Market Intermediation: Information, Computation, and Incentives

Matthias Gerstgrasser

University College

University of Oxford

Thesis submitted for the degree of

Doctor of Philosophy

September 2018

Abstract

Auctions are a major field of interest in game theory and in the wider microeconomics area, reflected by recognitions such as Nobel prizes to William Vickrey and Paul Milgrom. The algorithmic game theory literature too provides discussion of a wide range of different auction settings. But real-life markets are rarely comprised of a single monopolist facing buyers without alternative. We therefore explore market intermediation, in which we aim to match buyers and sellers to achieve some objective. While auctions have been well-explored in manifold variations, intermediation has received less attention in the literature.

We aim to move beyond the independent, single-unit case and explore the limits of what can be achieved in more complex scenarios. In the first part, we look at a correlated-priors setting. We show that the revenue-optimal mechanism for this can be computed using a polynomial-time algorithm for one buyer and one seller. For two or more buyers we show that this problem is NP-hard, in contrast, but that truthful-in-expectation mechanisms can be computed using an LP in polynomial time for fixed number of buyers and sellers. In this setting we further discuss how market intermediation relates to classical auctions, as well as reverse auctions. Further motivating our results, we show that our discussion of market intermediation can lead back to useful results for both of these settings, giving an improved algorithm for the optimal two-bidder auction, and showing for the first time that a reverse auction behaves differently than an auction.

In the second part, we consider an online intermediation setting, in which

the market maker encounters an unknown sequence of buyers and sellers one at a time, with knowledge of their independent priors. We explore this from the point of view of online algorithms and competitive analysis, comparing against an offline adversary who knows the buyer-seller sequence in advance. For the general case, we show that the competitive ratio of the intermediary's revenue grows as the square root of the number of buyers and sellers. In contrast, we consider two settings with natural restrictions; one in which the sequence is balanced, and one in which there is an upper limit on the number of items the intermediary is allowed to hold at any one time. For both these settings we show that the competitive ratio is constant.

Finally, in the third part we explore multi-unit intermediation. In this, we consider one seller and one buyer each having concave valuation of a number of items. The intermediary's aim will be to maximise welfare, while maintaining budget balance. This setting has been explored for the single-item case, along with simple reductions for divisible goods to that case. We will give a strong characterisation result as well as approximation guarantees for the multi-unit case.

Contents

Part 1

Introduction

1.1 Introduction

Auctions are a major field of interest in game theory and in the wider microeconomics area, reflected by recognitions such as Nobel prizes to William Vickrey and Paul Milgrom. With an increasing focus on decentralised systems, auctions and broader mechanism design questions have increasingly been of interest in Computer Science as well. The algorithmic game theory literature provides discussion of a wide range of different auction settings, with sponsored search auctions perhaps one of the most prominent examples. Mobile spectrum auctions as well have been studied in much detail, and more recently attention has been given to a variety of mechanism design problems related to cloud computing, shared economy settings, electronic commerce, and many others.

But the reality is rarely as one-sided as the plain auction model assumes. Few settings are imaginable where an auctioneer is naturally endowed with an item to begin with — usually, they have to acquire it themselves first. And even in situations where that is sometimes the case things don't always end there: In the aforementioned case of spectrum auctions, for instance, recent developments have moved toward reallocating spectrum from one licensee to another.

The aim of this thesis is therefore to explore a more general, two-sided market intermediation model, in which an auctioneer faces both buyers and sellers, and aims to facilitate trade. Examples of such settings are ample in the real world, with a wide range of subtly different variations. Consider an electronic marketplace such as eBay, or an online retailer such as Amazon. Both ultimately intermediate between suppliers and customers, though the mechanisms they employ are rather different, and our exposition will aim to capture several different aspects of intermediation. Amazon in particular is a curious case: It is both a classical online shop, buying stock and then later acting as seller toward its customers, as well as a market maker in its "Marketplace" incarnation, where it provides a platform for third-party sellers to interact with buyers. But both of these are intermediation: The difference merely being whether Amazon take ownership of the item first to sell on later, or whether supplier and customer are matched directly.

The simplest setting we will consider is the intermediation analogue of a single-item auction. This has been studied for independent priors [99], and we therefore focus on the correlated case. In this model, a market maker brings seller and buyer together, and runs a simultaneous double auction with both. We will later link this setting to the well-explored single-item auction, before we move on to more complex models. We then consider an online variant, where an intermediary faces a sequence of sellers and buyers, and must decide when to buy and sell items, and at which prices, not knowing the future sequence. Lastly we focus on a multi-unit setting, and welfare maximisation.

Together these will help to illuminate some of the key facets of market intermediation, as well as pose questions of broad relevance to the mechanism design field.

1.2 Background and General Literature Review

Background & Context

The rise of the Internet and the development of the World Wide Web have gone hand in hand with a fundamental transformation of computer science research in the past two decades. The most dramatic paradigm shift is that from monolithic systems designed, operated and controlled by single organisations to decentralised systems shared by many self-interested participants. This change has given rise to the field of Algorithmic Game Theory, the study of decision-making of strategic agents from a computational point of view. Applications of computer science arguments to game theoretic problems have led to a number of novel strands of research in this area, from questions of computational tractability to worst-case performance guarantees of equilibria concepts.

Within game theory, mechanism design is one of the most actively researched areas of enquiry, underscored by the award of the 2007 Nobel prize in economics to pioneers Myerson, Maskin and Hurwicz. Mechanism design deals with the development of robust social choice protocols involving strategic participants. In an era of decentralised routing, cloud computing and peer to peer systems, applications of algorithmic mechanism design results feature in many essential computing technologies. Relevant settings are manifold within modern society, and increasingly common in areas of inquest within computer science. They share in common a requirement for the establishment of a protocol of conduct that incentivises self-interested participants to act in a fashion that leads to a socially desirable outcome.

Applications of algorithmic game theory and mechanism design are diverse, and central to many novel developments underpinning the internet era. As one of the most famous examples, Google's rise to be one of the most popular search engines might have been based on the efficacy of their PageRank algorithm. However, the core of their business model of selling advertisement space is based on the Generalised Second-Price auction. This mechanism is unusual in that it is not truthful [53], and encourages bidders to strategise.

Auctions in general are a main theme within mechanism design, and a wealth of results in this area has focused for instance on approximately optimal mechanisms for combinatorial auctions. A major real-world application of these are allocations of wireless spectrum, as exemplified by frequency auctions for 3G and more recently 4G LTE mobile phone spectrum licenses [40, 72, 73]. The radio frequency spectrum is becoming increasingly contested, and efficient allocation of licenses will be a major mathematical challenge in the coming decades. More recent work [58, 59, 126, 127] has focused on dynamic spectrum auctions, a novel approach that would allow spectrum to be allocated as it was needed rather than ahead of time.

Another factor compounding the increased interest in algorithmic mechanism design is the shift from brick-and-mortar operations to electronic commerce in many industries as well as the rise in online social networking as a tool in everyday interaction. This has led to the creation of electronic markets in novel settings and is a major driver behind data-driven mechanism design.

Nisan et al. [105] provide a thorough introduction to a wide variety of applications of this field.

Algorithmic Game Theory

The increased interest in game theory within the computer science community led to a focus on algorithmic aspects of the field. A number of notable advances have been made using computer science techniques [105].

Firstly, a great deal of attention has been given to the computational complexity of equilibrium concepts. Nash's famous result [101, 100] is purely non-constructive - it guarantees the existence of Nash equilibria, but does not show how to reach or compute them. Daskalakis et al. [44] and Chen and Deng [28] have shown that indeed even computing a Nash equilibrium is computationally hard (it is PPAD-complete).

Secondly, in the spirit of traditional analysis of the performance of algorithms, Koutsoupias and Papadimitriou [76, 77] have introduced the concept now known as the *price of anarchy*, a worst-case performance guarantee of Nash equilibria solutions of games. This type of analysis has been particularly successful in its application to routing games such as in [113] and many others.

Thirdly, Nisan and Ronen [103, 104] have laid the foundations for algorithmic mechanism design.

Algorithmic Mechanism Design

The general problem in mechanism design is as follows: We are faced with a number of players, each of which has a piece of private information. As a mechanism designer, we want to pick one of several potential outcomes based on players' private information. For instance, we may wish to allocate items to those who value them most. However, players are self-interested, and prefer some outcomes over others. They may thus misrepresent their private information to us, in order to influence our decision making. (For instance, they may claim they value an item more than they do.)

In algorithmic mechanism design, we view the decision making process as an algorithm. We thus want to find an algorithm that takes as input the declared private information from participants and computes an outcome or an allocation $x(.)$, usually together with an algorithm that computes payments for each player $p_i(.)$. We may only care about the allocation, using payments to implement incentives for players. Or vice versa; In the theory of optimal auctions we are concerned with revenue maximisation, and use allocations purely to incentivise bidders.

One of the most important results in this context is the revelation principle, which shows that arbitrary dominant-strategy mechanisms can be reduced to truthful ones via a black-box construction [64] [105, chapter 9]. Hence we need only consider truthful mechanisms, that is mechanisms in which telling the truth is a dominant strategy for players. This greatly simplifies our task by narrowing down our search space; Truthful mechanisms furthermore are often desirable in practice, as they do not require players to strategise. There is several different angles from which this problem can be approached:

Firstly, given a social choice function or objective function, can we implement it in a (truthful) mechanism? If not, how close can we get? The major result in auction theory is that of the VCG mechanism developed by Vickrey, Clarke and Groves [119, 33, 66]. This mechanism allocates items to bidders who value them most, i.e. it implements the objective of social surplus (or welfare) maximisation. Affine maximisers are a generalisation of this and allow us to implement affine functions of bidders' valuations. In general, an exact implementation will not always be possible - in some settings we will loose some efficiency compared to the optimal solution if we knew players private information a priori. For truthful mechanisms, this leads to the question of an approximation ratio, that is the ratio of the value of the outcome achieved by a mechanism in the worst case to that of the optimal outcome.

The very question that started interest in algorithmic mechanism design, the approximation ratio of scheduling unrelated machines, remains largely unsolved to this day. This strategic version of the classical computer science problem was first posed by Nisan and Ronen [103, 104], who showed that the standard VCG mechanism could achieve an approximation ratio of n for n machines, and gave a lower bound of 2. A large body of research has since focussed on the deterministic [31, 80, 4, 5, 82, 83, 6], randomised [97, 90, 89, 88, 29] and fractional [31] cases, but a linear gap between upper and lower bounds on the approximation ratio remains.

Secondly, in a given domain, since we know we need only consider truthful mechanisms, a question of great interest is that of a characterisation of the class of truthful mechanisms. In other words, what do the truthful mechanisms in a given setting look like? For the unrestricted domain, that is, without any further structure on the space of outcomes, this has been answered: Roberts [108] gave a characterisation via a global condition, and showed that the implementable social choice rules are precisely the affine maximisers: That is, functions of the form $f(v) = \text{argmax}_j \sum_i \lambda_i v_{ij} + \beta_j$. This is a generalisation of the VCG mechanism, where we weight each of the players with a positive multiplier λ_i and add a constant β_j to each of the outcomes.

Myerson's [98] celebrated result showed that in single-parameter domains an allocation rule is implementable if and only if it is monotone, i.e. an agent's allocation can only increase if they increase their bid. Saks and Yu [115] later extended this result to any convex domain. Rochet [109] has generalised this further to a cyclic monotonicity condition. For cycles of length two this reduces to the previous result. However, all of these conditions are local to each of the players; Few results are known in terms of global characterisations similar to Roberts' result [108]. One notable exception worth mentioning is the work of Vidali who gave a characterisation of the geometry of truthful mechanisms [120] as well as some results on the aforementioned hierarchy of domains and the extension of truthful mechanisms from domain to superdomain [121]. In the case of scheduling, some progress has been made for the case of two machines [32, 51].

Lastly, how do we compute an optimal mechanism for a specific instance? In this view, we regard the mechanism design problem as a computational one. The difference for us as the algorithm designer, compared to classical algorithms questions, is that we do not know a priori the input. Again, we do not need to reason about agent's strategies, if we restrict ourselves to truthful mechanisms. In doing so, truthfulness becomes an additional constraint in the optimisation problem we are facing. This is similar to issues of computational complexity: There, polynomial runtime is often required as an additional constraint; here, it is truthfulness, or often both. Hence also the previous discussion on approximation ratio can be seen as a continuation of the traditional inquiry into the approximation ratio of polynomial time (or similar) algorithms. That computing the optimal mechanism in an unrestricted setting is NP-complete was first shown by Conitzer and Sandholm [36]. Much research has since focused on understanding the complexity of the optimal mechanism problem in more specified settings, for instance single-parameter domains [47, 48]. Myerson's result gives a computationally easy way of computing the optimal single-bidder auction for independent priors. More recent work has focused on more complex settings including multiple-goods auctions [44, 43, 61, 62] as well as correlated priors [106, 50].

Auctions and Markets

The concept of trade is one of the centrepieces of economic theory. From the perspective of a seller (or buyer) this has been treated extensively in auction theory, including the aforementioned seminal results by Myerson [98], Vickery, Clarke and Groves [119, 33, 66], and others. The study of trade from the perspective of the *market* perhaps goes back even further, with results such as those of Walras' tâtonnement process (or Walrasian auction) for price-finding in perfect competition markets [122, 123]. Modern game-theoretic research into this area started with Myerson and Satterthwaite [99], who give both a strong negative result – no truthful, individually rational mechanism can be both efficient and (even weakly) budget balanced – as well as positive results extending Myerson's result on auctions [98] to give a revenue-optimal double auction for independent priors.

Given their negative result on social welfare, research in this area has thus focused on mechanisms that satisfy some of the conditions above but not others, or make further assumptions. Rustichini et al. [114] show that the inefficiency disappears for many bidders. Cramton et al. [39] for instance have shown a positive result when initial endowments are symmetric. One of the key results on analysing market intermediation is by McAfee [92], who gives an approximately efficient truthful mechanism for many traders. This is further strengthened to a 2-approximation to the optimal gain from trade in the Bayesian setting in [93].

The importance of this setting is further underlined when looking at it in the larger context of market design. Research on supply chains [11, 12, 15, 16, 124] or spatially distributed markets [13, 14], for instance, use bilateral trade as a building block for more complex market design and analysis.

Our aim will be to investigate further several aspects of the market intermediation setting. We will discuss further literature relevant to those specific questions in the respective chapters of this thesis.

1.3 Preliminaries

With this background in mind, we now formally introduce game theory and mechanism design, and then the specific setting we are interested in.

Introduction & Basics of Game Theory

A general game theory setting has two basic ingredients: One, a set of **outcomes** $\omega \in \Omega$. Two, one or more **strategic, self-interested agents** $i \in I$. That is, each agent has a **preference** over the outcomes, and a set of available **strategies** or actions. Players each choose a strategy $s_i \in S_i$, and an outcome ω occurs as a function of the chosen strategies, i.e. $\omega = \omega(\mathbf{s})$. (Throughout this thesis a boldface variable $\mathbf{a}$ shall denote a vector made of components a_i.) A set of actions might be "rock, paper, scissors" for each of two players, and the set of outcomes could be "Player 1 wins", "Draw", "Player 2 wins". In this classic example, players' preference is simply to win. The rules of the game - paper beats rock, rock beats scissors, scissors beat paper - define a mapping from a pair of chosen actions to the space of outcomes. A **pure strategy** is simply a choice of one action; a **mixed** or **randomised** strategy is a probability distribution over the set of actions for a player. By abuse of notation we will write s_i for both pure and mixed strategies, unless the distinction is relevant. We will write $\mathbf{s}$ for the vector or profile of all players' strategies, and $\mathbf{s}_{-i}$ for the strategies of all players except i.

Preferences over outcomes could take several forms. In the most general, an ordering over the space of outcomes would be enough to formalise the notion of preference. This is generally thought to be the domain of social choice. In contrast, in game theory we will take preferences to be cardinal. That is, we assume if a particular outcome ω is realised, a player i will derive a real-valued **utility**, value or payoff $u_i(\omega)$. In words, we assume that agents not only know "Do I prefer outcome ω_1 over outcome ω_2?", but also "*How much* do I prefer outcome ω_1 over outcome ω_2?" This allows us to capture, for instance, the difference between a player who really wants to win, but doesn't mind too much between a draw and a loss; and one who really wants to avoid a loss, but cares less about winning or a draw. When mixed strategies or other sources of randomness are involved, we will generally assume players seek to maximise their expected utility. We will sometimes omit the expectation both in text and in formulas, when our intended meaning is clear.

The difficulty for players lies in choosing a strategy with the goal of achieving a preferred outcome (in expectation, in the worst case, etc.), knowing that their opponent(s) will do the same. This has given rise to a multitude of **solution concepts** for games. The most basic is that of dominant strategies. For our purposes, we define this as follows: A strategy s_i for player i is said to (weakly) **dominate** strategy s_i', if for any choice of strategies of the remaining players $\mathbf{s}_{-i}$, s_i gives a better or equal outcome to player i than strategy s_i':

$u_i(\omega(s_i, \mathbf{s}_{-i})) \geq u_i(\omega(s'_i, \mathbf{s}_{-i})) \quad \forall \mathbf{s}_{-i}$. A strategy s_i is said to be a **dominant strategy** if it dominates all other strategies of player i. A strategy profile $\mathbf{s}$ is a **dominant strategy equilibrium** if all players' strategies are dominant strategies.

Definition 1 (Dominant strategies)**.** *A strategy s_i is called a* ***dominant strategy*** *for player i, if for any choice of* $\mathbf{s}_{-i}$*,*

$$u_i(\omega(s_i, \mathbf{s}_{-i})) = \max_{s'_i} u_i(\omega(s'_i, \mathbf{s}_{-i}))$$

A strategy profile $\mathbf{s}$ *is called a* ***dominant strategy equilibrium****, if for every player i the strategy s_i is a dominant strategy for player i.*

A dominant strategy equilibrium is by far the strongest solution concept we will discuss. It has the clear advantage that it requires no counterfactual reasoning about other players' behaviour. But, it may not always be achievable. Not all games have a dominant strategy equilibrium. The most widely known, more general, solution concept is that of the **Nash equilibrium**. A strategy profile $\mathbf{s}$ is called a Nash equilibrium if for every player i, their strategy s_i is the best they could do *given all other players' strategies* $\mathbf{s}_{-i}$; i.e. if for every player i and every alternative strategy s'_i, $u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s'_i, \mathbf{s}_{-i})$. Where the strategies are randomised, the same inequality shall hold in expectation. We will also say that s_i is a best response for player i given $\mathbf{s}_{-i}$. For mixed strategies, this should hold for the respective expected utilities. Notice that unlike in the dominant strategy equilibrium the other players' strategies do matter in the Nash equilibrium. An advantage of the (possibly mixed) Nash equilibrium is that for a significant class of games it is guaranteed to exist - Nash's famous result [101]. A disadvantage is that more than one may exist. Further relaxations of this equilibrium concept are common in the literature, such as correlated and coarse correlated equilibria, but of less relevance for our purposes.

Definition 2 (Nash equilibrium)**.** *A strategy s_i is called a **best response strategy** for player i given $\mathbf{s}_{-i}$, if*

$$u_i(\omega(s_i, \mathbf{s}_{-i})) = \max_{s'_i} u_i(\omega(s'_i, \mathbf{s}_{-i}))$$

*A strategy profile $\mathbf{s}$ is called a **Nash equilibrium**, if for every player i the strategy s_i is a best response strategy for player i given $\mathbf{s}_{-i}$.*

If all of the information is public and known to all players, we have a so-called complete information game. More commonly we will encounter an incomplete information setting. In this, we assume players possess some private information, their **type** t_i. For instance, this may be their utility function. Players' strategies will now depend on their type: A player will play differently depending on which outcome they prefer.

In the most extreme case, we may know nothing about the players' private information. A middle ground is **Bayesian games**, in which we assume players' types are drawn from a known **prior distribution**. For our purposes, we will assume that this prior information is common knowledge. This gives rise to a new solution concept, the **Bayesian Nash equilibrium (BNE)**. A strategy profile $\mathbf{s}(\mathbf{t})$ together with a prior distribution ψ is a Bayesian Nash equilibrium, if for every player i, their strategy $s_i(t_i)$ maximises their expected utility with regard to other players' types drawn from the prior distribution.

Definition 3 (BNE)**.** *A strategy profile* $\mathbf{s}$ *is called a* ***Bayesian Nash equilibrium*** *(in a game with prior distribution* ψ*), if for every player* i *the strategy* s_i *maximises the expected utility for player* i *given* $\mathbf{t}_{-i}$ *is drawn from* ψ*:*

$$\mathbb{E}_{\mathbf{t}_{-i}\sim\psi_{-i}}[u_i(\omega(s_i(t_i), \mathbf{s}_{-i}(\mathbf{t}_{-i}))] = \max_{s_i'} \mathbb{E}_{\mathbf{t}_{-i}\sim\psi_{-i}}[u_i(\omega(s_i', \mathbf{s}_{-i}(\mathbf{t}_{-i})))]$$

Important Concepts in Mechanism Design & Auction Theory

In classical game theory, the focus is on reasoning about players' strategies $\mathbf{s}$, and what equilibria might arise as a result. A typical question one might ask is, for instance, "Given that everyone else will be acting strategically, which of my own strategies will achieve the outcome most desirable to me?". Or, "If all players behave strategically, what is the most likely outcome?"

In contrast, in **mechanism design**, we are concerned with the opposite question: "Given that players will act strategically, how can we design the rules of the game (the mapping from actions to outcomes), so that strategic behaviour will lead to an outcome that is optimal in some sense?" Generally, the difficulty for us as the mechanism designer will be that we do not know a priori the players' types, but want to optimise the outcome depending on that private information. Formally, we define a mechanism design problem as follows:

Definition 4 (Mechanism Design)**.** *Players each have a type* $t_i \in T_i$ *as before.*

This is private information of each player. As their action, players each send a message $\theta_i \in \Theta_i$ to the mechanism. The mechanism then selects an outcome $\omega \in \Omega$ as a function of the received messages θ. This function $\omega(\theta)$ is known to players before they send their messages. Generally, the mechanism will want to pick an outcome that optimises some function $f(\omega, \mathbf{t})$ that depends on the players' unknown types.

One important first distinction to make is the shape of the message space Θ_i. A priori messages could be of any form. A special case is if message and type space are identical, $T_i = \Theta_i$. In this case, players communicate a type to the mechanism directly. We call such a mechanism a **direct revelation** mechanism. Most of the mechanisms we will consider will be direct, but not all. As an example, if a player's type is how much they value an item for sale, then a direct revelation mechanism would ask them to submit a monetary bid for the item. Notice that that alone does not mean that they have to reveal their actual type — they are free to over- or underbid as they please. Direct revelation only means that messages and types come from the same space. An indirect mechanism in this example could be, for instance, to ask the player "Are you willing to buy the item for 10 pounds?" Their message would be a single-bit Yes or No answer, as opposed to their full type.

Definition 5 (Direct Revelation Mechanism)**.** *A mechanism in which type and message spaces are identical is called a* ***direct revelation mechanism****. Otherwise, a mechanism is called indirect.*

We have so far talked about players' types only as an abstract concept. In concrete terms, we will work in a private-values model, in which players' types are their valuation functions, $t_i = v_i \in \mathbb{R}^\Omega$. In words, each player knows how much they value each outcome, but the mechanism and other players do not. Other models exist; for instance, in a common value model it is assumed the value of outcomes are common to all players, but players only have partial information on these values (expressed in their type).

Definition 6 (Private Values)**.** *In the private values model, player's types are their valuation functions:* $t_i = v_i \in \mathbb{R}^\Omega$

Implementability

Just as game theory considers several different solution concepts, mechanism design considers several different ways of implementing a particular objective or **social choice function**. The most important implementation is that in dominant strategies. That is, we aim to construct an outcome function $\omega(\theta)$, so that each player has a dominant strategy for which message θ to send as a function of their type. This has the advantage that it is simple for players, who do not need to reason about other players' types or strategies. What is more, a key result in mechanism design tells us that if we restrict ourselves to mechanisms that have a dominant strategy equilibrium, we may as well restrict ourselves to only those mechanisms in which truth-telling is the dominant strategy for all players. This **revelation principle** is an incredibly powerful result:

Theorem 1 (Revelation Principle)**.** *For every mechanism in which all players have a dominant strategy, there is an equivalent direct-revelation mechanism in which truth-telling is a dominant strategy for all players.*

The proof of this well-known theorem is surprisingly easy: For an arbitrary dominant-strategy mechanism M, construct a direct revelation mechanism M' as follows: M' collects type messages θ from players. For each player i, M' computes what their optimal (dominant) strategy in M would be, if their type was indeed θ_i. M' then proceeds to feed those strategies into M; and picks whichever outcome M returned. Now, it is easy to see that it is indeed optimal for players to reveal their true type to M' constructed in this way; and indeed when players do so, M' will by definition always choose the same outcome that M would, thus showing the claim. See for instance [105] for more background information.

This result simplifies our task as a mechanism designer significantly, as we may restrict ourselves to mechanisms in which truth-telling is a dominant strategy. We also call such mechanisms **truthful**, or **dominant-strategy incentive compatible (DSIC)**.

Definition 7 (Truthfulness)**.** *A mechanism in which truth-telling is a dominant strategy for all players is called a truthful or dominant-strategy incentive compatible mechanism.*

This leads directly to one of the central questions in mechanism design: Which mechanisms have this property? And, which social choice functions (functions that map players' types to an outcome) can be implemented in this way? We call a social choice function **implementable** if it can be implemented in a DSIC mechanism. In the very general setting we have discussed so far, very few (interesting) social choice functions are implementable [64, 65, 116, 2, 3]; which is why we will discuss more restricted domains in the next subsection.

Other implementation concepts exist. The most important for us deal with randomness in the mechanism itself. So far we have assumed the function ω was deterministic, but this need not always be the case. In a **randomised mechanism**, the outcome function may depend not only on players' types but also on one or more random bits r coming from a known distribution. (Equivalently, it maps players' types to a probability distribution over outcomes.) There is two notions of truthfulness in this context. Firstly, in a **universally truthful** mechanism truth-telling is a dominant strategies for players ex post, or even if the random bits r were already known. This is only marginally more general than DSIC mechanisms, as it is easy to see that any such mechanism simply randomises between DSIC mechanisms. Secondly, in **truthful-in-expectation** mechanisms, truth-telling is a dominant strategy for players before the random coin tosses are known. This is strictly more powerful than deterministic DSIC mechanisms in some settings, and is usually what we are interested in in randomised mechanisms.

Definition 8 (Randomised Truthfulness)**.** *For a randomised mechanism, we define truthfulness as follows:*

1. *A randomised mechanism is called universally truthful if truth-telling is a dominant strategy for players ex post.*

2. *It is called truthful-in-expectation if truth-telling is a dominant strategy ex ante, i.e. if truth-telling maximises the expected utility of each player, independent of other players' strategies.*

Relaxations from dominant-strategy implementations are known. In an incomplete information setting as mechanism design usually is, Nash equilibria do not make sense. The most widely adopted solution concept outside of dominant strategies is thus that of **Bayes-Nash incentive compatibility** (BIC or BNIC). In this, truth-telling is a Bayesian Nash equilibrium, given the common prior information on players' types.

In addition to incentive compatibility, one usually needs to require a few more conditions. One is **individual rationality (IR)**, requiring that players never have negative utility from participating in a mechanism. As in most settings a mechanism could not force players to participate, this is already required for pragmatic reasons alone.

Definition 9 (Individual Rationality)**.** *A mechanism is called individually rational if player's utility is always nonnegative.*

Domains

With only definition 4 above, mechanism design is hard. Several impossibility results are known for the most general case [64, 65, 116, 2, 3]. While most of these deal with ordinal preferences, Gibbard's theorem [64, 65] also applies to cardinal preferences as we assume, and limits what is possible in this setting. Mechanism design therefore almost universally gives the mechanism additional power in the form of **payments**. That is, not only does the mechanism pick an outcome $\omega(\theta) \in \Omega$, it also computes payments $p_i(\theta)$ that it charges (or pays) each player. This will be used to incentivise players, usually to reveal their type truthfully. That of course works only if players care about payments. We will, as is standard in the literature, assume that players have **quasilinear** utility functions: If a player's utility for the original outcome

ω is $v_i(\omega)$, then their utility for the outcome, but when they are charged a payment p_i, will be $u_i = v_i(\omega) - p_i$. We will call the utility $v_i(\omega)$ for the original outcome (without a payment involved) the player's **valuation**.

Definition 10 (Quasilinear Utility)**.** *A quasilinear utility function is of the form* $u_i(\omega, p_i) = v_i(\omega) - p_i$*.*

Definition 11 (Mechanism Design with Money)**.** *Players each have a type* v_i*, which is private information of each player. As their action, players each send a message* $\theta_i \in \Theta_i$ *to the mechanism. The mechanism then picks an outcome* $\omega \in \Omega$*, and a vector of payments* $\mathbf{p}$*, as functions of the received messages* θ*. Each player* i *derives utility* $u_i(\theta) = v_i(\omega(\theta)) - p_i(\theta)$*. The functions* $\omega : \Theta \to \Omega$ *and* $\mathbf{p} : \Theta \to \mathbb{R}^n$ *are known to players before they submit their messages.*

In essence, we transform the original mechanism design problem without money into one with money, by defining a new outcome space $\hat{\Omega} = \Omega \times \mathbb{R}^n$ where n is the number of players; and by specifying player's utility function in the amended game. In theory, any mechanism design problem can be transformed in this way. Whether it makes sense depends, among other things, on whether quasilinear utility and monetary transfer are a sensible approach in a particular setting in practice. Most of mechanism design is with money, but there are some exceptions, see for instance kidney exchanges [1, 112].

It is important to note that usually we will only seek to optimise either the outcome **or** the payments, using the other purely as an incentive for players. It is precisely this relaxation that gives the mechanism extra power.

The chief example of a mechanism design problem is an auction. In an auction, the mechanism - now also called the auctioneer - wants to allocate one or more items to one or more bidders. The main difference compared to a more general mechanism design problem is that the outcome now takes the form of an allocation of available items to players; and that players only care about their own allocation, rather than the entire outcome space.

Definition 12 (Auction)**.** *In an auction with k items, players each have a private valuation $v_i \in \mathbb{R}^{2^k}$, i.e. a function from binary k-vectors to the real numbers. This is private information of each player. They each submit a message, often a bid $b_i \in \mathbb{R}^{2^k}$, to the mechanism. The mechanism computes for each player an allocation vector $x_i(\mathbf{b}) \in \{0,1\}^k$ and a payment $p_i(\mathbf{b}) \in \mathbb{R}$. Player i receives utility $u_i(\mathbf{b}) = v_i(x_i(\mathbf{b})) - p_i(\mathbf{b})$. Allocations must satisfy that $\sum_i x_{ij} \leq 1 \ \forall j$, i.e. that no item is allocated multiple times.*

If we are dealing with k distinct, different items, we call the auction a **multi-item auction**. If there is k identical copies of the same item, we call the setting a **multi-unit** auction, and simply write the allocation to each player as the number of units they receive. If there is a single item, we call it a **single-item** auction. Here, allocations to each player are a single bit, and valuations a real number indicating the value of the item to the player.

For multi-unit and multi-item settings, a further distinction is on the type of the valuation functions v_i: In one of the simplest cases, valuations are **additive**: The valuation for the union of two disjoint sets of items is the sum of the valuations for both sets, i.e. $v_i(X \cup Y) = v_i(X) + v_i(Y)$ $\forall X \cap Y = \emptyset$. In the multi-unit case this simplifies to a simple linear valuation,

$v_1(m) = m \cdot v_i(1)$. A generalisation of this is **submodular** valuations, which have the property that the marginal increase in value of adding more items to a players' allocation decreases the larger the allocation already is. Formally, if $X \subseteq Y$ and $y \notin Y$, then $v_i(X \cup y) - v_i(X) \geq v_i(Y \cup y) - v_i(Y)$, or one of a number of equivalent conditions. For multi-unit settings this simplifies to concave valuation functions. One further generalisation is **subadditive** valuations, which have the property that the value of the union of two sets is at most the sum of the values of the two sets: $v_i(X \cup Y) \leq v_i(X) + v_i(Y)$. With all of these, it is usually also assumed that the valuation functions are monotone increasing (though this does not follow from submodularity or subadditivity). A further generalisation is then that we do not assume even subadditivity, but still assume that valuations are **monotone increasing**: $v_i(X \cup Y) \geq v_i(X)$. In the most extreme case, we assume no knowledge at all about v_i. In the multi-item setting, this most general case is also called a **combinatorial auction**.

Definition 13 (Valuation Functions)**.** *A valuation function v is called*

1. *additive, if for all X, Y with $X \cap Y = \emptyset$, $v_i(X \cup Y) = v_i(X) + v_i(Y)$*

2. *submodular, if the following equivalent conditions hold:*

   (a) *for all $X \subseteq Y$ and $y \notin Y$, $v_i(X \cup y) - v_i(X) \geq v_i(Y \cup y) - v_i(Y)$*

   (b) *for all X, Y, $v_i(X \cup Y) \leq v_i(X) + v_i(Y) - v_i(X \cap Y)$*

3. *subadditive, if for all X, Y, $v_i(X \cup Y) \leq v_i(X) + v_i(Y)$*

4. *monotone increasing, if for all X, Y, $v_i(X \cup Y) \geq v_i(X)$*

5. *a general or unrestricted valuation function otherwise.*

All of the above types of valuations are commonly used in the literature. If monotonicity is assumed for additive, submodular and subadditive valuations, they form a hierarchy — every additive increasing function is an increasing submodular function; any such function is an increasing subadditive function in turn, etc. Subadditive and submodular valuations capture a situation where there is decreasing marginal returns, or where items are **substitutes** - players desire one or the other, but don't need both. The opposite is also possible, where items are **complements**, and neither one or the other are very useful on their own, but the bundle is. An extreme case of marginal returns is **unit-demand**, in which buyers only desire a single unit of the item, no more. (But the auctioneer may still have more than one unit for sale, making this different from the single-item case.)

Objective Functions

The final ingredient to our prototype mechanism design problem as outlined in definition 4 is that of the objective function the mechanism seeks to optimise or implement. Generally, we refer to a function $\omega : \mathbf{T} \to \Omega$ that maps types to outcomes as a **social choice** function. Some social choice functions can be implemented directly in a DSIC mechanism, usually with payments. Others, we cannot implement exactly. In those cases, we usually seek to approximate them, aiming to minimise the relative difference in performance either in the worst case, or – in the Bayesian setting – the expectation.

One of the main objectives studied, and perhaps the most straightforward, is that of (utilitarian) social welfare maximisation. **Social welfare** is the sum of the players' individual welfare, usually taken to be how much they value their allocation (or the outcome). Generally, players' payments are not taken to factor into their welfare.

Definition 14 (Social Welfare). *The social welfare of an outcome ω or an allocation $\mathbf{x}$ is defined as the sum of players' individual valuations:* $\mathrm{SW}(\omega) = \sum_i v_i(\omega)$ *and* $\mathrm{SW}(\mathbf{x}) = \sum_i v_i(\mathbf{x}_i)$

Social welfare is special because it can be implemented even in the most general setting of general valuation functions, and is essentially the only social choice function with that property [108]. The mechanism that achieves this is the **VCG mechanism**, named after its authors Vickrey, Clarke and Groves [119, 33, 66] This mechanism chooses the outcome that maximises social welfare, and charges each player i the externality they caused on the others; that is, the difference in utility that other players derive with i present compared to if i had not taken part.

Definition 15 (VCG mechanism). *The VCG mechanism implements the (or an) outcome ω that maximises social welfare $\omega(\mathbf{v}) = \mathrm{argmax}_\omega \sum_i v_i(\omega)$. Player i is charged their externality, so their payment is $p_i = -\sum_{j \neq i} v_j(\omega) + \max_{\hat{\omega}} \sum_{j \neq i} v_j(\hat{\omega})$.*

For the case of single-item auctions, this mechanism simplifies to the well-established format of a **second-price** or Vickrey auction.

Definition 16 (Second-Price Auction). *In a second-price auction, the highest bidder is awarded the item, and pays the second-highest bid.*

A key insight is that the VCG mechanism is truthful or DSIC, even for general valuation functions, and even without any structure on the outcome space.

Theorem 2 (VCG mechanism)**.** *The VCG mechanism is truthful / dominant strategy incentive compatible. [119, 33, 66]*

The second widely studied objective, generally in the context of auctions only, is that of **revenue maximisation**. In this, we seek to maximise the sum of the payments made by players, and use the allocation only as a means to incentivise players. Most often, we will still want this to be implemented in dominant strategies, and thus as a DSIC mechanism. Unlike for welfare maximisation, for revenue maximisation there is little we can do in terms of worst-case guarantees, or without any information on players' valuations. We therefore usually approach revenue in a Bayesian setting, assuming players' valuations $\mathbf{v}$ are drawn from a prior distribution ψ. We then seek to optimise the **expected revenue**.

Definition 17 (Revenue)**.** *The expected revenue of a DSIC mechanism given prior distribution ψ is $\mathbb{E}_{\mathbf{v}\sim\psi} \sum_i p_i(\mathbf{v})$.*

A further implication is that we generally use a different benchmark for revenue than for welfare or other objectives: In revenue, we compare the performance of a candidate mechanism to that of the best possible *truthful* mechanism. In welfare, on the other hand, we compare against the best possible outcome, without any game theoretic or strategic considerations.

One of the most important results in auction theory is that of Myerson [98], which resolves this problem for single-item auctions with independent priors. The theorem – in a nutshell – gives a transformation from valuations and priors to "virtual valuations", and shows that a Vickrey auction run on these virtual valuations is revenue-optimal. Assuming a reasonable "regularity" condition on the priors, the result is relatively straightforward; for the general case "ironing" is required.

Theorem 3 (Myerson Auction, part 1)**.** *For a single-item auction where players valuations are drawn from independent regular distributions, the expected revenue is equal to its expected virtual welfare surplus, where players' virtual valuation are defined as* $r(v_i) = v_i - \frac{1-\Psi_i(v_i)}{\psi_i(v_i)}$ *for prior distributions* ψ_i*. The expected revenue is thus maximised by awarding the item to the bidder with the highest nonnegative virtual valuation. A regular distribution is one for which the virtual valuation function* $r(v)$ *as just defined is nondecreasing. For irregular distributions a similar result holds.*

Where the prior distributions are identical (and thus i.i.d., as we assume independence to begin with), the virtual valuation function is also identical for all bidders, and monotone. The highest bidder will therefore also have the highest virtual valuation. The optimal auction is thus simply the Vickrey or VCG auction, but with a suitable minimum bid or reserve price, to avoid awarding the item if all bidders have negative virtual valuation.

But Myerson's result does much more than solve the optimal auction problem. It also gives a full characterisation of the class of DSIC mechanisms in this setting, showing that this is exactly the class of **monotone allocations**. It also shows that for every implementable social choice function, DSIC payments are uniquely determined. In deterministic mechanisms, a winning bidder will pay their **critical bid**.

Theorem 4 (Myerson Auction, part 2)**.** *In single-parameter (for our purposes: single-item) domains, DSIC-implementable social choice functions are exactly those that are monotone. Monotone allocation means that for $b_i \leq b_i'$ and any choice of $\mathbf{b}_{-i}$, $x_i(b_i, \mathbf{b}_{-i}) \leq x_i(b_i', \mathbf{b}_{-i})$. Furthermore, there is a unique payment function that makes $(\mathbf{x}, \mathbf{p})$ a DSIC mechanism, which is $p_i(\mathbf{b}) = \int_0^{b_i} z \cdot \frac{d}{dz} x_i(z, \mathbf{b}_{-i}) dz$. For $0-1$-allocations, this simplifies to a players' payment equalling their critical bid, that is the lowest bid for which they would still have won the item keeping $\mathbf{b}_{-i}$ constant.*

This is one of the key results in mechanism design, as it greatly aides mechanism design in any single-item or single-parameter setting, even for aims other than revenue maximisation, or for settings other than single-item auctions. For instance, it is easy to see that the welfare-maximising social choice function is monotone, and that VCG payments coincide precisely with critical bids; truthfulness of the Vickrey auction therefore already follows from Myerson's result.

Notice that for deterministic single-item domains, the allocation function is binary. Monotonicity therefore means that for any choice of $\mathbf{v}_{-i}$, there is a single jump from 0 to 1 in the output of $x_i(v_i, \mathbf{v}_{-i})$ as v_i increases, and that the location of this jump – the critical bid – determines i's payment if they do get allocated the item. Deterministic single-item mechanisms are therefore entirely determined by the critical bids; and DSIC on this domain is equivalent to i's critical bids depending only on $\mathbf{v}_{-i}$, not v_i. For single-bidder "auctions", this simplifies further, and tells us that the only truthful mechanisms are posted-price.

Myerson's result is also remarkable as monotone allocations on single-parameter domains, and the VCG mechanism / affine maximisers on the unrestricted domain, are two of very few complete characterisations of truthfulness that are known.

Some other objective functions have been discussed. Of particular interest is for instance the **makespan**, which is the maximum allocation to any one bidder. This is usually discussed in the setting of scheduling tasks to machines or workers, where the aim is to minimise the time until the last task is finished. This is closely related to the concept of **maximin fairness**, in which one seeks to maximise the minimum allocation that any player receives. (This could be seen as an egalitarian version of welfare maximisation.) While some of these – in particular makespan minimisation – have been studied extensively, relatively little is known about optimal mechanism design for objectives beyond welfare and revenue maximisation.

Market Intermediation

Most of classical mechanism design, and indeed auction theory, assumes a one-sided setting, in which there is one mechanism or auctioneer facing a set of relatively homogeneous bidders. Those bidders may differ in some way, for instance their types might be drawn from different priors, but the semantics of interaction tend to be uniform. The main aim of this thesis is to explore a more general setting of **market intermediation**, where the mechanism or auctioneer faces two or more classes of bidders, usually one or more sellers and one or more buyers, and aims to facilitate trade between them.

This setting is similar to an auction, in that the outcome consists of an allocation and a payment vector, and bidders only care about their own allocation and payment. For sellers, a positive allocation indicates that we *buy* an item from them, and the corresponding payment would be made *to* them. Myerson's monotonicity result still extends to this case, but care must be taken to interpret it correctly for sellers: The function $x_i(v_i, \mathbf{v}_{-i})$ now starts at 0 for a high value of v_i, and has a single jump to 1 as v_i decreases. Seller i will sell their item if their bid is *below* that critical bid.

We will consider several variations of this setting. One distinction is whether the auctioneer faces players one after another, or all simultaneously. The latter is called **offline** market intermediation, or sometimes also a **double auction**. The auctioneer collects bids from all buyers and sellers at the same time, and computes an allocation (which goods to transfer from which sellers to which buyers) and payments. On the other hand, if the auctioneer faces players one at a time, the setting is called **online market intermediation**. In this, the auctioneer trades which each trader separately; effectively conducting a sequence of single-bidder mechanisms, which are thus posted-price. The auctioneer may buy from a seller, and then later sell to a buyer.

Definition 18 (Market Intermediation)**.** *In market intermediation, a mechanism faces one or more sellers, and one or more buyers, and aims to facilitate trade.*

1. *If the intermediary faces all bidders at once, the setting is called a double auction.*

2. *If the intermediary faces them one at a time, it is called online market intermediation.*

In the auction setting, we usually have a straightforward feasibility constraint: The auctioneer can only allocate as much as they have. In market intermediation, we will generally require a similar condition, that the intermediary may only sell what they buy. This constraint is meaningful in both the offline and online versions. In the offline case, it is a restriction on the allocation vector. In the online version, it means that the intermediary must first buy an item, for them to be able to sell it later. In some cases, a stronger constraint is more appropriate, stating that the intermediary themselves may never own the item. This implies that they must sell exactly as many units of the item as they buy. This constraint is particularly justified in the context of welfare optimisation, as all items should be allocated in this case. We will only use this constraint in the offline version, as it does not make much sense in the online setting.

Definition 19 (Intermediation Feasibility Constraints)**.**

1. *An offline intermediation mechanism is said to be feasible or no-short-selling, if* $\sum_{i \in S} x_i \geq \sum_{j \in B} x_j$ *(ex post), where* S *denotes the set of sellers and* B *denotes the set of buyers.*

2. *An online intermediation mechanism is feasible if the intermediary's stock level is always nonnegative.*

3. *An offline intermediation mechanism is said to be inventory-balanced if* $\sum_{i \in S} x_i = \sum_{j \in B} x_j$ *(ex post).*

As in the auction setting, the auctioneer may care about social welfare maximisation, or revenue maximisation. The former goal at first might seem already solved, as clearly this setting is within the realm of the VCG mechanism. However, there is two twists specific to the intermediation domain:

Firstly, in a one-sided auction setting, payments only flow from buyers to the auctioneer. In the intermediation setting this is not true, as the intermediary will pay the seller(s). Crucially, VCG does not guarantee that the payment to the sellers be at most the payment from the buyers. The intermediary thus might have to subsidise trade with their own or outside money. Indeed, Myerson and Satterthwaite [99] showed that indeed the optimal social welfare cannot always be achieved even for a single buyer and single seller, without outside subsidies. But such subsidy may not always be feasible. It is thus customary in the intermediation setting to impose a **budget balance** constraint on the mechanism, requiring that revenue always be nonnegative. In some cases, an even stricter version is assumed, requiring that the intermediary always end up with no change in money at the end of trade.

Definition 20 (Budget Balance)**.**

1. *An intermediation mechanism is said to be (weakly) budget balanced if the intermediary never loses money:* $\sum_{i\in S} p_i \leq \sum_{j\in B} p_j$.
2. *An intermediation mechanism is said to be strongly budget balanced if the intermediary never loses nor gains money:* $\sum_{i\in S} p_i = \sum_{j\in B} p_j$.

Both of these conditions are usually taken to require their respective condition ex post.

Secondly, notice that since in our model sellers have a nonnegative valuation (and an endowment of items), social welfare is already nonzero before trade happens. It is therefore also common to focus not on the social welfare post trade, but on the increase in social welfare, also called **gain from trade**. This only makes sense in light of the previous comment; if we were to implement social welfare exactly, we would also implement gain from trade exactly. The distinction is thus in which of the two quantities we compare to its respective optimum when we reason about approximation mechanisms that also satisfy, e.g., budget balance.

For revenue maximisation, Myerson and Satterthwaite [99] give a direct extension of Myerson's single-item auction result (theorem 3). They show that for the setting with one seller, one buyer, a single item and independent priors, an approach based on virtual valuations gives the optimal double auction.

1.4 Outlook

Our main aim is to explore several open questions in this market intermediation setting further, as well as links to neighbouring fields. The remainder of this thesis will be structured into four main parts, each dealing with different questions. We will introduce the concepts and previous literature relevant to each of these topics separately in the corresponding chapter, where they go beyond the general literature review given in this introduction.

Firstly, we will focus on revenue optimisation in offline intermediation, when bidders' valuations are drawn from a correlated prior. This has not been looked at before for market intermediation. We show a polynomial-time algorithm for computing the optimal double auction with a single buyer and single seller. For two or more buyers, we show NP-hardness.

Secondly, our work on correlated priors links directly with previous results in auction theory, as well as to the less explored field of reverse auctions. Chapter 3 will detail these connections, and for the first time show an asymmetry between auctions and reverse auctions.

In Chapter 4 we will explore revenue-maximisation in an online intermediation setting. This is quite different in flavour from pure mechanism design results, and strongly connected to research on online algorithms. We give both upper and lower bounds on what is possible in this setting when competing against an offline adversary.

Lastly, we extend research on welfare maximisation with strong budget balance to a multi-unit setting. We give a full characterisation of the class of truthful strongly budget balanced mechanisms, and show several bounds on the approximation ratio.

Part 2

Single-item Double Auctions with Correlated Priors

2.1 Introduction

We first consider an offline double auction scenario, where the intermediary collects bids from one or more sellers and buyers and determines payments and allocations. Real-world instances of this are manifold, including in electronic markets. Companies such as eBay or Amazon (Marketplace) match sellers and buyers, and charge a fee for each successful transaction. Our aim is to maximize the intermediary's profit in such settings.

There is an extensive literature on this challenge, some of which is discussed below, but it mostly considers the case of many buyers and/or sellers with independent priors. Our interest here is different, in that we assume only a constant number of buyers and sellers (in the simplest version, just one of each), and the complexity arises from their joint probability distribution of valuations for the item. In the simplest version of this, where there is just one buyer and one seller, the intermediary can profit from buying the item from the seller and selling at a higher price to the buyer. We assume their valuations for the item come from a known joint distribution, which is the input to the problem. We consider two versions: the "no short selling" version, with the natural constraint that we cannot sell more items than we buy; and the more restrictive "balanced inventory" version, where in addition we must sell all the items we buy. For multiple buyers and sellers, we assume that the intermediary can buy and sell multiple items (but that each buyer/seller has unit demand/supply).

Related Work

The problem of optimal mechanisms in a market intermediation setting was first studied by Myerson and Satterthwaite [99]. In addition to an impossibility result for ex-post efficiency in a bilateral trade setting without an intermediary, they show optimal intermediation mechanisms for both social welfare as well as the intermediary's revenue in the case of one buyer and one seller, whose valuations are independent. Their revenue-maximization result is similar to Myerson's seminal single auction result [98] in that it, too, uses virtual valuation functions, for both buyer and seller. Welfare maximization for multiple buyers and sellers has been further studied for instance by McAfee [92] or more recently by [52, 57, 34].

Our own interest is chiefly in the complexity of computing optimal – that is, revenue-maximizing – or near-optimal mechanisms in the market intermediation setting. Prior work in this area has focused on the case where sellers' and buyers' valuations are independent. Deng et al. [45] show optimal and near-optimal mechanisms that can be computed in polynomial time for several variations of this setting, including continuous or discrete distributions and arbitrary or unlimited supply and demand. Niazadeh et al. [102] as well as Loertscher and Niedermayer [86, 87, 85] study a class of mechanisms called respectively *fee-setting mechanisms* or *affine fee schedules* in the independent setting. These are shown by Niazadeh et al. [102] to be able to extract a constant fraction of the optimum revenue in the worst case, under certain assumptions on the buyer's and seller's distribution. Deshmukh et al. [46] compare double auctions to a type of unlimited-supply, unit-demand auction, and give bounds on the worst-case revenue achievable by the intermediary.

Here we are interested in potentially correlated distribution over buyers' and sellers' valuations. The complexity of this has been studied for (non-double) auctions. Papadimitriou and Pierrakos [106] show that for two buyers, an optimal mechanism (for a discrete joint distribution) can be found in polynomial time via a reduction to finding a maximum-weight independent set on a bipartite graph. For continuous distributions they give a FPTAS. For the case of three buyers, in contrast, they show that it is NP-hard to approximate the optimal auction to within a certain constant fraction. Dobzinski et al. [50] show a polynomial-time algorithm for the two-buyer auction through derandomization and give polynomial-time approximation mechanisms for the many-buyers correlated single auction problem, building on previous work by Ronen [110].

2.2 Preliminaries

Definitions, Notation

We consider m buyers indexed by j, and k sellers indexed by i (where m, k are constants), each offering (respectively seeking) a single unit of an indivisible good. For fixed m, k, we use "$m \times k$" as shorthand for the m buyers, k sellers case. They cannot trade with each other directly, and can only trade with the intermediary. We assume each seller i has some valuation s_i and each buyer j has some valuation b_j for an item, and that these are drawn from a given joint probability distribution ψ. We focus on discrete distributions. For simplicity, we assume that the support of ψ is a grid of size n^{k+m}, with each player having possible valuations $\{i : 1 \leq i \leq n\}$ (as shown in Figure 2.1). The distribution ψ is assumed to be given explicitly as a $n \times n$ probability matrix.

In this chapter we focus on individually rational and incentive compatible mechanisms. Let $\mathbf{b}$ and $\mathbf{s}$ denote the vector of buyers' and sellers' bids received by the mechanism, and let $\mathbf{b}_{-j}$ and $\mathbf{s}_{-i}$ be (respectively) the bids of buyers other than j, and sellers other than i. Analogous to auctions, there are two equivalent ways in which we can define a deterministic, incentive compatible mechanism in this setting. Firstly, we can focus on allocations. For each seller / buyer we define a set S_i / B_j ($\subseteq$ supp(ψ)) of bid vectors in which we buy an item from seller i / sell an item to buyer j. Incentive compatibility means monotonicity of allocations, meaning for each seller i, if $(s, b) \in S_i$, and $s'_i < s_i$, then $(s'_i, s_{-i}, b) \in S_i$. In words, if given everyone else's bids s_{-i}, b, seller i's item would be bought by the intermediary if i bid

s_i, then it would also be bought for any lower bid s_i'. For short we will say that S_i is "downward-closed" in the direction of s_i. Similarly for each buyer j, B_j needs to be upward-closed in the direction of b_j.

Equivalently, we may think of a mechanism in terms of critical bids. Myerson tells us that the unique payments that make a monotone allocation rule (as just defined via the regions S_i and B_j) are precisely the critical bids. That is, the lowest (highest) bid for which a buyer (seller) would still be allocated the item (the sale of their item) if everyone else's bids remained fixed. We write $\sigma_i(\mathbf{b}, \mathbf{s}_{-i})$ (respectively, $\beta_j(\mathbf{b}_{-j}, \mathbf{s})$) for these critical bids. For simplicity sometimes just $\beta_j(\mathbf{b}, \mathbf{s})$ and $\sigma_i(\mathbf{b}, \mathbf{s})$. If $s_i \leq \sigma_i(\mathbf{b}, \mathbf{s}_{-i})$ we buy an item from seller i (paying $\sigma_i(\mathbf{b}, \mathbf{s}_{-i})$), and similarly if $b_j \geq \beta_j(\mathbf{b}_{-j}, \mathbf{s})$ we sell an item to buyer j (charging $\beta_j(\mathbf{b}_{-j}, \mathbf{s})$). We write $\beta_j(\mathbf{b}_{-j}, \mathbf{s}) = n + 1$ to indicate that a mechanism does not sell to buyer j at all for this combination of others' bids, independently of j's bid. Similarly $\sigma_i(\mathbf{b}, \mathbf{s}_{-i}) = 0$ to indicate not buying from seller i.

It is easy to see that these two yield equivalent definitions. Clearly S_i is simply the region "above" σ_i in the direction of s_i, (the graph of) which in turn is the boundary of S_i. Similarly B_j is the region below β_j in direction b_j. This is a slight generalization of the conceptually simpler picture in auctions, which was described in this geometric form by Papadimitriou and Pierrakos [106]. Here we have for each bidder a region B_j where they win the item, and a critical bid function β_j that gives their payment. If there is a single item to be sold, no two of the B_j may overlap. This constraint too generalizes to the market intermediation setting. Consider Figure 2.2 in contrast with Figure 2.1 to illustrate the difference. We will further discuss the auction setting in

Chapter 3, and direct the interested reader to Section 3.2 already for a more detailed description of the auction setting. As mentioned above, we consider two variants. In the "no short-selling" setting, we must buy at least as many items from sellers as we sell to buyers; in the "balanced inventory" variant we must buy exactly as many as we sell. Formally in terms of critical bids: (Again these can be expressed equivalently in terms of S_i and B_j.)

$$\begin{aligned}&\text{No Short-Selling:}\\ &\quad \forall(\mathbf{b},\mathbf{s}),\ |\{j : b_j \geq \beta_j(\mathbf{b},\mathbf{s})\}| \leq |\{i : s_i \leq \sigma_i(\mathbf{b},\mathbf{s})\}|\end{aligned} \tag{2.1a}$$

$$\begin{aligned}&\text{Balanced Inventory:}\\ &\quad \forall(\mathbf{b},\mathbf{s}),\ |\{j : b_j \geq \beta_j(\mathbf{b},\mathbf{s})\}| = |\{i : s_i \leq \sigma_i(\mathbf{b},\mathbf{s})\}|\end{aligned} \tag{2.1b}$$

The Geometry of Deterministic 1×1 Market Intermediation

If there is only a single buyer and a single seller, the constraints simplify significantly, most easily expressed in terms of now simply S and B. In the balanced-inventory case, constraint (2.1b) simplifies to $B = S$. That is, a mechanism in this setting, with only one buyer and seller each, is determined only by a single region of bid-combinations that yields a successful transaction. In the no-short-selling case, constraint (2.1a) simplifies to $B \subseteq S$. That is, S can potentially extend beyond B. However, we can say more, assuming optimality of the mechanism. Recall that by truthfulness, S is down-closed in the seller's direction and B is up-closed in the buyer's direction. If a mechanism is optimal, S must exactly be the down-closure (still in the seller's direction) of B. Firstly, it is easy ot see that the down-closure of B must be contained in S: B is contained in S, and S is down-closed. Secondly, if S extended beyond the down-closure of B, we could strictly improve our revenue by removing this protruding part of S. On the other hand, we may not elect to remove the part of $S - B$ that lies below any point in B due to truthfulness, i.e. down-closedness of S. Figures 2.1 and 2.3 illustrate this. More formally:

Lemma 1. *In the* 1×1 *market intermediation setting, an optimal mechanism will only buy the item from the seller if required to do so by the no-short-selling constraint, or monotonicity.*

Note the contrast with a standard 2-bidder auction, where the shape of the region in which we sell to one buyer does not fully determine the region in which we sell to the other. In a way, in the 1×1 market intermediation setting, we have fewer degrees of freedom to consider than in a two-buyer auction setting.

2.3 The Deterministic One Seller, One Buyer Case

For the case of one seller and one buyer, we show how to compute an optimal deterministic solution using a dynamic programming approach. In section 2.4 we show how to achieve this via modifications to known 2-bidder auctions in this setting, but the runtime guarantee of that approach, while still polynomial, is substantially worse. We represent a mechanism using values

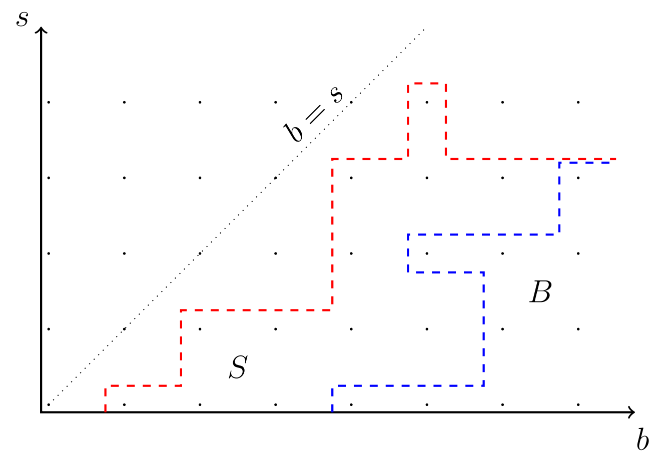

Figure 2.1: Example mechanism in the 1×1 case. Note that S contains B, to avoid short-selling. (In a balanced-inventory auction, B and S should coincide.) B and S lie below the diagonal $b = s$: any point above the diagonal is one where the buyer's bid is less than the seller's. The auction shown is suboptimal: in most of the S region, the item is being bought without being sold. Note that we draw the outline of the regions slightly away from the points on the prior support for easier readability.

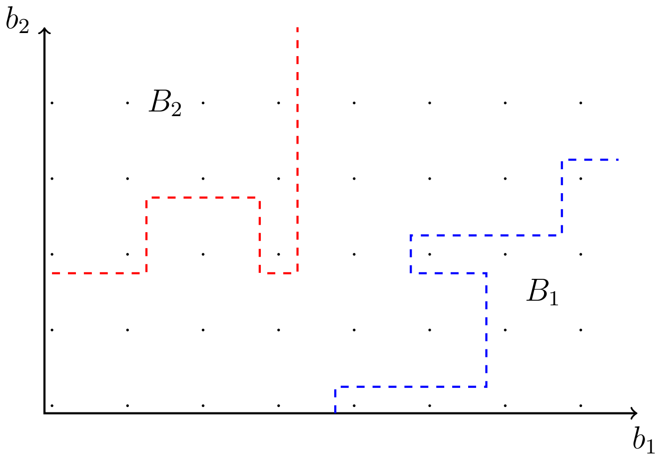

Figure 2.2: Compare this to a two-bidder auction. Here B_1 and B_2 indicate where we sell to each of the two buyers. In the two-bidder auction B_1 and B_2 must be disjoint, as we cannot sell the item twice. In the market intermediation setting, B must be contained in S. Note also that in this setting both B_1 and B_2 are upward-closed in the direction of the respective buyer's bid.

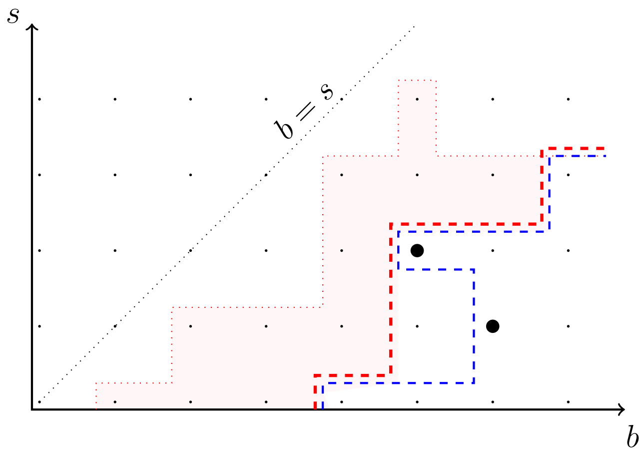

Figure 2.3: Removing the indicated area from S as in figure 1, the expected revenue of the mechanism cannot decrease. Below red line: Remaining region S, right of blue line: Region B. For the remaining part of S that is not also in B, we still buy but not sell the item. This can be optimal, e.g. if there is very high probability weight on the two points indicated. Crucially, if at a point $(b, s) \in S - B$ an optimal mechanism buys but not sells, then there must exist a point $(b, s') \in S \cap B$ with $s' > s$ where it buys and sells. Truthfulness then dictates that it also needs to buy at (b, s).

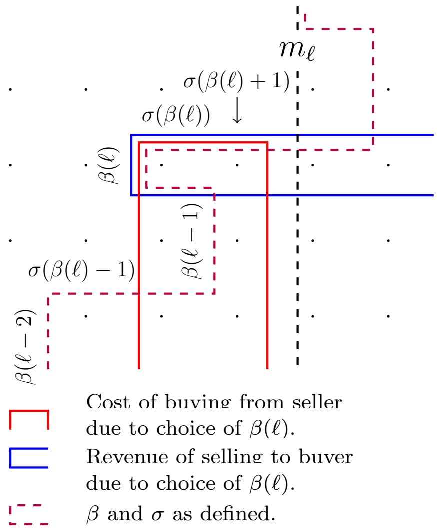

Figure 2.4: For fixed m_ℓ and $\beta(1), \cdots, \beta(\ell - 1)$, the choice of a particular $\beta(\ell)$ influences the expected revenue in two ways: On the one hand, the revenue from selling to the buyer at all points to the right of $\beta(\ell)$ in row ℓ. On the other hand, the cost of buying from the seller for all points below $\sigma(\beta(\ell)) = \ell$ in rows $\beta(\ell) \leq b < m_\ell$.

$\beta = (\beta(1), ..., \beta(n))$ for $\beta(s) \in \{1, \ldots, n+1\}$, where $\beta(s)$ signifies the left-most point in row s that is a member of B. If $\beta(s) = 1$ then the entire row s is in B. We set $\beta(s) = n + 1$ to signify that none of the points in row s are in B. We begin by noting that the contribution to the expected revenue that arises from the choice of one particular $\beta(s)$ does not depend on all the other β_{-s} simultaneously. Consider the expected revenue R for a given B, which is given by:

$$\mathbb{E}[R(\beta)] = \sum_{s=1}^{n} \beta(s) \sum_{b \geq \beta(s)} \psi_{bs} - \sum_{b=1}^{n} \sigma(b) \sum_{s \leq \sigma(b)} \psi_{bs} \tag{2.2}$$

That is, the first sum gives the expected profit from selling at points to the right of each $(\beta(s), s)$, while the second sum gives the cost of buying points that are below each $(b, \sigma(b))$. The contribution of a particular choice for one single $\beta(s)$ to the first of these sums is easily seen to be simply the profit of selling the points in row s to the right of and including $\beta(s)$. The impact of a particular $\beta(s)$ on the second of the sums is slightly more intricate. There are two ways in which the choice of $\beta(s)$ impacts the cost of buying. Firstly, we may have to buy the item at some points in row s, where we would not buy the item otherwise. Consider the minimum of $\beta(s+1), ..., \beta(n)$, say $\beta(t)$. We know that in row t, we buy and sell at points $(\beta(t), t)$,...,(n, t). So by truthfulness, we must also buy at points $(\beta(t), s)$,...,(n, s). This is regardless of our choice of $\beta(s)$. For points to the left of $(\beta(t), s)$, whether we buy the item or not does depend on $\beta(s)$. Secondly, in all those columns in which we buy due to $\beta(s)$, also affect the rows below s. We may increase the buying price from a lower value to s at those points, and we may have to buy the item (due to truthfulness) at points at which we would not otherwise buy it. The magnitude of this effect depends on all the $\beta(1), ..., \beta(s-1)$. This suggests a bottom-up dynamic programming approach, which we develop in this section.

Algorithm for the No Short-Selling Case

We next describe our dynamic programming algorithm first for the no short-selling setting. The idea is as follows: because the optimal choice of $\beta(1), ...,$ $\beta(\ell)$ depends only on the minimum of the $\beta(\ell+1), ..., \beta(n)$, we can iteratively compute all the potential optimal values for row 1 given values of $\min\{\beta(2), ..., \beta(n)\}$; then all optimal values of $\beta(1), \beta(2)$ given all possible values of $\min\{\beta(3), ..., \beta(n)\}$. We do not need to consider all n^2 combinations of $\beta(2)$ and $\beta(1)$. Since given $\beta(2)$ and $\min\{\beta(3), ..., \beta(n)\}$, we can immediately look up the best $\beta(1)$ using the information computed in the first step. We then proceed iteratively up the rows until we have computed tohe optimal values for β.

Let us start by defining $R(n, \beta(\ell), m_\ell)$ to be the expected revenue of the best deterministic mechanism that takes points $(\beta(\ell), \ell)$ and rightward in row ℓ, no points in rows $\ell+1$ and above, and does not have to pay for points in columns m_ℓ to n. We set $R(0, ., .) = 0$. The idea is that we want to capture the best possible revenue extractable from rows 1 to ℓ for a particular choice of $\beta(\ell)$, disregarding the cost of buying in columns m_ℓ to n. We can take $R(n, \beta(\ell), n+1)$ to denote the optimal revenue among mechanisms that have to pay for all rows. More precisely,

$$R(\ell,\beta(\ell),m_\ell) = \max_{\beta(1),\ldots,\beta(\ell-1)} \sum_{s=1}^{\ell} \beta(s) \sum_{b\geq\beta(s)} \psi_{bs} - \sum_{b=1}^{m_\ell-1} \sigma(b) \sum_{s\leq\sigma(b)} \psi_{bs} \tag{2.3}$$

It is easy to see that $\max_{\beta(n)} R(n,\beta(n),n+1)$ gives the revenue of the optimal auction. Indeed, by definition this is the maximum expected revenue extractable from all rows, if we have to pay in all columns. We can then show how to recursively compute the values of R, laying the groundwork for our dynamic programming algorithm.

Theorem 5 (Recursion for the no short-selling case)**.** *The $R(\ell,\beta(\ell),m_\ell)$ as defined above satisfy the following recursion:*

$$\begin{aligned} R(\ell,\beta(\ell),m_\ell) = \max_{\beta(\ell-1)} R\big(\ell-1,\beta(\ell-1),\min\{\beta(\ell),m_\ell\}\big)+ \\ \beta(\ell) \sum_{b\geq\beta(\ell)} \psi_{b\ell} - \ell \sum_{\beta(\ell)\leq b<m_\ell} \sum_{s\leq\ell} \psi_{bs} \end{aligned} \tag{2.4}$$

Proof. We can check this by splitting up the explicit formula for $R(\ell,\beta(\ell),m_\ell)$ into terms for rows below ℓ and row ℓ, and columns to the left of $\min(\beta(\ell),m_\ell)$ and those between the two.

$$\begin{aligned} R(\ell,\beta(\ell),m_\ell) = \sum_{s=1}^{\ell-1} \beta(s) \sum_{b\geq\beta(s)} \psi_{bs} + \beta(\ell) \sum_{b\geq\beta(\ell)} \psi_{b\ell} - \\ \sum_{b=1}^{\min(\beta(\ell),m_\ell)-1} \sigma(b) \sum_{s\leq\sigma(b)} \psi_{bs} - \sum_{b=\beta(\ell)}^{m_\ell-1} \sigma(b) \sum_{s\leq\sigma(b)} \psi_{bs} \end{aligned}$$

Observe that for $b \geq \beta(\ell)$, $\sigma(b)$ will be equal to ℓ (in the $(\ell,\beta(\ell))$-auction), so the last term in the above sum is precisely $\ell \sum_{b=\beta(\ell)}^{m_\ell-1} \sum_{s\leq\ell} \psi_{bs}$. Similarly,

$\min(\beta(\ell), m_\ell)$ is precisely the $m_{\ell-1}$ we used in the recursion, and therefore the first and third term are precisely $R(\ell-1, \beta(\ell-1), m_{\ell-1})$. Putting these together, we get that:

$$R(\ell, \beta(\ell), m_\ell) = R(\ell-1, \beta(\ell-1), m_{\ell-1}) + \beta(\ell) \sum_{b \geq \beta(\ell)} \psi_{b\ell} - \ell \sum_{b=\beta(\ell)}^{m_\ell - 1} \sum_{s \leq \ell} \psi_{bs} \tag{2.5}$$

i.e. precisely our claimed recursion. (The max follows from optimality of the auction.) The second term on the right hand side is the revenue from selling at points due to the choice of $\beta(\ell)$, while the third term accounts for the cost of buying at points due to this choice. Figure 2.4 illustrates these two terms. Note that if $\ell = 1$ then the first term vanishes since we defined $R(0, ., .) = 0$, and we are left with the explicit formula for $R(1, ., .)$. □

We can therefore compute the $R(\ell, \beta(\ell), m_\ell)$ recursively, as claimed. This suggests the following algorithm, listed below as Algorithm 1. This can easily be augmented to keep track of the values used for the $\beta(s)$, and to return the optimal β together with its expected revenue. Therefore, we can compute the optimal region B and thereby the optimal mechanism in the no short-selling setting in time $O(n^3)$.[1]

We can easily modify this algorithm to return the optimal mechanism that satisfies the balanced inventory property. We show the details in section 2.5. This modified algorithm also runs in time $\mathcal{O}(n^3)$, but with slightly better constant factors.

[1]Careful analysis of the algorithm presented shows that the last summand in the recursion for $R()$ has $(m_\ell - \beta(\ell)) \cdot \ell$ summands. It is easy to see however that we need not recompute the inner sum from scratch in each iteration. We can thus easily make the computation of the recursion run in linear time, giving the overall running time stated.

Algorithm 1 Optimal revenue in the no short-selling setting

for $\ell = 1, ..., n$ **do**
  for $\beta(\ell) = 1, .., n$ **do**
    for $m_\ell = 1, ..., n+1$ **do**
      if $\ell = 1$ **then**
        $R(1, \beta(1), m_1) \leftarrow \beta(1) \sum_{b \geq \beta(1)} \psi_{b1} - \sum_{\beta(1) \leq b < m_1} \psi_{b1}$
      else
        Compute $R(\ell, \beta(\ell), m_\ell)$ using the recursion in theorem 5.
return $\max_{\beta(n)} R(n, \beta(n), n+1)$.

2.4 Reducing the 1 $\times$ 1 Case to a 2-bidder Auction

For the two bidder single auction setting, Papadimitriou and Pierrakos [106] take a different approach. They reduce the problem of the optimal 2-bidder optimal auction to that of finding a maximum-weight independent set on a bipartite graph. We can similarly reduce the no short-selling market intermediation problem to this, with only a few steps, but at the cost of a larger polynomial runtime bound. Firstly, in the single item auction setting we have a strong exclusivity constraint: We cannot allocate the item to both bidders at the same time, so any valid mechanisms must satisfy $A_1 \cap A_2 = \emptyset$. In the intermediation setting we do not have such a constraint directly; rather we must have B contained in S. However, we can consider $\overline{S}$ the complement of S, and indeed have the same constraint $\overline{S} \cap B = \emptyset$. In an analogous fashion to their construction we can define marginal revenue contributions f_b for the buyer and f_s for the seller, the latter in our setting referring to the marginal profit contribution of *not buying* at point (i, j). It is easy to see that the maximum-weight independent set on the so constructed bipartite graph will give the optimal mechanism in the no short-selling setting. This approach runs in time $\mathcal{O}(n^6)$ for a support size of n^2.

It is furthermore possible to obtain an optimal deterministic mechanism via the derandomization approach described in Dobzinski et al. [50]. It is easy to see that the optimal truthful-in-expectation mechanism can be computed through an LP in either market intermediation setting. The derandomization applies to both without modification. A cursory analysis of this approach gives a runtime of $\mathcal{O}(n^7)$ for a support size of n^2. (The LP involves $2n^2$ variables. The exact time complexity depends on the LP solver chosen; Karmarkar's algorithm runs in time $\mathcal{O}(n^{3.5})$ for n variables.)

2.5 DP for the Balanced Inventory Case

For completeness, we show how to modify the dynamic programming approach for the balanced inventory setting, which we omitted in section 2.3. Let $R'(\ell, \beta(\ell), \beta(\ell+1))$ be the revenue of the optimal mechanism that satisfies the balanced inventory property, sells at points $(\beta(\ell), \ell)$ and rightward, and does not have to pay for columns $\beta(\ell+1)$ to n. Then the value of R' is given by the following explicit formula analogous to the no short-selling case.

$$R'(\ell, \beta(\ell), \beta(\ell+1)) = \max_{\beta(1)\leq\ldots\leq\beta(\ell-1)} \sum_{s=1}^{\ell} \beta(s) \sum_{b\geq\beta(s)} \psi_{bs} - \sum_{b=1}^{\beta(\ell+1)-1} \sigma(b) \sum_{s\leq\sigma(b)} \psi_{bs} \quad (2.6)$$

And we can prove a very similar recursion.

Theorem 6 (Recursion for the balanced inventory setting).
The $R'(\ell, \beta(\ell), \beta(\ell+1))$ *as defined above satisfy the following recursion:*

$$R'(\ell, \beta(\ell), \beta(\ell+1)) = \max_{\beta(\ell-1)\leq\beta(\ell)} R'(\ell-1, \beta(\ell-1), \beta(\ell)) + \beta(\ell) \sum_{b\geq\beta(\ell)} \psi_{b\ell} - \ell \sum_{\beta(\ell)\leq b<\beta(\ell+1)} \sum_{s\leq\ell} \psi_{bs} \quad (2.7)$$

Proof. The proof of theorem 5 applies mutatis mutandis. □

This suggests a slight modification of the no short-selling algorithm, listed as algorithm 2.

Algorithm 2 Optimal revenue in the balanced inventory setting

for $\ell = 1, ..., n$ **do**
  for $\beta(\ell) = 1, .., n$ **do**
    for $\beta(\ell+1) = \beta(\ell), ..., n+1$ **do**
      if $\ell = 1$ **then**
        $R'(1, \beta(1), \beta(2)) \leftarrow \beta(1) \sum_{b\geq\beta(1)} \psi_{b1} - \sum_{\beta(1)\leq b<\beta(2)} \psi_{b1}$
      else
        Compute $R'(\ell, \beta(\ell), \beta(\ell+1))$ using the recursion in theorem 6.
return $\max_{\beta(n)} R(n, \beta(n), n+1)$.

This algorithm obviously has slightly lower runtime than the one in the no short-selling setting, since we only need to loop through values $\beta(\ell+1) \geq \beta(\ell)$, and only consider values $\beta(\ell-1) \leq \beta(\ell)$ in the maximum in the recursion. This gives a constant-factor improvement in runtime.

2.6 Strong NP-hardness of the Multiple Buyers Case

For three or more buyers, it follows from Papadimitriou and Pierrakos [106] that computing the optimal mechanism is NP-hard. We show that this is also true for the 2×1 case (i.e. two buyers, one seller) in the no-short-selling setting by reducing from Maximum Independent Set. The idea here is to place high probability weight on high-revenue points along a diagonal in the $s = 1$ plane for each vertex of a given instance of Independent Set. We then use appropriately placed high-probability points for each of the edges to "force" a higher buying price for (at least) one of any two points corresponding to adjacent vertices. We can do this in a way that ensures that in the optimal mechanism the number of vertex points with a low buying price is maximized and corresponds to the maximum independent set.

Theorem 7 (NP-hardness)**.** *It is strongly NP-hard to compute the optimal mechanism in the 1 seller, 2 buyers setting with no short selling.*

Proof. In the following we construct a prior distribution in (b_1, b_2, s)-space. We will "choose" points, and place equal probability mass $\frac{1}{|V|+|E|}$ on all of these chosen points. In the analysis we will omit these weights to simplify the algebra. We place probability 0 on all other points in the prior support. We use K_1 and K_2 as constants whose values we define at the end of the proof.

The Construction

Given a graph (V, E) with $|V| = n$, pick any order of vertices and begin by placing probability weight $\frac{1}{|V|+|E|}$ on point $(K_1 + \lfloor \frac{n}{2} \rfloor - i, K_1 - \lfloor \frac{n}{2} \rfloor + i, 1)$ for each vertex $0 \leq i < n$. Next, enumerate the edges e_j, $0 \leq j < |E|$. We will write each edge as $e_j = (e_{j1}, e_{j2})$, where $e_{j1} < e_{j2}$ in the order of vertices just picked. For each edge put probability weight $\frac{1}{|V|+|E|}$ on point $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, K_2 + j)$. That is, we put probability weight for each edge on a point that has the same b_1-coordinate as the vertex point for its lower-numbered vertex and the same b_2-coordinate as its higher vertex. We choose these edge points with a different s-coordinate each, and all of them with a higher s-coordinate than the vertex points. It is clear that if the mechanism wants to buy and sell at an edge point $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, K_2 + j)$, it will also need to sell (and therefore buy by truthfulness) at one of the points $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j2}, K_2 + j)$ or $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j1}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, K_2 + j)$, when it sells to buyer 1 or buyer 2, respectively. But by truthfulness this entails a raised purchase price of $K_2 + j$ at the corresponding vertex points directly below $((K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j2}, 1)$ or $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j1}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}, 1))$ where it had otherwise been 1. Figure 2.5 illustrates this construction.

Reducing from Maximum Independent Set

Now, in order to ensure that the optimal mechanism raises the purchasing price at all vertex points except those that are in an independent set of maximum size, we need to pick constants K_1, K_2 in a way that ensures that:

1. The optimal mechanism always buys and sells at all the edge points.

2. The optimal mechanism raises the purchasing price at as few vertex points as possible.

From condition 1: The worst possible selling price at any edge point is given by $K_1 - \lfloor \frac{n}{2} \rfloor$, and the highest possible purchase price is $K_2 + |E| \leq K_2 + n^2$, for a revenue that is at least $K_1 - \lfloor \frac{n}{2} \rfloor - K_2 - n^2$. On the other hand, buying and selling at an edge point could necessitate a higher purchasing price at a vertex point, raising it by an amount that is bounded above by $K_2 + n^2$ as well. The profit obtained from the edge point must outweigh this. So in order to ensure that the optimal mechanisms buys and sells at all edge points, we need to ensure:

$$K_1 - \left\lfloor \frac{n}{2} \right\rfloor - 2K_2 - 2n^2 > 0$$

From condition 2: We need to ensure that if for an edge point $(K_1 + \lfloor \frac{n}{2} \rfloor - e_{j2}, K_1 - \lfloor \frac{n}{2} \rfloor + e_{j1}), K_2 + j)$, only one of the two corresponding vertex points already has a purchase price of at least K_2 due to another edge, but the other is still 1, the optimal mechanism will always prefer to sell to the buyer whose corresponding vertex point already has a high price. In other words, we need to ensure that the potential difference in revenue from selling to one buyer over the other is outweighed by the required raise in the purchase price by (at least) $K_2 - 1$. But the highest difference in selling price is bounded by n, and so we required that $K_2 > n + 1$.

Combining the two we get our desired result: Set $K_2 = 2n$ and $K_1 = 4n^2$ in the above construction for a given instance of Maximum Independent Set. Since the optimal mechanism will buy and sell at all the edge points, it is

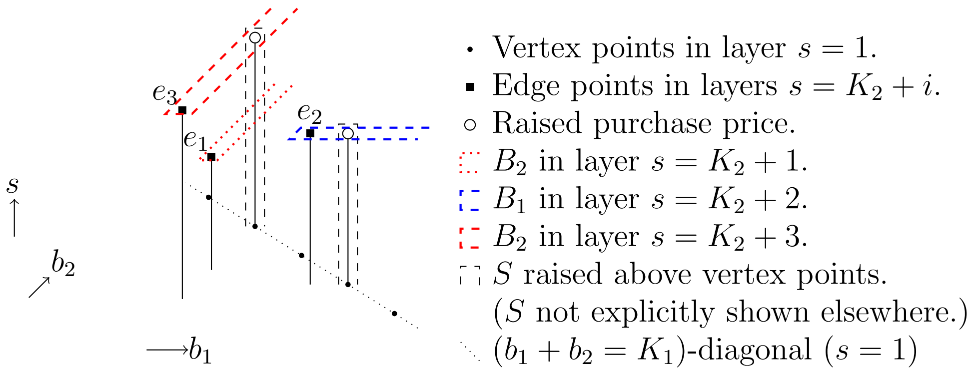

Figure 2.5: The construction for the reduction from Maximum Independent Set.

clear that at most one vertex point corresponding to two adjacent vertices can have a purchase price of 1. On the other hand, in the optimal mechanism the number of vertex points with a raised purchase price will be minimized. Therefore, the vertex points with purchase price 1 in the optimal mechanism correspond to the vertices of a maximum independent set in the graph. $\square$

2.7 Truthful-in-expectation Mechanisms

While in the preceding section we have shown that we cannot compute an optimal deterministic mechanism for the general case, we can however compute the optimal truthful-in-expectation mechanism for a fixed number of buyers and sellers. In single-item auctions a randomized mechanism is easily described by allocation probabilities $x_i(\mathbf{v})$ and expected payments $p_i(\mathbf{v})$ for all players for each possible bid vector. In the market intermediation setting with multiple buyers and sellers this is not obviously the case. For instance, there are many ways in which to make allocation probabilities of $\frac{1}{2}$ for each of two buyers and two sellers into a randomization over valid outcomes. The mechanism could flip a coin and buy from seller 1 and sell to buyer 1 on heads, seller 2 and buyer 2 on tails. It could not, however, independently flip four coins if we want to fulfill condition (2.1a) respectively (2.1b) ex-post. In the following we will consider the balanced inventory case. Our arguments easily extend to the no short-selling case. We first show that for our purposes, it is indeed sufficient to consider only the marginal allocation probabilities x_i, y_j and expected payments p_i, q_j. First, observe that any two randomized mechanisms that have the same marginal probabilities and expected payments will lead to identical expected utilities for players and expected revenue. It remains to show that any sensible vector of marginal allocation probabilities can be made into a probability distribution over valid outcomes (i.e. allocations which buy exactly as many items as they sell).

Theorem 8. *Let x, y be k-dimensional vectors of probabilities, i.e. $0 \leq x_i, y_i \leq 1$, with $\sum_i x_i = \sum_i y_i$. Then there exists a joint probability distribution over two k-dimensional $0/1$ vectors $\{\mathbf{a}, \mathbf{b} \in \{0,1\}^k | \sum_i a_i = \sum_i b_i\}$ which satisfies $\Pr(a_i = 1) = x_i$ and $\Pr(b_i = 1) = y_i$.*

Proof. Let $H_{2k} = [0,1]^{2k}$ be the $2k$-dimensional hypercube, and $H^*_{2k} = \{0,1\}^{2k}$ its vertices. Let $D_{2k} = \{(x,y) \in H_{2k} | \sum x_i = \sum y_i\}$ be the "generalized diagonal" of the hypercube. Let $D^*_{2k} = \{(x,y) \in \{0,1\}^{2k} | \sum x_i = \sum y_i\}$ be the vertices of H_{2k} with as many x-coordinates set to 1 as y-coordinates. That is, this is the set of valid (deterministic) allocation vectors for k buyers and sellers. Then our claim is equivalent to saying that D_{2k} is (in) the convex hull of D^*_{2k}. By the Krein-Milman theorem a convex set S is exactly the convex hull of its extreme points. An extreme point $s \in S$ is any point in S which can not be written as a convex combination of points in $S \setminus s$. Clearly D_{2k} is convex. It remains to show that the extreme points of D_{2k} are precisely D^*_{2k}. Clearly $D^*_{2k} \subseteq D_{2k}$. So let $(\mathbf{x}, \mathbf{y}) \in D_{2k} \setminus D^*_{2k}$ be a point in D_{2k} that does not have all elements equal to 0 or 1. We show that $(\mathbf{x}, \mathbf{y})$ is not an extreme point of D_{2k}.

If there is exactly one x_i with $0 < x_i < 1$, then there must be at least one y_j with $0 < y_j < 1$. (Otherwise $\sum x_i \notin \mathbb{N}$, but $\sum y_j \in \mathbb{N}$, which contradicts the assumption that $\sum x_i = \sum y_j$.) Then for $0 < \epsilon < \min\{x_i, 1 - x_i, y_j, 1 - y_j\}$, we have that $(x_i + \epsilon, y_j + \epsilon, \mathbf{x}_{-i}, \mathbf{y}_{-j}) \in D_{2k}$, and also $(x_i - \epsilon, y_j - \epsilon, \mathbf{x}_{-i}, \mathbf{y}_{-j}) \in D_{2k}$. Clearly $(\mathbf{x}, \mathbf{y})$ is a convex combination of these two. If there is at least two distinct $0 < x_i, x_\ell < 1, i \neq \ell$, then for $0 < \epsilon < \min\{x_i, 1 - x_i, x_\ell, 1 - x_\ell\}$, we have that $(x_i + \epsilon, x_\ell - \epsilon, \mathbf{x}_{-i\ell}, \mathbf{y}) \in D_{2k}$, and also $(x_i - \epsilon, x_\ell + \epsilon, \mathbf{x}_{-i\ell}, \mathbf{y}) \in D_{2k}$. Again, clearly $(\mathbf{x}, \mathbf{y})$ is a convex combination of these two. Similarly, if there is no $0 < x_i < 1$ there is at least two such y_j, y_ℓ. So D_{2k} is the convex hull of D^*_{2k}. This shows our claim. □

From this it follows immediately that we need only concern ourselves with the marginal allocation probabilities in computing an optimal randomized mechanism. Therefore we can write this as an LP following the approach of Dobzinski et al. [50] for auctions. We state the theorem for the balanced inventory case. The no-short-selling case follows immediately.

Theorem 9 (The optimal randomized mechanism as an LP).
For a fixed number of buyers and sellers, we can compute the optimal truthful-in-expectation mechanism in the balanced inventory setting using a linear program that is polynomial in the size of the prior.

Proof. Consider the following LP, where $\psi_{\mathbf{b},\mathbf{s}}$ denotes the prior probability of bid vectors $(\mathbf{b}, \mathbf{s})$.

$$\min \sum_{\mathbf{b},\mathbf{s}} \psi_{\mathbf{b},\mathbf{s}} \big(\sum_i -p_i(\mathbf{b},\mathbf{s}) + \sum_j q_j(\mathbf{b},\mathbf{s}) \big) \quad \text{s.t.} \tag{2.8a}$$

$$\sum_i x_i(\mathbf{b},\mathbf{s}) = \sum_j y_j(\mathbf{b},\mathbf{s}) \qquad \forall \mathbf{b},\mathbf{s} \tag{2.8b}$$

$$p_i(\mathbf{b},\mathbf{s}) \geq 0, \quad 0 \leq x_i(\mathbf{b},\mathbf{s}) \leq 1 \qquad \forall i,\mathbf{b},\mathbf{s} \tag{2.8c}$$

$$q_j(\mathbf{b},\mathbf{s}) \geq 0, \quad 0 \leq y_j(\mathbf{b},\mathbf{s}) \leq 1 \qquad \forall j,\mathbf{b},\mathbf{s} \tag{2.8d}$$

$$x_i(\mathbf{b}, s_i', \mathbf{s}_{-i}) s_i + p_i(\mathbf{b}, s_i', \mathbf{s}_{-i}) \leq$$

$$\leq -x_i(\mathbf{b},\mathbf{s}) s_i + p_i(\mathbf{b},\mathbf{s}) \qquad \forall i, s_i', \mathbf{b},\mathbf{s} \tag{2.8e}$$

$$y_j(b_j, \mathbf{b}_{-j}, \mathbf{s}) b_j + p_i(b_j, \mathbf{b}_{-j}, \mathbf{s}) \leq$$

$$\leq y_j(\mathbf{b},\mathbf{s}) b_j - q_j(\mathbf{b},\mathbf{s}) \qquad \forall j, b_j', \mathbf{b},\mathbf{s} \tag{2.8f}$$

$$x_i(\mathbf{b}, s_{\max}, \mathbf{s}_{-i}) = 0 \qquad \forall i, \mathbf{s}_{-i}, \mathbf{b} \tag{2.8g}$$

$$q_j(0, \mathbf{b}_{-j}, \mathbf{s}) = 0 \qquad \forall j, \mathbf{b}, \mathbf{s}_{-j} \tag{2.8h}$$

Here, conditions (b)-(d) ensure feasibility of the (class of) mechanisms. (e) and (f) ensure truthfulness on the sellers' and buyers' side, respectively. Together with (g) and (h) they give IR. The proof is the same as in Dobzinski et al. [50] for auctions. This gives us a non-empty class of valid mechanisms, all of which are optimal and have the same expected utilities for the players. Finding an arbitrary member of this class is easy.[2] □

[2]It is equivalent to finding for a given point in $(\mathbf{x}, \mathbf{y}) \in D_{2k}$ coefficients for a convex combination of vertices in D^*_{2k}. While the number of extreme points is exponential in k, by Carathéodory's theorem only $2k+1$ of the vertices are needed. To find them, we could simply pick any $2k$ vertices of D^*_{2k} to partition D_{2k} into two sets. Discard all the vertices not contained in the same partition as $(\mathbf{x}, \mathbf{y})$, and repeat with the remaining vertices.

2.8 Prior-independent Mechanisms with Approximation Guarantees

In the previous sections we have considered optimal mechanisms that take into account full information on the joint prior. An immediate follow-up question is what approximations are possible with less knowledge of the prior. In the following, we consider simple mechanisms that do not use any knowledge of the prior, other than that all prices lie in $[0, 1]$. We will show that while additive approximations are possible, essentially no multiplicative approximation of the optimal revenue can be guaranteed without knowledge of the prior.

Let D_x be the diagonal $b = s + x$, and let M_y be the mechanism that buys and sells iff the valuations are to the right of D_y, i.e. $b \geq s + y$. It is easy to see that for $x > y$, and any point on D_x, the revenue obtained by M_y is $2y - x$. (For point $(s, s + x)$, M_y buys from the seller for $(s + x) - y$ and sells to the buyer for $y + s$.)

Theorem 10 (Prior-independent mechanism for the 1×1 case)**.** *There is a prior-independent deterministic mechanism that extracts revenue at least* $\mathrm{OPT} - \frac{2}{3}n$ *in the single seller, single buyer setting with valuations on the unit interval (with valuations taking values* $\frac{i}{n}$*). Furthermore, this is optimal among prior-independent mechanisms in this setting.*

Proof. Consider the mechanism $M_{2n/3}$ that buys and sells iff $b \geq s + \frac{2n}{3}$. This achieves expected revenue of at least $\mathrm{OPT} - \frac{2n}{3}$ for any point (b, s): Clearly for any b and s, the optimal revenue will be at most n. On the other hand, if $s > b - \frac{2n}{3}$, then the revenue is bounded by $\frac{2n}{3}$. $M_{2n/3}$ will extract revenue at least $\frac{n}{3}$ for $s \leq b - \frac{2n}{3}$ and 0 otherwise.

Moreover, this is optimal among prior-independent mechanisms. Consider the points $(b, s) = (\frac{2n}{3}, 0)$, $(n, \frac{n}{3})$ and $(n, 0)$. If a mechanism does not buy and sell at $(\frac{2n}{3}, 0)$, then for the instance that has all the probability concentrated at this point, it generates revenue 0, while the optimal revenue would clearly be $\frac{2n}{3}$. Similarly for $(n, \frac{n}{3})$. Thus, we have to buy and sell at both those points if we want to generate profit at least $\mathrm{OPT} - \frac{2n}{3}$ in the worst case. On the other hand, if we buy and sell at both those points, it follows that by truthfulness we cannot charge more than $\frac{2n}{3}$ to the buyer nor offer less than $\frac{n}{3}$ to the seller at $(n, 0)$ either. Therefore we generate profit $\frac{n}{3}$ at $(n, 0)$, as opposed to a maximum revenue of n. Thus no deterministic prior-independent mechanisms can generate revenue greater than $\mathrm{OPT} - \frac{2n}{3}$ for all instances. $\square$

By allowing randomization we can do even better.

Theorem 11 (Prior-independent randomized mechanism for the 1×1 case)**.** *There is a prior-independent mechanism that extracts revenue at least* $\mathrm{OPT} - \frac{n}{2}$ *in the single seller, single buyer setting with valuations on the unit interval.*

Proof. Take the mechanism M that randomizes uniformly amongst $\{M_y : \frac{n}{2} \leq y \leq n\}$. Then for any point on a diagonal $D_{\frac{n}{2}+x}$, the expected revenue in M is given by $\int_0^x \frac{2}{n} \cdot [2(\frac{n}{2} + t) - (\frac{n}{2} + x)]dt = x$. But the optimal revenue for any point on $D_{\frac{n}{2}+x}$ is $\frac{n}{2} + x$. Clearly, for any diagonal to the left of $D_{\frac{n}{2}}$, the revenue is bounded by $\frac{n}{2}$. Thus, M gives expected revenue of at least $\mathrm{OPT} - \frac{n}{2}$.

□

On the other hand, for multiplicative guarantees knowledge of the prior is essential, as it is easy to see that we can give at best a $\frac{1}{n}$ worst-case guarantee, dependent on the granularity of the prior. We give a (weak) guarantee that depends on the granularity of the prior support, and show that this is the best we can do. It immediately follows that for arbitrary continuous distributions we could not give any guarantee.

Theorem 12 (Prior-independent randomized mechanism for the 1×1 case, multiplicative)**.** *There is a prior-independent mechanism that extracts revenue at least* $\frac{1}{n} \cdot \mathrm{OPT}$ *in the single seller, single buyer setting with valuations on* $\{i : 0 \leq i \leq n\}$*. Furthermore, this is optimal among prior-independent mechanisms in this setting.*

Proof. As before let D_i be the diagonal $s = b+i$, and let M_j be the mechanism the buys iff the valuations are to the right of D_j, i.e. $s \geq b + j$. It is easy to see that for $i > j$, and any point on D_j, the revenue obtained by M_i is $2j - i$.

Take the mechanism M that randomizes uniformly amongst $\{M_j : 1 \leq j \leq n\}$. Then for any point on a diagonal D_i it is easy to see that the expected revenue arising from the M_1,...,M_{i-1} equals 0. With probability $\frac{1}{n}$ the mechanism selects M_i, giving optimal revenue for any point on this diagonal. In expectation the mechanism thus generates revenue $\frac{1}{n} \cdot \mathrm{OPT}$.

To see that this is tight, suppose we want to extract a c-fraction of the optimal revenue. In order to generate a c-fraction of the optimal revenue for any point on D_1, we need to place at least c probability mass there. It is easy to see that this gives net revenue 0 for any point on D_2, so we need to place an additional c probability, for a total allocation probability of $2c$. By induction it follows that we need to buy and sell with probability nc at $(n, 0)$, and thus $c \leq \frac{1}{n}$.

□

2.9 Discussion and Further Work

One question raised by our results is that of the relation between single seller, single buyer market intermediation and two-bidder auctions. As discussed in section 2.4 the graph algorithm of Papadimitriou and Pierrakos [106] can be used to solve the no short-selling 1×1 market intermediation case, and the derandomization in Dobzinski et al. [50] applies immediately to both this and the 1×1 balanced inventory setting. These give running times of $\mathcal{O}(n^6)$ and $\mathcal{O}(n^7)$ in contrast to a running time of $\mathcal{O}(n^3)$ for our approach in the market intermediation setting. In the next chapter we show that with some additional work, this approach can also be made to work for 2-bidder auctions, although a constant-factor gap in complexity remains. What's more, we will show that the insights into the structure of the market intermediation setting we have gained here have direct implications for reverse or procurement auctions too. This is not entirely surprising if one thinks of market intermediation as a combination of auction and reverse auction into one setting; but it also shows that this line of inquiry has important theoretical implications, even in addition to direct applications.

An immediate question that follows from our hardness results is whether we can give good approximations in polynomial time. As we have shown in section 2.8, essentially nothing is possible without regard for the prior at all; but this does not preclude a middle ground that does take the prior into account, but does not give the exact optimal revenue. The approach in Dobzinski et al. [50] for n bidder auctions hinges on discarding all but the highest two bids. In a $k \times k$ market intermediation setting this can obviously guarantee at best a $\frac{1}{k}$ fraction of the optimal revenue in the worst case, as it is easy to construct instances in which the optimal mechanism conducts k transactions. For welfare maximization Dütting et al. [52] rank and pair the buyers and sellers; however to ensure truthfulness they need to charge the Vickrey price, giving no guarantee for revenue.

Part 3

Correlated Auctions and Reverse Auctions

3.1 Introduction

While market intermediation in itself is thoroughly motivated by its real-world applications already, this setting also stands out due to its close connection to more classical auction theory. In this chapter, we will explore this connection further, and present several novel results on auctions that follow directly from our investigation of intermediation.

Within mechanism design, auctions are a major field of interest [74, 72, 81]. In this, we consider a single auctioneer who wants to sell a single item to one of several bidders, each of whom has a private valuation for the item which the auctioneer does not know. The challenge is to allocate the item according to some measure of optimality based on the private valuations. In addition to social welfare optimisation, in which we aim to allocate the item to the bidder who values it most, revenue maximisation (where we aim to maximise the auctioneer's expected profit) has received major attention. Myerson's seminal result [98] showed that with independent priors, (revenue-) optimal single-item auctions have a closed-form solution: In the deterministic case, in which we are solely interested here, the item is sold to the bidder with the highest "virtual valuation" and their payment is their critical bid. This also gives the more general theory: For a specific kind of "truthful" allocation functions together with uniquely determined payments, bidders are incentivised to reveal their true valuations to the auctioneer. We may therefore regard the problem as one of finding allocation functions that satisfy this truthfulness constraint. For correlated priors, in contrast to the aforementioned independent-priors setting, this is an intricate computational problem. The case with three or more bidders has been shown to be intractable by Papadimitriou and Pierrakos [106]; but on the other hand both [106] as well as Dobzinski et al. [50] show that the optimal auction for *two* bidders can be computed in polynomial time; both approaches reduce the problem to generic known-polynomial problems.

In addition to selling an item, auctions may also be used by the auctioneer to buy an item or service from one of multiple sellers. These "reverse" or "procurement" auctions are widely used for instance to solicit bids for public projects. Many results from auctions carry over directly to the reverse auction case. For instance, the VCG mechanism for optimising social welfare works in a reverse auction, as do many other auction formats. So much do these cases appear to be mirror images of one another, that simple reverse counterparts of single-item auctions are rarely discussed explicitly in the literature; most of the published results on reverse auctions investigate more complex scenarios such as differing quality or service levels from different sellers, as in e.g. [91]. To our knowledge, a significant distinction between an auction and its direct reverse counterpart has not been discussed in the literature before.

Our main interest is in exploring the structural properties of correlated auctions further. For the two-bidder auction, this allows us to construct a $O(n^3)$ algorithm, which improves on $O(n^6)$ of previous approaches. Ours is the first algorithm to exploit directly structural properties of the problem. This algorithm is a direct extension of the dynamic programming approach of Section 2.3. For reverse auctions, we show that these behave differently than auctions, for any number of bidders; this raises interesting questions about their complexity.

Previous Work

Several variations of optimal correlated auctions have been investigated. Most relevant to our discussion is again the literature on the complexity of optimal correlated auctions in which the joint prior is given explicitly or as an oracle. Papadimitriou and Pierrakos [106] show that for two bidders, a (revenue-) optimal auction can be found in polynomial time. Their algorithm reduces the problem to finding a maximum-weight independent set on a bipartite graph, with edges encoding allocation constraints of the auction. This yields an algorithm that runs in time $O(n^6)$ for prior support size n^2 (each bidder's valuation taking one of n discrete values). For three or more bidders they show that it is NP-hard to approximate the optimal auction to within a factor of 1.0005. Dobzinski et al. [50] independently also give a polynomial algorithm for the two-bidder auction. They show that a truthful-in-expectation mechanism found via an LP can be derandomised. The runtime of this approach depends on the LP-solver chosen; standard interior point methods give $O(n^7)$. Furthermore they investigate k-lookahead auctions, in which an optimal auction is run on the highest k bidders' conditional distribution. They show that a polynomial-time algorithm for two bidders extends to a polynomial-time approximation algorithm for many bidders through the 2-lookahead auction. This builds on previous work by Ronen [110] and Ronen and Saberi [111]. Chen et al. [30] investigate the approximation ratio of the k-lookahead auction further. Caragiannis et al. [26] improve results by [50] on separation between deterministic vs. randomised expected revenue, and on the lower bound on the approximation ratio for the three-bidder auction by [106]. Diakonikolas et al. [49] show that an approximate trade-off between revenue and social welfare can be computed efficiently for two bidders. Esö [54] investigate optimal auctions for risk-averse buyers and sellers. A related setting with interdependent values has been investigated by several authors, see for instance [38, 84, 94]. Crémer and McLean [41, 42] discuss conditions for full surplus extraction with interim individual rationality.

Most of the literature on reverse or procurement auctions specifically seem to focus on more complex settings than the ones we are interested in. One major area of research is when sellers offer goods of differing qualities, see for instance Manelli and Vincent [91]. Chapter 13.5 by Hartline and Karlin [69] in Nisan et al. [105] discuss feasibility constraints in reverse auctions.

Several chapters in the same book briefly mention that they consider reverse auctions to be covered by the model they use or similar, e.g. pages 220, 269, 332 therein [105]. To the best of our knowledge, almost no literature looks specifically at the simple reverse auction setting we are interested in. The main exception to this we are aware of is a paper by Minooei and Swamy [95, 96], who discuss the more general setting of mechanism design for covering (as opposed to packing) problems. Conitzer and Sandholm [37] discuss collusion in combinatorial auctions and reverse auctions. They take the reverse setting to be a simple parallel of the forward case (as we do here), except for an explicit constraint on allowed allocations (in their case, for the VCG mechanism) which we also assume in this chapter. They consider among other results the complexity of computing whether collusion is possible in a (forward or reverse) auction, showing that this is NP-hard even for 2 colluders.

3.2 Preliminaries

We consider a single-item auction, in which an auctioneer wishes to sell one item to one of several bidders, numbered 1, …, k. We assume each bidder i has valuation v_i, which can take one of several discrete values. For ease of notation we take $v_i \in \{1, \ldots, n\} = [n]$; it is easily checked that none of our results depend on this. Let F denote the (joint) prior probability distribution over $\mathbf{v} = (v_1, \ldots, v_k)$. Our interest is only in deterministic mechanisms, which consist of allocation functions $x_i(\mathbf{v})$ together with payment functions $p_i(\mathbf{v})$ for each bidder. Let $x_i(\mathbf{v}) = 1$ if bidder i wins the item given bid vector $\mathbf{v}$, and $x_i(\mathbf{v}) = 0$ otherwise. Given that we assume the auctioneer only has a single copy of the item to sell, we require $\sum_i x_i(\mathbf{v}) \leq 1$ for all $\mathbf{v}$. We assume quasilinear utilities, and require the usual notions of truthfulness / DSIC and individual rationality, as defined formally below. We therefore can assume that players' bids are equal to their valuations. The auctioneer's aim will be to maximise their expected revenue $\mathbb{E}\left[p_i(\mathbf{v})\right]$.

$$\text{(Utilities)} \qquad u_i(\mathbf{v}) = v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \tag{3.1a}$$

$$\text{(DSIC)} \qquad v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i x_i(v_i', \mathbf{v}_{-i}) - p_i(v_i', \mathbf{v}_{-i}) \quad \forall i, \mathbf{v}, v_i' \tag{3.1b}$$

$$\text{(IR)} \qquad u_i(\mathbf{v}) \geq 0 \quad \forall i, \mathbf{v} \tag{3.1c}$$

$$\text{(1-item)} \qquad \sum_i x_i(\mathbf{v}) \leq 1 \quad \forall \mathbf{v} \tag{3.1d}$$

By Myerson [98], truthfulness in this domain for deterministic mechanisms is equivalent to monotone allocations, and the corresponding uniquely determined payments - the winner's critical bid. That is, if bidder i wins the auction given bid profile $\mathbf{v}$, then they also win the auction for bid profile $(v_i', \mathbf{v}_{-i})$, for any $v_i' > v_i$; and their payment will be the smallest $v_i' \leq v_i$ such that they would still win the auction given bid profile $(v_i', \mathbf{v}_{-i})$. If bidder i does not win they pay nothing (by IR (3.1c)).

$$x_i(\mathbf{v}) = 1 \Rightarrow \forall v_i' \geq v_i : x_i(v_i', \mathbf{v}_{-i}) = 1 \tag{3.2a}$$

$$p_i(\mathbf{v}) = \min\left\{v_i' : x_i(v_i', \mathbf{v}_{-i}) = 1\right\} \text{ if } x_i(\mathbf{v}) = 1, \text{ else } p_i(\mathbf{v}) = 0 \tag{3.2b}$$

Papadimitriou and Pierrakos [106] give a very elegant geometric representation of this condition: For each bidder i, their critical bid is given by a function $\alpha_i(\mathbf{v}_{-i})$ of the other bidders' bids; where $x_i(\mathbf{v}) = 1$ iff $v_i \geq \alpha_i(\mathbf{v}_{-i})$. Consider now for each bidder i the region $A_i = \{\mathbf{v} : v_i \geq \alpha_i(\mathbf{v}_{-i})\}$ of all bid vectors for which i wins the item. Clearly this is bordered by $\alpha_i(\mathbf{v}_{-i})$. Furthermore, if $(v_i, \mathbf{v}_{-i}) \in A_i$, then also $(v'_i, \mathbf{v}_{-i}) \in A_i$ for all $v'_i \geq v_i$. This follows both from the definition of A_i as the region bounded below by the graph of a function of v_{-i}, as well as directly from monotonicity. We will also say that A_i is "upward-closed in direction v_i" for this. The 1-item constraint (3.1d) entails that any two A_i must be disjoint. In summary, the picture we get is that looking for the optimal k-bidder auction is looking for a partition of the space of possible bid combinations into $k+1$ regions: k regions where the item is sold to each of the buyers (each upward-closed in the corresponding direction), and one where the item is not sold. Figure 3.1 shows this picture for the two-bidder case. Taking bidder 1's bid to be on the x-axis and bidder 2's on the y-axis, we are looking for A_1 to be rightward-closed, and A_2 to be upward-closed. There is a two-fold tradeoff: smaller $\alpha_i(u)$ means higher probability of drawing $v_i \geq \alpha_i(u)$, but selling at a lower price if so; smaller $\alpha_i(u)$ also means "blocking" more bid vectors for the other bidder. We will often identify a mechanism through either the regions A_i or the functions α_i; we will write $\alpha_i(\mathbf{v}_{-i}) = n+1$, if none of the bid vectors $(v_i, \mathbf{v}_{-i})$ are in A_i. When defining a mechanism through the α_i, DSIC and IR are automatic. The 1-item constraint (3.1d) for two bidders can be restated as a non-crossing property (3.3) [106]. The expected revenue has a simple closed form in terms of α_i given as equation (3.4) for two bidders.

$$\text{(Non-Crossing Property)} \qquad v_1 \geq \alpha_1(v_2) \Rightarrow v_2 < \alpha_2(v_1) \qquad \forall v_1, v_2 \qquad (3.3)$$

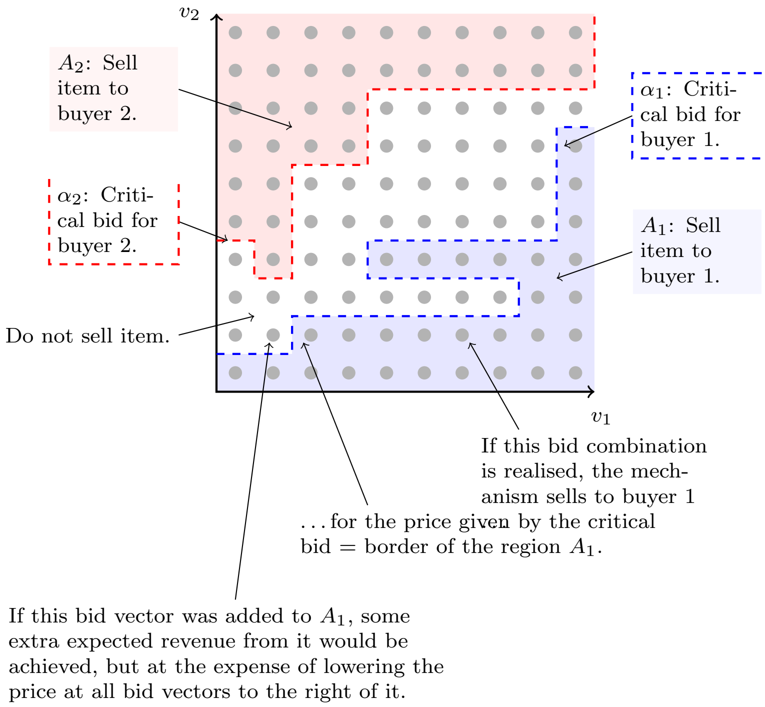

Figure 3.1: A mechanism as a partition of the bid space into regions of allocation, and the corresponding critical bid functions. Discrete prior support shown dotted, with critical bid functions and allocation regions drawn slightly larger for easier readability. (We are being slightly imprecise here: The graph of α_1 is the vertical part of the dashed blue line; The graph of α_2 is the horizontal part of the dashed red line.)

$$R = \sum_{j=1}^{n} \left[\alpha_1(j) \cdot \sum_{\ell=\alpha_1(j)}^{n} F(\ell, j) \right] + \sum_{\ell=1}^{n} \left[\alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j) \right] \qquad (3.4)$$

In a reverse auction, again a single auctioneer faces k bidders having their valuations drawn from a joint distribution F supported on $[n]^k$. We assume that each bidder holds one copy of a single type of item, each bidder's copy identical to all others', and that the auctioneer wishes to procure one copy. For simplicity we now write $x_i(\mathbf{v}) = -1$ if the mechanism buys a copy of the item from bidder i. This allows us to leave the definitions of utilities, DSIC and IR in equations (3.1a)-(3.1c) unchanged.[1] It is easy to see that in this context it makes little sense to require that the mechanism buys *at most* one copy of the item. Instead we require the mechanism to always buy *at least* one copy of the item, i.e. we require that $\sum_i x_i(\mathbf{v}) \leq -1$, replacing the corresponding constraint (3.1d).

Geometrically, we get a very similar picture of regions A_i in which the mechanism buys from bidder i. However, they now need to be downward-closed in direction of v_i. In the two-bidder case, A_1 ought to be leftward-closed, and A_2 downward-closed. Secondly, two or more of the A_i may now overlap (when the mechanism buys two or more copies); the constraint that $\sum_i x_i(\mathbf{v}) \leq -1$ means that the union of all A_i must cover all of the bid space.

[1] Taking instead v_i to be nonpositive, or adjusting equations (3.1a)-(3.1c) is equivalent.

3.3 The $O(n^3)$ Algorithm for the Two-Bidder Auction

We now show how to compute the optimal two-bidder auction in time $O(n^3)$. This immediately gives an $O(n^3)$ $\frac{5}{3}$-approximation algorithm for many bidders, as detailed in Dobzinski et al. [50]. We do this in three steps. First, we show that we only need to find the optimal allocation function for one of the bidders; finding the second bidder's allocation function is then very easy. Indeed, pre-computing all possible optimal allocations for the second bidder is easy. Second, we show that the optimal critical bid for bidder 1 in each $v_2 = c$ "row" only depends in a limited way on the critical bids in other rows. Third, we use these results to construct a very simple bottom-up dynamic programming algorithm.

Step 1: Disentangling the Two Bidders from One Another

One source of complexity in finding the optimal auction with correlated priors is that the two bidders' allocations interact: taking still v_1 to be the horizontal axis and v_2 the vertical one; if at a bid vector $\mathbf{w}$ we allocate the item to bidder 1 (i.e. $w \in A_1$), then by monotonicity we also do so at all bid vectors to its right. In turn, this means that we cannot sell to bidder 2 at any of the bid vectors $\mathbf{u} : u_1 \geq w_1 \wedge u_2 \leq w_2$ to the bottom right of $\mathbf{w}$. Vice versa, if at $\mathbf{w}$ we sell to bidder 2, we cannot sell to bidder 1 at any of the points to its top left. This argument applies repeatedly: Allocation to bidder 1 influences potential (and thus also optimal) allocation to bidder 2, which in turn influences potential & optimal allocation to bidder 1. A simple lemma shows how to disentangle the two bidders' allocation: Suppose (say) bidder 1's allocation is fixed. How does the choice of a value $\alpha_2(\ell)$ now influence the optimal choices of all other values of α_2? Simple: It does not. While the inclusion or exclusion of (ℓ, v_2) in A_2 influences other points (ℓ, v_2') in the same "column" through monotonicity for bidder 2; the only way it could influence a point $(\ell', v_2'), \ell' \neq \ell$, in another "column" is through monotonicity for bidder 1. But by assumption bidder 1's allocation is already fixed. Therefore, each choice of $\alpha_2(\ell)$ is independent of all the others. Furthermore, clearly in column $v_1 = \ell$ the optimal thing to do is to run an optimal single-bidder auction for bidder 2 on all the bid vectors not inside or below A_1. This principle holds for many bidders; We state it rigorously for two bidders:

Lemma 2. *In the optimal mechanism the following holds.*

$$\alpha_2(\ell) = \operatorname*{argmax}_{\alpha_2(\ell) > u(\ell)} \alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j)$$

$$u(\ell) = \max \{u : \alpha_1(u) \leq \ell\}$$

Proof. Let α_1 be fixed. Our aim is to find the α_2 so as to maximise the expected revenue (3.4) while maintaining the 1-item constraint (3.1d), which by Papadimitriou and Pierrakos [106] is equivalent to a non-crossing property of critical bid functions (3.3). Now, clearly the first sum in equation (3.4) is constant for fixed α_1. So, we are looking to solve

$$R_2 = \max_{\alpha_2} \sum_{\ell=1}^{n} \left[\alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j) \right] \tag{3.5}$$

Furthermore, from (3.3) it follows that we require

$$\alpha_2(\ell) > \max \{u : \alpha_1(u) \leq \ell\} =: u(\ell) \tag{3.6}$$

Since this is the only constraint on α_2, it follows that we may interchange the maximum and sum:

$$R_2 = \sum_{\ell=1}^{n} \left[\max_{\alpha_2(\ell) > u(\ell)} \alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j) \right] \tag{3.7}$$

And therefore, $\alpha_2(\ell)$ is as claimed.

$$\alpha_2(\ell) = \operatorname*{argmax}_{\alpha_2(\ell) > u(\ell)} \alpha_2(\ell) \cdot \sum_{j=\alpha_2(\ell)}^{n} F(\ell, j) \tag{3.8}$$

$\square$

Figure 3.2 illustrates this lemma. This result tells us that if we already knew one player's allocation & critical bid function, it would be easy to calculate the optimal allocation & critical bids for the second player. In fact, we can calculate all possible ones: Notice that the optimal $\alpha_2(\ell)$ depends only on $u(\ell)$, and no other information on α_1 or α_2. We can iterate through all (n^2) possible values of ℓ and u, and calculate (in linear time each) the optimal $a_2(\ell, u) = \operatorname{argmax}_{m > u} m \cdot \sum_{j=m}^{n} F(\ell, j)$, taking time $O(n^3)$ total.

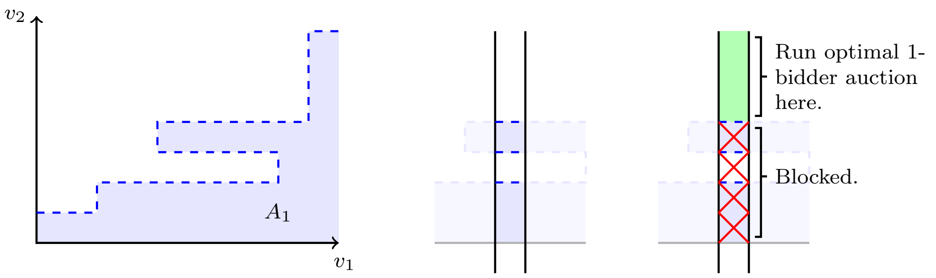

Figure 3.2: Finding the optimal allocation to bidder 2, given fixed allocation to bidder 1.
From left to right: (a) For a given fixed allocation to bidder 1 ...
(b) ... look at each $v_1 = c$-"column" separately ...
(c) ... and run the optimal single-bidder auction for bidder 2, on the part that is not blocked by allocation to buyer 1.

Step 2: Disentangling the Remaining Bidder's Allocation

Relying on the previous subsection, we can now extend an argument and algorithm that was first discussed by Gerstgrasser et al. [60]. Let us consider how the optimal choice for each $\alpha_1(j)$ depends on the value of α_1 in other rows. Let us assume that for some j the values of $\alpha_1(j+1)$, …, $\alpha_1(n)$ - i.e. the allocation to bidder 1 in rows above j - are fixed. A particular choice of $\alpha_1(j)$ contributes to the expected revenue of the mechanism in three ways:

- The contribution to expected revenue from the optimal $\alpha_1(1)$, …, $\alpha_1(j-1)$, which may depend on $\alpha_1(j)$.
- In *row* j, the expected revenue from selling at points $(\alpha_1(j), j)$, …, (n, j) at price $\alpha_1(j)$.
- For some *columns*, the choice of $\alpha_1(j)$ may entail that the mechanism may only sell to buyer 2 at points "above" row j. In particular, consider column $v_1 = \ell$. If $\alpha_1(j) \leq \ell$, then point $(\ell, j) \in A_1$, and thus (ℓ, j) and

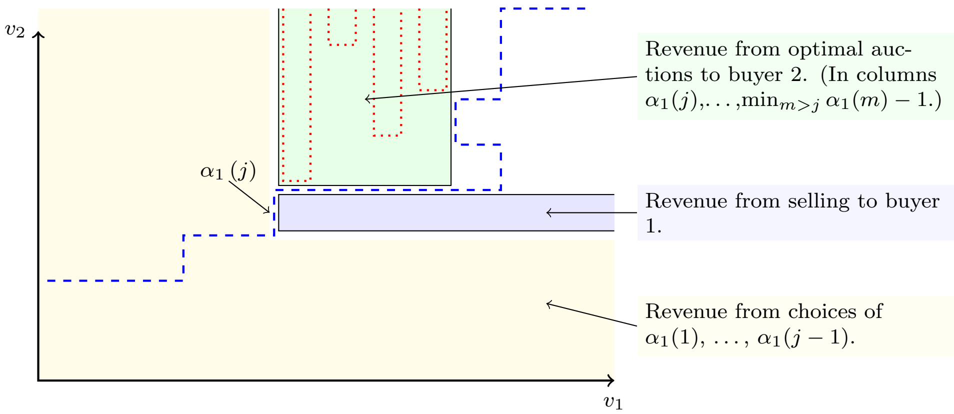

Figure 3.3: The contribution to the auctioneer's expected revenue due to the choice of $\alpha(j)$ as in Lemma 3. The blue area shows the bid profiles where we sell to buyer 1 due to this choice, the green area shows the bid profiles where we may sell to buyer 2 due to it (using the respective optimal single-bidder auctions, exemplified in red dotted lines).

all points directly below cannot allocate to buyer 2. On the other hand, this is not necessarily influenced by the choice of $\alpha_1(j)$: If for some $m > j$, (ℓ, m) is in A_1, i.e. $\alpha_1(m) \leq \ell$ (which is assumed to be fixed), then (ℓ, m),..., (ℓ, j),..., $(\ell, 1)$ could not allocate to bidder 2 irrespective of the choice of $\alpha_1(j)$. This means that this contribution to expected revenue occurs exactly for columns ℓ with $\alpha_1(j) \leq \ell < \min_{m>j} \alpha_1(m)$.

Figure 3.3 illustrates these three contributions. The crucial point here is in the last item, which we restate as a lemma due to its importance:

Lemma 3. *The optimal choice of $\alpha_1(1)$, ..., $\alpha_1(j)$ depends only on the* minimum *of $\alpha_1(j+1)$, ..., $\alpha_1(n)$, not all individual values.*

Step 3: The Dynamic Programming Algorithm

Using these two results, calculating the optimal auction is easy. First, let $r_2(\ell, u) = \max_{m>u}(m \cdot \sum_{j=m}^{n} F(\ell, j))$ be the expected revenue that can be obtained from bidder 2 on bid vectors $(\ell, u+1)$ and those directly above (cf. Lemma 2 and after). As discussed we can compute these in time $O(n^3)$, which our algorithm does as a first step. Now, by the discussion in the previous subsection, and Lemma 3, we only need to consider the minimum of α_1 in the top $n-j$ rows in order to calculate the optimal α_1 in the first j rows. We will write $\alpha_1(j) = n+1$ if none of the points in row j are in A_1. Our algorithm is as follows: for each possible value of $\min\{\alpha_1(2), \ldots, \alpha_1(n)\}$ we calculate the optimal value of $\alpha_1(1)$ and the associated expected revenue arising from this choice. We save these values as $a_1(1,m)$ and $r_1(1,m)$ for $m = 1, \ldots, n+1$. Then, we proceed upwards and for $j = 2, \ldots, n$ calculate the optimal values of $\alpha_1(j)$ for each possible value of $m = \min\{\alpha_1(j+1), \ldots, \alpha_1(n)\}$ using the already stored values of $r_1(j-1, .)$. In recursion form:

$$a_1(j,m) = \operatorname*{argmax}_{1 \le q \le n+1} \left\{ q \sum_{\ell=q}^{n} F(\ell, j) + \sum_{\ell=q}^{m-1} r_2(\ell, j) + r_1(j-1, \min\{q, m\}) \right\} \tag{3.9}$$

and $r_1(j,m)$ is the maximum attained by this expression. The first term is the expected revenue from selling to bidder 1 in this row (blue in Figure 3.3); the second is the revenue from auctioning to bidder 2 in the appropriate columns (green); the third is the expected revenue from rows below j (yellow). In the third summand we use $\min\{x_j, \ldots, x_n\} = \min\{x_j, \min\{x_{j+1}, \ldots, x_n\}\}$. It is easy to see that we do not need to calculate the sums inside the argmax from scratch for each value of q we consider; By memoising the partial sums we can evaluate the term inside the argmax in constant time. That then makes it possible to calculate all the a_1 and r_1 in time $O(n^3)$. $r_1(n, n+1)$ will then give the optimal expected revenue, and α_1 can be found by backtracking from

$a_1(n, n+1)$. See Algorithm 3 for the non-memoised version. Furthermore, this extends also to a $\frac{5}{3}$ approximation for many bidders: Dobzinski et al. [50] give this bound for the 2-lookahead auction, which runs the optimal auction for the two highest bidders, with their priors conditioned on the remaining $k-2$ bidders' bids. Thus, any algorithm for the two-bidder auction also can be used to solve the 2-lookahead auction for many bidders, if given access to the conditional prior. The improvement our algorithm gives over existing approaches thus transfers immediately also to approximation algorithms for many bidders.

Algorithm 3 Optimal auction with two correlated bidders.

for $\ell = 1, \ldots, n$ **do**
  for $u = 0, \ldots, n$ **do**
    $r_2(\ell, u) = \max_{m>u}(m \cdot \sum_{r=m}^{n} F(\ell, r))$
for $m = 1, \ldots, n+1$ **do**
  $r_1(0, m) = \sum_{q=1}^{m-1} r_2(0, q)$
for $j = 1, \ldots, n$ **do**
  for $m = 1, \ldots, n+1$ **do**
    $a_1(j, m) = \operatorname{argmax}_{1 \leq q \leq n+1}$
    $\left\{ q \sum_{\ell=q}^{n} F(\ell, j) + \sum_{\ell=q}^{m-1} r_2(\ell, j) + r_1(j-1, \min\{q, m\}) \right\}$
    $r_1(j, m) = \max_{1 \leq q \leq n+1}$
    $\left\{ q \sum_{\ell=q}^{n} F(\ell, j) + \sum_{\ell=q}^{m-1} r_2(\ell, j) + r_1(j-1, \min\{q, m\}) \right\}$
  return $r_1(n, n+1)$

Theorem 13. *It is possible to calculate the optimal two-bidder auction with correlated priors in time $O(n^3)$, where the prior support is of size n^2. This also gives a $O(n^3)$ $\frac{5}{3}$-approximation algorithm for many bidders.*

Proof. The memoised $O(n^3)$ algorithm follows easily from the conceptual prototype presented in Algorithm 3, by incrementally computing the sums inside the max and argmax. To verify that the algorithm returns the correct value, consider that we can write the expected revenue as follows:

$$R = \sum_{j=1}^{n}\left[\alpha_1(j)\cdot\sum_{\ell=\alpha_1(j)}^{n}F(\ell,j)\right]+\sum_{\ell=1}^{n}\left[\alpha_2(\ell)\cdot\sum_{j=\alpha_2(\ell)}^{n}F(\ell,j)\right] \tag{3.10}$$

Here the first term gives the expected revenue from selling to bidder 1, the second term gives the expected revenue from selling to bidder 2: for $v_2 = j$ we sell to bidder 1 at price $\alpha_1(j)$ for bid combinations $\mathbf{v} = (\alpha_1(j), j), \ldots, (n, j)$, giving probability $\sum_{\ell=\alpha_1(j)}^{n} F(\ell, j)$.

Now, as a first step we partition the range of the outer sum in the second term along those indices for which $\ell = \alpha_1(j)$ for an $\alpha_1(j)$ with $\alpha_1(j) < \min\{\alpha_1(j+1), \ldots, \alpha_1(n)\}$.

$$R = \sum_{j=0}^{n}\left[\alpha_1(j)\sum_{\ell=\alpha_1(j)}^{n}F(\ell,j)+\sum_{\ell=\alpha_1(j)}^{\min\{\alpha_1(j+1),\ldots\}-1}\alpha_2(\ell)\cdot\sum_{j=\alpha_2(\ell)}^{n}F(\ell,j)\right] \tag{3.11}$$

Note that $\sum_{\ell=\alpha_1(j)}^{\min\{\alpha_1(j+1),\ldots,\alpha_1(n)\}-1}[\ldots]$ is indexing over the empty set and we take it to equal 0, if $\alpha_1(j) \geq \min\{\alpha_1(j+1), \ldots, \alpha_1(n)\}$. We also take $\alpha_1(0) = 1$ and $F(., 0) = 0$. The outer sum now iterates over all the "rows" $v_2 = 1, \ldots, n$. Let now $r_2(\ell, u)$ be defined as before; $r_2(\ell, u) = \max_{m>u}(m \cdot \sum_{r=m}^{n} F(\ell, r))$, and $a_2(\ell, u) = \text{argmax}_{m>u}(m \cdot \sum_{r=m}^{n} F(\ell, r))$. By the discussion in Step 1 of this section, in the optimal auction $\alpha_2(\ell) = a_2(\ell, u(\ell))$ where $u(\ell) = \max\{u : \ell \leq \alpha_1(u)\}$ is the topmost point in column j that is in A_1. And similarly $r_2(\ell, u(\ell))$ is precisely the contribution to expected revenue from bidder 2 in column ℓ. It is easy to see that each column ℓ is counted in the outer sum precisely in the summand $j = u(\ell)$. We can therefore rewrite the expected revenue in terms of r_2 instead of α_2.

$$R = \sum_{j=0}^{n} \left[\alpha_1(j) \sum_{\ell=\alpha_1(j)}^{n} F(\ell, j) + \sum_{\ell=\alpha_1(j)}^{\min\{\alpha_1(j+1),\ldots,\alpha_1(n)\}-1} r_2(\ell, j) \right] \tag{3.12}$$

Let now $r_1'(j, m, q)$ be the j-th summand of this, with two free parameters:

$$r_1'(j, m, q) = q \sum_{\ell=q}^{n} F(\ell, j) + \sum_{\ell=q}^{m-1} r_2(\ell, j) \tag{3.13}$$

Then taking $\min \emptyset = n + 1$, we can write the optimal revenue using r_1'.

$$R = \sum_{j=0}^{n} \left[r_1' \Big(j, \min\{\alpha_1(j+1), \ldots, \alpha_1(n)\}, \alpha_1(j) \Big) \right] \tag{3.14}$$

So far we have only reasoned about the revenue given fixed α_1 and α_2, which when introducing r_2 we assumed to be optimal. Clearly for the optimal mechanism it holds, by definition, that that is optimal over all possibilities for α_1:

$$R = \max_{\substack{1 \le a_1(j) \le n+1; \\ a_1(0)=0}} \sum_{j=0}^{n} \left[r_1' \Big(j, \min\{a_1(j+1), \ldots, a_1(n)\}, a_1(j) \Big) \right] \tag{3.15}$$

Since not all of the summands depend on all of the $a_1(j)$, and since they are all non-negative, we can interchange summation and the maximum operators.[2]

$$R = \max_{a_1(n)} \Bigg[r_1'\big(n, \min \emptyset, a_1(n)\big) + \max_{a_1(n-1)} \Big[r_1'\big(n-1, \min\{a_1(n)\}, a_1(n-1)\big) + \max_{a_1(n-2)} \Big[r_1'\big(n-2, \min\{a_1(n), a_1(n-1)\}, a_1(n-2)\big) + \cdots \Big]\Big]\Bigg] \tag{3.16}$$

Now we can write R as a recursion in r_1. The following coincides with the definition of r_1 in equation (3.9):

$$\begin{aligned} r_1(j,m) &= \max_q \Big[r_1'\big(j,m,q\big) + r_1\big(j-1, \min\{q,m\}\big) \Big] \\ &= \max_q \Big[q \sum_{\ell=q}^{n} F(\ell, j) + \sum_{\ell=q}^{m-1} r_2(\ell, j) + r_1\big(j-1, \min\{q,m\}\big) \Big] \end{aligned} \tag{3.17}$$

It is easy to check that (for the optimal value of m), $r_1(j,m)$ is exactly the sum of the first j summands in the expression for the optimal revenue (equations (3.14), (3.15)).

$$r_1\Big(j, \min\{\alpha_1(j+1), \ldots, \alpha_1(n)\}\Big) = \sum_{s=0}^{j} \Bigg[r_1'\Big(s, \min\{\alpha_1(s+1), \ldots, \alpha_1(n)\}, \alpha_1(s)\Big)\Bigg] \tag{3.18}$$

From this it follows immediately that $R = r_1(n, n+1)$, as desired. Algorithm

[2] More generally, it is easy to check that for $f_i(x_i, \ldots, x_n) \geq 0$ the following holds.

$$\max_{\mathbf{x}} \left[\sum_{i=1}^{n} f_i \right] = \max_{x_n} \Big[f_n(x_n) + \max_{x_{n-1}} \Big[f_{n-1}(x_{n-1}, x_n) + \Big[\cdots + \max_{x_1} f_1(\mathbf{x}) \cdots \Big]\Big]\Big]$$

3 calculates this value $R = r_1(n, n+1)$ by construction. The algorithm also keeps track of the associated α_1. This can be retraced as $\alpha_1(n) = a_1(n, n+1)$, and $\alpha_1(j) = a_1\big(j, \min\{\alpha_1(j+1), \ldots, \alpha_1(n)\}\big)$. We can calculate α_2 as $\alpha_2(\ell) = a_2(\ell, \max\{j : \alpha_1(j) \leq \ell\})$. $\square$

3.4 The Reverse Auction Case

We now turn to the reverse auction case. This is in many settings equivalent to the auction case; However with correlated priors unexpected things happen. Recall that here we are looking for regions A_i which are downward-closed in direction v_i, may overlap, and must cover all of the bid space. In an auction, a large part of the auctioneer's power comes from the option of not selling the item; Indeed, reserve prices below which the item is not sold are at the heart of Myerson's seminal result [98]. For a single bidder, not selling is even *all* the power the auctioneer has to achieve any revenue. In a reverse auction, the equivalent of this is to buy multiple copies of the item from multiple sellers. In a way, both of these cases are suboptimal locally, but allow for higher expected revenue globally: If for a given bid vector $\mathbf{v}$ the auctioneer does not sell in the auction this clearly foregoes some potential contribution to expected revenue arising from selling at $\mathbf{v}$; but, it may allow the auctioneer to generate a higher contribution to expected revenue (through higher prices) at some other bid vectors. Similarly in the reverse auction, buying from multiple bidders at a bid vector $\mathbf{v}$ clearly incurs a double or multiple contribution to expected cost arising from $\mathbf{v}$; but, it may allow the auctioneer to achieve a lower expected cost elsewhere in the bid space as a result.

It is easy to see that the possibility of buying from multiple bidders generates a much richer space of potential outcomes than in the auction. Whereas in the auction there is $k + 1$ possible allocations for each bid vector (selling to each of the bidders, plus selling to none of them), in the reverse auction we potentially have to deal with $2^k - 1$ possible allocation (buying from any combination of bidders, except from none of them). The question we deal with in this section is whether all of these are actually relevant to the problem of finding the optimal reverse auction; that is, will all of these occur in an optimal mechanism? The answer is surprising: "Yes", for $k \geq 3$ bidders, so the reverse auction in these cases is clearly structurally different than the corresponding auction; but "No" for 2 bidders. The latter is surprising in itself, as a priori both the 2-bidder auction as well as the 2-bidder reverse auction potentially have three valid allocations. As it turns out, not even these two cases are structurally the same.

Theorem 14. *In the single-item reverse auction with two correlated sellers, the optimal mechanism will never buy from both bidders.*

Proof. Suppose for bid vector $\mathbf{x}$ we buy from both sellers. We consider two cases. Firstly, suppose there exists an i such that for no point $\mathbf{x}' = (x'_i, x_{-i})$ with $x'_i > x_i$ we buy only from seller i. Then we could strictly improve our cost if we did not buy from i at $\mathbf{x}$ and all those bid vectors $\mathbf{x}' = (x'_i, x_{-i})$ with $x'_i > x_i$. Thus the mechanism was not optimal. See Figure 3.4 (a) for an illustration.

So assume that for both i, there exists a bid vector $\mathbf{x}^{(\mathbf{i})} = (x'_i, x_{-i})$, with

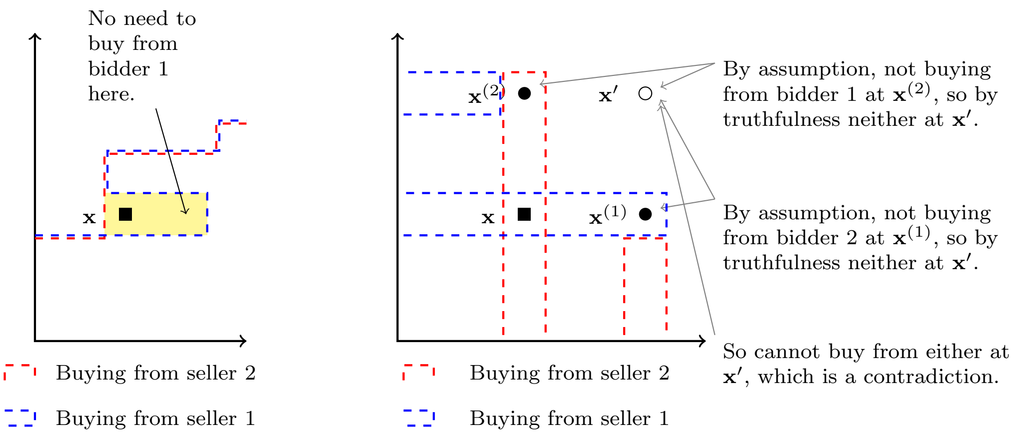

Figure 3.4: From left to right:
(a) In the first case of Theorem 14, if there is no point to the right of $\mathbf{x}$ in which we buy only from seller 1, we can improve our cost by not buying from seller 1 in the shaded region, i.e. moving the blue line so it coincides with the red one.
(b) Buying from either seller is blocked at $\mathbf{x}'$ due to truthfulness, in the second case of Theorem 14. By assumption we do not buy from seller 1 at $\mathbf{x}^{(2)}$, and thus by truthfulness cannot buy from seller 1 at $\mathbf{x}'$. Vice versa we assume we do not buy from seller 2 at $\mathbf{x}^{(1)}$, and so cannot buy from them at $\mathbf{x}'$ either. That leaves us with noone to buy the item from at $\mathbf{x}'$, violating our feasibility constraint.

$x_i' > x_i$, so that we buy only from seller i at $x^{(i)}$. But then by truthfulness it follows that at $\mathbf{x}' = (x_1', x_2')$ we cannot buy from either of the sellers. If we bought from seller 1 at $\mathbf{x}'$, we would also buy from seller 1 at $\mathbf{x}^{(2)}$ by truthfulness, but that contradicts our assumption. Vice versa for seller 2. But not buying at all at $\mathbf{x}'$ is not a valid mechanism by definition. Figure 3.4 (b) shows this case.

Thus, the optimal (valid) reverse auction can never buy from both bidders at once. $\square$

An immediate consequence of this result is that the optimal mechanism design problem in this setting is simpler than in the auction setting: We are now only looking for a partition of the bid space into *two* regions A_1 and $A_2 = A_1^c$. It is easy to check that this allows us to reduce the runtime of our two-bidder auction algorithm by a constant multiplicative factor, as we have fewer potential solutions to consider. We list this as Algorithm 4, again in non-memoised form for conciseness. To our knowledge this is the first algorithm specific to the reverse auction setting, exploiting structural arguments of this problem, and therefore also the first to show a lower runtime of this problem compared to the optimal correlated auction.

Algorithm 4 Optimal reverse auction with two correlated bidders.

$c_1(0, .) = 0$
for $j = 1, \ldots, n$ **do**
  for $m = 1, \ldots, n+1$ **do**
    $a_1(j, m) = \operatorname{argmin}_{1 \leq q \leq m}$
    $\left\{ q \sum_{\ell=1}^{q} F(\ell, j) + j \sum_{\ell=q}^{m-1} \sum_{s=1}^{j} F(\ell, s) + c_1(j-1, q) \right\}$
    $c_1(j, m) = \min_{1 \leq q \leq m}$
    $\left\{ q \sum_{\ell=1}^{q} F(\ell, j) + j \sum_{\ell=q}^{m-1} \sum_{s=1}^{j} F(\ell, s) + c_1(j-1, q) \right\}$
  return $c_1(n, n+1)$

Surprisingly, for $k \geq 3$ bidders the opposite holds: It is possible to construct instances in which it is optimal to buy all k copies of the item.

Theorem 15. *For three or more bidders, the optimal reverse auction may buy from all sellers.*

Proof. To show this, we will construct an instance. Our main gadget will be of the following form: Consider points $\mathbf{p}_1 = (c_{\mathrm{L}}, c_{\mathrm{M}}, c_{\mathrm{H}})$ and $\mathbf{p}_2 = (c_{\mathrm{M}}, c_{\mathrm{C}}, c_{\mathrm{H}})$ with high probability weight, and a third point $\mathbf{q} = (c_{\mathrm{M}}, c_{\mathrm{M}}, c_{\mathrm{H}})$ with very low probability weight, for some constants $c_{\mathrm{L}} << c_{\mathrm{M}} << c_{\mathrm{H}}$. We will want to make this so that the optimal mechanism will want to buy at the point $\mathbf{p}_1$ cheaply from seller 1 - and thus cannot buy at point $\mathbf{q}$ from seller 1, as by monotonicity that would also raise the purchase price at $\mathbf{p}_1$. Similarly for seller 2 and points $\mathbf{p}_2$ and $\mathbf{q}$. As a consequence, it will want to buy at $\mathbf{q}$ from

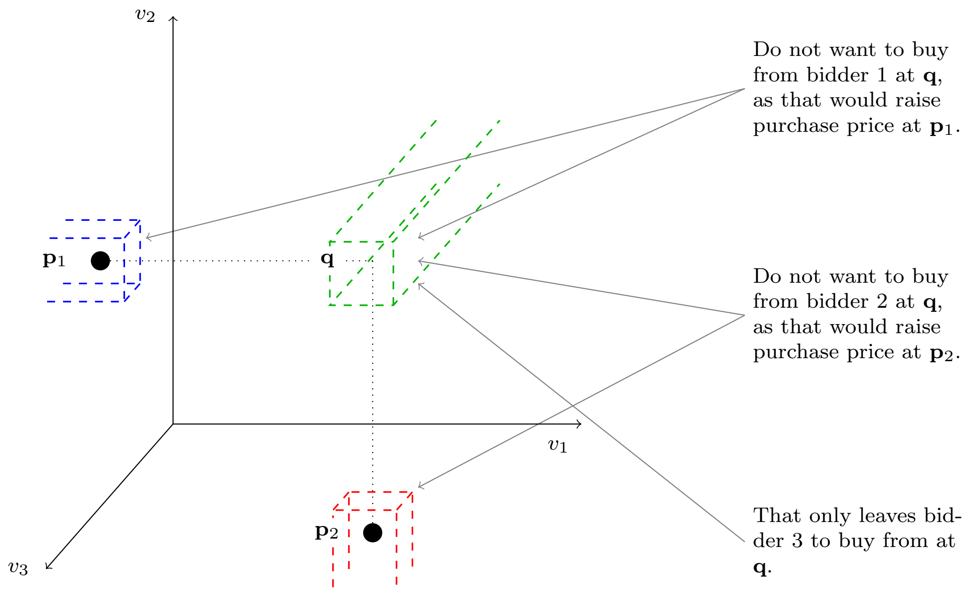

Figure 3.5: The gadget we will use in the proof of Theorem 15. High probability weight on points $\mathbf{p}_1$ and $\mathbf{p}_2$ makes it optimal to buy from seller 3 (in green) at point $\mathbf{q}$. Buying from either of the other sellers at $\mathbf{q}$ would raise the purchase price at either $\mathbf{p}_1$ or $\mathbf{p}_2$, thus raising the expected cost. By monotonicity, the mechanism then must also buy from seller 3 at all points behind $\mathbf{q}$ in this view.

seller 3. This will be at a very high purchase price, but if the probability weight on $\mathbf{q}$ is small enough, this will still be optimal in expectation. Figure 3.5 illustrates this construction. By monotonicity it follows that if the mechanism buys from seller 3 at $\mathbf{q} = (c_\mathrm{M}, c_\mathrm{M}, c_\mathrm{H})$, it must also buy from seller 3 at all points $(c_\mathrm{M}, c_\mathrm{M}, v_3), v_3 \leq c_\mathrm{H}$.

By creating three such gadgets in the right places and rotated appropriately, we can then make it optimal to buy from all three sellers at the intersection of these $\mathbf{q}$-segments. Consider the construction in Figure 3.6. In this we have one gadget consisting of $\mathbf{p}_{12} = (c_\mathrm{H}, c_\mathrm{L}, c_\mathrm{M})$, $\mathbf{p}_{13} = (c_\mathrm{H}, c_\mathrm{M}, c_\mathrm{L})$ and $\mathbf{q}_1 = (c_\mathrm{H}, c_\mathrm{M}, c_\mathrm{M})$ with the auctioneer buying from bidder 1 in the $\mathbf{q}_1$-segment; and similarly one gadget consisting of $\mathbf{p}_{21}$, $\mathbf{p}_{23}$ and $\mathbf{q}_2$ for bidder 2; and a third one comprising $\mathbf{p}_{31}$, $\mathbf{p}_{32}$ and $\mathbf{q}_3$ for bidder 3. Again let there be

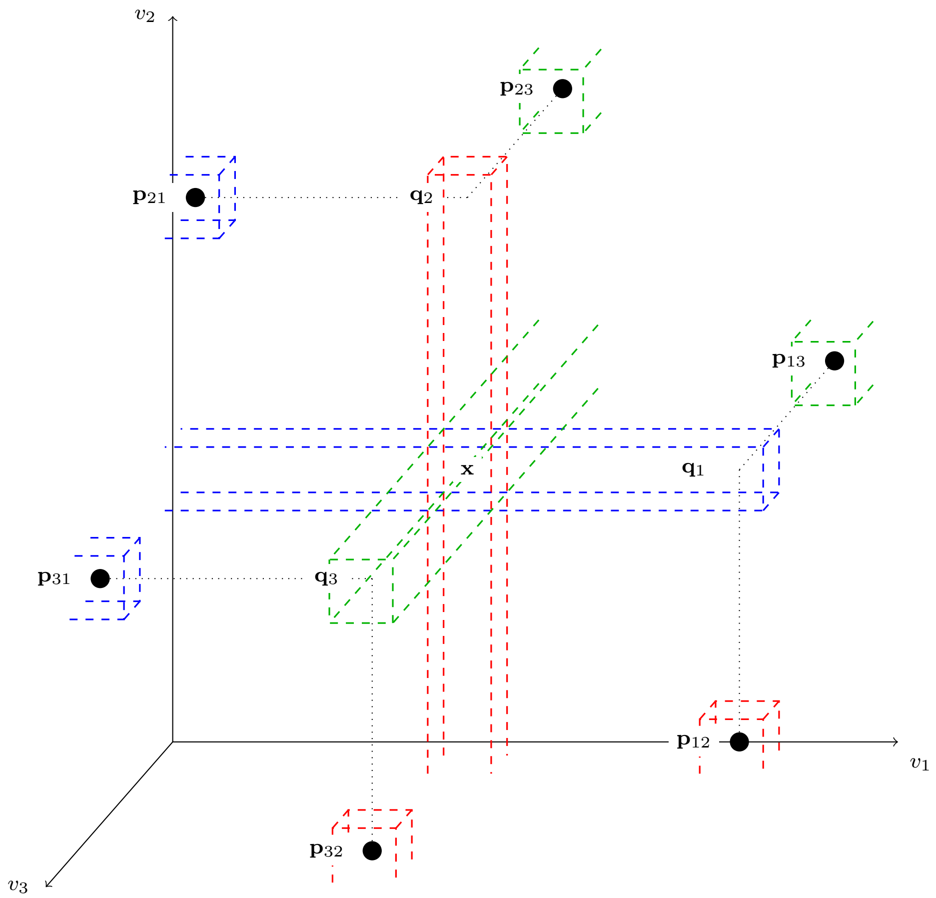

Figure 3.6: Three gadgets make up the construction used in the proof of Theorem 15. Notice that for bid vector $\mathbf{x}$ (in the centre at the intersection of the three $\mathbf{q}_i$-segments), the mechanism will buy from all three sellers, due to monotonicity and the allocation at the points $\mathbf{q}_i$.

very high probability weight on the $\mathbf{p}_{ij}$, and very small probability weight ϵ on the $\mathbf{q}_i$ (and everywhere else). The $\mathbf{q}_i$ are placed such that the $\mathbf{q}_i$-segments intersect at the point $\mathbf{x} = (c_{\mathrm{M}}, c_{\mathrm{M}}, c_{\mathrm{M}})$. It is easy to check that the optimal mechanism will indeed buy from bidder i in each $\mathbf{q}_i$-segment, for ϵ small enough; it will thus buy from all three bidders at $\mathbf{x}$.

To show this formally, wlog we take the prior support to be $[3]^3$, and $c_L = 1, c_M = 2, c_H = 3$; it is easy to see that the following arguments work for any other choice of these constants. Let there be probability weight $\frac{1-\epsilon}{6}$

on points $\mathbf{p}_{12} = (3, 1, 2)$, $\mathbf{p}_{13} = (3, 2, 1)$, $\mathbf{p}_{21} = (1, 3, 2)$, $\mathbf{p}_{23} = (2, 3, 1)$, $\mathbf{p}_{31} = (1, 2, 3)$, $\mathbf{p}_{32} = (2, 1, 3)$, and probability weight $\frac{\epsilon}{21}$ on each of the remaining 21 points of the prior support. We will denote by $\mathbf{q}_1 = (3, 2, 2)$, $\mathbf{q}_2 = (2, 3, 2)$, $\mathbf{q}_3 = (2, 2, 3)$ among these. Notice how for $i = 1, 2, 3$ each of these sets of two $\mathbf{p}_{ij}$ and one $\mathbf{q}_i$ forms one of the gadgets discussed in the main text. To show this we will proceed as follows. First, we show that the optimal mechanism has the property that the auctioneer buys each of the points $\mathbf{p}_{ij}$ for price 1 from seller j (and only seller j). There are two things to check here; step 1(a), we show that a valid allocation exists that has this property. Step 1(b), we show that any allocation with this property has lower expected cost than any allocation without this property. Step 2, we deduce from this that the optimal mechanism buys from all three sellers at point $\mathbf{x} = (2, 2, 2)$.

Step 1(a): There is a valid mechanism that buys from seller j (and only seller j) at each point $\mathbf{p}_{ij}$, for price 1: This is easy to see. We show one such mechanism in Figure 3.7. Each cell lists the bidder(s) that the item is bought from for this given bid vector; with the high probability bid vectors shown in bold face.

Step 1(b): Any mechanism that allocates at the points $\mathbf{p}_{ij}$ in this manner has lower expected cost than any mechanism that does not. We show this by giving first an upper bound on the expected cost of any mechanism with this property. This consists of an exact expression for the contribution to expected cost incurred at the $\mathbf{p}_{ij}$, plus an upper bound on the contribution at all the other points. Second, we give a lower bound of the expected cost of any mechanism which does not have this property; for this it suffices to

$v_3 = 1$	$v_1 = 1$	$v_1 = 2$	$v_1 = 3$
$v_2 = 3$	1	**3**	3
$v_2 = 2$	3	3	**3**
$v_2 = 1$	2	2	2

$v_3 = 2$	$v_1 = 1$	$v_1 = 2$	$v_1 = 3$
$v_2 = 3$	**1**	2	3
$v_2 = 2$	1	1,2,3	1
$v_2 = 1$	2	2	**2**

$v_3 = 3$	$v_1 = 1$	$v_1 = 2$	$v_1 = 3$
$v_2 = 3$	1	1	3
$v_2 = 2$	**1**	3	2
$v_2 = 1$	2	**2**	2

Figure 3.7: The full allocation for the instance in Theorem 15. Each cell shows the bidder(s) the mechanism buys from at the given bid vector. High probability points are shown in bold face.

lower bound the expected cost incurred at the $\mathbf{p}_{ij}$.

So, assume a mechanism allocates at the $\mathbf{p}_{ij}$ in the manner claimed; then the expected cost can be (very crudely) bounded above by $6 \cdot 1 \cdot 1 \cdot (\frac{1-\epsilon}{6}) + 21 \cdot 3 \cdot 3 \cdot (\frac{\epsilon}{21}) = 1 + 8\epsilon$. The first term is the contribution from the 6 points $\mathbf{p}_{ij}$ where we buy at price 1 from exactly 1 seller with probability $\frac{1-\epsilon}{6}$ each, the second term a bound from the 21 remaining points, where we buy from at most from 3 sellers, for at most a price of 3, with probability $(\frac{\epsilon}{21})$ each.

On the other hand, if a mechanism allocated at any of the $\mathbf{p}_{ij}$ differently (while maintaining monotonicity), that would mean either raising the purchase price to at least 2 at a $\mathbf{p}_{ij}$ (either due to buying from the same bidder at a higher price, or buying from a different bidder at price ≥ 2), or buying from more than one buyer at a $\mathbf{p}_{ij}$. Either way we would incur at least an extra $(\frac{1-\epsilon}{6})$ expected cost at one of the p_{ij}. The resulting total expected cost of the mechanism would thus also be at least $7 \cdot (\frac{1-\epsilon}{6})$.

It is easy to check that $1 + 8\epsilon$ is less than $7(\frac{1-\epsilon}{6})$ if $\epsilon < \frac{1}{55}$. So, for any such ϵ the optimal mechanism will have the property that the auctioneer buys each of the points $\mathbf{p}_{ij}$ for price 1 from seller j (and only seller j).

Step 2: Since the optimal mechanism buys from seller j for price 1 at each $\mathbf{p}_{ij}$, it follows that it buys from bidder i at each $\mathbf{q}_i$, as buying from either of the other bidders would contradict the low buying price at a $\mathbf{p}_{ij}$. Therefore by monotonicity, it will buy from all three sellers at $\mathbf{x} = (2, 2, 2)$. For k bidders, this construction easily generalises. Use k gadgets, each with $k - 1$ points $\mathbf{p}_{ij}$ forcing the mechanism to buy point $\mathbf{q}_i$ from the remaining bidder.

$\square$

A "May-Buy" Variant of the Reverse Auction Setting

As we discuss in section 3.2, in the reverse auction setting it makes little sense to require the mechanism to buy at most one copy of the item, i.e. to require $\forall \mathbf{v} : \sum_i x_i(\mathbf{v}) \geq -1$: even for the simple goal of maximising the social welfare $\sum_i v_i x_i(\mathbf{v})$ it would be optimal to simply set $x \equiv 0$; that is, to not buy the item at all. For cost minimisation, obviously the same holds - if allowed, it would be optimal to never buy the item, and thus incur cost 0. We therefore require that the mechanism always buys *at least* one copy of the item. We see no reason to prohibit it from buying more than one copy and discarding the extra copies; indeed requiring exactly one copy to be bought seems very restrictive. Hence, we require that $\forall \mathbf{v} : \sum_i x_i(\mathbf{v}) \leq -1$ as mentioned in the main text. This requirement of buying at least one copy isn't often discussed explicitly in the literature, but it is strictly needed. Conitzer and Sandholm [37] is the only paper we are aware of that mentions this specifically (and they use the same requirement we do). We believe that this definition of the reverse auction is most relevant and appropriate, as it most closely mirrors the auction setting.

However, we briefly mention one conceivable alternative to this, and show that this too is distinct from the auction setting. That is, our claim that reverse auctions behave differently than auctions is not an artefact of our definition of the reverse auction. Consider the following reverse auction variant: Instead of requiring the mechanism *must* buy the item, require that it *may* buy a copy of the item, and derive some fixed utility u^* from it if they do. Showing that this is different from the auction is easy. For instance, even for two bidders, it is possible for the optimal mechanism to buy from both or neither of the bidders. We only give a sketch proof; it can be made rigorous in the same manner as the proof of Theorem 15.

Theorem 16 (In the may-buy setting we may buy from all sellers.)**.** *In the may-buy setting, it is possible to construct an instance in which it is optimal to buy from all sellers for some bid combinations.*

Proof. We assume u^* to be large relative to all coordinates of all points in the discussion below. For three or more bidders, an instance like the one used in the "must-buy" setting (Theorem 15 and figure 3.6) works in the may-buy setting, with some extra probability on the $\mathbf{q}_i$. For two bidders, a similar construction works. Consider figure 3.8. Sufficiently high probability weight on bid profile $\mathbf{p}_1$ means the optimal mechanism will want to buy from seller 2 there for as low a prices as possible. This in turn means that it will not want to buy from seller 2 at any of the bid profiles (v_1, v_2') with $v_2' \geq v_2$. In particular it will want to avoid buying from seller 2 at bid profile $\mathbf{q}_1$.

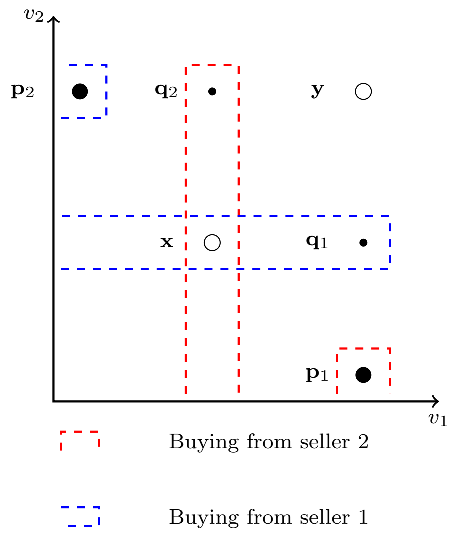

Figure 3.8: The two-bidder case in Theorem 16. In the optimal mechanism, we want to ensure a low price for (high-probability) bid vector $\mathbf{p}_1$ by buying from seller 2. We also want to ensure we get positive utility from bid vector $\mathbf{q}_1$, thus buying from seller 1. Vice verse for points $\mathbf{p}_2$ and $\mathbf{q}_2$. As a result, we may buy from both sellers at (very low probability) point $\mathbf{x}$. Note that in the may-buy setting, we can choose not to buy at all for bid vector $\mathbf{y}$.

However, it will still want to buy from seller 1 at $\mathbf{q}_1$, to get (lower) revenue $u^* - (\mathbf{q}_1)_1$. A similar reasoning shows that it will want to buy from seller 2 at $\mathbf{q}_2$. By truthfulness it follows that the optimal mechanism will then have to buy from both sellers at bid profile $\mathbf{x} = ((\mathbf{q}_2)_1, (\mathbf{q}_1)_2)$. (If the probability weight on $\mathbf{x}$ is small enough this will still be optimal.)

On the other hand, there are situations in which the optimal mechanism will want to forego buying at a particular bid profile. For instance, in figure 3.8 if there is sufficiently little prior probability at bid profile $\mathbf{y}$, it will be optimal to not buy there at all, to avoid raising the purchase price at $\mathbf{q}_1$ or $\mathbf{q}_2$. That is to say, in the may-buy setting all 2^k valid outcomes may occur in an optimal mechanism. $\square$

3.5 Auctions, Reverse Auctions, and Double Auctions

One way to visualise the relation of the discussed reverse auction setting to forward auctions is by looking at the allowed allocations in each setting. We begin with the two-bidder case. For a two-bidder auction, the allowed allocations are $(0, 0)$, $(1, 0)$ and $(0, 1)$ in the deterministic case, and any point in the triangle created by taking the convex hull of these in the randomised case. Three constraints work together to create this feasible region: Firstly for the two bidders, $0 \leq x_1 \leq 1$ and $0 \leq x_2 \leq 1$, i.e. each bidder may not receive the item, or receive one copy, or in the randomised case any expected fraction in between. In other words, they do not want less than 0 (i.e. they won't sell a copy of the item), and they do not want to buy more than 1 copy. The third constraint comes from the auctioneer, stating that at most one copy of the item may be allocated: $0 \leq x_1 + x_2 \leq 1$. Figure 3.9 (a) and (b) show this.

Correspondingly in a reverse auction, the allowed allocations are $(-1, 0)$, $(0, -1)$ and $(-1, -1)$, or their convex hull in the randomised case, as in figure 3.9 (c). It is worth pointing out that a market intermediation setting as discussed in Gerstgrasser et al.[60] also fits into this picture, with allowable allocations $(0, 0)$ (neither buy nor sell the item), $(0, -1)$ (buy the item but do not sell it), and $(1, -1)$ (buy and sell the item). This is with bidder 1 being the buyer and bidder 2 being the seller, but here the opposite case is indeed entirely symmetric. Figure 3.9 (d) shows this case.

In this framework, it is also easy to visualise the meaning of lemma 14: It

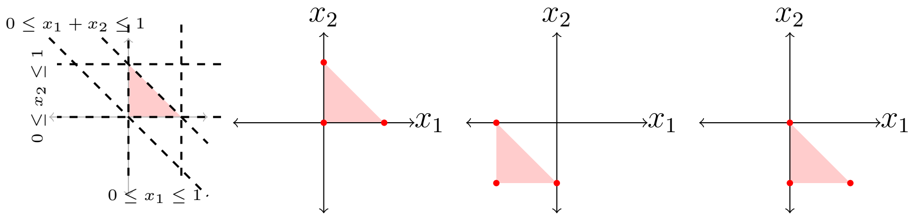

Figure 3.9: From left to right:
(a) The constraints in the two-bidder auction, with the intersection shaded in red.
(b) The resulting allowed allocations, with integer allocations shown bold.
(c) The corresponding picture in the reverse auctions setting.
(d) The feasible allocations in a market intermediation setting. (Top left quadrant would be the symmetric case with the roles of buyer and seller reversed.)

says that an optimal mechanism will never use allocations in the region greyed out in figure 3.10 (a), but rather will only use allocations shown in bold red points (in the deterministic case) respectively bold red line (in the randomised case).[3] A very similar result holds in the case of market intermediation, as detailed in Chapter 2. In this setting, the corresponding region shown in pale red in figure 3.10 (b) may be utilised by an optimal mechanism, but only when this is required due to truthfulness. This is essentially lemma 1: In one-seller-one-buyer market intermediation the region where the auctioneer buys the item from the seller is the smallest region which contains the region where the item is sold to the buyer and is closed under decreasing seller's bid.

For three or more bidders, the picture gets more complex, and again serves to visualise the difference between auctions and reverse auctions. In this case, it is stark. Geometrically, in both settings the auctioneer's feasibility constraint slices through a cube induced by the bidder's feasibility constraint, but in different parallel planes. In the case of the three-bidder auction it

[3]Notice that the proof of lemma 14 extends immediately to the truthful-in-expectation case.

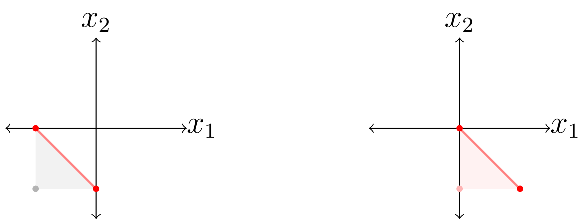

Figure 3.10: From left to right:
(a) Lemma 14 states that an optimal reverse auction will never utilise any allocation vector within the greyed-out region.
(b) The corresponding result for market intermediation states that the shaded red region will only be used when required by monotonicity.

gives the pyramid spanned by the origin and the unit vectors. In the three-bidder reverse auction the feasible region extends "one layer further", being spanned by the point $(-1,-1,-1)$ where the auctioneer buys from all three sellers, similar to the origin in the forward auction case; a "first layer up" similarly to an auction, consisting of the three points $(0,-1,-1)$, $(-1,0,-1)$ and $(-1,-1,0)$ where they buy from two of them; and "second layer up" comprising the three points $(0,0,-1)$, $(-1,0,0,)$ and $(0,-1,0)$ where they buy from one seller. Figure 3.11 shows these two cases. Similarly to Theorem 14 for the two bidders, a more general statement holds here stating that in the optimal mechanism the allocations below the "top" $\sum x_i = -1$ layer will only occur if needed due to truthfulness.

For the two possible market intermediation scenarios with three bidders we get similar pictures. Interestingly, the one-seller-two-buyers case is close to the auction, while the two-sellers-one-buyer case is closer to the reverse auction. In the 1×2 case, taking bidder 1 to be the seller, we get allowed allocations $(-1,0,0)$ - buy, but don't sell; $(0,0,0)$ - neither buy nor sell; and $(-1,1,0)$ and $(-1,0,1)$ - buy and sell to either buyer. Similarly to the auction

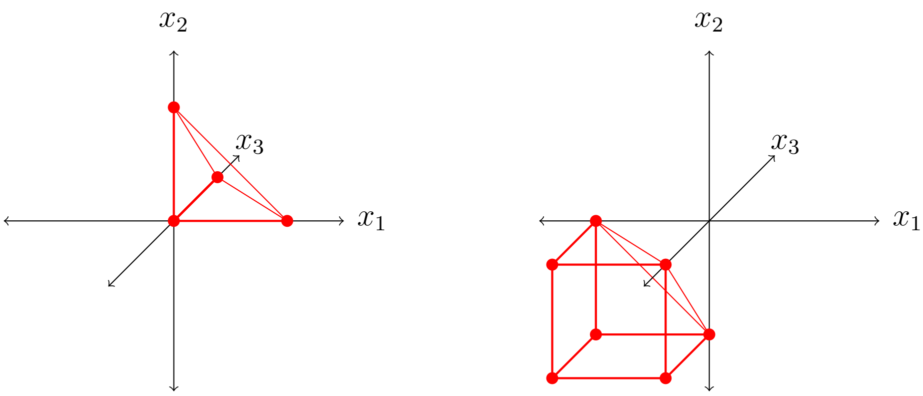

Figure 3.11: From left to right:
(a) The allowed allocations in a three-bidder auction. Integer allocations shown as bold dots, randomised allocations in outline.
(b) The corresponding allowed allocations in a three-bidder reverse auction.

case, this forms a pyramid, in one of the adjacent octants with one negative coordinate. In the 2×1 case we get a "pyramid plus extra layer" shape, similar to the reverse auction case but in one of the octants with two negative coordinates. It is worth noting that Papadimitriou and Pierrakos [106] show that the three-bidder auction is NP-hard, while Gerstgrasser et al. [60] show the same for the one-seller-two-buyers market intermediation. In other words, the two "pyramid shaped" cases are resolved. On the other hand, both "pyramid plus extra layer" cases are still open. It is unclear whether the larger space of allocations adds extra complexity to these problems, or if the additional degrees of freedom might make them easier.

It is also easy to see how this would generalise further to more complex cases. E.g. imagine an auction, with the following twist: Not only does the auctioneer have a copy of the item, some or all of the bidders also already have a copy. They still each have demand for an additional copy, so the auctioneer could just consider this a plain single-item auction. Or, the auctioneer could sell their copy, but also buy another copy from one bidder, and sell that to another bidder at a profit. In other words, an auction with potential for intermediation. Or, market intermediation where the intermediary starts out with a copy of the item.

The obvious question here is if there is a more comprehensive theory unifying these different cases. Or indeed, covering also more complex scenarios within this framework of a "general single-item market model". Notice for instance that a three-bidder unit-demand auction with two copies of the item for sale would also lead to a "pyramid plus extra layer" shape of allowed allocations, similar to the reverse auction.

3.6 Discussion and Future Work

Our $O(n^3)$ algorithm for the two-bidder auction is interesting not only in its own right - presenting a substantial improvement of the previously known $O(n^6)$ approaches, both for the exact solution for two bidders but also for a $\frac{5}{3}$-approximation for many bidders; but it is significant also because of the techniques used in arriving at it. Our algorithm is the first to exploit structural properties of the 2-bidder auction design problem, whereas previous approaches reduce the problem to a generic graph algorithm respectively a derandomised LP. In contrast, structural insights allow us to decompose the problem into a simple recursion. Generalised versions of these observations hold for many settings, and could potentially be useful there.

Lemma 2 gives a useful characterisation of the optimal two-bidder auction, and indeed this generalises to virtually all correlated mechanism design settings we can think of. For two-bidder settings, this gives a useful decomposition into optimal single-bidder mechanisms, that may be tractable in many cases. For more general cases, an analogue of Lemma 2 states — informally — that in an optimal mechanism, for any partition of bidders into two sets S_1, S_2, for a fixed bid vector $\mathbf{v}_1$ of bidders in S_1, the allocation to bidders in S_2 is simply that of the optimal $|S_2|$-bidder mechanism taking into account the (fixed) allocations to bidders in S_1. This may be a useful aid in exploring multi-bidder settings.

It is interesting to consider why this generalisation cannot be used to show a polynomial-time algorithm for the three-bidder auction. For this setting, we get two characterisations depending on whether we partition into $|S_1| = 2$, $|S_2| = 1$, or vice versa. If we take $S_1 = \{b_1, b_2\}$, $S_2 = \{b_3\}$, we get that for any fixed v_1, v_2, the allocation to bidder 3 in the induced affine linear v_3-subspace (i.e. in each "vertical line" in the bid space) is that of the optimal single-bidder auction to bidder 3 on all bid vectors that are above any allocated to bidders 1 and 2 (i.e. on those not blocked by the allocation to bidders 1 and 2). Vice versa, if we take $S_1 = \{b_3\}$, $S_2 = \{b_1, b_2\}$, we get that for fixed v_3, the allocation to bidders 1 and 2 on the induced v_1, v_2-subspace (each horizontal plane in the bid space) is that of the optimal two-bidder mechanism on all bid vectors not blocked by the allocation to bidder 3. While these are interesting properties, what is missing is an analogue of Lemma 3, which allows us to rewrite the optimal revenue as a tractable recursion. In the three-bidder case, we do not have an efficient way of computing the optimal allocation to two or even one bidder. Even in such settings where this result does not lead to an efficient algorithm, it may still be useful in characterising the optimal mechanism. One open question that is of particular interest here is whether Lemma 2 and its generalisation also hold as an "if and only if" statement. That is, it would be nice if they would give sufficient conditions for a mechanism to be optimal.

Our results on correlated reverse auctions for the first time (to our knowledge) show an asymmetry between auctions and reverse auctions. For two bidders, a further structural analysis allows us to show a small reduction in complexity compared to the auction, and to devise the first algorithm specific to the correlated reverse auction setting. Our result for three or more bidders is surprising, as it shows a much higher dimensional space of possible outcomes - exponential (in the number of bidders) compared to linear in an auction. We take this as evidence that the reverse auction case is interesting to consider as a separate problem from the standard auction model.

Part 4

Online Market Intermediation

4.1 Introduction

We next focus on a setting similar to the offline double auction, but in which buyers and sellers appear one-by-one. As buyers and sellers interact with the intermediary one-on-one, the intermediary is essentially limited to posted-price mechanisms. It is therefore not so much the individual interaction (and each trader's valuation and prior distribution) that is challenging, but rather the sequence in which they appear. We therefore analyse this setting in the tradition of worst-case analysis online algorithms, in which we assume a candidate online algorithm does not know the sequence of buyers and sellers in advance, and compare its performance against an offline adversary (which knows the sequence of traders, but not the draws from the prior, in advance).

A similar setting to ours has been studied before by Blum et al. [20]. They consider multiple buyers and sellers trading identical copies of an item. Arrivals and departures of traders are revealed to the intermediary as they happen, and may overlap. Most relevant for us, they give results on the intermediary's revenue compared to that of an offline adversary. The key difference to our work is the overlapping of traders. This gives rise to an analysis based on online weighted matchings, as the intermediary themselves never holds on to copies of the item, which are traded directly between seller and buyer. In contrast, in our work the intermediary must make purchasing decisions without knowledge of possible future buyers, and maintains a stock of items themselves. We consider both settings to capture different scenarios: Compare eBay, where seller and buyer are matched directly, to Amazon, where the intermediary themselves maintains a stock of items. Clearly the challenges that both face are very different from one another. In the overlapping setting, the aim is to match buyers and sellers. In the one-by-one setting, particular matchings are not important (except as a proof technique, as we shall see), but great difficulty arises from the unknown future arrivals.

Our interest is in studying this online intermediation setting, in which the intermediary encounters buyers and sellers as a one-by-one sequence. The intermediary's aim will be to maximise their revenue, by purchasing from sellers and selling on to buyers at a profit. Our analysis will be two-fold: As regards the utilities and traders' valuations, we treat them in a Bayesian fashion, assuming they are drawn from (independent) priors. In contrast, we will analyse the sequence of arrivals using competitive analysis, where arrivals are chosen adversarially, and are unknown to the online intermediary a priori. We study the case of arbitrary sequences of buyers and sellers and show that the competitive ratio – the ratio of the optimal offline profit over the profit obtained by the online algorithm – is $\Theta(n)$, where n is the number of buyers and sellers. The Θ-notation hides important factors that measure how much the probability distributions differ from the uniform distribution. The upper bound is obtained by detail-free mechanism: only very simple nuggets of information are required from those valuations, in order to produce asymptotically optimal algorithm. This pessimistic analysis is matched by two positive results. First, when the intermediary can hold only up to k items for some fixed k, the competitive ratio is constant; the explanation is that the restriction of the k items reduces significantly the freedom of the offline algorithm. Second, when we consider balanced sequences of buyers and sellers— in particular, when there is equal number of buyers and sellers and the i-th buyer is preceded by at least i sellers—the competitive ratio drops to a small constant. In all these cases, we study truthful mechanisms, which are posted price mechanisms, interesting in their own right, due to achieving respectable performance and robust simplicity in a variety of auctions, despite their constrained nature.

Prior Work

Our line of inquiry relates to a number of active research areas within mechanism design. Most obvious is the connection to the already discussed literature on (offline) market intermediation and double auctions. Myerson and Sattherthwaite [99] laid the groundwork for this, later extended to many bidders by Deng et al. [45], and to budget constraints by Blumrosen and Dobzinski [21]. The assumptions we make in this chapter are essentially in line with this tradition of research.

In sequential auctions, a first step was a comparison of simple posted-price mechanisms to (revenue-) optimal mechanisms, including a proof of equivalence under some assumptions [27, 22, 125]. An approach based on no-regret learning was adopted by [19, 20] in the setting of digital goods auctions, in which the supply is unlimited. Online mechanism design is also intimately connected to secretary problems, which has been explored in several publications [68, 67, 10]. For a general review of online mechanisms, we refer the reader to [107].

A slightly different line of research has been pursued on prior-free online auctions. In this, the exact prior distribution is not known, but some larger class of distributions that restricts the prior is known, e.g. MHR or regular distributions. A single-item independent-priors setting was explored in [7] and extended to multiple items in [9]. Digital goods auctions in this analysis have also received some attention in the literature [71, 78, 79], as have procurement auctions [17].

4.2 Preliminaries

We consider the problem faced by a broker who faces a sequence $\sigma = (a_1, \ldots, a_n)$ of buyers and sellers, each trading a single item. Each trader a_i can be either a seller or a buyer that has a valuation for the item v_i. A seller is willing to sell the item at any price higher than v_i and a buyer is willing to buy at any price lower than v_i. We assume that the valuation v_i of every seller is drawn independently from a known probability distribution with probability density function (pdf) and cumulative density function (cdf)

$s(.)$ and $S(.)$ respectively; similarly, there are pdf and cdf $b(.)$ and $B(.)$ for every buyer. We make the simplifying assumptions that the distributions are continuous and regular with support $[0, 1]$ (although our results hold by a simple scaling argument for $[0, D]$, for any $D > 0$). Regularity means that their virtual values $v + \frac{S(v)}{s(v)}$ for selling and $v + \frac{1-B(v)}{b(v)}$ for buying are strictly increasing; this restriction can be removed if we allow randomized algorithms. We also require the standard assumption that $s(.)$ and $b(.)$ are strictly positive in their domain.

The intermediary interacts with the buyers and sellers by posting prices in the following online manner: when agent a_i arrives, the intermediary posts a price depending only on the previous observations: a_j for $1 \leq j \leq i$, the previous prices set p_j for $1 \leq j < i$ as well as the history of the intermediary's internal state (the stock level) up to that point. If the trader is a seller and the posted price is higher than his valuation, he sells the item and the intermediary pays the posted price. If the trader is a buyer and the intermediary has at least one item in stock, one item is transferred to the buyer who pays the intermediary the posted price. This asymmetry between the buyers and sellers is crucial in designing competitive online algorithms.

Our goal is, given the distributions $S(\Delta)$ and $B(\Delta)$, to maximise the profit of the intermediary, relative to that achieved by the optimal offline algorithm which knows the sequence (but not the realizations of the valuations) in advance. We use a standard metric from online algorithms, the competitive ratio. An algorithm A is c-competitive if for any input σ:

$$c(\mathcal{R}_\sigma(A) + \alpha) \geq \mathcal{R}_\sigma(\text{OPT}) \qquad (4.1)$$

where $\mathcal{R}_\sigma(A)$ and $\mathcal{R}_\sigma(\text{OPT})$ are the expected revenue achieved by algorithms A and the optimal offline algorithm OPT for the sequence σ of buyers and sellers. The additive constant α is necessary as we have to allow the online algorithm at least a constant loss margin in the worst case.

We study three variations of our setting. The first and most natural is the general case, where the sequence of buyers and sellers is unrestricted. In this case, the adversary – who designs the worst-case sequence – is really powerful and there is not much to do. For example, how can any online algorithm do well against both sequences S^n and S^nB^n when it has to be the same in the common prefix? However, it is possible to find reasonable bounds on the competitive ratio when we restrict the class of the probability distributions.

In Section 4.3, we consider the classes of distributions parametrised by two parameters: M and N. Specifically, we will assume that these parameters bound the density from above, and the rate of vanishing near zero from below:

$$b(v), s(v) \leq M \tag{4.2}$$

$$\lim_{v \to 0} \frac{S(v)}{v} \geq N \tag{4.3}$$

On the other hand, we consider (later in the same section) the case in which the intermediary can only hold up to K items, where K is a constant. This curtails so much the power of the offline algorithm that the competitive ratio drops to a constant, which depends on the probability distributions.

In the second variation (Section 4.4), we consider the natural restriction in which the sequences of buyers and sellers are balanced, i.e., they contain the same number of buyers and sellers and, for every i, buyer i is preceded by at least i sellers. In this case we analyse the profit obtained by the optimal offline algorithm and design a simple competitive online algorithm.

4.3 General Setting

In this section we study the simplest setting with arbitrary sequences of buyers and sellers. When we impose conditions on the distributions, we construct a lower bound which outlines the nature of the uncertainty in this problem. This is complemented by designing a natural algorithm avoiding those pitfalls.

For the presentation of the lower bound we make the simplifying assumption that the distributions S and B for the sellers and buyers are the same, whose cdf and pdf we denote by F and f. This can only strengthen the lower bound.

The two conditions on the distributions 4.2, 4.3 will be used and warrant a short discussion to provide intuition. Without loss of generality we will assume that $M \geq 1$ (otherwise just take $\max\{1, M\}$). One immediate consequence is that for all $v \in [0, 1]$

$$F(v) = \int_0^v f(t)dt \leq \int_0^v M dt = M \cdot v \tag{4.4}$$

For example, the uniform distribution satisfies this property with $M = 1$ and the truncated exponential over the unit interval with $M = \frac{e}{e-1} \simeq 1.582$. We mention that it is inequality (4.4) that we need in the following and not the supremum norm bound per se. Also, all the results in this section can be adapted to discrete distributions; although there is no density in this case, we can just use 4.4 as $\Pr[X \leq t] \leq M \cdot t$ for all $t \in [0, 1]$, where X is the random variable of an item value.

Lower Bound

This lower bound quantifies the notion that due to uncertainty, one cannot be too aggressive in buying many items as there is no guarantee that enough buyers will appear later. It does so by considering long sequences of sellers, where the online algorithm does not know if they'll be followed by equally long segments of buyers.

Theorem 17. *For any continuous distribution of item values, no (even randomised) online market intermediator can achieve a competitive ratio better than $\Omega(\sqrt{n})$, where n is the length of the sequence.*

Proof. Let X_i denote the value agent i has for the item, and assume these values are i.i.d. from a continuous distribution over $[0, 1]$ with cdf F, for both buyers and sellers. First consider a sequence of of n consecutive sellers S^n. Let P_i be the random variable of the price that the online algorithm posts to the i-th seller. Notice that this may include any dependence on the previous values $X_1, \ldots, X_{i-1}$ that we have observed in the past, as well as any internal randomisation of the algorithm itself. Also, let Y_i be an indicator variable denoting whether we bought an item at step i, i.e.

$$Y_i = \begin{cases} 1, & X_i \leq P_i \\ 0, & \text{otherwise.} \end{cases}$$

Then, the expected total number of items bought Y is

$$\begin{aligned} \mathbb{E}[Y] = \sum_{i=1}^{n} \mathbb{E}[Y_i] &= \sum_{i=1}^{n} \mathbb{E}_{P_i}[\mathbb{E}_{Y_i}[Y_i|P_i]] \\ &= \sum_{i=1}^{n} \mathbb{E}_{P_i}[\Pr[X_i \leq P_i]] \\ &= \sum_{i=1}^{n} \mathbb{E}[F(P_i)]. \end{aligned}$$

Also, if we define R_i to be the random variable representing the amount the online algorithm pays to the i-th seller, then the expected total expenses R of the algorithm are

$$\begin{aligned} \mathbb{E}[R] = \sum_{i=1}^{n} \mathbb{E}[R_i] &= \sum_{i=1}^{n} \mathbb{E}_{P_i}[\mathbb{E}_{R_i}[R_i|P_i]] \\ &= \sum_{i=1}^{n} \mathbb{E}_{P_i}[P_i \Pr[X_i \leq P_i]] \\ &= \sum_{i=1}^{n} \mathbb{E}[P_i F(P_i)]. \end{aligned}$$

Now notice that, since the algorithm has only a constant available loss margin (see 4.1), there must be a constant α such that that

$$\mathbb{E}[R] \leq \alpha. \tag{4.5}$$

Next, by applying repeatedly Jensen's inequality for the convex function $t \mapsto t^2$ we can get that

$$\mathbb{E}[Y]^2 = \left(\mathbb{E}\left[\sum_{i=1}^{n} F(P_i)\right]\right)^2 \leq \left(\mathbb{E}\left[\sum_{i=1}^{n} F(P_i)\right]^2\right) \leq n \cdot \mathbb{E}\left[\sum_{i=1}^{n} F(P_i)^2\right].$$

But from (4.4), for all $i = 1, 2, \ldots, n$ it is

$$F(P_i)^2 \leq M \cdot P_i \cdot F(P_i)$$

so from the above chain of inequalities we have

$$\mathbb{E}[Y]^2 \leq nM\mathbb{E}\left[\sum_{i=1}^{n} P_i F(P_i)\right] = nM\mathbb{E}[R].$$

Combining this with 4.5 finally gives

$$\mathbb{E}[Y] \leq \sqrt{\alpha M} \cdot \sqrt{n}. \tag{4.6}$$

Now consider the addition of n consecutive buyers, that is the sequence $S^n B^n$. By the previous analysis, even if the online algorithm manages to sell all the items acquired during the initial stage S^n, by 4.6 its expected total profit cannot exceed

$$\mathbb{E}[Y] \cdot 1 = \sqrt{\alpha M} \cdot \sqrt{n} = O(\sqrt{n}).$$

On the other hand, consider the offline algorithm that posts a fixed price of $q = \frac{1}{8M}$ to all sellers and a price of $p = \frac{1}{2M}$ to all buyers. Then, if N_s is the random variable representing the number of items the offline algorithm bought during the initial S^n part of the input, it is $\mathbb{E}[N_s] = nF(q)$. Furthermore, by assuming that the offline algorithm will only try to sell items to the first N_s buyers, that is buyers $i = n+1, n+2, \ldots, n+N_s$, the expected number of items sold is $\mathbb{E}[N_s](1-F(p)) \geq \mathbb{E}[N_s](1-Mq) = \frac{1}{2}nF(q)$. Thus, the offline profit is lower-bounded by

$$\frac{1}{2}nF(q)\cdot p - nF(q)\cdot q = nF(q)(\frac{p}{2} - q) - nF\left(\frac{1}{8M}\right)\left(\frac{1}{4M} - \frac{1}{8M}\right) = \Omega(n).$$

□

This shows us that candidate algorithms should, if they encounter a long sequence of sellers, buy a superconstant number of items for at most a constant price. We will leverage this in the following subsection.

Upper Bound

The structure of the lower bound can serve as a guide to a simple algorithm, with a competitive ratio containing $\sqrt{n}$ times a factor dependent on the constants N, M.

Algorithm for the general case: The algorithm contains just two rules for posting prices to sellers and buyers:

- The price posted to the i-th seller encountered so far (which is not the same as the seller encountered at step i) is $q(i) = \frac{1}{i^{1/2+\epsilon}}$ for $\epsilon > 0$ arbitrarily small.

- For any buyer the posted price is $p = \frac{1}{2M}$ (if the intermediary has items available for sale).

Note that only M needs to be known instead of the actual distribution, making this mechanism mostly detail-free. Its analysis can be broken down in three parts. We show that there is an upper bound on the amount of money the algorithm can spend, which is absorbed by the additive constant in 4.1. Afterwards, we study a mechanism that is more restricted, but easier to analyse, to bound the probability that the online algorithm will succeed at each the offline adversary attempts to trade. Combining this with a lower bound on the income attained from each sale, we get competitiveness of the online algorithm.

In this section we use the extra assumption on the distributions of item values: we want the ratio $\frac{S(v)}{v}$ not to vanish to zero as v becomes arbitrarily small. Essentially, we want to ensure that the agents' valuations span the support of the distribution: any posted price has a positive probability of leading to a successful trade. Formally, we assume that there is a universal constant $N > 0$ (not controlled by the adversary) such that

$$\lim_{v \to 0} \frac{S(v)}{v} \geq N \tag{4.7}$$

Various stronger conditions exist that ensure that 4.7 holds, for example that the density function is non-zero on the left boundary, $s(0) \geq N$, or even that $s(v) \geq N$ for all $v \in [0, 1]$. Notice for example that the uniform distribution satisfies (4.7) with $N = 1$. and so does the exponential, while for the truncated exponential over $[0, 1]$, $N = \frac{e}{e-1} \simeq 1.582$.

Theorem 18. *The above algorithm is $O(\frac{M}{N} n^{1/2+\epsilon})$-competitive, where $\epsilon > 0$ is arbitrarily small.*

Proof. First notice that for any input sequence, the expected total amount spent by the algorithm is upper-bounded by

$$\sum_{i=1}^{\infty} S(q(i))q(i) \overset{(4.4)}{\leq} M \cdot \sum_{i=1}^{\infty} q^2(i) = M \sum_{i=1}^{\infty} \left(\frac{1}{i^{1/2+\epsilon}} \right)^2 \leq M \sum_{i=1}^{\infty} \frac{1}{i^{1+2\epsilon}} \equiv c < \infty$$

since the exponent of the denominator in the series is greater than 1. It follows that the profit is at least $\frac{k}{2M} - c$, where k is the number of items sold. Thus, when analysing the expected profit of this algorithm asymptotically as $k \to \infty$, we can get a lower bound by focusing just on the term involving k, and absorbing the rest with the additive term in (4.1). In order to simplify the analysis, we will consider a more restricted algorithm that ignores most of the potential trades it could attempt.

We start by bounding the profit that the optimal offline algorithm could achieve. Imagine for a moment the offline algorithm could ignore the drawn valuations, and always buy for free and sell at a maximum price of 1. The profit achievable in this way depends only on the maximum number of trades $m(\sigma)$ possible for a sequence σ of sellers and buyers (and it clearly upper bounds the optimal offline profit.) The maximum number of trades is equivalent to the size of a maximum matching in the bipartite graph $G(\sigma) = G(\sigma, E)$ defined as follows: the vertices are the sellers and the buyers, and there is an edge $(s, b) \in E$ if and only of it is possible to sell an item bought from seller s to buyer b or, equivalently, if s appears before b in σ. Then, clearly $\mathbb{E}[\mathcal{R}^{\mathrm{OPT}}] \leq m(\sigma)$.

It is possible to compute a maximum matching online: keep track of sellers encountered in a FIFO queue. Whenever a buyer is encountered, match it to the oldest seller in the queue and remove that seller from the queue. (Do not match a buyer if the queue is empty.) This computes the lexicographically first maximum matching (this is Lemma 4, page 118, for $K \to \infty$). It is clear that exactly the first $m(\sigma)$ sellers will be included in this matching. On the buyers side, the first $m(\sigma)$ "feasible" buyers will be included. The edges of this matching are "ordered": if (s, b) and (s', b') are both in the matching and $s <_\sigma s'$, then $b <_\sigma b'$.

Next we move on into lower-bounding the expected online profit. Our restricted algorithm is as follows: post price $q(i)$ as above to sellers. If buyer j is matched to seller i (in the offline lexicographically first maximum matching described above), and the algorithm has successfully bought an item from i, then post price p to j. If the sale is unsuccessful, throw away the item bought from seller i and move on. The algorithm will never have spare items to sell to unmatched buyers. It will throw away items from unmatched sellers at the end. It is clear that for the matched first $m(\sigma)$ sellers, this algorithm will buy an item with probability $S(q(i))$ from the i-th seller and, in the event of succeeding, sell the same item to the matched buyer $\mu(i)$ with probability

$$1 - B(p) \overset{4.2}{\geq} \left(1 - M\frac{1}{2M}\right) = \frac{1}{2}.$$

Thus, for each of these sellers, the probability of a successful transaction (purchase and sale) occurring is at least $S(q(i)) \cdot (1 - B(p)) \geq \frac{S(q(i))}{2}$. This gives us an expected income of at least

$$\begin{aligned}
\sum_{i=1}^{m(\sigma)} \frac{S(q(i))}{2} \cdot p = \frac{p}{2} \sum_{i=1}^{m(\sigma)} S(q(i)) & \qquad (4.8) \\
\geq \frac{\frac{1}{2M}}{2} \cdot m(\sigma) \cdot \min_{i=1,\ldots,m(\sigma)} S(q(i)) & \qquad (4.9) \\
\geq \frac{m(\sigma)}{4M} S\left(\frac{1}{m(\sigma)^{1/2+\epsilon}}\right) & \qquad (4.10)
\end{aligned}$$

since prices $Q(i)$ decrease as i increases. But $\frac{1}{m(\sigma)^{1/2+\epsilon}} \to 0$ as $m(\sigma) \to \infty$, and so by property 4.7 we can ensure that for long enough sequences

$$S(q(m(\sigma))) \geq \frac{N}{2} q(m(\sigma)) = \frac{N}{2} \frac{1}{m(\sigma)^{1/2+\epsilon}}.$$

So, using 4.10 and the fact that the algorithm's losses are a constant c, the expected online profit is lower bounded by

$$\frac{N}{4M} m(\sigma)^{1/2-\epsilon} - c = \Omega(\frac{N}{M} m(\sigma)^{1/2-\epsilon}).$$

This results in a competitive ratio of

$$O(\frac{M}{N} m(\sigma)^{1/2+\epsilon})$$

since the size of the maximum matching can be at most $m(\sigma) \leq \frac{n}{2}$. This shows our claim. $\square$

Limiting the Number of Unsold Items

Taking a closer look, the root of inefficiency seems to be the long sequences of sellers the online algorithm faces. Limiting the number of unsold goods to be less than or equal to K at any given time, imposes a similar constraint to having a full warehouse, while at the same time limiting the profit the offline algorithm can generate from those long sequences. Essentially, this constraint restricts the adversary far more than the online algorithm.

Algorithm for K items: The algorithm consists of two prices q and p, one for sellers and one for buyers, and attempts to buy items only if the current stock is $k_i < K$ for a constant K. Prices are selected so as to maximise the profit for the "sequence" SB. Formally, select q, p that maximize: $-qS(q) + p(1 - B(p))S(q)$.

Theorem 19. *This algorithm is $(-qS(q) + p(1 - B(p))S(q))$-competitive for any fixed K.*

Proof. Similarly to the unrestricted case, we begin by bounding the profit achievable by the offline optimal. This is again bounded by the maximum number of possible transactions given the buyer-seller sequence. Construct a directed bipartite graph as before. The maximum number $m_K(\sigma)$ of possible transaction holding up to K items corresponds to the size of a maximum matching with no "temporal" cut of size greater than K. Here, we define "temporal" cut to mean any cut in the graph that separates earlier vertices (buyers and sellers) $1, \ldots, i$ from later vertices $i+1, \ldots$. This is precisely the condition that we cannot hold more than K items at a time, corresponding to matching more than K sellers from an initial segment to buyers later in the sequence.

Again, we can compute one such matching online, using a FIFO queue of length K, adding sellers to the queue while it is not full and match buyers greedily as before. Consider again a restricted algorithm that posts prices q to those sellers that can be added to the queue. To bound our cost, we need to avoid buying from sellers we already know cannot be matched to buyers due to the K-limit. When a buyer is encountered and is matched to a seller, post price p if an item was successfully bought from the matched seller. Attempt no other sales. As before, we use $\mu(i)$ to denote the buyer matched to seller i.

Using a similar method to the analysis of the offline algorithm at the end of Theorem 17, we show that the online algorithm generates a constant amount of profit for any matching attempted by the offline. Observe that for any pair of seller i and buyer $\mu(i)$in the maximum matching, we extract an expected profit of $-qS(q) + p(1 - B(p))S(q)$. We will also buy items from sellers that go unmatched. However, by the construction of the matching, there will be at most K of these, leading to losses of at most Kq, which can be captured by the additive constant. (To see this: We only buy from sellers that entered the FIFO queue. If a seller enters the queue, the only way it will not be matched to a buyer is if it is in the queue when we reach the end of the sequence. There can be at most K sellers in the queue at any time, and in particular then.)

Therefore, our online algorithm's profit is at least

$$m_K(\sigma)\left(-qS(q) + p(1 - B(p))S(q)\right) - Kq,$$

while the optimal offline algorithm can achieve at most $m_K(\sigma)$. This shows our claim.

$\square$

To conclude the analysis, we present a lemma showing that the matching of possible trades computed by the restricted online algorithm is indeed maximum, thus upper bounding the offline profit.

Lemma 4 (FIFO Maximum Matching)**.** *The matching computed using a FIFO queue of size K as in the preceding proof is a maximum matching.*

Proof. We show this for the K-items case. The general case works similarly, or follows by setting K large enough. Let F be the matching computed by our FIFO algorithm, and let M be any arbitrary maximum matching in $G(\sigma)$. We will show that we can transform M into F using a series of changes that do not reduce its size.

Let i be the index of the first vertex that is not matched in the same way in F and M. That is, all edges in F and M that are either between vertices before i, or originate at a vertex before i, are identical in both matchings. (There cannot be any matchings that terminate in a vertex smaller than i but originate after i due to the construction of the graph.) We will show using a case-by-case analysis that we can change M into M' so that i is matched the same way as in F, without changing any edges originating before i, and with $|M| = |M'|$. It follows that we can repeat this procedure until M is transformed into F, and thus $|M| = |F|$, i.e. F is a maximum matching.

1. If i is a buyer: This is not possible. If i is matched in either matching,

the edge is originating from a vertex before i, and thus must be the same in both matchings by our hypothesis.

2. So i is a seller.

   (a) If i is matched in both matchings. Let j_F be its match in F, and j_M in M.

      i. $j_F < j_M$

         A. j_F unmatched in M. Make edge ij_M into ij_F. Can't violate K-limit this way, as we're making the edge shorter.

         B. j_F matched in M. Make edge ij_M into ij_F, and match the seller originally matched to j_F in M with j_M. We can't violate the K-limit this way.

      ii. $j_M < j_F$ - This is not possible.

         A. It is not possible that j_M is unmatched in F, as we encounter it before j_F, and would have matched i to it.

         B. It is not possible that j_M is matched to a seller other than i in F. Not to one before i by hypothesis, and not to one after i by construction of the FIFO algorithm.

   (b) If i is matched in F but not in M. Let j_F be i's match in M.

      i. j_F unmatched in M. This cannot happen. Notice that we cannot have any buyers between i and j_F that are unmatched in F, nor can we have any that are matched to sellers after i. Thus, all buyers between i and j_F are matched to sellers before i in both F and M. There can be at most $K - 1$ of them, as there is one more edge originating from i in F, and a cut between i and $i+1$ has size at most K in F. Therefore, we could add the edge ij_F to M without violating the K-limit. Thus M was not maximum, contradicting out assumption.

ii. j_F matched in M. Let s_M be the seller matched to j_F in M. Again, all buyers between i and j_F are matched to sellers before i in both matchings. So we can replace $s_M j_F$ with $s j_F$ without violating the K- limit.

(c) If i is matched in M but not in F. Let j_M be its match in M.

i. j_M is matched in F. This cannot happen due to the FIFO construction.

ii. j_M is unmatched in F. This cannot happen. If i were to enter the FIFO queue, it would be matched to j_M (or an earlier available buyer) in F. If i does not enter the FIFO queue this can only be because the queue was full. But if the queue was full, this means that K sellers before i were matched to buyers between i and j_M (otherwise j_M would be matched to one of them if F). So there is K edges going from sellers before i to buyers between i and j_M in F. So there is also K edges going that way in M, as they are identical on vertices before i. So there is $K+1$ edges going from nodes before and including i to vertices after i in M, violating the K- limit.

□

4.4 Balanced Sequences

In this section we analyze the optimal profit obtained by offline and online algorithms on balanced sequences.

Definition 21. *A sequence of buyers and sellers is called balanced if it contains the same number of buyers and sellers and, for every i, the i-th buyer is preceded by at least i sellers.*

For example, the sequence $SBSSBSBB$ is balanced, but the sequence $SBBSSB$ is not.

Since the buyers and sellers appear one-by-one, the only truthful mechanisms are those that post a price (or a distribution/menu of prices) and the trader decides whether to accept it or not. We make here a distinction between two types of mechanisms:

- *Non-adaptive* mechanisms compute all buying and selling (posted) prices in advance.

- *Adaptive* mechanisms may compute a price at step i based on the previous history of interaction with agents $1, \ldots, i-1$. In particular, they may take into account how many items have been bought and sold before.

The optimal profit and the posted prices of an adaptive mechanism can be calculated by a simple dynamic programming approach. Let $\mathcal{R}^k(a_1, \ldots, a_n)$ denote the optimal profit when the intermediary starts with k items and faces a sequence $a_1, \ldots, a_n$ of traders, where $a \in \{S, B\}$. Optimal revenue and prices can be calculated (at each step) by the following recurrence in polynomial time:

$$\mathcal{R}^0(\epsilon) = 0 \tag{4.11}$$

$$\mathcal{R}^0(B, \sigma) = \mathcal{R}^0(\sigma) \tag{4.12}$$

$$\mathcal{R}^k(B, \sigma) = \max_p \left\{ \left(p + \mathcal{R}^{k-1}(\sigma)\right)(1 - F(p)) + \mathcal{R}^k(\sigma)F(p) \right\} \tag{4.13}$$

$$\mathcal{R}^k(S, \sigma) = \max_p \left\{ \left(\mathcal{R}^{k-1}(\sigma) - p\right)(1 - F(p)) + \mathcal{R}^k(\sigma)F(p) \right\} \tag{4.14}$$

In this section, we show that the optimal online mechanism for balanced sequences has constant competitive ratio. In our analysis, we will employ a (fictional) class of fractional mechanisms. We allow these to buy fractions of items, and assume that a posted price of p would buy a $S(p)$-fraction of an item respectively sell $B(p)$ of an item. This simplifies our analysis: in a fractional mechanism, we know in advance the exact quantity of stock held at any point, rather than merely an expected number of items. In particular, we will know whether the intermediary has sufficient stock levels for a sale at any point. A consequence of this is that adaptive and non-adaptive fractional mechanisms are identical.

The remainder of this chapter establishes the following relations of optimal revenue for balanced $2n$-sequences:

$$\begin{aligned} \text{adaptive}(\sigma) &\le \text{fractional}(\sigma) \\ &\le \text{fractional}(S^n B^n) \\ &\simeq \text{non-adaptive}(\sigma) \end{aligned}$$

$$\leq 2 \cdot \text{non-adaptive online}(\sigma)$$

It follows that, since even a non-adaptive online mechanism is 2-competitive, for balanced sequences the competitive ratio of the optimal online mechanism is constant.

Theorem 20. *The per-item revenue of an optimal fractional mechanism on the sequence S^nB^n is given by the following optimisation problem:*

$$\begin{aligned} \max \quad & p(1 - B(p)) - qS(q) && (4.15) \\ s.t. \quad & 1 - B(p) = S(q) && (4.16) \\ & p, q \in [0, 1] && (4.17) \end{aligned}$$

Proof. It is easy to see that the optimal posted prices are given by the following optimisation problem:

$$\begin{aligned} \max \quad & \sum_1^n p_i(1 - B(p_i)) - q_iS(q_i) && (4.18) \\ \text{s.t.} \quad & \sum_1^n (1 - B(p_i)) \leq \sum_1^n S(q_i) && (4.19) \\ & p_i, q_i \in [0, 1] && (4.20) \end{aligned}$$

Here, q_i and p_i are the posted prices for sellers and buyers, respectively. It is easy to see that these are indeed the optimal posted prices, by definition – the constraints are exactly the constraints on feasibility (buying at most as much as selling), and the objective function gives the mechanism's revenue. The first constraint must be tight at the (an) optimum, as otherwise the mechanism is paying more than necessary to the sellers. (This follows since we are dealing with fractions, not expectations.) We can solve this optimisation through Lagrangian methods, where the Lagrangian is as follows:

$$\mathcal{L} = \left(\sum_1^n p_i(1 - B(p_i)) - q_i S(q_i) \right) - \lambda \cdot \left(\sum_1^n (1 - B(p_i)) - \sum_1^n S(q_i) \right)$$

We are interested in the points where its derivatives with respect to p_i vanish:

$$(1 - B(p_i)) - p_i b(p_i) = -\lambda b(p_i) \qquad \Leftrightarrow \qquad p_i - \frac{1 - B(p_i)}{b(p_i)} = \lambda$$

This has at most one solution for λ fixed by regularity of the distribution. Similarly the q_i will have a unique solution. Therefore all $p_i = p$ and all $q_i = q$ for global prices p, q. Dividing the objective by n gives us the stated result. □

The following lemma allows us to relate this to arbitrary sequences of n sellers and n buyers.

Lemma 5. *Amongst all sequences with n sellers and n buyers, $S^n B^n$ obtains the maximum optimal profit with fractional mechanisms.*

Proof. Let σ be an arbitrary sequence with n buyers and n sellers, and let q_i, p_i be the posted prices of an optimal fractional mechanism for this sequence. We have that the p_i and q_i have to satisfy the following condition:

$$\sum_1^n (1 - B(p_i)) \leq \sum_1^n S(q_i)$$

Otherwise, the mechanism would sell more than it buys. So the p_i, q_i satisfy the constraints on the prices given in the optimisation problem in 4.18 - 4.20. Since they are a therefore a feasible solution to that optimisation problem, the mechanism's revenue can clearly be at most as the solution to that optimisation problem, which is precisely our claim. □

Theorem 21. *For any sequence σ, the profit obtained by any adaptive posted price mechanism is at most as much as the profit of the optimal fractional mechanism.*

Proof. We will show that the revenue of an optimal adaptive mechanism is bounded from above by the revenue of the optimal fractional adaptive profit. Since in fractional mechanisms there is no difference between adaptive and non-adaptive, our claim follows.

We begin by analysing a fixed adaptive mechanism. For seller i, let Q_i denote the price posted to seller i, and $\tilde{Q}_i$ be the probability of a purchase happening at that price. In an adaptive mechanism, Q_i and $\tilde{Q}_i$ are random variables, depending on the previous history of trades. We do the same for P_j and $\tilde{P}_j$ for buyer j. We have the following (assuming strictly monotone cdf):

$$\mathbb{E}[Q_i S(Q_i)] = \mathbb{E}[S^{-1}(\tilde{Q}_i)\tilde{Q}_i]$$
$$\mathbb{E}[P_j(1 - B(P_j))] = \mathbb{E}[B^{-1}(1 - \tilde{P}_j)\tilde{P}_j]$$

Let X_i and Y_j be the payment given to or received by seller i or buyer j. The profit of the mechanism is $\sum_j Y_j - \sum_i X_i$. Thus, we have:

$$
\begin{aligned}
\mathbb{E}\left[\sum_{j=1}^{n} Y_j - \sum_{i=1}^{n} X_i\right] &= \sum_j \mathbb{E}\left[\mathbb{E}\left[Y_j | P_j\right]\right] - \sum_i \mathbb{E}\left[\mathbb{E}\left[X_i | Q_i\right]\right] && (4.21)\\
&= \sum_j \mathbb{E}\left[P_j\left(1 - B(P_j)\right)\right] - \sum_i \mathbb{E}\left[Q_i S(Q_i)\right] && (4.22)\\
&= \sum_j \mathbb{E}[B^{-1}(1 - \tilde{P}_j)\tilde{P}_j] - \sum_i \mathbb{E}[S^{-1}(\tilde{Q}_i)\tilde{Q}_i] && (4.23)\\
&\leq \sum_j \mathbb{E}[\tilde{P}_j] B^{-1}(1 - \mathbb{E}[\tilde{P}_j]) - \sum_i \mathbb{E}[\tilde{Q}_i] S^{-1}(\mathbb{E}[\tilde{Q}_i]), && (4.24)
\end{aligned}
$$

where the last inequality follows from the assumption that both distributions are regular. Observe that in the last line, $\mathbb{E}[\tilde{P}_j]$ and $E[\tilde{Q}_i]$ give the expected number of items bought, which would translate into fractions in a fractional mechanism; and $B^{-1}(1-\mathbb{E}[\tilde{P}_j])$ and $S^{-1}(\mathbb{E}[\tilde{Q}_i])$ could be interpreted as prices in a fractional mechanism.

So, this is essentially the objective of the fractional mechanism optimisation problem. Next, we show that the constraints translate as well. Let $\mathcal{S}_t$ be the (indices of the) sellers in the initial t-segment of the sequence, and similarly $\mathcal{B}_t$ the buyers. Let Z_t be the fraction if items bought from or sold to the trader at step t. The the current stock level of the intermediary is given by $\sum_{i \in \mathcal{S}_t} \mathbb{E}[Z_i] - \sum_{j \in \mathcal{B}_t} \mathbb{E}[Z_j]$, and is non-negative by the constraint that the intermediary cannot sell what they do not possess. Therefore, for any t, the following holds:

$$
0 \leq \sum_{i \in \mathcal{S}_t} \mathbb{E}[Z_i] - \sum_{j \in \mathcal{B}_t} \mathbb{E}[Z_j] \qquad (4.25)
$$

$$= \sum_{i \in \mathcal{S}_t} \mathbb{E}[\mathbb{E}[Z_i | \tilde{Q}_i]] - \sum_{j \in \mathcal{B}_t} \mathbb{E}[\mathbb{E}[Z_j | \tilde{P}_j]] \tag{4.26}$$

$$= \sum_{i \in \mathcal{S}_t} \mathbb{E}[\tilde{Q}_i] - \sum_{j \in \mathcal{B}_t} \mathbb{E}[\tilde{P}_j] \tag{4.27}$$

But this is again precisely the constraint 4.19, and so together with 4.24 we get the optimisation problem for the optimal fractional mechanism, which shows our claim. □

To complete the analysis, we will provide matching lower bounds from non-adaptive posted prices mechanisms that asymptotically obtain the same profit as the fractional mechanisms for balanced sequences of the form $S^n B^n$. Having that, we will show that posted price mechanisms can obtain a constant fraction of $S^n B^n$'s profit from any balanced sequence containing n buyers and sellers.

Fortunately, there is a simple non-adaptive posted-price mechanism which achieves asymptotically the same profit as the optimal fractional mechanism. As we will show, just setting the prices q and p to sellers and buyers respectively as computed in Theorem 20 is good enough.

Theorem 22. *The mechanism that buys and sells with the prices q and p as computed in Theorem 20 achieves profit asymptotically equal to the optimal fractional mechanism for the sequence $S^n B^n$.*

Proof. Let $Y \sim \mathrm{Bin}(n, S(q))$ be the number of items bought from S^n and let X be the number of items sold to B^n. The number of items sold is not a simple binomial distribution, but rather the minimum of two random variables: the number if items bought, and the number of items that could be sold if there was enough stock. That is, $X \sim \min\{Y, X'\}$, where $X' \sim \text{Bin}(n, 1 - B(p)) = \text{Bin}(n, S(q))$. (Due to the condition on p, q in Theorem 20.) The expected value of Y is easy enough to compute, however X is more tricky. A simple solution is to consider

$$\Pr\left[\left(\frac{1}{n}|Y - \mathbb{E}[Y]| \geq \epsilon\right) \wedge \left(\frac{1}{n}|X' - \mathbb{E}[X']| \geq \epsilon\right)\right] \leq 2\Pr\left[\frac{1}{n}|Y - \mathbb{E}[Y]| \geq \epsilon\right] \tag{4.28}$$

$$\leq 4e^{-2n\epsilon^2} \tag{4.29}$$

from a Hoeffding bound. Thus with probability at least $1 - 4e^{-2n\epsilon^2}$ we have that $Y \leq nS(q) + n\epsilon$ and both X' and $X \geq nS(q) - n\epsilon$. Taking a worst case approach, for the remaining $4e^{-2n\epsilon^2}$ we can lose at most n. So the expected profit is at least:

$$\mathbb{E}[X \cdot p - Y \cdot q] \geq \left(1 - 4e^{-2n\epsilon^2}\right) \cdot n\left((S(q) - \epsilon)p - (S(q) + \epsilon)q\right) - 4ne^{-2n\epsilon^2} \tag{4.30}$$

This holds for any $\epsilon > 0$. Setting $\epsilon = \frac{1}{\log n}$, as $n \to \infty$ the profit converges to

$$p(1 - B(p)) - qS(q) + O(\frac{n}{\log n}) = p(1 - B(p)) - qS(q) + o(n).$$

This is asymptotically the same as the optimal fractional mechanism. □

The following trivial claim (for balanced sequences) allows us to compare online non-adaptive and, therefore, offline non-adaptive mechanisms to adaptive mechanisms.

Lemma 6. *Every balanced sequence with n buyers and n sellers contains a* *(non-contiguous) subsequence of the form $S^k B^k$ for some $k \geq \lceil \frac{n}{2} \rceil$.*

Proof. There are at least $\lceil \frac{n}{2} \rceil$ sellers before the $\lceil \frac{n}{2} \rceil$th buyer, and exactly $\lceil \frac{n}{2} \rceil$ buyers after them (including the $\lceil \frac{n}{2} \rceil$th buyer themselves). □

We are now ready to define the online non-adaptive algorithm.

Online non-adaptive algorithm: The algorithm buys only from the first $\lceil \frac{n}{2} \rceil$ buyers at price q and sells at price p to the first $\lceil \frac{n}{2} \rceil$ buyers which have a matching number of sellers before them , where p and q are again as defined in in Theorem 20 for balanced sequences of n sellers and n buyers.

Theorem 23. *For every balanced sequence σ containing n buyers and n sellers, the above online non-adaptive algorithm obtains at least half the profit of the optimal adaptive offline algorithm.*

Proof. The proof is immediate: The profit of the online algorithm is at least half the optimal fractional profit of $S^n B^n$, which is at least equal to the optimal adaptive profit for any balanced sequence σ containing n buyers and n sellers. □

Note that the algorithm must know in advance the number n of buyers and sellers. If the number n of sellers and buyers is not known, the risk is that the algorithm will buy more items than it will manage to sell. Compare, for example, the sequences $(SB)^n$ and $S^n B^n$; the probability that a buyer will appear when there are no available items to be sold is higher in the first sequence than the second. An immediate follow-up question then is if the mechanism can be adapted to unknown n. From our analysis, it seems plausible that we could aim for a weaker version of Theorem 22, by ignoring a fraction of sellers to guarantee that there will be no large surplus of items at the end of the sequence. Indeed, Giannakopoulos et al. [63] show this for a more general setting.

Note also that we could use the same idea to improve the competitive ratio of the online algorithm but, in order to keep the presentation clear, we made no effort to optimize the ratios.

One final comment on the profit of adaptive mechanisms: Among all balanced sequences, the sequence that gives the maximum profit is not the sequence S^nB^n; intuitively, by moving some buyers earlier in the sequence, we obtain an improved profit by adapting the remaining buying prices to the outcome of these potential trades. For example, it should be intuitively clear that the sequence $S^{n/2}BS^{n/2}B^{n-1}$ has (slightly) better adaptive profit than the sequence S^nB^n for large n. Our work above shows that the difference is asymptotically insignificant, but it remains an intriguing question to determine the balanced sequence with the maximum profit.

4.5 Discussion

We have shown that even in a simple online intermediation setting, we cannot hope to achieve a constant competitive ratio against an offline adversary. On the other hand, we show that our $\sqrt{n}$ bound is tight, and a matching online algorithm that achieves this approximation ratio. More importantly, even for very mild, natural restrictions of either limited stock or balanced sequences, we can indeed achieve a constant approximation ratio with an online mechanism.

This illuminates a more general question inherent in any online intermediation setting: How much stock do you buy, not knowing future sales levels? Our results suggest that this question is answerable even in the worst-case. The proof techniques employed also highlight a fundamental distinction of the online setting compared to its offline counterpart, or indeed to overlapping arrivals and departures: We employ techniques from online algorithms rather than matchings (online or offline).

Many interesting open questions remain. First and foremost, we make relatively strong assumptions on the prior distributions. The eventually published paper [63], which contains significantly reworked but similar results to those in this chapter, gives some bounds for more general classes of distributions. A follow-up paper [75] considers the setting where the sequence is drawn from a distribution, but valuations are chosen adversarially.

A more general question would be to compare our results, which are developed within the tradition of competitive analysis, to bounds on the worst-case regret of an online (learning or adaptive) mechanism for such an intermediation setting.

Part 5

Multi-Unit Intermediation

5.1 Introduction

In the previous chapters, we have focused on the intermediary's revenue. In this final part, we will discuss social welfare, and return again to the offline setting. Welfare in the context of intermediation mechanisms is a very different beast from welfare in auctions. In the latter, we generally employ the VCG mechanism to pick the optimal outcome or allocation, and don't concern ourselves too much with the payments which are used only to implement the social choice function truthfully - that the auctioneer ends up with a sum of money at the end is either irrelevant or a nice bonus. In two-sided mechanisms, we cannot afford ourselves this luxury: Notice that payments are made from the buyer(s) to the intermediary, but *from the intermediary to the seller(s).* If we do not take care, those payments may easily exceed those from the buyers, meaning the intermediary will be left with a net loss. In many settings, that is infeasible or undesirable.

For this reason in the context of welfare maximisation with intermediation mechanisms, it is usually assumed that an additional **budget balance** constraint is imposed on the mechanism. In its less stringent form, weak budget balance requires that the intermediary never (ex post) is left with a monetary deficit, meaning that payments to sellers may never exceed payments from buyers. A more restrictive version, strong budget balance requires that the intermediary never ends up with any non-zero amount of money, or that the total net flow of payments is always zero. The latter condition is also motivated partly by practical concerns. For instance, a mechanism may leave the intermediary with payments equalling the entire increase in welfare. Clearly, this is not ideal, as it means that the net utility of players (other than the intermediary) has not changed at all.

In this chapter, then, we will require our usual notions of truthfulness and individual rationality, and in addition strong budget balance:

- *Incentive Compatibility ((DS)IC)*: It must be a dominant strategy for the agents (buyers and sellers) to behave truthfully, hence not "lie" about their valuations for the items in the market. This enables the market mechanism to make an informed decision about the trades to be made.

- *Individual Rationality (IR)*: It must not harm the utility of an agent to participate in the mechanism.

- *Strong Budget Balance (SBB)*: All monetary transfers that the mechanism executes are among participating agents only. That is, no money is injected into the market, and no money is burnt or transferred to any agent outside of the market.

Such mechanisms were first studied in the seminal paper by Myerson and Satterthwaite [99] and since then in various other publications in the economics literature. Unfortunately, this first foray into this area immediately delivered a strong impossibility result: No mechanism that satisfies the above conditions can implement the optimal social welfare exactly. (In fact, not even a weakly budget balanced mechanism could.)

This chapter studies a generalisation of the classical bilateral trade setting by allowing the seller to hold multiple units initially. These units are assumed to be of a single resource, so that agents only express valuations in terms of how many units they have in possession. The final utility of an agent (buyer or seller) is then determined by their valuation and the payment they paid or received. We focus our study on characterising which mechanisms satisfy the above three properties and which of these feasible mechanisms achieve a good social welfare (i.e., total utility of buyer and seller combined).

Due to its simplicity, our setting is fundamental to any strategic setting where items are to be redistributed or reallocated. Our characterisation efforts show that all feasible mechanisms must belong to a very restricted class, already for this very simple setting with one buyer, one seller, and a relatively simple valuation structure. The specific mechanisms we develop are very simple, and suitable for implementation with very little communication complexity.

Our Contribution

Our first main contribution is a full characterisation of the class of truthful, individually rational and strongly budget balanced mechanisms in this setting. We do this separately for two classes of valuation functions: submodular valuations and general non-decreasing valuations. Section 5.3 presents a high-level argument for the submodular case. A full and rigorous formal proof for both settings is given in Appendix 5.4. Essentially, for the general case, any mechanism that aims to be truthful, strongly budget balanced and individually rational can only allow the agents to trade a single quantity of items at a predetermined price. The trade then only occurs if both the seller and buyer agree to it. This leads to a very clean characterization and has the added benefit of giving a robust, simple to understand mechanism: the agents do not have to disclose their entire valuation to the mechanism, and only have to communicate whether they agree to trade one specific quantity at one specific price.

Secondly, we give approximation mechanisms for the social welfare objective in the Bayesian setting in Section 5.5, for the case of submodular valuations. Theorem 29 presents a 2-approximate deterministic mechanism. For randomised mechanisms, we show a $e/(e-1)$-approximation in Theorem 30.

Prior Work

The first approximation result for bilateral trade was presented in McAfee [93], where for the Bayesian single-item case the author proves that the optimal *gain from trade* can be 2-approximated by the *median mechanism*, which is a mechanism that sets the seller's median valuation as a fixed price for the item, and trade occurs if and only if p lies in between the buyer's and seller's valuation and the buyer's valuation exceeds p. The analysis in [93] is done under the assumption that the seller's median valuation does not exceed the median valuation of the buyer. The gain from trade is defined as the increase in social welfare as a result of trading the item. Blumrosen and Dobzinski [21] extended the analysis of this mechanism by showing that it also 2-approximates the social welfare without the latter assumption on the medians.

Blumrosen and Dobzinski [21] furthermore consider the classical bilateral trade setting (with a single item) and present various mechanisms for it that approximate the optimal social welfare. Their best mechanism achieves an approximation factor of $e/(e-1)$. As in the present chapter, there are prior distributions on the traders' valuations, and the quantity being approximated is the expectation over the priors, of the optimal allocation of the item.

The weaker notion of Bayesian incentive compatibility is considered by Blumrosen and Mizrahi [23], who propose a mechanism in which the seller offers a take-it-or-leave-it price to the buyer. They prove that this mechanism approximates the harder *gain from trade* objective within a factor of $1/e$ under a technical albeit often reasonable *monotone hazard rate condition* on the buyer's distribution.

The class of DSIC, IR, and SBB mechanisms for bilateral trade was characterised by Colini et al. [34] to be the class of *fixed price mechanisms*. In the present work, we characterise this set of mechanisms for the more general multi-unit bilateral trade setting, thereby extending their result.

Various recent papers analyse more general two-sided markets, where there are multiple buyers and sellers, who hold possibly complex valuations over the items in the market. Colini et al. [35] analyse a more general scenario with multiple buyers, sellers, and multiple distinct items, and use the same feasibility requirements as ours (DSIC, IR, and SBB). Segal-Halevi at al. [117] have considered a similar setting but focus on *gains from trade (GFT)* (i.e., the increase in social welfare resulting from reallocation of the items) instead of welfare. They initially considered a multi-unit setting like ours (albeit with multiple buyers and sellers), and they extend their work in [118] to allow multiple types of goods. They present a mechanism that approximates the optimal GFT asymptotically in large markets. Balseiro et al. [18] design two-sided market mechanisms for one seller and multiple buyers with a temporal component, where valuations are correlated between buyers but independent across time steps. A good approximation (of factor 1/2) of the social welfare using the more permissive notion of *Bayesian* Incentive Compatibility (BIC) was achieved by Brustle et al. [25]. Their optimality benchmark is different from the one we consider as they compare their mechanism to the best possible BIC, IR, and SBB mechanism. A very recent work by Babaioff et al. [8], proposes mechanisms that achieve social welfare guarantees for both optimality benchmarks. Feldman and Gonen [56] consider optimizing the gains from trade in a two-sided market setting tailored to online advertising platforms, and the authors extend this idea further in [55] by considering two-sided markets in an online setting.

5.2 Preliminaries

In a *multi-unit bilateral trade* instance there is a buyer and seller, where the seller holds a number of units of an item. This number will be denoted by k. The buyer and seller each have a *valuation function* representing how much they value having any number of units in possession. These valuation functions are denoted by v and w, respectively. Precisely stated, a valuation function is a function $v : [k] \cup \{0\} \to \mathbb{R}_{\geq 0}$ where $v(0) = 0$. Note that we use the standard notation $[a]$, for a natural number a, to denote the set $\{1, \dots, a\}$. We denote by v the valuation function of the buyer, drawn from f, and we denote by w the valuation function of the seller, drawn from g. For $q \in [k]$, the valuation $v(q)$ or $w(q)$ of an agent (i.e., buyer or seller) expresses in the form of a number the extent to which he would like to have q units in his possession.

A mechanism M interacts with the buyer and the seller and decides, based on this interaction, on an *outcome*. An outcome is defined as a quadruple (q_B, q_S, p_B, p_S), where q_B and q_S denote the numbers of items allocated to the buyer and the seller respectively, such that $q_B + q_S = k$. Moreover, p_B and p_S denote the payments that the mechanism charges to the buyer and seller respectively. Note that typically the payment of the seller is negative since he will get money in return for losing some items, while the payment of the buyer is positive since he will pay money in return for obtaining some items. Let Ω be the set of all outcomes. For brevity we will often refer to an outcome simply by the number of units traded q_B; and by slight abuse of notation we will sometimes use p to donate a constant per-unit price.

Formally, a mechanism is a function $M : \Sigma_B \times \Sigma_S \to \Omega$, where Σ_B and Σ_S denote strategy sets for the buyer and seller. A *direct revelation* mechanism is a mechanism for which Σ_B and Σ_S consists of the class of valuation functions that we want to consider. That is, in such mechanisms, the buyer and seller directly report their valuation function to the mechanism, and the mechanism decides an outcome based on these reports. We want to define our mechanism in such a way that there is a dominant strategy for the buyer and seller, under the assumption that their valuation functions are in a given class $\mathcal{V}$. It is well known (see e.g. [24]) that then we may restrict our attention to direct revelation mechanisms in which the dominant strategy for the buyer and seller is to report the valuation functions that they hold. Such mechanisms are called *dominant strategy incentive compatible (DSIC)* for $\mathcal{V}$. We will assume from now on that M is a direct revelation mechanism. In this chapter, we consider for $\mathcal{V}$ two natural classes of valuation functions:

- Monotonically increasing submodular functions, i.e., valuation functions v such that for all $x, y \in [k]$ where $x < y$ it holds that $v(x) - v(x-1) \geq v(y) - v(y-1)$ and $v(x) \leq v(y)$. This reflects a common phenomenon observed in many economic settings involving identical goods: Possessing more of a good is never undesirable, but the increase in valuation still goes down as the held amount increases. For a monotonically increasing submodular function v and number of units $x \in [k]$, we denote by $\tilde{v}(x)$ the *marginal* valuation $v(x) - v(x-1)$. Thus, it holds that $\tilde{v}(x) \geq \tilde{v}(y)$ when $x < y$.

- Monotonically increasing functions, i.e., valuation functions v such that

$v(x) < v(y)$ for all $x < y$, where $x, y \in [k]$.

Besides the DSIC requirement, there are various additional properties that we would like our mechanism to satisfy.

- Ideally, our mechanism should be *strongly budget balanced (SBB)*, which means that for any outcome (q_B, q_S, p_B, p_S) that the mechanism may output it holds that $p_B = -p_S$. This requirement essentially states that all money transferred is between the buyer and the seller only.

- Additionally, we want that running the mechanism never harms the buyer and the seller. This requirement is known as *(ex-post) individual rationality (IR)*. Note that when v and w are the valuation functions of the buyer and the seller, then the initial utility of the buyer is 0 and the initial utility of the seller is $w(k)$. Thus, a mechanism M is individually rational if for the outcome $M(v, w) = (q_B, q_S, p_B, p_S)$ it always holds that $v(q_B) - p_B \geq 0$ and $w(q_S) - p_S \geq w(k)$.

- We would like the mechanism to return an outcome for which the total utility is high. That is, we want the mechanism to maximise the sum of the buyer's and seller's utility, which is equivalent to maximizing the sum of valuations $v + w$ when strong budget balance holds.

We characterise in Section 5.3 the class of DSIC, SBB, IR mechanisms for both valuation classes. In Section 5.5, we subsequently provide various approximation results on the quality of the solution output by some of these mechanisms. For these results, we assume the standard *Bayesian setting*: The mechanism has no knowledge of the buyer's and seller's precise valuation, but knows that these valuations are drawn from known probability distributions over valuation functions. Our approximation results provide mechanisms that guarantee a certain outcome quality (which is measured in terms of *social welfare*, defined in Section 5.5) for arbitrary distributions on the valuation functions.

Formally, in the Bayesian setting, a multi-unit bilateral trade instance is a pair (f, g, k), where $k \in \mathbb{N}$ is the total number of units that the seller initially has in his possession, and f and g are probability distributions over valuation functions of the buyer and the seller respectively. Note that we do not impose any further assumptions on these probability distributions.

5.3 Characterisation

In [34] the authors prove that every DSIC, IR, SBB mechanism for classical bilateral trade (i.e. the case where $k = 1$) is a *fixed price mechanism*: That is, the mechanism is parametrised by a price $p \in \mathbb{R}_{\geq 0}$ such that the buyer and seller trade if and only if the buyer's valuation exceeds the price and the price exceeds the seller's valuation. Moreover, in case trade happens, the buyer pays p to the seller. In this chapter we characterise the set of DSIC, IR, and SBB mechanisms for multi-unit bilateral trade, and we thereby generalise the characterisation of [34].

Definition 22. *A* multi-unit fixed price mechanism *is a sequential posted price mechanism with a fixed per-unit price p, potentially with bundling. Such a mechanism iteratively proposes a quantity q of units to both the buyer and seller simultaneously, which the seller and buyer can choose to either accept or reject. If both agents accept, q additional units are reallocated from the* *seller to the buyer, the buyer pays pq to the seller, and the mechanism may then either proceed to the next iteration or terminate. If one of the two agents rejects, the mechanism terminates. Quantity q may vary among iterations, but must be pre-determined prior to execution of the mechanism.*

Theorem 24. *Any mechanism that satisfies DSIC, IR and SBB must be (equivalent to) a* multi-unit fixed price mechanism*.*

For increasing submodular valuations, any number of iterations is allowed. For general increasing valuations, the mechanism is further restricted to execute only one iteration (or equivalently, it may only offer one bundle for a fixed price).

In simple terms, our result states that for the submodular valuations case, the only thing to be done truthfully in this setting is to set a fixed per-unit price p, and ask the buyer and seller if they want to trade one or several units of the good at per-unit price p. This repeats until one agent rejects. In the general monotone case this is further restricted to a single such proposed trade. The following is a high-level argument of the proof of Theorem 24 for the submodular setting. We refer the reader to Appendix 5.4 for a full formal proof, which also includes the general case.

Lemma 7. *All prices must be fixed in advance, and cannot depend on the bid / valuation of neither the seller nor the buyer.*

Proof. This follows immediately from DSIC and SBB: By DSIC, for any outcome, the price charged to the buyer can't depend on the buyer's bid, otherwise one can construct scenarios in which the price charged to the buyer could be manipulated to the buyer's benefit by misreporting the bid. The same holds for the seller. By SBB the payment of the buyer completely determines the payment of the seller (the payment is simply negated) so neither payment can depend on either's bid. $\square$

Theorem 25. *Suppose in a DSIC, SBB, IR mechanism the price for the outcome in which q units are traded is qp for a fixed per-unit price for all potential outcomes. Then the allocation chosen for a given pair of valuation functions is the one arising when asking bidders sequentially if they want to trade one unit (or a bundle of units), until one rejects.*

Proof. To see this, consider the seller's utility function $u_s(q) = q \cdot p + w(k - q)$ and the buyer's utility function $u_b(q) = v(q) - q \cdot p$, if q units would be traded at unit price p. Since both valuation functions are concave, it is easy to see that both utility functions are concave, and each has a single peak (one or more equal adjacent maxima, and no further local maxima). Furthermore they both start at 0, and once either of them becomes negative, it stays negative. Suppose we sequentially ask both bidders if they want to trade one unit for price p, until one rejects. Then the quantity traded is $\min(\operatorname{argmax}(u_s), \operatorname{argmax}(u_b))$, i.e. the first of the two peaks. If the mechanism iteratively proposes them bundles $q_1, q_2, \ldots$, then the same expression on the traded quantity would apply, but with the utility functions restricted to the domain $\{0, q_1, q_1 + q_2, \ldots\}$. If we ask them about the big all-k-item bundle, we would choose the bundle outcome iff $u(k) > u(0)$, for both, and 0 if for either of them $u(0) > u(k)$, i.e. if one (the first) of the peaks of the two utility functions restricted to $\{0, k\}$ is at 0.

Now, DSIC means that for any bid of the opposing agent, the agent cannot get anything better than what she gets by telling the truth. If the quantity traded by the mechanism would be larger than $\min(\text{argmax}(u_s), \text{argmax}\,(u_b))$, then the bidder with the lowest peak could improve her utility by claiming that all outcomes higher than her peak are wholly unacceptable (utility less than 0) to them; by IR, the mechanism would then be limited to trading the quantity at the first peak. If, on the other hand, the traded quantity would be less than the quantity of the first peak, then both players would gain by lying, in order to make the mechanism choose to trade a higher quantity (if such a quantity is at all present in the mechanism's set of tradeable quantities.) □

Theorem 26. *In a DSIC, SBB, IR mechanism, all potential outcomes, i.e., (quantity,price)-pairs, must have the same per-unit price.*

Proof. Suppose two outcomes have different per-unit prices. W.l.o.g. suppose for $q_1 < q_2$, $p_1/q_1 < p_2/q_2$, i.e. the per-unit price is higher in the larger allocation. Then there exists a valuation function v_{s1} for the seller in which the seller prefers outcome q_2 over q_1, but both give positive utility; and there exists another valuation function v_{s2} that gives negative utility for q_1, but the same utility for q_2. I.e. $0 < u_{s1}(q_1) < u_{s2}(q_2)$ but $u_{s2}(q_1) < 0 < u_{s2}(q_2) = u_{s1}(q_2)$. Now if for a given buyer's valuation, the chosen outcome given v_{s1} is q_1, then the seller would have an incentive to misreport v_{s2}, making outcome q_1 unavailable to the mechanism due to IR, thus making it choose q_2. Vice versa, if per-unit prices are decreasing, the same argument works for the buyer. □

Together, these three results give a full characterisation of the class of

DSIC, IR, SBB mechanisms in this setting. We give a full proof in the following section, including also the case of general monotone valuations.

5.4 Detailed Proof of Theorem 24

We denote the class of monotonically increasing submodular functions with domain $[k]$ by $\mathcal{S}_k$. We denote the class of monotonically increasing functions with domain $[k]$ by $\mathcal{I}_k$.

The definition below defines multi-unit fixed price mechanisms as a direct revelation mechanism. From the point of view of providing a rigorous proof, this is more convenient to work with than the sequential posted price definition given in the shorter version.

Definition 23. *Let $p \in \mathbb{R}_{\geq 0}$, let $S \subseteq [k]$, and let $\tau = (\tau_B, \tau_S, \tau_\cap)$ be a vector of three tie-breaking functions specified below. The* multi-unit fixed price mechanism *$M_{p,S,\tau}$ is the direct revelation mechanism that returns for a multi-unit bilateral trade instance (f, g, k) an outcome $M_{p,S,\tau}(v, w) = (q_B, q_S, p_B, p_S)$ on reported valuation functions v and w, where*

- $\tau_B(v) \subseteq \arg_q \max\{v(q) - qp : q \in S \cup \{0\}\}$ *and* $\tau_B(v) \neq \varnothing$,
- $\tau_S(w) \subseteq \arg_q \max\{w(k - q) + qp : q \in S \cup \{0\}\}$ *and* $\tau_S(w) \neq \varnothing$,
- $\tau_\cap(v, w)$ *is a tie-breaking function that selects an element in* $\tau_B(v) \cap \tau_S(w)$ *in case this intersection is non-empty,*
- $q_B = k - q_S = \begin{cases} \min\{\max \tau_B(v), \max \tau_S(w)\} & \text{if } d_B \cap d_S = \varnothing, \\ \tau_\cap(v, w) & \text{otherwise.} \end{cases}$,

- $p_B = -p_S = q_B p$.

Informally stated, the mechanism offers the buyer and seller a fixed unit price p and a set of quantities S. It then asks the buyer and seller which quantity in $S \cup \{0\}$ they would like to trade when for each unit the buyer would pay p to the seller. The mechanism then makes the buyer and seller trade the minimum of these two demanded numbers at a unit price of p. Typically the preferred quantity is unique for both the buyer and the seller, but in case of indifferences the buyer and seller will specify a set of multiple preferred quantities. In such cases, the tie-breaking functions τ_B, τ_S determine which quantities among the sets of indifferences are considered for trade, and the tie-breaking function $\tau_\cap$ is finally used to determine the traded quantity in case the sets selected by τ_B and τ_S intersect. Otherwise, the minimum of the maximum quantities of τ_B and τ_S is traded.

It turns out that multi-unit fixed price mechanisms characterise the set of all DSIC, IR, and SBB mechanisms with respect to monotonically increasing submodular valuation functions. Moreover, with the additional restriction that S is a singleton set, they characterise the set of all DSIC, IR, and SBB mechanisms with respect to monotonically valuation functions.

We first prove sufficiency.

Theorem 27. *For all $p \in \mathbb{R}_{\geq 0}$ and $S \subseteq [k]$, the mechanism is IR, SBB, and DSIC with respect to the class of monotonically increasing submodular valuation functions. Moreover, if $|S| = 1$, then $M_{p,S,\tau}$ is IR, SBB, and DSIC with respect to the class of monotonically increasing valuation functions.*

Proof. First we prove the statement for the class of monotonically increasing submodular valuation functions. The SBB property holds trivially by definition of the mechanism, $p_B = -p_S$.

Let v and w be increasing submodular valuation functions of the buyer and seller. Let q_B and p_B be the quantity given to the buyer and payment made by the buyer under the outcome $M_{p,S,\tau}(v, w)$. If $\tau_B(v) \cap \tau_S(w)$ is non-empty, then the function $\tau_\cap$ selects a utility maximizing quantity for both the buyer and seller, so IR obviously holds in that case. If $\tau_B(v) \cap \tau_S(w) = \varnothing$, the mechanism $M_{p,S,\tau}$ is IR for the buyer: his utility is $v(q_B) - p_B = v(q_B) - q_B p = v(\min\{\max \tau_B(v), \max \tau_S(w)\}) - \min\{\max \tau_B(v), \max \tau_S(w)\}p$. The value $\max \tau_B(v)$ is defined as a utility-maximizing quantity in S for the buyer, given that the buyer pays p for each unit. If the buyer's valuation function v is submodular, getting any quantity less than $\max \tau_B(v)$ at a price of p per unit will yield the buyer a non-negative utility. Therefore, the buyer's utility is non-negative. For the seller, the argument to establish the IR property is similar: His utility is $w(k - q_B) + p_B = w(k - q_B) + q_B p = w(k - \min\{\max \tau_B(v), \max \tau_S(w)\}) + \min\{\max \tau_B(v), \max \tau_S(w)\}p$. The value $\max \tau_S(w)$ is defined as a utility maximizing quantity in S for the seller to give to the buyer, given that the seller receives a payment of p for each unit. As the buyer's valuation function v is submodular, giving any quantity less than $\max \tau_S(w)$ to the buyer at a price of p per unit will yield the seller a non-negative increase utility. Therefore, the seller's utility increase is non-negative.

For the DSIC property, observe that if the mechanism sets $q_B \in \tau_S(w)$, then the mechanism chooses the outcome that is the utility-maximizing one for the seller among all outcomes in the range of the mechanism. On the other hand, if $q_B \in \tau_B(v) \setminus \tau_S(w)$ then the seller can only manipulate the outcome by misreporting a valuation that causes q_B to attain a smaller value, and hence in this case the mechanism will select an outcome where a smaller quantity is traded against a price of p per unit. By increasingness and submodularity of the seller's valuation function, this will result in a lower utility for the seller. Hence, it is a dominant strategy for the seller to not misreport his valuation function. For the buyer, the argument is similar: If the mechanism sets $q_B \in \tau_B(v)$, then the mechanism chooses the outcome that is the utility-maximizing one for the buyer among all outcomes in the range of the mechanism. On the other hand, if $q_B \in \tau_S(w) \setminus \tau_B(v)$ then the buyer can only manipulate the outcome by misreporting a valuation that causes q_B to attain a smaller value, and hence in this case the mechanism will select an outcome where a smaller quantity is traded against a price of p per unit. By increasingness and submodularity of the buyer's valuation function, this will result in a lower utility for the buyer. Hence, it is a dominant strategy for the buyer to not misreport his valuation function.

Next, we prove the statement for the larger class of monotonically increasing valuation functions. Again, the SBB property holds trivially.

As we now work under the assumption that $|S| = 1$, let q be the quantity such that $S = \{q\}$. Let v and w be increasing valuation functions for the buyer and seller respectively. By definition of the mechanism and the increasingness of the valuation functions, it holds that $\tau_B(v) \in \{\{q\}, \{0\}, \{q, 0\}\}$. Likewise, $\tau_S(w) \in \{\{q\}, \{0\}, \{q, 0\}\}$. Therefore, for both the buyer and seller, the traded quantity is 0 or the unique positive quantity q in case he prefers trading q units at least as much as trading 0 units. Hence the buyer and seller both experience a non-negative increase in utility for the outcome decided by the mechanism. This establishes IR. For DSIC, observe that if a positive quantity is traded in the selected outcome under truthful reporting, then the only effect that misreporting can achieve is that a quantity of 0 at a price of 0 is traded instead, which would leave both the buyer and the seller with a 0 increase in utility, hence this will not increase either player's utility. If on the other hand a quantity of 0 is traded at a price of 0, then $\{0\} = \tau_B(v), 0 \notin \tau_S(w)$ or $\{0\} = \tau_S(w), 0 \notin \tau_B(v)$ or $0 \in \tau_B(v), 0 \in \tau_S(w)$. In the first case, clearly the buyer is not incentivised to manipulate the mechanism into producing the alternative outcome where q units are traded, and the seller is unable to manipulate the mechanism into producing that outcome as it selects the minimum of $\tau_S(w)$ and $\tau_B(w)$, where the latter equals $\{0\}$ regardless of the sellers report. For the second case, symmetric reasoning can be applied to conclude that none of the two agents are incentivised to misreport. For the third case, it trivially holds that none of the agents are incentivised to manipulate the mechanism into trading q instead of 0 units. This establishes DSIC. □

Next, we show necessity, i.e., all DSIC, IR, and SBB direct revelation mechanisms are multi-unit fixed price mechanisms.

Theorem 28. *Let M be a multi-unit bilateral trade mechanism that is IR, SBB, and DSIC with respect to the class of monotonically increasing submodular valuation functions. Then, there exist $p \in \mathbb{R}_{\geq 0}$, $S \subseteq [k]$, and τ such that $M = M_{p,S,\tau}$. Moreover, if M is also IR, SBB, and DSIC with respect to the* *bigger class of monotonically increasing valuation functions, then* $|S| = 1$.

We divide this proof up into several lemmas. We start by proving the theorem for the smaller class of monotonically increasing submodular valuation functions. First, we show that whenever M trades the same number of items for two distinct pairs of valuation functions, then it must charge the same payments. Second, we extend this by showing that whenever the mechanism trades distinct numbers of items for any two distinct pairs of valuation functions, then the mechanism must charge the same price proportional to the number of items traded. It follows that we may associate to M a unit price p such that the payment from the buyer to the seller is always $q_B p$, where q_B is the traded quantity. Lastly, we show that there is a set S such that the range of quantities that the seller may let the mechanism trade from (by means of reporting a valuation function to the mechanism), is equal to $(S \cup \{0\}) \cap [\arg_q \max\{v(q) - pq : q \in S \cup \{0\}\}]$. By the fact that the valuation functions are increasing and submodular, and by the fact that M is DSIC, it follows that truthful reporting of the seller will result in the mechanism trading

$$\arg_q \max\{w(k-q) + pq : q \in (S \cup \{0\}) \cap [\arg_q \max\{v(q) - pq : q \in S \cup \{0\}\}]\}$$

units. This expression is equal to $\min\{\max d_B, \max d_S\}$ if $d_B \cap d_S = \varnothing$ (where d_B and d_S are defined as in Definition 23) , and otherwise it is a set from which an arbitrary quantity $\tau(v, w)$ may selected. This implies that $M = M_{p,S,\tau}$ for the appropriate choices of p, S, and τ.

With respect to the larger class of monotonically increasing valuation functions, the set of DSIC, IR, and SBB mechanisms must be smaller. We prove for this class that whenever the set S consists of more than one quantity, then there must be a pair of valuation functions in which either the buyer or seller is better off by not truthfully reporting his valuation function.

We now proceed by stating and proving formally the claims sketched above. In the proofs of the claims below, we use the following terminology and notation. For ease of exposition, we denote from now on an outcome by a pair (q, p) where q is the traded number of units (i.e., the quantity that the buyer gets assigned) and p is the payment of the buyer, which is equal to the negated payment of the seller by the SBB property. For a reported valuation v of the buyer, let $M_v = \{\omega \in \Omega \mid \exists w : M(v, w) = \omega\}$ be the menu of outcomes offered to the seller when the buyer reports v. That is, when the buyer reports v, the seller can select one of the outcomes ω in M_v by reporting (not necessarily truthfully) some valuation in reply to v. Likewise, we let $N_w = \{\omega \in \Omega \mid \exists v : M(v, w) = \omega\}$ be the menu of outcomes offered to the buyer when the seller reports v. We let $M = \bigcup_v M_v = \bigcup_w N_w$ be the set of all outcomes that the mechanism can produce, and we let S be the projection of M on the quantity obtainable by the buyer (i.e., the set S consists of all quantities that the mechanism can possibly trade).

The next lemma shows that there is a unique payment that the mechanism charges for every quantity in S, which implies that M consists of at most k outcomes.

Lemma 8. *Let M be IR, SBB, and DSIC with respect to the class of valuation functions $\mathcal{C}$, where $\mathcal{C}$ is either $\mathcal{S}_k$ or $\mathcal{I}_k$. Let (v, w) and (v', w') be two pairs in $\mathcal{C}^2$. Let $M(v, w) = (q_B, p_B)$ and $M(v', w') = (q'_B, p'_B)$. If $q_B = q'_B$,* *then* $p_B = p'_B$.

Proof. As M is DSIC, it is immediate that for every v'' and for every quantity q it holds that there are no two distinct payments p, p' such that (q, p) and (q, p') are both in $M_{v''}$. Also, let (q, p) and (q', p') be in $M_{v''}$, where $q < q'$. Then $p \leq p'$, as otherwise there are valuations of the seller where misreporting results in trading less items at a higher price, which would violate the DSIC property.

Let (v, w) and (v', w') be as in the statement of the lemma, i.e., such that $q_B = q'_B$. When the buyer reports v and the seller reports w, by assumption (q_B, p_B) is the outcome, so $(q_B, p_B) \in M_v$. Define the valuation function w^* as the function that grows linearly, extremely steeply up to the quantity $k - q_B = k - q'_B$, and grows extremely slowly at a rate of $\epsilon > 0$ from $k - q_B$ onward. We define the function v^* similarly: It grows at an extremely high rate up to the quantity $q_B = q'_B$ and grows at extremely slow rate ϵ from q'_B onward.

We first consider a deviation by the seller from (v, w) to (v, w^*). Let $(q, p) \in M_v$ be the outcome of the mechanism on report (v, w^*). As we have chosen the valuation w^* to be sufficiently steep up to $k - q_B$ items, IR would be violated for a seller with valuation w^* if more than q_B items are traded.

Suppose now that upon report (v, w^*) the mechanism trades strictly less than q_B items. i.e., $q < q_B$. We prove that then, $p = p_B$: If we would assume $p < p_B$, then a seller with valuation w^* would misreport w, as in terms of valuation he is practically indifferent between trading q and q_B items (ϵ needs to be chosen small enough for this), and his received payment would increase from p to p_B. If on the other hand we would assume that $p > p_B$, then a seller with valuation w would misreport w^* as he would retain more items, and receive a higher payment. Thus $p = p_B$.

By entirely analogous reasoning, when (v', w^*) is reported to the mechanism, the mechanism also trades $q'_B = q_B$ items or less. Let $q' < q_B$ be the number of traded items under (v', w^*). The payment is equal to p'_B, and if less than q_B items are traded, then $w'(k - q') = w'(k - q_B)$.

Next, we define from w^* a valuation function w^{**} for which it holds that under (v, w^{**}) and (v', w^{**}) the same number of items is traded at prices p_B and p'_B respectively. We do this as follows: If $q' = q$ then we simply let w^{**} be w^*. Otherwise, if $q' \neq q$, assume without loss of generality that $q' > q$ and let $\bar{w}^*$ be the valuation function that grows extremely steeply up to $k - q$ units and increases extremely slowly after $k-q$ units. By considering the deviation by the seller from profile (v, w^*) to profile $(v, \bar{w}^*)$, we see that under $(v, \bar{w}^*)$ at most q' units are traded and the payment is still equal to q_B. Likewise, under $(v', \bar{w}^*)$ at most q' items are traded and the payment is q'_B. Thus, the minimum number of traded items among the pair of strategy profiles $(v, \bar{w}^*)$ and $(v', \bar{w}^*)$ is larger than the minimum number of traded items among the pair $(v, \bar{w}^*)$ and $(v', \bar{w}^*)$. Repeating this operation will thus eventually yield a strategy profile w^{**} such that under (v, w^{**}) and (v', w^{**}) the same number of items is traded at prices p_B and p'_B respectively.

Now, if we would suppose for contradiction that $p_B \neq p'_B$, then we may assume without loss of generality that $p_B < p'_B$. When the seller reports w^{**}, a buyer with valuation function v would now be incentivised to report the valuation function v' instead of v, since then his payment decreases, and he still receives the same number of items. This is a contradiction to the DSIC property. Therefore, $p_B = p'_B$ which proves our claim. $\square$

By Lemma 8 there is a unique payment for each quantity $q \in S$, and we denote this payment by $p(q)$. The next lemma extends the previous lemma by stating essentially that payments must grow linearly with the number of allocated items, when one of the players changes his reported valuation.

Lemma 9. *Let M be IR, SBB, and DSIC with respect to the class $\mathcal{C}$, where $\mathcal{C}$ is either $\mathcal{S}_k$ or $\mathcal{I}_k$. Let (v, w) and (v', w') be two pairs in $\mathcal{C}^2$. Let $M(v, w) = (q_B, q_S, p_B, p_S)$ and $M(v', w') = (q'_B, q'_S, p'_B, p'_S)$. If $q_B > 0$, then $p'_B = (q'_B/q_B)p_B$.*

Proof. First we show that $p(\cdot)$ is a non-decreasing function. Suppose that this is not true, and assume that $(q, p(q))$ and $(q', p(q))$ is the pair of outcomes in M such that (i) $q' > q$ and $p(q) > p(q')$, (ii) there is no $q'' \in S$ such that $q' > q'' > q$ and (iii) q is minimal. Let (v, w) and (v', w') be two valuation profiles that result in these two respective outcomes $(q, p(q))$ and $(q', p(q))$. Let w^* be a valuation function that increases linearly at an extremely high rate up to $k - q'$ and increses extremely slowly afterward. We now see that when (v, w^*) is reported, q' units or less are traded due to IR, and in fact at most q units are traded due to DSIC (because, if q' units were traded, a seller with valuation w^* would misreport w to trade less units for more money), and no less than q'' units are traded due to DSIC, where $q'' \leq q$ is the least quantity such that $p(q'') = p(q)$. (Otherwise the seller with valuation w^* could misreport w and trade more units at almost the same valuation, for significantly less money. We use here that w^* increases sufficiently slowly on the interval $[k - q', k]$). We also observe that when (v', w^*) is reported,

(i) q' units or less are traded due to IR, (ii) the traded quantity cannot be any quantity with a higher payment than $p(q')$ since otherwise a seller with valuation w' would misreport w^* if the buyer reports v', and (iii) the traded quantity cannot be any quantity with a lower payment than $p(q')$ since otherwise a seller with valuation w^* would misreport w'. Thus, exactly q' units are traded when (v', w^*) is reported. We conclude that $(q, p(q))$ and $(q', p(q'))$ are both in N_{w^*}, and therefore a buyer with valuation v would have an incentive to misreport v' if the seller reports w^*, which violates DSIC and yields a contradiction. We conclude that the payment function $p(\cdot)$ is non-decreasing.

Also, note that $0 \in S$ and $p(0) = 0$ as otherwise the IR property would be violated for a buyer whose valuation is identically 0 and a seller whose valuation is strictly increasing: When such a valuation profile is reported, the seller's valuation implies that no positive number of items can be traded for payment 0; the buyer's valuation implies that no positive number of items can be traded for a positive payment; so 0 items must be traded for a payment that is not positive (due to IR of the buyer) and not negative (due to IR of the seller).

The claim of this lemma is equivalent to the claim that there exists a unique value p such that $p(q) = pq$ for all $q \in S$. We will demonstrate this by means of contradiction: Suppose that there is no such value p. Let q be the lowest quantity in S such that $p(q'') \neq p(q)q''/q$ for all $q'' \in S, q'' < q$. Let q' be the highest quantity in S such that $p(q'') = p(q')q''/q'$ for all $q'' \in S, q'' < q'$. Note that there is no $q'' \in S$ such that $q' < q'' < q$, and that $p(\cdot)$ behaves linearly up to q', and that q essentially serves as the least witness for the non-linearity of $p(\cdot)$.

We distinguish two cases: The case where $p(q)q'/q > p(q')$, and the case where $p(q)q'/q < p(q')$. First let us assume that $p(q)q'/q > p(q')$. We will derive a contradiction by constructing valuation functions (v^*, w^*) for the buyer and seller such that the following properties are satisfied: (i) outcomes $(q, p(q))$ and $(q', p(q'))$ are in M_{v^*} and N_{w^*}; (ii) the buyer with valuation v^* strictly prefers outcome $(q', p(q'))$ over all outcomes in $N_{w^*} \setminus \{(q', p(q'))\}$; and (iii) the seller with valuation w^* strictly prefers outcome $(q, p(q))$ over all outcomes in $M_{v^*} \setminus \{(q, p(q))\}$. This is a contradiction because (ii) requires that the mechanism outputs $(q', p(q'))$ (otherwise the buyer with valuation v^* would be incentivised to misreport so that $(q', p(q'))$ is output) and (iii) requires that the mechanism outputs $(q, p(q))$ (otherwise the seller with valuation w^* would be incentivised to misreport so that $(q, p(q))$ is output).

Therefore, we will now define the appropriate valuations v^* and w^*. Let v^* be a valuation function that grows linearly at an extremely high rate up to quantity q' and increases extremely slowly afterward. This causes all outcomes where a positive quantity is traded a positive utility for the buyer with valuation v^*, moreover, the maximum utility for such a buyer is achieved at outcome $(q', p(q'))$. This already establishes property (ii). To see that property (i) holds, let w be any seller's valuation so that $(q', p(q')) \in N_w$ (which must exist because $q' \in S$). By the definition of v^*, the mechanism selects outcome $(q', p(q'))$ on report (v^*, w) and therefore $(q', p(q')) \in M_{v^*}$. Now, consider any valuation profile (v, w) that results in outcome $(q, p(q))$, so that $(q, p(q)) \in M_v$. Let w' be the valuation function that grows linearly at an extremely high rate up to quantity $k - q$ and after that point grows linearly at a rate of $p(q)/q - \epsilon/q$ up to quantity k. The initial increase up to point $k - q$ is so steep that the seller can never experience a utility above $w'(k)$ when any quantity higher than q is traded. The value $\epsilon > 0$ is chosen to be so small that the only outcome at which a seller with valuation w' has a positive utility is $(q, p(q))$. Therefore, upon report (v, w') the mechanism outputs $(q, p(q))$ and we may infer that $N_{w'} = \{(q, p(q)), (0, 0)\}$, from which it follows that $(q, p(q)) \in M_{v^*}$, as $(q, p(q))$ must be the selected outcome upon report (v^*, w) (by the DSIC property). This establishes property (i) for v^*.

For valuation function w^*, let w^* increase linearly at an extremely high rate up to quantity $k - q$, and increase extremely slowly afterwards. Clearly, the seller with valuation w^* prefers the outcome $(q, p(q))$ among all outcomes in M, which establishes property (iii). Let (v, w) be any report upon which the mechanism outputs $(q, p(q))$, so that $(q, p(q)) \in M_v$. Then $(q, p(q))$ is also output upon report (v, w^*) which establishes that $(q, p(q)) \in N_{w^*}$. Next, let (v, w) be any report upon which the mechanism outputs $(q', p(q'))$, so that $(q', p(q')) \in N_w$. Let v' be a function that increases at rate $p(q')/q' + \epsilon/q'$ for sufficiently small $\epsilon > 0$ up to quantity q', and increases extremely slowly afterward. Then $(q', p(q'))$ is output when (v', w) is reported, so that $(q', p(q')) \in M_{v'}$. Moreover, trading any quantity higher than q' would yield a negative utility for a buyer with valuation v' (because ϵ is extremely small). Therefore, for $q'' > q'$ it holds that $(q'', p(q'')) \notin M_{v'}$ and in particular $(q, p(q)) \notin M_{v'}$. Thus, when (v', w^*) is reported, the outcome $(q', p(q'))$ is output by the mechanism, and this establishes that $(q', p(q')) \in N_{w^*}$. Note that here we need that $p(q') > 0$, which is the case as the outcome the mechanism returns on (v', w') is IR by assumption and the valuation functions are monotonically increasing. This completes the proof for the case where $p(q)q'/q > p(q')$.

For the case where $p(q)q'/q < p(q')$ we proceed in a similar fashion: Again, we will derive a contradiction by constructing valuation functions (v^*, w^*) for the buyer and seller such that (i) outcomes $(q, p(q))$ and $(q', p(q'))$ are in M_{v^*} and N_{w^*}; (ii) the buyer with valuation v^* strictly prefers outcome $(q, p(q))$ over all options in $N_{w^*} \setminus \{(q, p(q))\}$; and (iii) the seller with valuation w^* strictly prefers outcome $(q', p(q'))$ over all options in $M_{v^*} \setminus \{(q, p(q))\}$. This is a contradiction because (ii) requires that the mechanism outputs $(q, p(q))$ (otherwise the buyer with valuation v^* would be incentivised to misreport so that $(q, p(q))$ is output) and (iii) requires that the mechanism outputs $(q', p(q'))$ (otherwise the seller with valuation w^* would be incentivised to misreport so that $(q', p(q'))$ is output). Note that the difference with the previous case is that here we construct v^* such that the higher of the two quantities q and q' is preferred, instead of the lower one. Likewise, w^* is now constructed such that the lower of the two quantities is preferred instead of the higher one.

We start in this case with the construction of w^*. Let w^* be a valuation function that increases linearly at an extremely high rate up to quantity $k-q$. From $k-q$ to $k-q'$, valuation w^* increases by an amount of $p(q)-p(q')+\epsilon$, where $\epsilon > 0$ is sufficiently small, and w^* increases extremely slowly from $k-q'$ onward. The increase in valuation from quantities $k-q$ to $k-q'$ is slightly higher than the amount by which the payment changes among the quantities q and q', this causes the seller with valuation w^* to encounter a slightly lower (but positive) increase in utility when quantity q is traded instead of quantity q'. Moreover, among all quantities in S up to q', the maximum utility for a seller with valuation w^* is achieved at quantity q', which already establishes property (iii). Lastly, note that due to the extreme steepness of w^* up to $k-q$, the utility of the seller is lower than $w^*(k)$ when any quantity higher than q is traded, so the mechanism will never do so by the IR constraint. It remains to establish property (i). Let (v, w) be any report where the mechanism selects outcome $(q', p(q'))$. It follows by DSIC that outcome $(q', p(q'))$ will also be selected on report (v, w^*), so that $(q', p(q')) \in N_{w^*}$. Next, let (v, w) be any report where the mechanism selects outcome $(q, p(q))$. Let w' be a valuation function that increases extremely steeply up to $k-q$ and increases extremely slowly afterwards, so that the report (v, w') results in $(q, p(q))$ and hence $(q, p(q)) \in N_{w'}$, and because of IR we also infer that $(q'', p(q'')) \notin M_{w'}$ when $q'' > q$. Let v' be a valuation function that increases linearly up to quantity q and increases extremely slowly afterwards, where $v(q) = p(q) + \epsilon$, and $\epsilon > 0$ is sufficiently small. Note that trading a positive quantity lower than q would result in a negative utility for a buyer with valuation v', so that such outcomes are not in $M_{v'}$. Therefore, when (v', w') is reported the outcome selected by the mechanism must be $(q, p(q))$, which shows that $(q, p(q)) \in M_{v'}$. It follows now that the selected outcome upon report (v', w^*) must be $(q, p(q))$ which yields $(q, p(q)) \in N_{w^*}$ and establishes property (i) for w^*.

Lastly, we design v^*. Let v^* simply increase extremely steeply up to the quantity q, and increase extremely slowly afterwards, so that a buyer with valuation v^* experiences positive utility for all outcomes in M, and maximum utility when outcome $(q, p(q))$ is selected. This straightforwardly establishes property (ii). For property (i), let (v, w) be any profile where outcome $(q, p(q))$ results, so that $(q, p(q)) \in N_w$. By DSIC, outcome $(q, p(q))$

is also selected when (v^*, w) is reported, so $(q, p(q)) \in M_{v^*}$. Next, let (v, w) be any profile where outcome $(q', p(q'))$ results, so $(q', p(q')) \in M_v$. Let w' be a function that increases extremely steeply up to quantity $k - q'$, and increases extremely slowly afterwards, so that trading any quantity higher than q' would result in a decrease in utility for a seller with valuation w' (hence the mechanism cannot trade such quantities when w' is reported, by the IR property), and the maximum increase in utility is achieved when $(q', p(q'))$ is chosen. Therefore reporting (v, w') results in outcome $(q', p(q'))$, thus $(q', p(q')) \in N_{w'}$ and $(q'', p(q'')) \notin N_{w'}$ for all $q'' > q'$. Therefore, when (v^*, w') is reported, outcome $(q', p(q'))$ is selected, which establishes property (i) for v^* and completes the proof for the case $p(q)q'/q > p(q')$.

$\square$

Let M be IR, SBB, and DSIC with respect to the class of monotonically increasing submodular valuation functions. From the above it follows that for a mechanism that is IR, SBB, and DSIC with respect to $\mathcal{S}_k$ or $\mathcal{I}_k$, there exists a price $p \in \mathbb{R}_{\geq 0}$ such that for all pairs (v, w) of monotonically increasing submodular valuation functions, the payment charged to the buyer is $q_B p$ (and the payment charged to the seller is $-q_B m$ by SBB). We will refer to p as the *unit price*.

The above corollary establishes the needed properties on the payments of the mechanism. The remaining lemmas use Lemma 9 by implicitly assuming the existence of the unit price p in their statement, and they characterise the quantities S tradable by the mechanism and the quantities that appear in the menus M_v and N_w. The next lemma states that the utility maximizing outcome in S for a buyer with any valuation function v is always in M_v.

Lemma 10. *If M is SBB, IR, and DSIC with respect to $\mathcal{S}_k$, and suppose that unit price p is positive. Then, for all $v \in \mathcal{S}_k$ it holds that $(q, p(q)) \in M_v$ for the lowest q in the set $\arg_q \max\{v(q) - p(q) : q \in S\}$.*

Proof. Let q be the lowest quantity in $\arg_q \max\{v(q) - p(q) : q \in S\}$. Let (v', w') be any report that results in outcome $(q, p(q))$, so that $(q, p(q)) \in N_{v'}$. Let w^* be a valuation function that increases extremely steeply up to the quantity $k-q$ and increases extremely slowly afterwards. Observe that by our assumption that $p > 0$, a seller with valuation w^* strongly prefers outcome $(q, p(q))$ over all other outcomes in S, and trading any quantity larger than q would violate IR. by DSIC, outcome $(q, p(q))$ is thus selected when (v', w^*) is reported, hence $(q, p(q)) \in N_{w^*}$ and $(q', p(q')) \notin N_{w^*}$ for all $q' > q$. So when (v, w^*) is reported, an outcome is selected from N_{w^*} that maximises the utility of the buyer with valuation v, and this outcome is $(q, p(q))$.

□

The following lemma strengthens the previous.

Lemma 11. *Suppose M is SBB, IR, and DSIC with respect to $\mathcal{S}_k$ and suppose that unit price p is positive. Let $v \in \mathcal{S}_k$ and let $q \leq \min \arg_{q'} \max\{v(q') - p(q') : q' \in S\}$ be a quantity not exceeding the least utility-maximizing quantity for a buyer with valuation v. It holds that $\arg_{q'} \max\{v(q') - p(q') : q' \in S, q' \leq q\}$ is the singleton set containing the quantity $q' = \max S \cap [q]$, and that $q \in M_v$.*

Proof. Note that the existence of the unit price p implies that the utility function of the buyer is a submodular function of the traded quantity. Therefore, the utility function for a buyer with valuation v is increasing up to the least utility-maximizing outcome in S, after which it stays constant up to the highest utility-maximizing outcome in S, after which it starts decreasing. Let q be any quantity less than or equal to the least utility-maximizing quantity, i.e., less than $\min \arg\max\{v(q') - p(q') : q' \in S\}$. Then, the utility-maximizing quantity q' in $S \cap [q]$ for a buyer with valuation v is $\max S \cap [q]$. It remains to prove that $(q', p(q'))$ is in M_v. Let (v', w') be any report resulting in outcome $(q', p(q'))$, so that $(q', p(q')) \in N_{v'}$. Let w'' be any function increasing extremely steeply up to quantity $k - q'$, after which it increases extremely slowly. Then $(q', p(q'))$ is the result of report (v', w''), and note that it is not IR to trade a quantity exceeding q' when w'' is reported, so $(q'', p(q'')) \notin N_{w''}$ for $q'' > q'$, and $(q', p(q')) \in N_{w''}$. Therefore, when (v, w'') is reported, a quantity of q' is traded, and no higher quantity. (Note that we use positivity of p here.) Thus, $(q', p(q'))$ is in M_v, which proves the claim. $\square$

The above lemma shows that for a mechanism M that is SBB, IR, and DSIC with respect to $\mathcal{S}_k$, if $p > 0$, then for any $v \in \mathcal{S}_k$, the menu M_v that the buyer presents to the seller includes the outcomes $(q, p(q))$ such that q is in the subset of S obtained by truncating S at the buyer's least-quantity utility-maximizing outcome.

We can prove the following symmetric lemma for the seller.

Lemma 12. *Suppose a mechanism M is SBB, IR, and DSIC with respect to $\mathcal{S}_k$ and suppose that the unit price p is positive. Let $w \in \mathcal{S}_k$ and let $q \leq \min \arg_{q'} \max\{w(k - q') + p(q') : q' \in S\}$ be a quantity not exceeding the least utility-maximizing quantity for a seller with valuation w. It holds that*

$\arg_{q'} \max\{w(k-q')+p(q') : q' \in S, q' \leq q\}$ *is the singleton set containing quantity* $q' = \max S \cap [q]$*, and that* $q \in N_w$*.*

Proof. Note that the existence of the unit price p implies that the increase in utility of the seller is a submodular function of the traded quantity q. Therefore, the function for a buyer with valuation v is increasing up to the least utility-maximizing outcome in S, after which it stays constant up to the highest utility-maximizing outcome in S, after which it starts decreasing. Let q be any quantity less than or equal to the least utility-maximizing quantity, i.e., less than $\min \arg\max\{w(k-q')+p(q') : q' \in S\}$. Then, the utility-maximizing quantity q' in $S \cap [q]$ for a seller with valuation w is $\max S \cap [q]$. It remains to prove that $(q', p(q'))$ is in N_w. Let (v', w') be any report resulting in outcome $(q', p(q'))$, so that $(q', p(q')) \in N_{w'}$. Let v'' be any function increasing at a rate $p+\epsilon$ up to quantity q', for a sufficiently small $\epsilon > 0$, after which it increases extremely slowly. Then $(q', p(q'))$ is the result of report (v'', w'), and note that it is not IR to trade a quantity exceeding q' when v'' is reported, so $(q'', p(q'')) \notin M_{v''}$ for $q'' > q'$, and $(q', p(q')) \in N_{v''}$. Therefore, when (v'', w) is reported, a quantity of q' is traded, and no higher quantity. Thus, $(q', p(q'))$ is in M_w, which proves the claim. □

The above lemma shows that for any unit price mechanism M that is SBB, IR, and DSIC with respect to $\mathcal{S}_k$, for any $w \in \mathcal{S}_k$, the menu M_w that the seller presents to the buyer includes the outcomes $(q, p(q))$ such that q is in the subset of S obtained by truncating S at the buyer's least-quantity utility-maximizing outcome.

The last two lemmas combined imply that the menu of buyer consist of the outcomes $(q, p(q))$ in S where the quantity does not exceed the least utility-maximizing outcome, plus an additional arbitrary subset of utility-maximizing outcomes; and the same holds for the seller. We will show that next.

Lemma 13. *Suppose a mechanism M is SBB, IR, and DSIC with respect to $\mathcal{S}_k$ and suppose that the unit price p is positive. Let $v, w \in \mathcal{S}_k$. Let $q = \min \arg_{q''} \max\{v(q'') - p(q'') : q'' \in S\}$ be the least utility maximizing quantity for the buyer with valuation v, then $M_v = \{(q'', p(q'')) : q'' \in S \cap [q]\} \cup T$ where $T \subseteq \arg_{q''} \max\{v(q'') - p(q'') : q'' \in S\}$. Similarly let $q' = \min \arg_{q''} \max\{w(k - q'') + p(q'') : q'' \in S\}$ be the least utility maximizing quantity for the seller with valuation w, then $N_w = \{(q'', p(q'')) : q'' \in S \cap [q']\} \cup T'$ where $T' \subseteq \arg_{q''} \max\{w(k - q'') + p(q'') : q'' \in S\}$.*

Proof. Lemma 11 and 12 show that $M_v \supseteq \{(q'', p(q'')) : q'' \in S \cap [q]\}$ and $N_w \supseteq \{(q'', p(q'')) : q'' \in S \cap [q']\}$. Let $\hat{q} \in S$ such that $\hat{q} > \max \arg_{q''} \max\{v(q'') - p(q'') : q'' \in S\}$ and let $\check{q} \in S$ such that $\check{q} > \max \arg_{q''} \max\{w(k - q'') + p(q'') : q'' \in S\}$. It suffices to show that $(\hat{q}, p(\hat{q})) \notin M_v$ and that $(\check{q}, p(\check{q})) \notin N_w$.

Suppose $(\hat{q}, p(\hat{q})) \in M_v$. Let w^* be a valuation function increasing extremely steeply up to quantity $k - \hat{q}$, and increases extremely slowly afterwards. By DSIC, the outcome $(\hat{q}, p(\hat{q}))$ is selected on report (v, w^*), where we use that $p > 0$. However, by Lemma 12 it holds that $(q, p(q)) \in N_{w^*}$, so that it also must hold by the DSIC property that outcome $(q, p(q))$ is selected, which is a contradiction.

Suppose $(\check{q}, p(\check{q})) \in N_w$. Let v^* be a valuation increasing at extremely high rate up to quantity $\check{q}$, that increases extremely slowly afterwards. By

DSIC, the outcome $(\check{q}, p(\check{q}))$ is selected on report (v^*, w). However, by Lemma 11 it holds that $(q', p(q')) \in N_{v^*}$, so that it also must hold by the DSIC property that outcome $(q, p(q))$ is selected, which is a contradiction. $\square$

We are now finally ready to prove the necessity-part of our characterisation of IR, DSIC, SBB multi-unit bilateral trade mechanisms.

Proof of Theorem 24. By Lemma 8, for all $q \in S$ there is a price $p(q)$ such that a payment of $p(q)$ is charged whenever q units are traded. By Lemma 9, there is a unit price p such that $p(q) = p \cdot q$ for all $q \in S$. This establishes already that the payment function of any IR, DSIC, and SBB mechanism is in accordance with Definition 23, hence it remains to establish that the traded quantity is also prescribed by Definition 23.

First we consider the special case $p = 0$. By increasingness of the valuation function of the seller, it follows that the mechanism can only trade a quantity of 0 units in order to satisfy IR. Subsequently it follows by IR and SBB that the mechanism is required to charge a payment of 0. A mechanism that always trades 0 units at price 0 is by definition equal to a mechanism $M_{0,\varnothing,\tau}$, where τ is arbitrary and irrelevant as there is only a single outcome that the mechanism outputs.

Next, assume that $p > 0$. We prove the claim separately for $\mathcal{S}_k$ and $\mathcal{I}_k$, and we start with $\mathcal{S}_k$. By Lemma 13, for every pair of valuations (v, w) it holds that $M_v = \{q \in S : q \leq \min \arg'_q \max\{v(q') + p(q')\}\} \cup T$ where T is an arbitrary set of utility-maximizing quantities in S for a buyer with valuation v, and $N_w = \{q \in S : q \leq \min \arg_{q'} \max\{w(k - q') + w(q)\}\} \cup T'$ where T' is an arbitrary set of utility maximizing quantities in S for a seller with valuation w. Let $\tau_S(w)$ be the seller's utility maximizing quantities in N_w and let $\tau_B(v)$ be the buyer's utility-maximizing quantities in M_v. If $\tau_S(w)$ and $\tau_B(v)$ intersect, then by DSIC, mechanism must output any quantity in $\tau_S(w) \cap \tau_B(v)$: call this quantity $\tau_\cap(v, w)$. Otherwise, if $\tau_S(w) \cap \tau_B(v) = \varnothing$, the mechanism must output $\min\{\max \tau_S(w), \max \tau_B(v)\}$, in order to satisfy the DSIC property: Assume that and $\max \tau_B(v) > \max \tau_S(w)$ (the other case is symmetric) and suppose that the mechanism trades any quantity $q \neq \max \tau_S(w)$. Since the traded quantity q must lie in the intersection of M_v and N_v and since $\tau_B(v)$ and $\tau_S(w)$ are sets of highest quantities in M_v and N_w respectively, we have $q < \max \tau_S(w)$. Hence, among the quantities in N_w a quantity less than $\max \tau_S(w)$ is traded, but the buyer prefers quantity $\max \tau_S(w)$ because $\max \tau_S(w)$ is closer to the buyer's set $\max \tau_B(v)$ of utility-maximizing quantities in M_v, which would give the buyer a higher utility due to submodularity. The buyer would thus misreport such that $\max \tau_S(w)$ is output instead. Note the tie-breaking functions $\tau = (\tau_B, \tau_S, \tau_\cap)$ we just established, as well as the derived traded quantity q given to the buyer, agree precisely with those of Definition 23. We complete the equivalence by noting that $0 \in S$ as we can define a seller's utility function that grows extremely steeply up to quantity k, so that $(0, 0)$ is the only IR outcome. This implies that $M = M_{p, S \setminus \{0\}, \tau}$.

It remains to prove the claim for $\mathcal{I}$. Suppose for contradiction that there are at least two positive quantities q, q' in S, where $0 < q < q'$. We apply the same technique as in Lemma 9. Let (v, w) be a valuation profile such that the mechanism selects $(q, p(q))$ when (v, w) is reported, and let (v', w') be a valuation profile such that the mechanism selects $(q', p(q'))$ when (v', w') is reported.

Let v^* be a valuation function that increases extremely slowly up to quantity $q-1$, then jumps to a value of $pq+2\epsilon$ at quantity q and proceeds again to grow extremely slowly up to quantity $q'-1$, and finally jumps to a value of $pq'+\epsilon$ at quantity q' after which it grows extremely slowly onward. The only IR quantities that the mechanism can trade when a buyer reports v^* are 0, q, and q'. As $(q, p(q)) \in N_w$, the mechanism must select the outcome $(q, p(q))$ when (v^*, w) is reported, because of DSIC. So, $(q, p(q)) \in M_{v^*}$. Next, we construct a function w'' for which it holds that $(q, p(q)) \notin N_{w''}$ and $(q', p(q')) \in N_{w''}$: Valuation w'' is defined such that it increases extremely steeply up to the quantity $k-q'$. Subsequently it increases by an amount of $p+(q'+1)\epsilon$ to quantity $k-q'+1$, and it increases at a rate of $p-\epsilon$ afterward. Note that the only IR quantities that can be traded under w'' are 0 and q'. We thus have that $(q, p(q)) \notin N_{w''}$ and $(q', p(q')) \in N_{w''}$ because $(q', p(q'))$ in $M_{v'}$ and by DSIC the mechanism must select $(q', p(q'))$ when (v', w'') is reported. Therefore, when (v^*, w') is reported, $(q', p(q'))$ is selected so we see that $(q', p(q')) \in M_{v^*}$.

Let w^* be a valuation function defined as follows. Let $\epsilon > 0$ be sufficiently small. We let $w^*(k) = kp$, and for all $q'' > 0$ not equal to q or q', We let $w^*(k-q'') = p \cdot (k-q'') - \epsilon$, so that the seller's increase in utility for trading q'' units is $pq'' - (w^*(k) - w^*(k-q'')) = pq'' - (pk - p(k-q'') + \epsilon) = -\epsilon$, so when w^* is the valuation of the seller, the mechanism cannot trade q'' items as that would violate IR. Moreover, we define $w^*(k-q) = p \cdot (k-q) + \epsilon$ and $w^*(k-q') = p \cdot (k-q') + 2\epsilon$, so that trading q or q' units leads to an increase in utility for a seller with valuation w^* and so that trading q'' units is the preferred quantity to trade for a seller with valuation w^*. We now see that $(q', p(q')) \in N_{w^*}$, because $(q', p(q'))$ is in $M_{v'}$ so that by DSIC the mechanism outputs $(q', p(q'))$ on report (v', w^*). Next, we construct a function v'' for which it holds that $(q', p(q')) \notin M_{v''}$ and $(q, p(q)) \in M_{v''}$. This function is defined as follows: $v(q'') = pq'' - \epsilon$ for all q'' except q, where $v(q) = pq + \epsilon$. When a buyer reports v'', by IR the mechanism can either trade 0 or q units and no other quantity. On report (v'', w) the mechanism must output $(q, p(q))$ due to DSIC and because $(q, p(q)) \in N_w$ by assumption. Thus $(q, p(q)) \in N_{v''}$ and $(q', p(q')) \notin M_{v''}$, hence when (w^*, v') is reported the mechanism outputs $(q, p(q))$ because of DSIC. This establishes $(p, p(q)) \in N_{w^*}$.

We thus have constructed two functions v^* and w^* for which it holds that both $(q, p(q))$ and $(q', p(q'))$ are in both M_{v^*} and N_{w^*}. Moreover, a buyer with valuation v^* strictly prefers $(q, p(q))$ over $(q', p(q'))$, so by DSIC the mechanism must output the outcome $(q, p(q))$ when (v^*, w^*) is reported. However, a seller with valuation w^* strictly prefers $(q', p(q'))$ over $(q, p(q))$, so by DSIC the mechanism must output the outcome $(q', p(q'))$ when (v^*, w^*) is reported, which is a contradiction. So, we must refute the assumption that there are at least 2 quantities that the mechanism can trade.

Hence, either $(0, 0)$ is always output, in which case the claim is trivial (the mechanism is equal to $M_{0,\varnothing,\tau}$, where τ is not relevant), or there is a unique positive quantity q such that the mechanism selects either $(q, p(q))$ or $(0, 0)$ and outputs $(q, p(q))$ on at least one valuation profile (v, w). It now suffices to prove, by definition of the mechanism (Definition 23), that outcome $(q, p(q))$ is selected if both players experience an increase in utility from this outcome. Let (v', w') be an arbitrary valuation profile for which the latter holds. As $(q, p(q)) \in M_v$, we infer that $(q, p(q))$ is output on report (v, w') so that $(p, p(q)) \in M'_w$. Thus, by DSIC, the mechanism must select $(q, p(q))$ on report (v', w') as otherwise a buyer with valuation v' would report v instead. This proves that the mechanism equals $M_{p,\{(q,p(q))\},\tau}$. □

5.5 Approximation Mechanisms

In this section we study the design of DSIC, IR, SBB mechanisms that optimise the social welfare, i.e., the sum of the buyer's and seller's valuation. From Theorem 24, our characterization states that such a mechanism needs to be a multi-unit fixed price mechanism, so that the design challenge lies in an appropriate choice of unit-price p and quantities offered at each iteration of the mechanism.

We focus on the case of increasing submodular valuations. Obviously, every item traded can only increase the social welfare. Therefore, given that the objective is to maximise it, we repeatedly offer a single item for trade.[1] The challenge lies thus in determining the right unit price p. It is easy to see that no sensible analysis can be done if absolutely nothing is known about the valuation functions of the buyer and seller. Therefore, we assume a *Bayesian setting*, as introduced in Section 5.2 in order to model that the mechanism designer has statistical knowledge about the valuations of the two agents: The buyer's (and seller's) valuation is assumed to be unknown to the mechanism, but is assumed to be drawn from a probability distribution f

[1] Also, with respect to our tie-breaking rule mentioned at the end of the last section: We simply employ the tie breaking rule that favours the highest quantity to trade, which is the dominant choice when it comes to maximising social welfare.

(and g) which is public knowledge. We show that we can now determine a unit price that leads to a good social welfare in expectation.

For a valuation function v of the buyer, we write $\hat{v}$ to denote the *marginal increase function* of v: $\hat{v}(q) = v(q) - v(q-1)$ for $q \in [k]$. Thus, $\hat{v}$ is a non-increasing function. Similarly, for a valuation function w of the seller, we write $\breve{w}$ to denote the *marginal decrease function* of w: $\breve{w}(q) = w(k-q+1) - w(k-q)$, for $q \in [k]$, so that $\breve{w}$ is a non-decreasing function. Thus, for all $q \in [k]$, the increase in social welfare as a result of trading q items as opposed to $q-1$ items is $\hat{v}(q) - \breve{w}(q)$. Note that therefore if v and w are increasing submodular valuation functions of the buyer and seller respectively, then the social welfare is maximised by trading the maximum number of units q such that $\hat{v}(q) > \breve{w}(q)$. We measure the quality of a mechanism on a bilateral trade instance (f, g, k) as the factor by which its expected social welfare is removed from the expected optimal social welfare $OPT(f, g, k)$ that would be attained if the buyer and seller would always trade the maximum profitable amount:

$$
\begin{aligned}
&OPT(f, g, k) = \\
&= \mathop{\mathbf{E}}_{v \sim f, w \sim g}\left[w(k) + \sum_{q=1}^{\max\{q' : \hat{v}(q') > \breve{w}(q')\}} (\hat{v}(q) - \breve{w}(q))\right] \\
&= \mathop{\mathbf{E}}_{v \sim f, w \sim g}\left[\sum_{q=1}^{k} \breve{w}(q) + \sum_{q=1}^{\max\{q' : \hat{v}(q') \geq \breve{w}(q')\}} (\hat{v}(q) - \breve{w}(q))\right]
\end{aligned}
$$

For $q \in [k]$ and a buyer's valuation function w, we denote by $GFT(v, w, q)$

the value $\max\{0, \hat{v}(q) - \check{w}(q)\}$ (where "GFT" is intended to stand for "Gain From Trade"). Note that $GFT(v, w, q)$ is non-increasing in q and that $OPT(f, g, k)$ can be written as

$$OPT(f, g, k) = \sum_{q=1}^{k} \mathbb{E}_{w\sim g}[\check{w}(q) + GFT(v, w, q)].$$

Note that a social welfare as high as $OPT(f, g, k)$ can typically not be attained by any DSIC, IR, SBB mechanism. However, it is still a natural benchmark for measuring the performance of such a mechanism, and we will see next that there exists such a mechanism that achieves a social welfare that is guaranteed to approximate $OPT(f, g, k)$ to within a constant factor. In particular, for a mechanism M, let $q_M(v, w)$ be the number of items that M trades on reported valuation profiles (v, w), and define

$$SW(M,(f, g, k)) = \mathbb{E}_{v\sim f, w\sim g}[v(q_M(v, w)) + w(k - q_{M(v,w)})]$$

as the expected social welfare of mechanism M. We say that the mechanism M achieves an *α-approximation to the optimal social welfare*, for $\alpha > 1$, iff $OPT(f, g, k)/SW(M, (f, g, k)) < \alpha$.

We show next that the multi-unit fixed price mechanism where p is set so that $\sum_{q=1}^{k} \mathbf{Pr}_{w\sim g}[\check{w}(q) \leq p] = k/2$ achieves a 2-approximation to the optimal social welfare.

Theorem 29. *Let (f, g, k) be a multi-unit bilateral trade instance where the supports of f and g contain only increasing submodular functions. Let M be the multi-unit bilateral trade mechanism where at each step one item is offered* *for trade at price p until either agent reject the offer, where p is chosen so that $\sum_{q=1}^{k} \mathbf{Pr}_{w \sim g}[\check{w}(q) \leq p] = k/2$. (Informally: p is the price such that the seller is expected to accept to trade half of his units at price p.) Mechanism M achieves a 2-approximation to the optimal social welfare.*

Proof. Let v be an arbitrary buyer's valuation function. We show that the mechanism achieves a 2-approximation if f is the distribution having only v in its support, and hence v is the buyer's valuation with probability 1. It suffices to prove the claim under this assumption, because the unit-price p depends on distribution g only. Hence, if M achieves the claimed social welfare guarantee for every fixed buyer's valuation function, then it also achieves this guarantee for every distribution on the buyer's valuation. For ease of notation, we will abbreviate $SW(M, (f, g, k))$ to simply SW and we let $\ell = \max\{q : \hat{v}_k(q) \geq p\}$ be the highest quantity that the buyer would like to trade at unit-price p. In the remainder of the proof, we will omit the subscript $w \sim g$ from the expected value operator.

We first observe that SW can be written as follows, where we write $\mathbf{1}[\cdot]$ to denote the indicator function and E_q for the event that $\hat{v}(q) \geq p \geq \check{w}(q)$.

$$\begin{aligned} SW &= \mathbb{E}\left[\sum_{q=1}^{k} (\check{w}(q) + \mathbf{1}[E_q] GFT(v, w, q))\right] \\ &= \mathbb{E}\left[\sum_{q=1}^{\ell} (\check{w}(q) + \mathbf{1}[E_q] GFT(v, w, q))\right] + \mathbb{E}\left[\sum_{q=\ell+1}^{k} \check{w}(q)\right] \qquad (5.1) \end{aligned}$$

We will bound these last two expected values separately in terms of $OPT(f, g, k)$, and subsequently we will combine the two bounds to obtain the desired approximation factor.

We start with the quantities up to ℓ, for which first rewrite the expression as follows.

$$\begin{aligned}
&\mathbf{E}\left[\sum_{q=1}^{\ell}(\breve{w}(q) + \mathbf{1}[E_q]GFT(v,w,q))\right] \\
&= \sum_{q=1}^{\ell}\mathbb{E}[\breve{w}(q)] + \sum_{q=1}^{\ell}\mathbf{Pr}[E_q]\mathbf{E}[GFT(v,w,q)) \mid E_q] \\
&= \sum_{q=1}^{\ell}\mathbb{E}[\breve{w}(q)] + \sum_{q=1}^{\ell}\mathbf{Pr}[E_q]\mathbb{E}[GFT(v,w,q)) \mid E_q].
\end{aligned}$$

Now, observe that $\mathbf{Pr}[E_q] = \mathbf{Pr}[p \geq \breve{w}(q)]$ for quantities $q \leq \ell$. Since $\sum_{q=1}^{k}\mathbf{Pr}[p \geq \breve{w}(q)] = k/2$ and $\mathbf{Pr}[p \geq \breve{w}(q)]$ is decreasing in q, this implies that $\sum_{q=1}^{\ell}\mathbf{Pr}[E_q] = \sum_{q=1}^{\ell}\mathbf{Pr}[p \geq \breve{w}(q)] \geq \ell/2$. Using additionally the fact that $\mathbf{E}[GFT(v,w,q)) \mid E_q]$ is also non-increasing in q, we obtain the following bound.

$$\begin{aligned}
&\mathbb{E}\left[\sum_{q=1}^{\ell}(\breve{w}(q) + \mathbf{1}[E_q]GFT(v,w,q))\right] && (5.2)\\
&\quad\geq \sum_{q=1}^{\ell}\mathbb{E}[\breve{w}(q)] + \frac{\sum_{q=1}^{\ell}\mathbf{Pr}[E_q]}{\ell}\sum_{q=1}^{\ell}\mathbb{E}[GFT(v,w,q)) \mid E_q] \\
&\quad\geq \sum_{q=1}^{\ell}\mathbb{E}[\breve{w}(q)] + \frac{1}{2}\sum_{q=1}^{\ell}\mathbb{E}[GFT(v,w,q)) \mid E_q] \\
&\quad\geq \sum_{q=1}^{\ell}\mathbb{E}[\breve{w}(q)] + \frac{1}{2}\sum_{q=1}^{\ell}\mathbb{E}[GFT(v,w,q))] \\
&\quad\geq \frac{1}{2}\sum_{q=1}^{\ell}\mathbb{E}[\breve{w}(q) + GFT(v,w,q)] && (5.3)
\end{aligned}$$

For the quantities higher than ℓ, we first observe that non-increasingness of $\mathbf{Pr}[\breve{w}(q) < p]$ in the quantity q implies that $\mathbf{Pr}[\breve{w}(q) > p]$ is non-decreasing in q. Moreover, $\sum_{q=1}^{k} \mathbf{Pr}[\check{w}(q) \leq p] = k/2$ means that $\sum_{q=1}^{k} \mathbf{Pr}[\check{w}(q) > p] = \sum_{q=1}^{k} \mathbf{Pr}[\check{w}(q) \leq p]$, hence it holds that $\sum_{q=\ell+1}^{k} \mathbf{Pr}[\check{w}(q) > p] \geq \sum_{q=1}^{k} \mathbf{Pr}[\check{w}(q) \leq p]$. Therefore, we derive

$$\begin{aligned}
\mathbb{E}\left[\sum_{q=\ell+1}^{k} \check{w}(q)\right] &= \frac{1}{2}\sum_{q=\ell+1}^{k} \mathbb{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=\ell+1}^{k} \mathbb{E}[\check{w}(q)] \\
&\geq \frac{1}{2}\sum_{q=\ell+1}^{k} \mathbb{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=\ell+1}^{k} \mathbb{E}[\check{w}(q) \mid \check{w}(q) > p]\mathbf{Pr}[\check{w}(q) > p] \\
&\geq \frac{1}{2}\sum_{q=\ell+1}^{k} \mathbb{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=\ell+1}^{k} \hat{v}(q)\mathbf{Pr}[\check{w}(q) > p] \\
&\geq \frac{1}{2}\sum_{q=\ell+1}^{k} \mathbb{E}[\check{w}(q)] + \frac{1}{2}\sum_{q=\ell+1}^{k} \mathbb{E}[GFT(v,w,q)] \\
&\geq \frac{1}{2}\sum_{q=\ell+1}^{k} \mathbb{E}[\check{w}(q) + GFT(v,w,q)], \qquad (5.4)
\end{aligned}$$

where the second inequality holds because $\check{w}(q)$ conditioned on $\check{w}(q) > p$ is always higher than $\hat{v}(q)$ which does not exceed p. Moreover, the third inequality follows because $\mathbf{E}[GFT(v,w,q)] = \mathbb{E}[(\hat{v}(q) - \check{w}(q))\mathbf{1}(\hat{v}(q) > \check{w}(q))] \leq \mathbb{E}[\hat{v}(q)\mathbf{1}(\hat{v}(q) > \check{w}(q))] \leq \mathbb{E}[\hat{v}(q)\mathbf{1}(p > \check{w}(q))] = \hat{v}(q)\mathbf{Pr}[p > \check{w}(q)]$.

We now use (5.3) and (5.4) to bound (5.1) and obtain the desired inequality

$$SW \geq \frac{1}{2}\sum_{q=1}^{k} \mathbb{E}[\check{w}(q) + GFT(v,w,q)] = \frac{OPT(f,g,k)}{2},$$

which proves the claim. □

The above 2-approximation mechanism is deterministic. We show next that we can do better if we allow randomisation: Consider the *Generalized Random Quantile Mechanism*, or M_G, which draws a number x in the interval $[1/e, 1]$ where the CDF is $\ln(ex)$ for $x \in [1/e, 1]$. The mechanism then sets a unit price $p(x)$ such that $\mathbf{E}_w[\max\{q : w(q) \geq qp(x)\}] = \sum_{q=1}^{k} \mathbf{Pr}_w[\breve{w}(q) \leq p(x)] = xk$, repeatedly offering single item trades as before. In words, the price is set such that the expected number of units that the seller is willing to sell, is an x fraction of the total supply, where x is randomly drawn according to the probability distribution just defined. This randomised mechanism satisfies DSIC, IR, and SBB, because it is simply a distribution over multi-unit fixed price mechanisms. Note that this mechanism is also a generalisation of a previously proposed mechanism: In [21], the authors define the special case of this mechanism for a single item, and call it the *Random Quantile Mechanism.* They show that it achieves a $e/(e-1)$-approximation to the social welfare, and we will prove next that this generalisation preserves the approximation factor, although the proof we provide for it is substantially more complicated and requires various additional technical insights.

Theorem 30. *Let (f, g, k) be a multi-unit bilateral trade instance where the supports of f and g contain only increasing submodular functions. The Generalised Random Quantile Mechanism M_G achieves a $e/(e-1)$-approximation to the optimal social welfare.*

Proof. As in the proof of Theorem 29, we fix a valuation function v for the buyer. It suffices to prove the claim under this assumption, because the unit-price p depends on distribution g only. For ease of notation, we will again abbreviate $SW(M_G, (f, g, k))$ to simply SW.

We first rewrite $OPT(f,g,k)$ as follows:

$$OPT(f,g,k) = \sum_{q=1}^{k} \mathbb{E}_w[\max\{\hat{v}(q), \breve{w}(q)\}]$$
$$= \sum_{q=1}^{k} \mathbb{E}_w[\hat{v}(q)] + \sum_{q=1}^{k} \mathbb{E}_w[(\breve{w}(q) - \hat{v}(q))\mathbf{1}[\breve{w}(q) \geq \hat{v}(q)]]$$
$$= \sum_{q=1}^{k} \hat{v}(q) + \sum_{q=1}^{k} \mathbb{E}_w[\breve{w}(q) - \hat{v}(q) \mid \breve{w}(q) \geq \hat{v}(q)]$$
$$\cdot \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)]$$
$$= \sum_{q=1}^{k} \hat{v}(q) + \sum_{q=1}^{k} (\mathbf{E}_w[\breve{w}(q) \mid \breve{w}(q) \geq \hat{v}(q)] - \hat{v}(q))$$
$$\cdot \mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(b)]. \tag{5.5}$$

In the remainder of the proof, we will derive a lower bound of $(1 - 1/e)$ times the expression (5.5) on SW, which implies our claim. We first observe that SW can be bounded and rewritten as follows.

$$SW = \sum_{q=1}^{k} \mathbb{E}_w[\breve{w}(q)\mathbf{1}[\breve{w}(q) \geq \hat{v}(q)]]$$
$$+ \sum_{q=1}^{k} \mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)]\mathbb{E}_{w,x}[\hat{v}(q)\mathbf{1}[p(x) \in [\breve{w}(q), \hat{v}(q)]]$$
$$+ \breve{w}(q)\mathbf{1}[p(x) \notin [\breve{w}(q), \hat{v}(q)]] \mid \breve{w}(q) < \hat{v}(q)]$$
$$\geq \sum_{q=1}^{k} \mathbb{E}_w[\hat{v}(q)\mathbf{1}[\breve{w}(q) \geq \hat{v}(q)] + (\breve{w}(q) - \hat{v}(q))\mathbf{1}[\breve{w}(q) \geq \hat{v}(q)]]$$
$$+ \sum_{q=1}^{k} \mathbb{E}_{w,x}[\hat{v}(q)\mathbf{1}[p(x) \in [\breve{w}(q), \hat{v}(q)]] \mid \breve{w}(q) < \hat{v}(q)]]$$
$$\cdot \mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)]$$

$$
\begin{aligned}
&= \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] \\
&\quad + \sum_{q=1}^{k} (\mathbf{E}_w[\breve{w}(q) \mid \breve{w}(q) \geq \hat{v}(q)] - \hat{v}(q))\mathbf{Pr}[\breve{w}(q) \geq \hat{v}(q)] \\
&\quad + \sum_{q=1}^{k} \mathbb{E}_{w,x}[\hat{v}(q)\mathbf{1}[p(x) \in [\breve{w}(q), \hat{v}(q)] \mid \breve{w}(q) < \hat{v}(q)]] \\
&\quad \cdot \mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)] \\
&= \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] \\
&\quad + \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_{w,x}[p(x) \in [\breve{w}(q), \hat{v}(q)] \mid \breve{w}(q) < \hat{v}(q)]] \\
&\quad \cdot \mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)] \qquad (5.6) \\
&+ \sum_{q=1}^{k} (\mathbf{E}_w[\breve{w}(q) \mid \breve{w}(q) \geq \hat{v}(q)] - \hat{v}(q))\mathbf{Pr}[\breve{w}(q) \geq \hat{v}(q)].
\end{aligned}
$$

Next, we bound the first part (5.6) of the last expression, i.e., excluding the last summation.

$$
\begin{aligned}
(6) &\leq \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)] \\
&\quad \cdot \frac{\int_{1/e}^{z:p(z)=\hat{v}(q)} \mathbf{Pr}_w[\breve{w}(q) \leq p(x)]\frac{1}{x}dx}{\mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)]} \\
&= \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] \\
&\quad + \int_{1/e}^{z:p(z)=\hat{v}(q)} \left(\sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) \leq p(x)] \right) \frac{1}{x}dx \\
&\geq \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \int_{1/e}^{z:p(z)=\hat{v}(q)} \sum_{q=1}^{k} \hat{v}(q)\frac{kx}{k}\frac{1}{x}dx
\end{aligned}
$$

$$= \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \sum_{q=1}^{k} \hat{v}(q) \int_{1/e}^{z:p(z)=\hat{v}(q)} 1dx$$

$$= \sum_{q=1}^{k} \hat{v}(q)\mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)] + \sum_{q=1}^{k} \hat{v}(q)(\mathbf{Pr}[\breve{w}(q) < \hat{v}(q)] - \frac{1}{e})$$

$$= (1 - 1/e) \sum_{q=1}^{k} \hat{v}(q), \tag{5.7}$$

where for the inequality we used that both $\hat{v}(q)$ and $\mathbf{Pr}_w[\breve{w}(q) < \hat{v}(q)]$ are non-increasing in q, so that replacing all the probabilities by the average probability xk/k yields a lower value. Substituting (5.6) by (5.7) and using the expression (5.5) for OPT then yields the desired bound.

$$\begin{aligned} SW \geq (1 - 1/e)\Big(\sum_{q=1}^{k} \hat{v}(q) + \sum_{q=1}^{k} (\mathbf{E}_w[\breve{w}(q) \mid \breve{w}(q) \geq \hat{v}(q)] - \\ \hat{v}(q))\mathbf{Pr}_w[\breve{w}(q) \geq \hat{v}(q)]\Big) \\ = (1 - 1/e)OPT(f, g, k). \end{aligned}$$

□

Currently we have no non-trivial lower bound on the best approximation factor achievable by a DSIC, IR, SBB mechanism, and we believe that the approximation factor of $e/(e-1)$ achieved by our second mechanism is not the best possible. For our first mechanism, it is rather easy to see that the analysis of the approximation factor of our first mechanism is tight, and that it is a direct extension of the median mechanism of [93], for which it was already shown in [21] that it does not achieve an approximation factor better than 2: The authors show that 2 is the best approximation factor possible for any deterministic mechanism for which the choice of p does not depend on the buyer's distribution.

For the more general class of increasing valuation functions, an approximation factor of $(2e-1)/(e-1) \approx 2.582$ to the optimal social welfare is achieved by a mechanism of [21]: They use a $e/(e-1)$-approximation mechanism for the single-item setting, which yields a $(2e-1)/(e-1)$ approximation mechanism for the multi-unit setting through a conversion theorem which they prove. We note that their conversion theorem is more precisely presented for the setting with a buyer and a seller who holds one *divisible* item. However, their proof straightforwardly carries over to the multi-unit setting. It would be an interesting open challenge to improve this currently best-known bound of $(2e-1)/(e-1)$ for general increasing valuations.

5.6 Discussion

Our results give a full characterisation of truthful mechanisms for the multi-unit bilateral trade setting. This is of importance not only for theoretical considerations, but also due to its practical consequences: We have shown that the class of truthful mechanisms in this setting consists only of very simple constant unit-price, sequential posted price mechanisms. These are not only obviously truthful, but also very easy to implement. They require little computation on the participants side, and the communication complexity of such a protocol is minimal.

Many interesting open questions remain in this area. In the simple setting we consider, we do not know matching upper and lower bounds on the approximation ratio. For the multi-unit setting studied here, the next step would be to generalize first to markets with multiple buyers and sellers and then to indistinguishable agents, each entering the market with an endowment of items. In an orthogonal direction, it is interesting to consider the case with a single pair of a seller and a buyer, but multiple item types that are substitutes of each other, with more complex valuation functions as described in e.g. [70].

Part 6

Conclusion

In this thesis, we have investigated several different key aspects of market intermediation. Firstly, we showed how a simple unit-supply, unit-demand offline setting differs from a common auction setting (and how they are similar), once one looks at correlated priors. We made explicit several implications our research has on the theory of auctions and reverse auctions, showing for the first time an asymmetry between the two. Next, we looked at the challenges that appear when moving to an online setting, and gave tight bounds on the approximation ratio. Lastly, we considered welfare in a multi-unit setting, and gave both a complete characterisation of the design space of truthful strongly budget balanced mechanisms, as well as strong approximation results.

Several common threads have run through our treatment of these different facets of the intermediation setting. The most important one is that of prior information: In all of our results, our mechanisms hinge on some partial knowledge of bidders' valuations. Indeed, for none of the questions we discussed, one could hope for meaningful worst-case results. At the end of Section 2.8 we explicitly discussed this for the correlated priors case, but it is easy to see for our aims in later chapters as well.

Second, we care about the computational tractability of the mechanisms we propose. We make this most explicit in the earlier chapters, but also in the online setting many of the mechanisms discussed are carefully stated as simple optimisation problems. In multi-unit intermediation, one of the main points of our characterisation result is the low computational and communications complexity of the resulting class of truthful budget balanced mechanisms.

Third, the concept of truthfulness has been a major theme in our inquiries. In the context of correlated priors, we derived powerful results on revenue optimal mechanisms through a geometric characterisation of truthfulness. We further explored the implications of feasibility constraints in auctions and reverse auctions, when coupled with the usual truthfulness requirement. In multi-unit intermediation, we gave a complete characterisation of all truthful budget balanced mechanisms. Such a full description of that class of mechanism greatly aids any further questions one might ask about them, and indeed is not known in many settings.

These lines of inquiry have led us to a number of meaningful discoveries about the market intermediation problem, as well as about its relation to classical auction settings. One of the key insights perhaps is that sellers are different than buyers. We have discussed this in most detail in the context of auctions and reverse auctions, but the implications are clear also for intermediation: Two-sided does not imply symmetric! It is perhaps impossible to consider market intermediation without also thinking about auctions and reverse auctions, given their close connection. Some recent work [52] has tried to decompose intermediation into mechanisms for each side and a ranking for matching buyers and sellers; our discussion suggests that such an approach ought to take into account differences between buying and selling sides in its design.

Indeed, we propose that this relationship goes both ways: Our research into market intermediation has led us directly to novel and surprising results about auctions and reverse auctions. Intermediation is thus not only a worthwhile field in itself, but also for its implications to already well studied one-sided auction settings.

Our results also yield interesting insight in how intermediation *differs* from classical auction theory. One distinction that stands out throughout our endeavours is the close relation that intermediation bears to matching algorithms. In some cases, this is an implicit constraint, when the market maker is not allowed to own items themselves; while our discussion on multi-unit intermediation in Chapter 5 restricted itself to single-seller and single-buyer cases, even the multi-unit aspect already exhibited this. Our characterisation result crucially hinges on matching units bought and units sold (and payments). But even where matching transactions is not a requirement for the intermediary, our discussion in Chapter 4 has shown that the theory of matchings may be indispensable as a tool for the analysis of (non-matching) mechanisms!

Many exciting open problems remain in this area, and we believe we have shown not only our hard results, but also that these further questions are worth investigating. Each of the particular areas we looked at has direct follow-ups open. In Chapter 2, the complexity of the two-sellers, one-buyer case is still unclear, and correspondingly in Chapter 3 that of the three-bidder reverse auction. In Chapter 4, many more complex settings immediately follow from the groundwork of our simple distilled-to-its-core model, as does regret-minimisation analysis of the problem. Chapter 5 similarly has a multitude of more complex settings one is left wondering about: What about dropping the predetermined distinction between buyer and seller, instead giving non-zero initial endowment to both? What about many traders instead of two? Multiple types of items?

Beyond that, our results and the immediate connections to other mechanism design settings pose more general questions: How does intermediation relate to auctions and reverse auctions, beyond the correlated priors aspect? We discussed briefly in section 3.5 some initial thoughts on a more comprehensive framework. This could clearly be extended even further, for instance to include general reallocation mechanisms. Should there be another, orthogonal, dimension along which we classify mechanism design problems, in addition to the traditional hierarchy (see section 1.3) increasingly complex valuation functions? These questions touch the very core of how we think about mechanism design, and again underline the importance of intermediation and more general market models as an area of study within the field.

Bibliography

[1] Abraham, D. J., Blum, A., and Sandholm, T. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proc. 8th ACM Conference on Electronic Commerce (EC 2007)* (2007).

[2] Arrow, K. J. A difficulty in the concept of social welfare. *Journal of Political Economy 58*, 4 (1950), 328–346.

[3] Arrow, K. J. *Social choice and individual values.* Yale university press, 1951.

[4] Ashlagi, I., Dobzinski, S., and Lavi, R. An optimal lower bound for anonymous scheduling mechanisms. In *Proc. 10th ACM Conference on Electronic Commerce (EC 2009)* (2009), pp. 169–176.

[5] Ashlagi, I., Dobzinski, S., and Lavi, R. Optimal lower bounds for anonymous scheduling mechanisms. *Mathematics of Operations Research 37*, 2 (2012), 244–258.

[6] Auletta, V., Christodoulou, G., and Penna, P. Mechanisms for scheduling with single-bit private values. In *Proc. 5th International*

Symposium on Algorithmic Game Theory (SAGT 2012) (2012), pp. 25–36.

[7] Babaioff, M., Blumrosen, L., Dughmi, S., and Singer, Y. Posting prices with unknown distributions. In *Proc. 2nd Symposium on Innovations in Computer Science (ICS 2011)* (2011), pp. 166–178.

[8] Babaioff, M., Cai, Y., Gonczarowski, Y. A., and Zhao, M. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Proc. 2018 ACM Conference on Economics and Computation (EC 2018)* (2018), pp. 373–373.

[9] Babaioff, M., Dughmi, S., Kleinberg, R., and Slivkins, A. Dynamic pricing with limited supply. In *Proc. 13th ACM Conference on Electronic Commerce (EC 2012)* (2012), pp. 74–91.

[10] Babaioff, M., Immorlica, N., Kempe, D., and Kleinberg, R. Online auctions and generalized secretary problems. *ACM SIGecom Exchanges 7*, 2 (2008), 7.

[11] Babaioff, M., and Nisan, N. Concurrent auctions across the supply chain. In *Proc. 3rd ACM Conference on Electronic Commerce (EC 2001)* (2001), pp. 1–10.

[12] Babaioff, M., and Nisan, N. Concurrent auctions across the supply chain. *Journal of Artificial Intelligence Research 21* (2004), 595–629.

[13] Babaioff, M., Nisan, N., and Pavlov, E. Mechanisms for a spatially distributed market. In *Proc. 5th ACM Conference on Electronic Commerce (EC 2004)* (2004), pp. 9–20.

[14] Babaioff, M., Nisan, N., and Pavlov, E. Mechanisms for a spatially distributed market. *Games and Economic Behavior 66*, 2 (2009), 660–684.

[15] Babaioff, M., and Walsh, W. E. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. In *Proc. 4th ACM Conference on Electronic Commerce (EC 2003)* (2003), pp. 64–75.

[16] Babaioff, M., and Walsh, W. E. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. *Decision Support Systems 39*, 1 (2005), 123–149.

[17] Badanidiyuru, A., Kleinberg, R., and Singer, Y. Learning on a budget: posted price mechanisms for online procurement. In *Proc. 13th ACM Conference on Electronic Commerce (EC 2012)* (2012), pp. 128–145.

[18] Balseiro, S., Mirrokni, V., Leme, R. P., and Zuo, S. Dynamic double auctions: Towards first best. *Columbia Business School Research Paper No. 18-64. SSRN 3213460* (2018).

[19] Bar-Yossef, Z., Hildrum, K., and Wu, F. Incentive-compatible online auctions for digital goods. In *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2002)* (2002), pp. 964–970.

[20] Blum, A., and Hartline, J. D. Near-optimal online auctions. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)* (2005), pp. 1156–1163.

[21] Blumrosen, L., and Dobzinski, S. (almost) efficient mechanisms for bilateral trading. *arXiv preprint arXiv:1604.04876* (2016).

[22] Blumrosen, L., and Holenstein, T. Posted prices vs. negotiations: an asymptotic analysis. In *Proc. 9th ACM Conference on Electronic Commerce (EC 2008)* (2008), pp. 49–49.

[23] Blumrosen, L., and Mizrahi, Y. Approximating gains-from-trade in bilateral trading. In *Proc. 12th Conference on Web and Internet Economics (WINE 2016)* (2016), pp. 400–413.

[24] Börgers, T. *An Introduction to the Theory of Mechanism Design.* Oxford University Press, 2015.

[25] Brustle, J., Cai, Y., Wu, F., and Zhao, M. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proc. 2017 ACM Conference on Economics and Computation (EC 2017)* (2017), pp. 589–590.

[26] Caragiannis, I., Kaklamanis, C., and Kyropoulou, M. Limitations of deterministic auction design for correlated bidders. *ACM Transactions on Computation Theory (TOCT) 8*, 4 (2016), 13.

[27] Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. Multi-parameter mechanism design and sequential posted pricing. In

Proc. 42nd Annual ACM symposium on Theory of computing (STOC 2010) (2010), pp. 311–320.

[28] Chen, X., and Deng, X. Settling the complexity of two-player nash equilibrium. In *Proc. 47th IEEE Symposium on Foundations of Computer Science (FOCS 2006)* (2006), pp. 261–272.

[29] Chen, X., Du, D., and Zuluaga, L. F. Copula-based randomized mechanisms for truthful scheduling on two unrelated machines. In *Proc. 6th International Symposium on Algorithmic Game Theory (SAGT 2013)* (2013), pp. 231–242.

[30] Chen, X., Hu, G., Lu, P., and Wang, L. On the approximation ratio of k-lookahead auction. In *Proc. 7th Workshop on Internet and Network Economics (WINE 2011)* (2011), pp. 61–71.

[31] Christodoulou, G., Koutsoupias, E., and Vidali, A. A lower bound for scheduling mechanisms. In *Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)* (2007), pp. 1163–1170.

[32] Christodoulou, G., Koutsoupias, E., and Vidali, A. A characterization of 2-player mechanisms for scheduling. In *Proc. 16th Annual European Symposium on Algorithms (ESA 2008)* (2008), pp. 297–307.

[33] Clarke, E. H. Multipart pricing of public goods. *Public Choice 11*, 1 (1971), 17–33.

[34] Colini-Baldeschi, R., de Keijzer, B., Leonardi, S., and Turchetta, S. Approximately efficient double auctions with strong budget balance. In *Proc. 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016)* (2016).

[35] Colini-Baldeschi, R., Goldberg, P. W., de Keijzer, B., Leonardi, S., Roughgarden, T., and Turchetta, S. Approximately efficient two-sided combinatorial auctions. In *Proc. 2017 ACM Conference on Economics and Computation (EC 2017)* (2017), pp. 591–608.

[36] Conitzer, V., and Sandholm, T. Complexity of mechanism design. In *Proc. 18th Conference on Uncertainty in Artificial Intelligence (UAI 2002)* (2002), pp. 103–110.

[37] Conitzer, V., and Sandholm, T. Failures of the VCG mechanism in combinatorial auctions and exchanges. In *Proc. 5th international Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)* (2006), pp. 521–528.

[38] Constantin, F., Ito, T., and Parkes, D. C. Online auctions for bidders with interdependent values. In *Proc. 6th international Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)* (2007), p. 110.

[39] Cramton, P., Gibbons, R., and Klemperer, P. Dissolving a partnership efficiently. *Econometrica: Journal of the Econometric Society* (1987), 615–632.

[40] Cramton, P., Shoham, Y., and Steinberg, R. *Combinatorial auctions.* MIT Press, 2006.

[41] Crémer, J., and McLean, R. P. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica: Journal of the Econometric Society 53*, 2 (1985), 345–361.

[42] Crémer, J., and McLean, R. P. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society 56*, 6 (1988), 1247–1257.

[43] Daskalakis, C., Deckelbaum, A., and Tzamos, C. The complexity of optimal mechanism design. In *Proc. 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014)* (2014), pp. 1302–1318.

[44] Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. The complexity of computing a nash equilibrium. *SIAM Journal on Computing 39*, 1 (2009), 195–259.

[45] Deng, X., Goldberg, P., Tang, B., and Zhang, J. Revenue maximization in a bayesian double auction market. *Theoretical Computer Science 539* (2014), 1–12.

[46] Deshmukh, K., Goldberg, A. V., Hartline, J. D., and Karlin, A. R. Truthful and competitive double auctions. In *Proc. 10th Annual European Symposium on Algorithms (ESA 2002)* (2002), pp. 361–373.

[47] Dhangwatnotai, P., Dobzinski, S., Dughmi, S., and Roughgarden, T. Truthful approximation schemes for single-parameter agents. In *Proc. 49th IEEE Symposium on Foundations of Computer Science (FOCS 2008)* (2008), pp. 15–24.

[48] Dhangwatnotai, P., Dobzinski, S., Dughmi, S., and Roughgarden, T. Truthful approximation schemes for single-parameter agents. *SIAM Journal on Computing 40*, 3 (2011), 915–933.

[49] Diakonikolas, I., Papadimitriou, C., Pierrakos, G., and Singer, Y. Efficiency-revenue trade-offs in auctions. In *Proc. 39th International Colloquium on Automata, Languages, and Programming (ICALP 2012)* (2012), pp. 488–499.

[50] Dobzinski, S., Fu, H., and Kleinberg, R. D. Optimal auctions with correlated bidders are easy. In *Proc. 43rd Annual ACM symposium on Theory of computing (STOC 2011)* (2011), pp. 129–138.

[51] Dobzinski, S., and Sundararajan, M. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *Proc. 9th ACM Conference on Electronic Commerce (EC 2008)* (2008), pp. 38–47.

[52] Dütting, P., Roughgarden, T., and Talgam-Cohen, I. Modularity and greed in double auctions. In *Proc. 15th ACM Conference on Electronic Commerce (EC 2014)* (2014), pp. 241–258.

[53] Edelman, B., Ostrovsky, M., and Schwarz, M. Internet advertising and the generalized second price auction: Selling billions of

dollars worth of keywords. Tech. rep., National Bureau of Economic Research, 2005.

[54] Eső, P. An optimal auction with correlated values and risk aversion. *Journal of Economic Theory 125*, 1 (2005), 78–89.

[55] Feldman, M., and Gonen, R. Online truthful mechanisms for multi-sided markets. *arXiv preprint arXiv: 1604.04859* (2016).

[56] Feldman, M., and Gonen, R. Removal and threshold pricing: Truthful two-sided markets with multi-dimensional participants. In *Proc. 11th International Symposium on Algorithmic Game Theory (SAGT 2018)* (2018), pp. 163–175.

[57] Feldman, M., Gravin, N., and Lucier, B. Combinatorial auctions via posted prices. In *Proc. 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2015)* (2015), pp. 123–135.

[58] Gandhi, S., Buragohain, C., Cao, L., Zheng, H., and Suri, S. A general framework for wireless spectrum auctions. In *Proc. 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN 2007)* (2007), IEEE, pp. 22–33.

[59] Gandhi, S., Buragohain, C., Cao, L., Zheng, H., and Suri, S. Towards real-time dynamic spectrum auctions. *Computer Networks 52*, 4 (2008), 879–897.

[60] Gerstgrasser, M., Goldberg, P. W., and Koutsoupias, E. Revenue maximization for market intermediation with correlated pri-

ors. In *Proc. 9th International Symposium on Algorithmic Game Theory (SAGT 2016)* (2016), pp. 273–285.

[61] Giannakopoulos, Y., and Koutsoupias, E. Duality and optimality of auctions for uniform distributions. *arXiv preprint arXiv:1404.2329* (2014).

[62] Giannakopoulos, Y., and Koutsoupias, E. Duality and optimality of auctions for uniform distributions. In *Proc. 15th ACM Conference on Electronic Commerce (EC 2014)* (2014).

[63] Giannakopoulos, Y., Koutsoupias, E., and Lazos, P. Online market intermediation. In *Proc. 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)* (2017).

[64] Gibbard, A. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society* (1973), 587–601.

[65] Gibbard, A. Straightforwardness of game forms with lotteries as outcomes. *Econometrica: Journal of the Econometric Society* (1978), 595–614.

[66] Groves, T. Incentives in teams. *Econometrica: Journal of the Econometric Society* (1973), 617–631.

[67] Hajiaghayi, M. T. Online auctions with re-usable goods. In *Proc. 6th ACM Conference on Electronic Commerce (EC 2005)* (2005), pp. 165–174.

[68] Hajiaghayi, M. T., Kleinberg, R., and Parkes, D. C. Adaptive limited-supply online auctions. In *Proc. 5th ACM Conference on Electronic Commerce (EC 2004)* (2004), pp. 71–80.

[69] Hartline, J., and Karlin, A. Profit maximization in mechanism design. In Nisan et al. [105], ch. 13.

[70] Kelso Jr, A. S., and Crawford, V. P. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society* (1982), 1483–1504.

[71] Kleinberg, R., and Leighton, T. The value of knowing a demand curve: bounds on regret for online posted-price auctions. In *Proc. 44th IEEE Symposium on Foundations of Computer Science (FOCS 2003)* (2003), pp. 594–605.

[72] Klemperer, P. Auction theory: A guide to the literature. *Journal of Economic Surveys 13*, 3 (1999), 227–286.

[73] Klemperer, P., Ed. *The economic theory of auctions*, vol. 1. Edward Elgar Pub, 2000.

[74] Klemperer, P. *Auctions: Theory and Practice*. Princeton University Press, 2004.

[75] Koutsoupias, E., and Lazos, P. Online trading as a secretary problem. In *Proc. 11th International Symposium on Algorithmic Game Theory (SAGT 2018)* (2018), pp. 201–212.

[76] Koutsoupias, E., and Papadimitriou, C. Worst-case equilibria. In *Proc. 16th Symposium on Theoretical Aspects in Computer Science (STACS 1999)* (1999), pp. 404–413.

[77] Koutsoupias, E., and Papadimitriou, C. Worst-case equilibria. *Computer Science Review 3*, 2 (2009), 65–69.

[78] Koutsoupias, E., and Pierrakos, G. On the competitive ratio of online sampling auctions. In *Proc. 6th Workshop on Internet and Network Economics (WINE 2010)* (2010), pp. 327–338.

[79] Koutsoupias, E., and Pierrakos, G. On the competitive ratio of online sampling auctions. *ACM Transactions on Economics and Computation (TEAC) 1*, 2 (2013), 10.

[80] Koutsoupias, E., and Vidali, A. A lower bound of 1+ φ for truthful scheduling mechanisms. In *Proc. 32nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2007)* (2007), pp. 454–464.

[81] Krishna, V. *Auction theory.* Academic Press, 2009.

[82] Lavi, R., and Swamy, C. Truthful mechanism design for multidimensional scheduling via cycle monotonicity. In *Proc. 8th ACM Conference on Electronic Commerce (EC 2007)* (New York, NY, USA, 2007), pp. 252–261.

[83] Lavi, R., and Swamy, C. Truthful mechanism design for multidimensional scheduling via cycle monotonicity. *Games and Economic Behavior 67*, 1 (2009), 99–124.

[84] Levin, D., and Smith, J. L. Optimal reservation prices in auctions. *The Economic Journal* (1996), 1271–1283.

[85] Loertscher, S., and Niedermayer, A. When is seller price setting with linear fees optimal for intermediaries? Tech. rep., Discussion Papers, Department of Economics, Universität Bern, 2007.

[86] Loertscher, S., and Niedermayer, A. Fee setting intermediaries: on real estate agents, stock brokers, and auction houses. Tech. rep., Discussion paper//Center for Mathematical Studies in Economics and Management Science, 2008.

[87] Loertscher, S., and Niedermayer, A. Fee-setting mechanisms: On optimal pricing by intermediaries and indirect taxation. Tech. rep., SFB/TR 15 Discussion Paper, 2012.

[88] Lu, P. On 2-player randomized mechanisms for scheduling. In *Proc. 5th International Workshop on Internet and Network Economics (WINE 2009)* (2009), pp. 30–41.

[89] Lu, P., and Yu, C. Randomized truthful mechanisms for scheduling unrelated machines. In *Proc. 4th International Workshop on Internet and Network Economics (WINE 2008)* (2008), pp. 402–413.

[90] Lu, P., Yu, C., et al. An improved randomized truthful mechanism for scheduling unrelated machines. In *Proc. 25th Symposium on Theoretical Aspects in Computer Science (STACS 2008)* (2008), pp. 527–538.

[91] Manelli, A. M., and Vincent, D. R. Optimal procurement mechanisms. *Econometrica: Journal of the Econometric Society* (1995), 591–620.

[92] McAfee, R. P. A dominant strategy double auction. *Journal of Economic Theory 56*, 2 (1992), 434–450.

[93] McAfee, R. P. The gains from trade under fixed price mechanisms. *Applied Economics Research Bulletin 1*, 1 (2008), 1–10.

[94] Milgrom, P. R., and Weber, R. J. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society* (1982), 1089–1122.

[95] Minooei, H., and Swamy, C. Near-optimal and robust mechanism design for covering problems with correlated players. In *Proc. 9th Conference on Web and Internet Economics (WINE 2013)* (2013), pp. 377–390.

[96] Minooei, H., and Swamy, C. Near-optimal and robust mechanism design for covering problems with correlated players. *ACM Transactions on Economics and Computation (TEAC) 3*, 4 (2015), 19.

[97] Mu'alem, A., and Schapira, M. Setting lower bounds on truthfulness. In *Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)* (2007), pp. 1143–1152.

[98] Myerson, R. B. Optimal auction design. *Mathematics of Operations Research 6*, 1 (1981), 58–73.

[99] Myerson, R. B., and Satterthwaite, M. A. Efficient mechanisms for bilateral trading. *Journal of Economic Theory 29*, 2 (1983), 265–281.

[100] Nash, J. Non-cooperative games. *Annals of Mathematics* (1951), 286–295.

[101] Nash, J. F., et al. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences (PNAS) 36*, 1 (1950), 48–49.

[102] Niazadeh, R., Yuan, Y., and Kleinberg, R. Simple and near-optimal mechanisms for market intermediation. In *Proc. 10th Conference on Web and Internet Economics (WINE 2014)* (2014), pp. 386–399.

[103] Nisan, N., and Ronen, A. Algorithmic mechanism design. In *Proc. 31st Annual ACM symposium on Theory of computing (STOC 1999)* (1999), pp. 129–140.

[104] Nisan, N., and Ronen, A. Algorithmic mechanism design. *Games and Economic Behavior 35* (2001), 166–196.

[105] Nisan, N., Roughgarden, T., Tardos, É., and Vazirani, V., Eds. *Algorithmic Game Theory.* Cambridge University Press, 2007.

[106] Papadimitriou, C. H., and Pierrakos, G. On optimal single-item auctions. In *Proc. 43rd Annual ACM symposium on Theory of computing (STOC 2011)* (2011), pp. 119–128.

[107] Parkes, D. Online mechanisms. In Nisan et al. [105], ch. 16.

[108] Roberts, K. The characterization of implementable choice rules. In *Aggregation and Revelation of Preferences. Papers presented at the 1st European Summer Workshop of the Econometric Society* (1979), J.-J. Laffon, Ed., pp. 321–348.

[109] Rochet, J.-C. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics 16*, 2 (1987), 191–200.

[110] Ronen, A. On approximating optimal auctions. In *Proc. 3rd ACM Conference on Electronic Commerce (EC 2001)* (2001), pp. 11–17.

[111] Ronen, A., and Saberi, A. On the hardness of optimal auctions. In *Proc. 43rd IEEE Symposium on Foundations of Computer Science (FOCS 2002)* (2002), pp. 396–405.

[112] Roth, A. E., Sönmez, T., and Ünver, M. U. Kidney exchange. *The Quarterly Journal of Economics 119*, 2 (2004), 457–488.

[113] Roughgarden, T., and Tardos, É. How bad is selfish routing? *Journal of the ACM (JACM) 49*, 2 (2002), 236–259.

[114] Rustichini, A., Satterthwaite, M. A., and Williams, S. R. Convergence to efficiency in a simple market with incomplete information. *Econometrica: Journal of the Econometric Society* (1994), 1041–1063.

[115] Saks, M., and Yu, L. Weak monotonicity suffices for truthfulness on convex domains. In *Proc. 6th ACM Conference on Electronic Commerce (EC 2005)* (2005), pp. 286–293.

[116] Satterthwaite, M. A. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory 10*, 2 (1975), 187–217.

[117] Segal-Halevi, E., Hassidim, A., and Aumann, Y. Double auctions in markets for multiple kinds of goods. In *Proc. 27th International Joint Conference on Artificial Intelligence (IJCAI 2018* (2018), pp. 489–497.

[118] Segal-Halevi, E., Hassidim, A., and Aumann, Y. Muda: A truthful multi-unit double-auction mechanism. In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI 2018)* (2018), pp. 1193–1201.

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

[120] Vidali, A. The geometry of truthfulness. In *Proc. 5th International Workshop on Internet and Network Economics (WINE 2009)* (2009), pp. 340–350.

[121] Vidali, A. Extending characterizations of truthful mechanisms from subdomains to domains. In *Proc. 7th Workshop on Internet and Network Economics (WINE 2011)* (2011), pp. 408–414.

[122] Walras, L. *Éléments d'économie politique pure, ou, Théorie de la richesse sociale.* F. Rouge, 1896.

[123] Walras, L. *Elements of pure economics.* Routledge, 2013.

[124] Walsh, W. E., Wellman, M. P., and Ygge, F. Combinatorial auctions for supply chain formation. In *Proc. 2nd ACM Conference on Electronic Commerce (EC 2000)* (2000), pp. 260–269.

[125] Yan, Q. Mechanism design via correlation gap. In *Proc. 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)* (2011), pp. 710–719.

[126] Zhou, X., Gandhi, S., Suri, S., and Zheng, H. ebay in the sky: strategy-proof wireless spectrum auctions. In *Proc. 14th ACM International Conference on Mobile Computing and Networking (MOBICOM 2008)* (2008), ACM, pp. 2–13.

[127] Zhou, X., and Zheng, H. Trust: A general framework for truthful double spectrum auctions. In *Proc. 28th IEEE Conference on Computer Communications (IEEE INFOCOM 2009)* (2009), pp. 999–1007.