation function, in particular, additive valuations, and *subadditive* (complement-free) valuations where $v(S \cup T) \leq v(S) + v(T)$ for all $S, T \subseteq M$. We use $[n]$ to denote the set $\{1, \cdots, n\}$.

5.1.1 Fairness and Welfare

An allocation $A = (A_1, A_2, \cdots, A_n)$ is a partition of the items into n disjoint subsets, where A_i is the set of items allocated to agent i. Our aim is to obtain envy-free allocations with high welfare. As before, we say an allocation with payments (A, p) is *envy-free* if for each $i, j \in N$, $v_i(A_i) + p_i \geq v_i(A_j) + p_j$, and an allocation envy-freeable if there exist payments p such that (A, p) is envy-free. There are two natural types of payment. First, we have *subsidy payments* if $p_i \geq 0$. Second, we have *transfer payments* if $\sum_{i \in N} p_i = 0$, To distinguish these, we denote a subsidy payment to agent i by s_i and a transfer payment by t_i. We define the *total transfer* of an allocation as the sum $\sum_i |t_i|$.

We measure the welfare of an allocation A using the general concept of *ρ-mean welfare*, $\mathrm{W}^\rho(A) = \left(\frac{1}{n}\sum_{i\in N} v_i(A_i)^\rho\right)^{\frac{1}{\rho}}$. This class of welfare functions, introduced by Arunachaleswaran et al. [6], encompasses a range of welfare functions including the two most important cases: $\rho \to 0$, the *Nash social welfare*, is the geometric mean of the values of the agents, denoted by $\mathrm{NSW}(A) = \left(\prod_{i\in N} v_i(A_i)\right)^{\frac{1}{n}}$, and $\rho = 1$, the *utilitarian social welfare* or simply *social welfare* (scaling by the number of agents), denoted by $\mathrm{SW}(A) = \sum_{i\in N} v_i(A_i)$. With transfer payments, our interest lies in utilities rather than simply valuations. In particular, the ρ-mean welfare of an allocation with transfers (A, t) is $\mathrm{W}^\rho(A, t) = \left(\frac{1}{n}\sum_{i\in N}(v_i(A_i) + t_i)^\rho\right)^{\frac{1}{\rho}}$.

5.1.2 Fair Division With Transfer Payments

In this chapter, we study the following question.

Is there an allocation with transfers that simultaneously satisfies
(i) envy-freeness, (ii) high welfare, and (iii) a negligible total transfer?

We have seen that envy-freeable allocations always exist. Thus, with transfer payments, we can obtain the property of envy-freeness. The reader may ask whether transfers are necessary. Specifically, given the guaranteed existence of envy-freeable allocation, can such allocations provide high welfare? The answer is *no*. Even worse, no positive guarantee on welfare can be obtained without transfers. This is true even for the case of additive valuations. To see this, consider the following simple example for Nash social welfare.

Example 5.1.1. *Take two agents and two items $\{a, b\}$. Let the valuation functions be additive with $v_{1,a} = 1, v_{1,b} = \frac{1}{2}$ for agent 1 and $v_{2,a} = \frac{1}{2}, v_{2,b} = \epsilon$ for agent 2. Observe there are only two envy-freeable allocations: either agent 1 gets both items or agent 1 gets item a and agent 2 gets item b. For both these envy-freeable allocations the corresponding Nash social welfare is at most $\sqrt{\epsilon}$. In contrast, the optimal Nash social welfare is $\frac{1}{2}$ when agent 1 gets b and agent 2 gets a.*

It follows that to find envy-free solutions with non-zero approximation guarantees for welfare we must have transfer payments. At the outset, if we restrict ρ to be equal to 1, the result of Halpern and Shah [44] implies that the allocation that maximizes utilitarian welfare can be made envy-free with transfer payments. However, we show that this allocation can require arbitrarily large transfers relative to the number of agents. The main point of concern in using transfer payments to achieve envy-freeness is that it may be difficult for the participants to include a substantial quantity of money in the system in order to implement this solution. Consequently, this creates a third requirement, i.e. to bound the total transfers. Thus the holy grail here is to obtain high welfare using only *negligible transfers*: formally, we desire transfers whose sum (of absolute values) is independent of the number of items m. In particular, we want an allocation with transfers (A, t) such that the welfare of A is at least α times the welfare of the welfare-maximizing allocation A^* (for some large $\alpha \in [0, 1]$) and $\sum_{i \in N} |t_i| = O(f(n))$ for some function f. Specifically, the payments are negligible in the number of items (and thus in the total welfare) as m grows.

At first glance, this task seems impossible. If envy-freeable solutions cannot themselves ensure non-zero welfare guarantees, how could negligible transfer payments induce high welfare? Surprisingly, this is possible for some important classes of valuation functions and types of welfare. However, it is indeed not possible for other classes and types. Investigating the boundary of this dichotomy is the purpose of this chapter.

5.2 Transfer Payments and ρ-Mean Welfare

In this section we familiarize the reader with the structure of envy-freeable allocations and transfer payments, and introduce our preliminary results. We begin with the general case of ρ-mean welfare.

Lemma 5.2.1. *For subadditive valuations, any envy-free allocation with transfers (A, t) satisfies $\mathrm{W}^\rho(A, t) \geq \frac{1}{n} \mathrm{W}^\rho(A^*)$.*

Proof. By the envy-freeness property $v_i(A_i) + t_i \geq v_i(A_j) + t_j$. Thus

$$v_i(A_i) + t_i \geq \frac{1}{n}\left(\sum_j v_i(A_j) + t_j\right) \geq \frac{1}{n} v_i(M) \geq \frac{1}{n} v_i(A_i^*)$$

Here the second inequality follows by subadditivity. Hence

$$\begin{aligned} \mathrm{W}^\rho(A, t) &= \left(\frac{1}{n}\sum_{i\in N}(v_i(A_i) + t_i)^\rho\right)^{\frac{1}{\rho}} \geq \left(\frac{1}{n}\sum_{i\in N}\left(\frac{v_i(A_i^*)}{n}\right)^\rho\right)^{\frac{1}{\rho}} \\ &= \frac{1}{n}\left(\frac{1}{n}\sum_{i\in N}\left(v_i(A_i^*)\right)^\rho\right)^{\frac{1}{\rho}} = \frac{1}{n}\,\mathrm{W}^\rho(A^*) \end{aligned}$$

as desired. □

The resultant welfare guarantee of $\alpha = \frac{1}{n}$ is not particularly impressive. But it is a strictly positive guarantee, which was unachievable without transfer payments. The bound is also tight as shown by the following simple example.

Example 5.2.2. *Take $m = n$ items and n agents. Let the valuation functions be additive with $v_{ii} = 1$ and $v_{ij} = 0$ for $j \neq i$. Consider the allocation assigning the grand bundle to agent 1. This is envy-freeable with transfer payments $t_1 = -\frac{n-1}{n}$ and $t_i = \frac{1}{n}$, for any agent $i \neq 1$. For social welfare ($\rho = 1$) the corresponding welfare guarantee is $\alpha = \frac{1}{n}$.*

But how expensive is it to obtain this welfare guarantee? To answer this, we provide a short review concerning the computation of transfer payments. Recall that an allocation A is envy-freeable if there exist payments p such that (A, p) is envy-free. The characterization due to Halpern and Shah [44] described in the previous chapter shows us how to find, for any envy-freeable allocation A, the minimum *subsidy* payments s such that (A, s) is envy-free. Let $l(i)$ be weight of a maximum weight path from node i to any other node in G_A. Setting $s_i = l(i)$ for each agent i gives an envy-free allocation with minimum subsidy payments. We do not wish to subsidize the mechanism, so we convert these subsidies into transfer payments. To do this, let $\bar{s} = \frac{1}{n}\sum_{i\in N} s_i$ be the average subsidy. Then setting $t_i = s_i - \bar{s}$ for each agent gives a valid set of transfer payments, which we dub the *natural* transfer payments. We remark that the natural transfer payments do not always minimize the total transfer, but they will be sufficient for our purposes. We are now ready to compute transfer payments for subadditive valuations in the ρ-mean welfare setting. We begin by restating a theorem from the previous chapter.

Theorem 4.0.4. *For monotone valuations there is a polytime algorithm to find an envy-free allocation with subsidies (A, s) with $s_i \leq 2(n-1)$ for all i.*

Observe that any bound on the maximum subsidy for each agent also applies to the maximum natural transfer for each agent. Combining this observation with the previous result gives us the following corollary.

Corollary 5.2.3. *For subadditive valuations, there exists an envy-free allocation with transfers (A, t) such that $\frac{\mathrm{W}^{\rho}(A,t)}{\mathrm{W}^{\rho}(A^*)} \geq \frac{1}{n}$ and with a total transfer $\sum_i |t_i|$ of at most $2n^2$. This allocation can be computed in polytime.*

Thus, we can quickly obtain an envy-free allocation with transfers whose total transfer is negligible, i.e., independent of m. But, as stated, we only have a low welfare guarantee for this general ρ-mean welfare class. In the next section, we will show that high welfare and negligible transfers are achievable for the special case of $\rho = 0$, that is, Nash Social Werlfare. First, we conclude this section by presenting a generalization of our earlier theorem that will later be useful. We say that an allocation B has *b-bounded envy* if $v_i(B_j) - v_i(B_i) \leq b$ for every pair $i, j \in N$.

Lemma 5.2.4. *Given an allocation B with b-bounded envy there is a polytime algorithm to find an envy-free allocation with transfers (A, t) with $\sum_{i \in N} |t_i| \leq 2bn^2$.*

Proof. Let $B = \{B_1, B_2, \ldots, B_n\}$ be an allocation with b-bounded envy. Consider the envy-freeable allocation $A = (B_{\pi(1)}, \cdots, B_{\pi(n)})$ obtained by computing a maximum-weight matching between the bundles in B and the agents. Applying an approach from the previous chapter, let P be a path of maximum weight in the envy-graph G_A. Without loss of generality, $P = (1, \cdots, r)$. By definition of the envy-graph, we then have

$$\begin{aligned} w_A(P) &= \sum_{i=1}^{r-1} v_i(A_{i+1}) - v_i(A_i) \\ &= \sum_{i=1}^{r-1} v_i(A_{i+1}) - v_i(B_i) + v_i(B_i) - v_i(A_i) \\ &= \sum_{i=1}^{r-1} v_i(B_{\pi(i+1)}) - v_i(B_i) + \sum_{i=1}^{r-1} v_i(B_i) - v_i(A_i) \\ &\leq b(n-1) + \sum_{i=1}^{r-1} v_i(B_i) - v_i(A_i) \end{aligned} \tag{5.1}$$

Here the inequality holds as $v_i(B_{\pi(i+1)}) - v_i(B_i) \leq b$ for each agent i, and $r < n$. We have

$$\begin{aligned} \sum_{i=1}^{r-1} v_i(B_i) - v_i(A_i) &\leq \sum_{i:v_i(B_i)\geq v_i(A_i)} v_i(B_i) - v_i(A_i) \\ &\leq - \sum_{i:v_i(B_i)<v_i(A_i)} v_i(B_i) - v_i(A_i) \\ &= \sum_{i:v_i(B_i)<v_i(A_i)} v_i(A_i) - v_i(B_i) \\ &= \sum_{i:v_i(B_i)<v_i(A_i)} v_i(B_{\pi(i)}) - v_i(B_i) \\ &\leq b(n-1). \end{aligned} \tag{5.2}$$

Above the second inequality holds as the social welfare of A is the maximum over all allocations of the bundles in B; in particular, $\sum_i v_i(A_i) \geq \sum_i v_i(B_i)$. The last inequality again follows as B has b-bounded envy.

Together (5.1) and (5.2) give $w_A(P) \leq 2b(n-1)$. This implies that A can be made envy-free with a subsidy $s_i \leq 2b(n-1)$ to each agent i. Hence, setting $t_i = s_i - \bar{s}$, we have that (A, t) is envy-free with a total transfer payment of at most $\sum_{i \in N} |t_i| \leq 2bn^2$. $\square$

5.3 Transfer Payments and Nash Social Welfare

In the following two sections, we present our main results concerning Nash social welfare and utilitarian social welfare. Here we show that, with transfers, excellent welfare guarantees can be obtained for NSW. Conversely, in Section 5.4, we will see that only much weaker guarantees can be obtained for utilitarian welfare.

5.3.1 NSW with General Valuation Functions

Now, recall from Example 5.1.1 that no positive welfare guarantee can be obtained in the case of Nash social welfare for even the basic case of additive valuations. Our first result for Nash social welfare is therefore somewhat surprising. With transfer payments, constant factor welfare guarantees can be obtained for general valuations. That is, envy-freeness and high welfare are simultaneously achievable.

Theorem 5.3.1. *For general valuations, there exists an envy-free allocation with transfers* (A, t) *such that* $\frac{\text{NSW}(A,t)}{\text{NSW}(A^*)} \geq e^{-1/e}$.

Proof. Let A^* be an allocation that maximizes Nash social welfare. Now, let A be an envy-freeable allocation induced by reallocating the bundles in A^* to maximize utilitarian social welfare. Recall this can be found by taking a maximum weight matching between the agents and the bundles of A^*; let $\pi(i)$ be the agent who receives bundle A_i^* in the allocation A. By Theorem 4.2.1, this allocation is envy-freeable. So let t be any valid set of transfer payments such that (A, t) is envy-free.

By definition we have that $v_i(A_i^*) = v_i(A_{\pi(i)})$, for all $i \in N$. Then, by envy-freeness, we have $v_i(A_i) + t_i \geq v_i(A_{\pi(i)}) + t_{\pi(i)} = v_i(A_i^*) + t_{\pi(i)}$. Denote by $t_{\max}$ the maximum positive transfer payment, i.e. $t_{\max} = \max_i t_i$, and let m be an agent whose transfer t_m is equal to $t_{\max}$. By envy-freeness, no agent envies agent m, so $v_i(A_i) + t_i \geq t_{\max}$ for all i. Putting this all together, we have

$$\frac{\prod_{i=1}^{n} v_i(A_i) + t_i}{\prod_{i=1}^{n} v_i(A_i^*)} \geq \prod_{i=1}^{n} \frac{\max\left[v_i(A_i^*) + t_{\pi(i)}, t_{\max}\right]}{v_i(A_i^*)}$$

Now define $N^+ = \{i \mid t_{\pi(i)} \geq 0\}$ and $N^- = N \setminus N^+$.

$$\begin{aligned}
\prod_{i=1}^{n} \frac{\max\left[v_i(A_i^*) + t_{\pi(i)}, t_{\max}\right]}{v_i(A_i^*)} &\geq \prod_{i \in N^+} \frac{v_i(A_i^*) + t_{\pi(i)}}{v_i(A_i^*)} \cdot \prod_{i \in N^-} \frac{\max\left[v_i(A_i^*) + t_{\pi(i)}, t_{\max}\right]}{v_i(A_i^*)} \\
&\geq \prod_{i \in N^-} \frac{\max\left[v_i(A_i^*) + t_{\pi(i)}, t_{\max}\right]}{v_i(A_i^*)}
\end{aligned}$$

Next let N_1^- be the indices corresponding to negative transfers that also satisfy $t_{\max} \leq v_i(A_i^*)+t_{\pi(i)}$, and let N_2^- be the indices corresponding to negative transfers that also satisfy $t_{\max} > v_i(A_i^*) + t_{\pi(i)}$. Furthermore, set $v_i(A_i^*) + t_{\pi(i)} = t_{\max} + \alpha_i$. Observe that, for $i \in N_1^-$, we have $\alpha_i \geq 0$, but for $i \in N_2^-$, we have $\alpha_i < 0$. Applying this gives

$$\begin{aligned}
\prod_{i \in N^-} \frac{\max\left[v_i(A_i^*) + t_{\pi(i)}, t_{\max}\right]}{v_i(A_i^*)} &\geq \left(\prod_{i \in N_1^-} \frac{t_{\max} + \alpha_i}{t_{\max} + \alpha_i - t_{\pi(i)}}\right) \cdot \left(\prod_{i \in N_2^-} \frac{t_{\max}}{t_{\max} + \alpha_i - t_{\pi(i)}}\right) \\
&\geq \left(\prod_{i \in N_1^-} \frac{t_{\max}}{t_{\max} - t_{\pi(i)}}\right) \cdot \left(\prod_{i \in N_2^-} \frac{t_{\max}}{t_{\max} - |\alpha_i| - t_{\pi(i)}}\right) \\
&\geq \left(\prod_{i \in N_1^-} \frac{t_{\max}}{t_{\max} - t_{\pi(i)}}\right) \cdot \left(\prod_{i \in N_2^-} \frac{t_{\max}}{t_{\max} - t_{\pi(i)}}\right) \\
&= \left(\prod_{i \in N^-} \frac{t_{\max}}{t_{\max} - t_{\pi(i)}}\right)
\end{aligned}$$

Now for, $i \in N^-$, let $k_i = |t_{\pi(i)}|$. Since $\sum_{i \in N} t_i = 0$ we have $\sum_{i \in N^+} t_i = \sum_{i \in N^-} t_i := T$. Thus

$$\left(\frac{\prod_{i=1}^{n} v_i(A_i) + t_i}{\prod_{i=1}^{n} v_i(A_i^*)}\right)^{1/n} \geq \left(\prod_{i \in N^-} \frac{t_{\max}}{t_{\max} - t_{\pi(i)}}\right)^{1/n} = \left(\prod_{i \in N^-} \frac{t_{\max}}{t_{\max} + k_i}\right)^{1/n}$$

Observe, by the *arithmetic-geometric mean inequality*, that $\prod_{i \in N^-}(t_{\max} + k_i)$ is maximized when $k_i = k_j = T/|N^-|$. In addition, $t_{\max} \geq T/|N^+|$. So

$$\left(\prod_{i \in N^-} \frac{t_{\max}}{t_{\max} + k_i} \right)^{1/n} \geq \left(\frac{\frac{T}{|N^+|}}{\frac{T}{|N^+|} + \frac{T}{|N^-|}} \right)^{|N^-|/n} = \left(\frac{n - |N^+|}{n} \right)^{\frac{n-|N^+|}{n}}$$
$$\geq \min_x \left(\frac{1}{x} \right)^{\frac{1}{x}} \geq e^{-1/e} \qquad \square$$

This theorem is noteworthy; for general valuation functions, with transfers, it allows us to simultaneously obtain high Nash social welfare and envy-freeness. But what of our third objective, negligible transfer payments? The approach applied in the proof of Theorem 5.3.1 cannot guarantee negligible transfers. Specifically, simply reallocating the bundles of the allocation A^* that maximizes Nash social welfare can require large transfers. In particular, the following example shows this method may require transfers as large as $\Omega(\sqrt{m})$.

Example 5.3.2. *Take an instance with two agents and m items. Assume the first agent has a valuation function given by $v_1(S) = |S|$, for each $S \subseteq M$; assume the second agent has a valuation function given by $v_2(S) = \sqrt{|S|}$, for each $S \subseteq M$. The reader may verify that the Nash welfare maximizing allocation A^* is to give the first agent $\frac{2m}{3}$ items and the second agent $\frac{m}{3}$ items. This allocation is also the allocation that maximizes utilitarian social welfare by reassigning the bundles of A^*. Thus $A = A^*$. However, to make the allocation envy-free requires a minimum transfer payment of $\Omega(\sqrt{m})$, from the first agent to the second agent.*

Of course, this example does not rule out the possibility that, for general valuation functions, an envy-free allocation with transfers that has high welfare and negligible payments exists. In particular, simply allocating each agent half the items requires no transfer payments at all, and gives high Nash social welfare. So simultaneously obtaining high Nash social welfare and envy-freeness via negligible transfers for general valuation functions remains an open question. Fortunately, we can show that these three properties are simultaneously achievable for important special classes of valuation function.

5.3.2 NSW Guarantees with Negligible Transfers

Here we prove that for (i) additive valuations, and (ii) matroid rank valuations, it is always possible to obtain envy-free allocations with high Nash social welfare and negligible transfers. Furthermore, for additive valuations we can do this using polynomial time algorithms.

Theorem 5.3.3. *For additive valuations, given an α-approximate allocation to maximum Nash social welfare, there exists a polynomial time computable envy-free allocation with transfers (A, t) such that $\frac{\text{NSW}(A,t)}{\text{NSW}(A^*)} \geq \frac{1}{2}\alpha \cdot e^{-1/e}$ with a total transfer $\sum_i |t_i|$ of at most $2n^2$.*

Proof. Let B be the α-approximate allocation to the maximum Nash social welfare; that is $\frac{\text{NSW}(B)}{\text{NSW}(A^*)} \geq \alpha$. Now Caragiannis et al. [26] gave a polytime algorithm which, given input B, outputs an EF1 allocation B' with a Nash social welfare guarantee of $\frac{\alpha}{2}$.

Next, recall the proof of Theorem 5.3.1. Observe that, during the proof, we did not use the fact that A^* maximizes Nash social welfare. Thus the $e^{-1/e}$ approximation ratio holds if we start with any other allocation $\hat{A}$ instead of A^*. That is by reallocation the bundles of $\hat{A}$ we obtain an envy-freeable allocation A whose Nash social welfare is that least a factor $e^{-1/e}$ of that of A^*. In particular, we can do this for the allocation $\hat{A} = B'$ given by Caragiannis et al [26]. So, by Theorem 5.3.1, there exists an envy-free allocation with transfers (A, t) such that $\frac{\text{NSW}(A,t)}{\text{NSW}(B')} \geq e^{-1/e}$. Now

$$\frac{\text{NSW}(A,t)}{\text{NSW}(A^*)} = \frac{\text{NSW}(B')}{\text{NSW}(A^*)} \cdot \frac{\text{NSW}(A,t)}{\text{NSW}(B')} \geq \frac{1}{2}\alpha \cdot e^{-1/e}$$

Furthermore, because B' is EF1 and A is obtained by the same procedure as in Lemma 5.2.4, we obtain transfer payments with $\sum_i |t_i| \leq 2n^2$. □

We remark that, for additive valuations, polytime algorithms to find allocations that α-approximate the maximum NSW do exist. Specifically, Barman et al. [11] present an algorithm with an approximation guarantee of $\alpha = \frac{1}{1.45}$. Together with Theorem 5.3.3, we thus obtain in polytime an envy-free allocation with negligible transfers and a Nash social welfare guarentee of $\frac{1}{2.9}e^{-1/e}$.

Better existence bounds can be obtained for the additive case if we remove the requirement of a polytime algorithm. A well-known result of Caragiannis et al. [27] states that for additive valuations, the Nash welfare maximizing allocation is EF1. In fact, a recent result of Benabbou et al. [17] provides a similar statement for the case of *matroid rank* valuation functions, a sub-class of submodular functions. A valuation function is matroid rank if it is submodular, and the marginal value of any item is binary (i.e. for any set S of items and any item x not in S, $v_i(S \cup \{x\}) - v_i(S) \in \{0,1\}$). Here, an NSW-maximizing allocation is EF1[17]. Combining this with Lemma 5.2.4, the corresponding envy-free allocation with transfers (A,t) has transfers satisfying $\sum_i |t_i| \leq 2n^2$. Further, by Theorem 5.3.1, we have $\frac{\mathrm{NSW}(A,t)}{\mathrm{NSW}(A^*)} \geq e^{-1/e}$ as desired.

Theorem 5.3.4. *For matroid rank valuations, there exists an envy-free allocation with transfers* (A,t) *with* $\frac{\mathrm{NSW}(A,t)}{\mathrm{NSW}(A^*)} \geq e^{-1/e}$ *and* $\sum_i |t_i| \leq 2n^2$. □

5.4 Transfer Payments and Utilitarian Social Welfare

In this section we present our results on utilitarian social welfare, which differ markedly from those in the previous section. To begin, recall that an allocation B has *b-bounded envy* if $v_i(B_j) - v_i(B_i) \leq b$ for every pair of agents $i, j \in N$. Without transfers, allocations with *b-bounded envy* may have very low welfare.

Lemma 5.4.1. *For utilitarian social welfare, there exist instances with additive valuation functions such that any allocation with b-bounded envy has a welfare guarantee of at most* $2\sqrt{\frac{b}{m}} + \frac{1}{n}$.

Proof. Consider the following instance with additive valuations. Let $v_{n,j} = 1$ for each $j \in M$ and let $v_{ij} = \epsilon$ for all $i \neq n$ and all $j \in M$. Evidently, to maximize utilitarian social welfare we simply give all the items to agent n. So $\mathrm{SW}(A^*) = m$. Because the items are interchangeable for every agent, any allocation A can be described as $(y_1, \cdots, y_{n-1}, y_n =$

x), where y_i is the fraction of items allocated to agent i. to the agent. Since every item must be allocated, we have $\sum_{i=1}^{n-1} y_i = (1-x)$. The corresponding welfare guarantee for the allocation A is then $\frac{\text{SW}(A)}{\text{SW}(A^*)} = (1-\epsilon)x + \epsilon$.

Now suppose A has b-bounded envy. Therefore, $v_i(A_j) - v_i(A_i) \leq b$, for any pair of agents $i, j \in N$. In particular, $m(y_i - x) \leq b$ since agent n cannot envy agent i too much and $\epsilon m(x - y_i) \leq b$ since agent i cannot envy agent n too much. Summing the later inequality over all agents i gives $\epsilon m((n-1)x - (1-x)) \leq (n-1)b$. This implies $x \leq (1 - \frac{1}{n})\frac{b}{\epsilon m} + \frac{1}{n}$. Thus

$$
\begin{aligned}
\frac{\text{SW}(A)}{\text{SW}(A^*)} &= (1-\epsilon)x + \epsilon \\
&\leq (1-\epsilon)\left(\left(1 - \frac{1}{n}\right)\frac{b}{\epsilon m} + \frac{1}{n}\right) + \epsilon \\
&\leq 2\sqrt{\frac{b}{m}} - \frac{b}{m} + \frac{1}{n} \\
&\leq 2\sqrt{\frac{b}{m}} + \frac{1}{n}
\end{aligned}
$$

Here the second inequality holds by setting $\epsilon = \sqrt{\frac{b}{m}}$. □

Lemma 5.4.1 implies that any EFk allocation in the given example, with $k = o(m)$, cannot provide a welfare guarantee that is significantly higher than $\frac{1}{n}$. The natural question to ask, now, is whether the problem inherent in Lemma 5.4.1 can be rectified with a small quantity of transfers. On the positive side, the result from the previous chapter shows that a small quantity of subsidy independent of the number of items is always sufficient to eliminate envy. A similar result also extends to the corresponding natural transfer payments. Combining this result with Lemma 5.2.1 tells us that a utilitarian welfare guarantee of $\frac{1}{n}$ can be achieved alongside envy-freeness with a negligible total transfer. Unfortunately, for the above example, the Iterated Matching Algorithm returns an allocation whose social welfare is only a $\frac{1}{n}$-fraction of the optimal welfare. The following corollary shows that this was inevitable: unlike for NSW, in order to make any improvement above this threshold, non-negligible transfers are required.

Corollary 5.4.2. *For any $\alpha \in \left[\frac{1}{n}, 1\right]$, there exists an instance with additive valuations such that any envy-free allocation with transfers (A, t) satisfying $\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \geq \alpha$ requires a total transfer $\sum_{i \in N} |t_i| \geq \frac{1}{4}\left(\alpha - \frac{1}{n}\right)^2 m$.*

Proof. Take the same instance as in Lemma 5.4.1. Now for utilitarian social welfare, we have $\mathrm{SW}(A) = \mathrm{SW}(A, t)$ as $\sum_{i \in N}(v_i(A_i) + t_i) = \sum_{i \in N} v_i(A_i) + \sum_{i \in N} t_i = \sum_{i \in N} v_i(A_i)$. Let A^* be the allocation that maximizes the social welfare. Thus

$$\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \; = \; (1-\epsilon)x + \epsilon \; \geq \; \alpha \tag{5.3}$$

Next, observe that $x = y_n \geq y_i$, for each i otherwise the allocation A is not envy-freeable. Thus $t_n \leq 0$. Then, by envy-freeness of (A, t), we must have $mx + t_n \geq my_i + t_i$ and $\epsilon m y_i + t_i \geq \epsilon m x + t_n$. It follows that

$$(n-1) \cdot (\epsilon m x + t_n) \; \leq \; \epsilon m \cdot \sum_{i=1}^{n-1} y_i + \sum_{i=1}^{n-1} t_i \; = \; \epsilon m \cdot (1 - x) - t_n$$

Rearranging we obtain $n \cdot (\epsilon m x + t_n) \leq \epsilon m$. In particular,

$$-t_n \; \geq \; \epsilon m \cdot \left(x - \frac{1}{n}\right) \tag{5.4}$$

Combining (5.3) and (5.4) we get

$$\sum_{i \in N} |t_i| \; \geq \; |t_n| \; \geq \; -t_n \; \geq \; \epsilon m \cdot \left(x - \frac{1}{n}\right) \; \geq \; \epsilon m \cdot \left(\frac{\alpha - \epsilon}{1 - \epsilon} - \frac{1}{n}\right)$$

Finally, choosing $\epsilon = 1 - \sqrt{\frac{1-\alpha}{1-\frac{1}{n}}}$ gives the desired bound. $\square$

So, for utilitarian social welfare, non-negligible transfers are required to ensure both envy-freeness and high welfare. Recall, though, that balancing constraints on the valuation functions have been used in the literature to sidestep impossibility bounds on welfare. The reader may wonder if such constraints could be used to bypass the result in Corollary 5.4.2: are negligible transfer payments sufficient to obtain high welfare when the valuation functions are constant-sum? The answer is NO, as we shall see in the subsequent theorem.

In recent work, Barman et al. [12] considered the case of subadditive valuations with the constant-sum condition, and gave a polynomial-time algorithm that finds an EF1 allocation with social welfare at least $\Omega(\frac{1}{\sqrt{n}})$ of the optimal welfare. Applying the algorithm of Lemma 5.2.4 to the resulting allocation gives us an envy-free allocation with negligible transfers and welfare ratio $\Omega(\frac{1}{\sqrt{n}})$. Once again, we show that this threshold cannot be crossed without non-negligible transfers.

Theorem 5.4.3. *There exist instances with constant-sum additive valuations such that any envy-free allocation with transfers (A, t) satisfying $\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \geq \alpha$ has a total transfer $\sum_{i \in N} |t_i| \geq (\alpha - \frac{2}{\sqrt{n}})\frac{m}{\sqrt{n}}$, for any $\alpha \in [\frac{2}{\sqrt{n}}, 1]$.*

Proof. Consider an instance with m items and n agents. Divide the items into $\sqrt{n}$ sets, each of cardinality $\frac{m}{\sqrt{n}}$. Let B_ℓ be the set of items $\{(\ell-1)\frac{m}{\sqrt{n}}+1, (\ell-1)\frac{m}{\sqrt{n}}+2, \cdots, \ell\frac{m}{\sqrt{n}}\}$. We now define a collection of constant-sum additive valuation functions. We partition the set of agents into two parts; agents in the set $H = \{1, \cdots, \sqrt{n}\}$ have high value for a small number of items, and agents in the set $L = \{\sqrt{n}+1, \cdots, n\}$ have low value for a large number of items. A high value agent i has valuations $v_{ij} = 1$ for $j \in B_i$ and zero otherwise. Thus for each agent $i \in H$ there is a corresponding set B_i which it values. Each low value agent has a uniform valuation of $v_{ij} = \frac{1}{\sqrt{n}}$ for all $j \in M$. Observe that the value each agent has for the grand bundle is exactly $\frac{m}{\sqrt{n}}$, that is, constant-sum. Note that any allocation to a high value agent i can be described by the fraction of B_i which it receives. Consider an envy-freeable allocation A that assigns an x_i-fraction and a y_{ki}-fraction of B_i to $i \in H$ and $k \in L$ respectively. By envy-freeability we must have $x_i \geq y_{ki}$ for all $i \in H$ and $k \in L$. We also have that $x_i + \sum_{k \in L} y_{ki} = 1$. Observe that the utilitarian social welfare is maximized by allocating B_i to the high value agent i; this allocation satisfies

$\mathrm{SW}(A^*) = m$. We then have

$$\begin{aligned}
\frac{\mathrm{SW}(A)}{\mathrm{SW}(A^*)} &= \frac{1}{m}\left(\frac{m}{\sqrt{n}}\sum_{i\in H} x_i + \frac{m}{\sqrt{n}}\sum_{i\in H}\sum_{k\in L}\frac{y_{ki}}{\sqrt{n}}\right) \\
&= \frac{1}{m}\left(\frac{m}{\sqrt{n}}\sum_{i\in H} x_i + \frac{m}{n}\sum_{i\in H}(1-x_i)\right) \\
&= \frac{1}{\sqrt{n}}\left(\left(1-\frac{1}{\sqrt{n}}\right)\sum_{i\in H} x_i + 1\right) \\
&\geq \alpha
\end{aligned}$$

From this we can infer that $\sum_{i\in H} x_i \geq \sqrt{n}\alpha - 1$. Now, let t be valid transfer payments. By envy-freeness, we see that for any i, k

$$\frac{m}{\sqrt{n}} \cdot \sum_{j\in H} \frac{y_{kj}}{\sqrt{n}} + t_k \geq \frac{m}{\sqrt{n}} \cdot \frac{x_i}{\sqrt{n}} + t_i$$

First summing over $i \in H$ and then summing over $k \in L$ gives

$$\frac{m}{\sqrt{n}}\sum_{k\in L}\sum_{j\in H} y_{kj} + \sqrt{n}\sum_{k\in L} t_k \;\geq\; m\left(1-\frac{1}{\sqrt{n}}\right)\sum_{i\in H} x_i + (n-\sqrt{n})\sum_{i\in H} t_i$$

Sum $\sum_{i\in N} t_i = 0$, rearranging gives

$$0 \;\geq\; m\sum_{i\in H} x_i - \frac{m}{\sqrt{n}}\sum_{i\in H}\left(x_i + \sum_{k\in L} y_{ki}\right) + n\sum_{i\in H} t_i \;\geq\; m\left(\sum_{i\in H} x_i - 1\right) + n\sum_{i\in H} t_i$$

In particular,

$$-\sum_{i\in H} t_i \;\geq\; \frac{m}{n}\left(\sum_{i\in H} x_i - 1\right)$$

Now recall that $\sum_{i \in H} x_i \geq \sqrt{n}\alpha - 1$. Thus

$$\sum_{i \in N} |t_i| \;\geq\; -\sum_{i \in H} t_i \;\geq\; \frac{m}{\sqrt{n}}\left(\alpha - \frac{2}{\sqrt{n}}\right) \;\geq\; 0 \qquad \square$$

So non-negligible transfer payments are required even under constant-sum valuations. This adds to our collection of negative results for utilitarian social welfare. Are any positive results possible? Specifically, can we at least match the lower bounds on transfer payments inherent in the these negative results?

To conclude this chapter, we present results that upper bound the total transfer required to obtain an envy-free allocation with a utilitarian social welfare guarantee. We give upper bounds for additive and general valuation functions. In both cases, the bound we obtain is a function of the maximum value that an agent receives in the welfare-optimal allocation. In particular, while the lower bounds are obtained as functions of m, the upper bounds we get are functions of the product of n and $\max_i v_i(A_i^*)$. In allocations that distribute utility uniformly among the agents these expressions are comparable; even in the worst case, since $v_i(A_i^*) \leq m$ for any i, they differ by some function of only n, and this difference is independent of the number of items. We begin with the additive case.

Theorem 5.4.4. *For additive valuations, for any $\alpha \in (0, 1]$, there is an envy-free allocation with transfers (A, t) such that $\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \geq \alpha$ with total transfer $\sum_{i \in N} |t_i| \leq n(\alpha \max_i v_i(A_i^*) + 2)$.*

Proof. We prove this result with a simple polytime algorithm (Algorithm 2) that outputs the desired allocation with transfers (A, t).

By additivity, the optimal allocation A^* assigns each item in M to an agent with the greatest valuation for that item. Consequently, $X = (X_1, \cdots, X_n)$ maximizes welfare among all reassignments of its bundles, so X is an envy-freeable allocation of the items $\bigcup_{i \in [n]} X_i$. By construction, we have $\mathrm{SW}(X) \geq \alpha\, \mathrm{SW}(A^*)$. Now let P be any path in the

Algorithm 2: Envy-free allocation with high welfare and small transfers for additive valuations

$A_i \leftarrow \emptyset$ for all $i \in N$;
$A^* = (A_1^*, \cdots, A_n^*) \leftarrow$ Welfare-Maximizing Allocation;
for $i = 1$ *to* n **do**
  $A_i \leftarrow$ minimal set $X_i \subseteq A_i^*$ with $v_i(X_i) \geq \alpha \cdot v_i(A_i^*)$
end
Use the Iterated Matching Algorithm to allocate $M \setminus \bigcup_{i \in N} A_i$;
Compute the natural transfers $(t_1, \cdots, t_n)$;
return (A, t)

envy-graph G_X, without loss of generality, $P = (1, 2, \ldots, r)$. Then

$$\begin{aligned} w_X(P) &= \sum_{i=1}^{r-1} v_i(X_{i+1}) - v_i(X_i) \\ &\leq -\left(v_r(X_1) - v_r(X_r)\right) \\ &\leq \max_i v_i(X_i) \\ &\leq \alpha \max_i v_i(A_i^*) + 1 \end{aligned}$$

Here the first inequality holds by Theorem 4.2.1 as the envy-graph contains no positive cycle. The last inequality holds by the minimality of X_i. Next, let $(Y_1, \cdots, Y_n)$ be the allocation of the remaining items $\bigcup_i (A_i^* \setminus X_i)$ given by the Iterated Matching Algorithm. The key properties we require from this algorithm are that the allocation Y is envy-freeable and that, for any path P, weight of the path $w_Y(P) \leq 1$. But Algorithm 2 simply outputs the allocation $A_i = X_i \cup Y_i$ for each agent i. Hence

$$w_A(P) \ = \ w_X(P) + w_Y(P) \ \leq \ \alpha \max_i v_i(A_i^*) + 2$$

Now if we take s to be the minimum subsidy payments required for envy-freeness then $s_i \leq \alpha \max_i v_i(A_i^*) + 2$, for each agent i. Using the transfer payments $t_i = s_i - \bar{s}$, we have that $\sum_i |t_i| \leq n(\alpha \max_i v_i(A_i^*) + 2)$, as claimed. □

Finally, we show how to upper bound the transfer payments in the case of general valuation functions. Here, the welfare target is limited to the constant factor $\frac{1}{3}$, and the gap between our lower and upper bounds widens by a factor of n, but once again, this gap is independent of m.

Theorem 5.4.5. *For general valuations, for any $\alpha \in \left(0, \frac{1}{3}\right]$, there is an envy-free allocation with transfers (A, t) such that $\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \geq \alpha$ with total transfer $\sum_{i \in N} |t_i| \leq 2n^2 \left(3\alpha \max_i v_i\left(A_i^*\right) + 2\right)$.*

Proof. We prove this result using an algorithm (see Algorithm 3) that outputs the desired allocation with transfers (A, t).

Algorithm 3: Envy-free allocation with high welfare and small transfers for general valuations

```
B_i ← ∅ for all i ∈ N;
A* = (A*_1, ··· , A*_n) ← Welfare-Maximizing Allocation;
Let π be an ordering of the agents with A*_π(1) ≥ A*_π(2) ≥ ··· ≥ A*_π(n);
for k = 1 to n do
    Let S be the set of all X ⊆ M such that:
        • v_j(X) ≥ 3α v_π(k)(A*_π(k)) for some j with B_j = ∅,
        • X ⊆ A*_i for some i, and
        • X ∩ B_ℓ = ∅ for all ℓ
    if |S| ≠ ∅ then
        X̂_k ← argmin_{X∈S} |X|;
        B_j ← X̂_k;
    end
end
Apply the envy-cycles procedure of Lipton et al. [55] to allocate the items in
  M \ ⋃_{i∈N} B_i;
Reassign bundles B = (B_1, ··· , B_n) to the agents to maximize the sum of utilities.
  Call this allocation A;
Compute the natural transfers (t_1, ··· , t_n);
return (A, t)
```

We first show the bound on the transfer payments. Let $X = (X_1, \cdots, X_n)$ be the partial allocation obtained when the `for` loop finishes in Algorithm 3. Note that by the ordering of the optimal allocation, and by minimality of the allocated sets, we have, for any pair i, j of agents, $v_i(X_j) \leq 3\alpha \max_i v_i(A_i^*) + 1$. Thus $v_i(X_j) - v_i(X_i) \leq 3\alpha \max_i v_i\left(A_i^*\right) + 1$. At this stage, applying the envy-cycles procedure of [55] does not increase the envy by more than one. Let $B = (B_1, \cdots, B_n)$ be the partial allocation obtained after this step. We therefore have $v_i(B_j) - v_i(B_i) \leq 3\alpha \max_i v_i(A_i^*) + 2$. Now, by Lemma 5.2.4, we have that (A, t) is envy-free and $\sum_{i \in N} |t_i| \leq 2n^2 \left(3\alpha \max_i v_i\left(A_i^*\right) + 2\right)$.

In order to show that $\frac{\mathrm{SW}(A,t)}{\mathrm{SW}(A^*)} \geq \alpha$, it suffices to show $\frac{\mathrm{SW}(X)}{\mathrm{SW}(A^*)} \geq \alpha$: since we add items to X to obtain A, we have $\mathrm{SW}(X) \leq \mathrm{SW}(A)$, and since introducing transfers does not affect utilitarian welfare, we have $\mathrm{SW}(A, t) = \mathrm{SW}(A)$. Let $S \subseteq N$ be the set of time steps in which a set was allocated during the `for` loop . The welfare $\mathrm{SW}(X)$ then satisfies

$$\mathrm{SW}(X) \;\geq\; \sum_{k \in S} 3\alpha \cdot v_{\pi(k)}(A^*_{\pi(k)}) \tag{5.5}$$

Next consider rounds $k \in N \setminus S$, that is, the rounds when a bundle is not allocated. Since agent $\pi(k)$ can otherwise be allocated the set $A^*_{\pi(k)}$, if no set is allocated in round k then either agent $\pi(k)$ has already received a set $X_{\pi(k)}$ of value at least $3\alpha \cdot v_{\pi(k)}(A^*_{\pi(k)})$ or some other agent who came before her received a set $X_{f(k)} \subseteq A^*_{\pi(k)}$ of value at least $3\alpha \cdot v_{\pi(k)}(A^*_{\pi(k)})$. Thus $\max\left[v_{f(k)}(X_{f(k)}), v_{\pi(k)}(X_{\pi(k)})\right] \geq \alpha \cdot v_{\pi(k)}(A^*_{\pi(k)})$ and so

$$\begin{aligned}
\sum_{k \in N \setminus S} \alpha \cdot v_{\pi(k)}(A^*_{\pi(k)}) \;&\leq\; \sum_{k \in N \setminus S} \max\left[v_{f(k)}(X_{f(k)}), v_{\pi(k)}(X_{\pi(k)})\right] \\
&\leq\; \sum_{k \in N \setminus S} v_{f(k)}(X_{f(k)}) + v_{\pi(k)}(X_{\pi(k)}) \\
&\leq\; 2 \cdot \mathrm{SW}(X)
\end{aligned} \tag{5.6}$$

Summing (5.5) and (5.6) immediately gives the utilitarian welfare guarantee.

$$\begin{aligned}
3\,\mathrm{SW}(X) \;&\geq\; \sum_{k \in S} 3\alpha \cdot v_{\pi(k)}(A^*_{\pi(k)}) + \sum_{k \in N \setminus S} 3\alpha \cdot v_{\pi(k)}(A^*_{\pi(k)}) \\
&=\; 3\alpha \sum_{k \in N} v_{\pi(k)}(A^*_{\pi(k)}) \\
&=\; 3\alpha \cdot \mathrm{SW}(A^*)
\end{aligned}$$

□

Chapter 6

Summary and Conclusions

To summarize, Part I of this thesis studies two popular multi-item auctions: the sequential auction and the simultaneous auction. We first show that both standard types (first-price and second-price) of sequential auctions have a focal subgame-perfect equilibrium, and that the price can (very rarely) be non-monotone at this equilibrium. Thus the declining price anomaly is not guaranteed to hold in the equilibria of full-information sequential auctions with three or more buyers.

The main motivation behind the seminal work of Paes Leme et al. [64] is to study the quality (i.e., the welfare) of equilibrium outcomes in first-price sequential auctions. They show that for *unit-demand* bidders, the price of anarchy (the welfare ratio between an optimal outcome and the worst-case subgame perfect equilibrium outcome) is at most 2. They also show that this ratio is unbounded for submodular valuations when there are at least four bidders. Recently, Ahunbay and Vetta [2] studied the two-bidder, identical-item setting, and showed that the price of anarchy is bounded by a constant for two bidders with decreasing marginals, and is exactly m for two bidders with general monotone valuations. Since their bounds apply to the two-bidder identical-items setting, and the bounds in the opposite direction apply only with four or more bidders and more general valuation forms, some parts of this picture remain to be filled in. Some first steps towards completing this picture appear in Ahunbay's recent thesis [1].

We then study the risk-free profitability of a bidder in a multi-item auction. Aside from resolving the remaining gaps in this work, one could also study the performance of other strategies of interest (besides the risk-free strategy) in these auctions from a theoretical perspective, expanding on the large body of research in empirical economics on this topic.

Part II of this thesis studies the fair division problem. We show that with the standard normalizing assumption, one dollar per agent is sufficient to guarantee an envy-free allocation with payments when the valuations are additive. We also study the same parameter for the case of general monotone valuations. Finally, we study the tradeoffs between fairness, welfare and the total quantity of payments in this setting. Several important problems in fair division remain open. For instance, it is not known for many important valuation classes (additive, submodular, XOS, subadditive) exactly what constant fraction of the MMS value can be guaranteed to all agents. We also do not know the valuation classes for which an *EFX* allocation exists (an allocation $A_1, \ldots, A_n$ is EFX if, for any pair i, k of agents, either $v_i(A_i) \geq v_i(A_k)$, or $v_i(A_i) \geq v_i(A_k \setminus \{g\})$ for *every* item g in A_k). For the fair division with subsidy problem, Barman et al. [13] show a bound of $n-1$ dollars for *dichotomous* valuations, a class that is neither contained in nor contains the additive valuations. The gap between the lower bound of $n-1$ and the upper bound of $O(n^2)$ remains open for all other non-trivial valuation classes that aren't subsumed by additive valuations. The idea of using subsidies or transfers between agents in order to achieve fairness in indivisible item-allocation is an incredibly natural and easily implemented one. Despite this, the problem remains relatively unsolved from both the existential and the mechanism-design perspectives. A compelling direction for future research is to comprehensively analyze the exact power and limitations of envy-free item allocations with payments.

Bibliography

[1] Mete Ahunbay. *Equilibria and efficiency of sequential multiunit auctions*. PhD thesis, McGill University, 2021.

[2] Mete Şeref Ahunbay and Adrian Vetta. The price of anarchy of two-buyer sequential multiunit auctions. In *Web and Internet Economics*, pages 147–161, 2020.

[3] A. Alkan, G. Demange, and D. Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4):1023–1039, 1991.

[4] N. Alon. Splitting necklaces. *Advances in Mathematics*, 63:241–253, 1987.

[5] E. Aragones. A derivation of the money Rawlsian solution. *Social Choice and Welfare*, 12(3):267–276, 1995.

[6] E. Arunachaleswaran, S. Barman, R. Kumar, and N. Rathi. Fair and efficient cake division with connected pieces. In *Web and Internet Economics*, pages 57–70, 2019.

[7] O. Ashenfelter. How auctions work for wine and art. *The Journal of Economics Perspectives*, 3(3):23–36, 1989.

[8] O. Ashenfelter and D. Genesove. Legal negotiation and settlement. *American Economic Review*, 82:501–505, 1992.

[9] A. Ashta. Wine auctions: More explanations for the declining price anomaly. *Journal of Wine Research*, 17(1):53–62, 2006.

[10] J. Bae, E. Beigman, R. Berry, M. Honig, and R. Vohra. Sequential bandwidth and power auctions for distributed spectrum sharing. *Journal on Selected Areas in Communications*, 26(7):1193–1203, 2008.

[11] S. Barman, S. Krishnamurthy, and R. Vaish. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 557–574, 2018.

[12] S. Barman, U. Bhaskar, and N. Shah. Optimal bounds on the price of fairness for indivisible goods. In *Web and Internet Economics*, pages 356–369, 2020.

[13] Siddharth Barman, Anand Krishna, Y. Narahari, and Soumyarup Sadhukhan. Achieving envy-freeness with limited subsidies under dichotomous valuations. *CoRR*, abs/2201.07419, 2022. URL `https://arxiv.org/abs/2201.07419`.

[14] A. Beggs and K. Graddy. Declining values and the afternoon effect: evidence from art auctions. *Rand Journal of Economics*, 28:544–565, 1997.

[15] X. Bei, N. Chen, X. Hua, Tao B., and E. Yang. Optimal proportional cake cutting with connected pieces. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.

[16] X. Bei, X. Lu, P. Manurangsi, and W. Suksompong. The price of fairness for indivisible goods. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 81–87, 2019.

[17] N. Benabbou, M. Chakraborty, A. Igarashi, and Y. Zick. Finding fair and efficient allocations when valuations don't add up. In *Algorithmic Game Theory*, pages 32–46, 2020.

[18] D. Bertsimas, V. Farias, and N. Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, 2011.

[19] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Symposium on Discrete Algorithms (SODA)*, pages 700–709, 2011.

[20] J. Black and D. de Meza. Systematic price differences between successive auctions are no anomaly. *Journal of Economics and Management Strategy*, 1(4):607–628, 1992.

[21] S. Buccola. Price trends at livestock auctions. *American Journal of Agricultural Economics*, 64:63–69, 1982.

[22] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[23] P. Burns. Experience and decision making: a comparison of students and businessmen in a simulated progressive auction. In V. Smith, editor, *Research in Experimental Economics*, volume 3, pages 139–157. JAI Press, 1985.

[24] Yang Cai and Christos Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *Proceedings of 15th ACM Conference on Economics and Computation (EC)*, pages 895–910, 2014.

[25] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. The efficiency of fair division. In *Internet and Network Economics*, pages 475–482, 2009.

[26] I. Caragiannis, N. Gravin, and X. Huang. Envy-freeness up to any item with high nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 527–545, 2019.

[27] I. Caragiannis, D. Kurokawa, H. Moulin, A. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Trans. Econ. Comput.*, 7(3): 12:1–32, 2019.

[28] O. Chanel, L. Gérard-Varet, and S. Vincent. Auction theory and practice: Evidence from the market for jewellery. In V. Ginsburgh and P. Menger, editors, *Economics of the Arts: Selected Essays*, pages 135–149. North-Holland, 1996.

[29] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. *Journal of the ACM*, 63(2):11:1–11:19, 2016.

[30] Y. Cohler, J. Lai, D. Parkes, and A. Procaccia. Optimal envy-free cake cutting. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11, 2011*, 2011.

[31] L. E. Dubins and E. H. Spanier. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1):1–17, 1961.

[32] U. Feige. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009.

[33] M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of 49th Symposium on Theory of Computing (STOC)*, pages 201–210, 2013.

[34] M. Feldman, B. Lucier, and V. Syrgkanis. Limits of efficiency in sequential auctions. In *Proceedings of the 9th Conference on Web and Internet Economics (WINE)*, pages 160–173, 2013.

[35] D. Foley. Resource allocation and the public sector. *Yale Econ Essays*, 7(1):45–98, 1967.

[36] P. Funk. Auctions with interdependent valuations. *International Journal of Game Theory*, 25:51–64, 1996.

[37] I. Gale and M. Stegeman. Sequential auctions of endogenously valued objects. *Games and Economic Behavior*, 36(1):74–103, 2001.

[38] M. Gallegati, G. Giulioni, A. Kirman, and A. Palestrini. What's that got to do with the price of fish?: Buyer's behavior on the Ancona fish market. *Journal Econ. Behav. Organ.*, 80(1):20–33, 2011.

[39] G. Gamow and M. Stern. *Puzzle-Math*. Viking Press, 1958.

[40] M. Ghodsi, M. Taghi Hajiaghayi, M. Seddighin, S. Seddighin, and H. Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 539–556, 2018.

[41] V. Ginsburgh. Absentee bidders and the declining price anomaly in wine auctions. *Journal of Political Economy*, 106(6):319–335, 1998.

[42] V. Ginsburgh and J. van Ours. On organizing a sequential auction: Results from a natural experiment by Christie's. *Oxford Economic Papers*, 59(1):1–15, 2007.

[43] C-J. Haake, M. Raith, and F. Su. Bidding for envy-freeness: A procedural approach to n-player fair-division problems. *Social Choice and Welfare*, 19(4):723–749, 2002.

[44] D. Halpern and N. Shah. Fair division with subsidy. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, pages 374–389, 2019.

[45] A. Hassidim, H. Kaplan, Y. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of 12th Conference on Electronic Commerce (EC)*, pages 295–296, 2011.

[46] A. Hu and L. Zou. Sequential auctions, price trends, and risk preferences. *Journal of Economic Theory*, 158:319–335, 2015.

[47] P. Jehiel and B. Moldovanu. Strategic nonparticipation. *Rand Journal of Economics*, 27 (1):84–98, 1996.

[48] Philippe Jehiel, Benny Moldovanu, and Ennio Stacchetti. How (not) to sell nuclear weapons. *American Economic Review*, 86(4):814–29, 1996.

[49] Paul Klemperer. *Auctions: Theory and Practice*. Princeton University Press, 2004. ISBN 9780691119250.

[50] F. Klijn. An algorithm for envy-free allocations in an economy with indivisible objects and money. *Social Choice and Welfare*, 17:201–215, 2000.

[51] Vijay Krishna. *Auction Theory*. Academic Press, 2 edition, 2009. ISBN 9780123745071.

[52] D. Kurokawa, A. Procaccia, and J. Wang. Fair enough: Guaranteeing approximate maximin shares. *J. ACM*, 65(2), 2018.

[53] V. Lambson and N. Thurston. Sequential auctions: Theory and evidence from the Seattle Fur Exchange. *Games and Economic Behavior*, 37(1):70–80, 2006.

[54] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006.

[55] R. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.

[56] K. Lusht. Order and price in a sequential auction. *Journal of Real Estate Finance and Economics*, 8:259–266, 1994.

[57] Eric S. Maskin. *On the Fair Allocation of Indivisible Goods*, pages 341–349. Palgrave Macmillan UK, London, 1987.

[58] P. McAfee and D. Vincent. The declining price anomaly. *Journal of Economic Theory*, 60:191–212, 1993.

[59] C. Mezzetti. Sequential auctions with informational externalities and aversion to price risk: decreasing and increasing price sequences. *The Economic Journal*, 121(555): 990–1016, 2011.

[60] P. Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.

[61] P. Milgrom and R. Weber. A theory of auctions and competitive bidding II. In P. Klemperer, editor, *The Economic Theory of Auctions*. Edward Elgar, 2000.

[62] Roger B. Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, 1981.

[63] Vishnu V. Narayan, Enguerrand Prebet, and Adrian Vetta. The declining price anomaly is not universal in multi-buyer sequential auctions (but almost is). In Dimitris Fotakis and Evangelos Markakis, editors, *Algorithmic Game Theory*, pages 109–122, Cham, 2019. Springer International Publishing.

[64] R. Paes Leme, V. Syrgkanis, and E. Tardos. Sequential auctions and externalities. In *Proceedings of 23rd Symposium on Discrete Algorithms*, pages 869–886. Society for Industrial and Applied Mathematics, 2012.

[65] J. Pesando and P. Shum. Price anomalies at auction: Evidence from the market for modern prints. In V. Ginsburgh and P. Menger, editors, *Economics of the Arts: Selected Essays*, pages 113–134. North-Holland, 1996.

[66] G. Rodriguez. Sequential auctions with multi-unit demands. *The B.E. Journal of Theoretical Economics*, 2009.

[67] F. Salladarre, P. Guilloteau, P. Loisel, and P. Ollivier. The declining price anomaly in sequential auctions of identical commodities with asymmetric bidders: empirical evidence from the Nephrops norvegicus market in France. *Agricultural Economics*, 48:731–741, 2017.

[68] S. Sosnick. Bidding strategy at ordinary auctions. *Journal of Farm Economics*, 45(1): 8–37, 1961.

[69] H. Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.

[70] W. Stromquist. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8): 640–644, 1980.

[71] L-G. Svensson. Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica*, 51(4):939–954, 1983.

[72] K. Tadenuma and W. Thomson. The fair allocation of an indivisible good when monetary compensations are possible. *Mathematical Social Sciences*, 25(2):117–132, 1993.

[73] S. Thiel and G. Petry. Bidding behaviour in second-price auctions: Rare stamp sales, 1923-1937. *Applied Economics*, 27(1):11–16, 1995.

[74] Z. Tu. A resale explanation for the declining price anomaly in sequential auctions. *Review of Applied Economics*, 6(1-2):113–127, 2010.

[75] G. van den Berg, J. van Ours, and M. Pradhan. The declining price anomaly in Dutch Dutch rose auctions. *American Economic Review*, 91:1055–1062, 2001.

[76] H. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63 – 91, 1974.

[77] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

[78] R. Weber. Multiple object auctions. In R. Engelbrecht-Wiggans, M. Shubik, and R. Stark, editors, *Auctions, Bidding and Contracting: Use and Theory*, pages 165–191. New York University Press, 1983.