ROBUST MARKET DESIGN:
INFORMATION AND COMPUTATION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Inbal Talgam-Cohen
December 2015

Abstract

A fundamental problem in economics is how to allocate precious and scarce resources, such as radio spectrum or the attention of online consumers, to the benefit of society. The vibrant research area of market design, recognized by the 2012 Nobel Prize in economics, aims to develop an engineering science of allocation mechanisms based on sound theoretical foundations.

Two central assumptions are at the heart of much of the classic theory on resource allocation: the *common knowledge* and *substitutability* assumptions. Relaxing these is a prerequisite for many real-life applications, but involves significant informational and computational challenges. The starting point of this dissertation is that the computational paradigm offers an ideal toolbox for overcoming these challenges in order to achieve a robust and applicable theory of market design.

We use tools and techniques from combinatorial optimization, randomized algorithms and computational complexity to make contributions on both the informational and computational fronts:

1. We design simple mechanisms for maximizing seller revenue that do not rely on common knowledge of buyers' willingness to pay. First we show that across many different markets – including notoriously challenging ones in which the goods are heterogeneous – the optimal revenue benchmark can be surpassed or approximated by adding buyers or limiting supplies, and then applying the standard Vickrey (second-price) mechanism. We also show how, by removing the common knowledge assumption, the classic theory of revenue maximization expands to encompass the realistic but complex case in which buyers are interdependent in their willingness to pay.

2. We prove positive and negative results for maximizing social welfare without substitutability, i.e., without the convexity property known to drive economic efficiency. On the positive side, we design natural greedy mechanisms for two-sided markets with strong incentive properties, whose welfare performance depends on the market's "distance" from substitutability. On the negative side, we show how computational challenges related to complementarity lead to the economic failure of competitive markets, in the sense that there do not exist simple prices that guide such a market to an efficient allocation.

These results carry implications for the practice of market design, both for revenue-maximizing sellers such as Internet companies running online auctions, and for welfare-maximizing policy makers such as governments running spectrum auctions.

In loving memory of Daphna Talgam

Acknowledgements

Although I spent many hours of solitary work towards this dissertation, the most meaningful and memorable parts were the interaction and collaboration with my advisors, mentors, colleagues and friends.

First and foremost I would like to thank my advisor Tim Roughgarden. Working under his guidance has been an enormous privilege, and he will always be an inspiration for me in the depth and breadth of his knowledge and insight, his passion for teaching and mentoring, his integrity, curiosity and daring intellectual, his leadership and generosity to his students.

I also wish to thank my mentor Paul Milgrom. Taking his course my second year at Stanford was a turning point in my life, and like so many others I have been profoundly influenced by his genius (as is apparent throughout this dissertation). I am deeply grateful for his support and for all that I have learned through our conversations.

I thank Ramesh Johari for kindly serving on my reading committee and for his excellent guidance in writing this dissertation, Robert Wilson for doing me the honor of chairing my dissertation committee, and Ashish Goel for kindly serving on the committee.

Thanks to my wonderful co-authors Qiqi Yan and Paul Duetting, without whom parts of this dissertation would not have come to existence. I was fortunate to learn from Qiqi's originality and modesty, and Paul's clarity of thought and productivity.

I thank the Hsieh family for honoring me with the title of Hsieh Fellow and for generously funding my research.

I wish to acknowledge those without whom I would not have begun my PhD studies: Special thanks to Uriel Feige, my MSc advisor at the Weizmann Institute, whose profound wisdom and concise, on-spot advice have guided me throughout my PhD studies. I thank Justice Esther Hayut for allowing me the honor and privilege of interning with her; witnessing firsthand how her deep thinking and economic reasoning drove fairness and justice inspired me to pursue this PhD.

I owe a great deal to many researchers at Stanford who were always available to provide invaluable advice and guidance: Itai Ashlagi, Yonatan Gur, Fuhito Kojima, Uri Nadav, Serge Plotkin, Al Roth, Amin Saberi, Yoav Shoham, Gregory Valiant, Ryan Williams and Virginia Vassilevska Williams. Special thanks to Vasilis Gkatzelis for his generosity and encouragement.

Heartfelt thanks to the mentors, collaborators and co-authors with whom I was fortunate to connect through my internships and research projects: Noga Alon, Moshe Babaioff, Shahar Dobzinski, Michal Feldman, Prabhakar Krishnamurthy, Mohammad Mahdian, Yishai Mansour, Mukund Sundararajan, Moshe Tennenholtz, Sergei Vassilvitskii, László Végh and Omri Weinstein.

Thanks to those who made Stanford such a wonderful place for me: Tim's group members (Peerapong, Shaddin, Rishi, Zhiyi, Anthony, Kostas, Jamie, Okke and Joshua), the members of the market design coffee, and admins Lynda Harris and Ruth Harris.

I am grateful to the folks at ACM for entrusting me with co-leading XRDS magazine, an experience which introduced me to many brilliant peers, enriched my knowledge of the computer science community and made me a better community member.

Thanks to my dear friends in the Bay Area, for helping me get through the difficult times and celebrate the good ones: Ayelet, Elinor, Gili, Ifat, Neta and Shiri, as well as Ankita, Erika, Kahye, Kshipra, Irene and Leor, Keren and Orit, and finally Dana.

I am grateful beyond words to my parents Yoav and Daphna, brothers Erez and Tomer and family members Ruti, Efrat, Avi and Dorit for their love and support. And to Noy, thank you for standing by me always and for sharing this incredible journey with me, which you and Dory have made so special.

Contents

List of Tables

List of Figures

List of Algorithms

1

Introduction

"[W]e may have reason to suspect that the future demand will come from a distribution which differs from that governing past history in an unpredictable way."

- Herbert Scarf

"Using a term like nonlinear science is like referring to the bulk of zoology as the study of non-elephant animals." - Stanislaw Ulam

---

1.1 Overview

Robust Market Design. Market design is the theory and practice of designing mechanisms for resource allocation among selfish agents. The goal of such mechanisms is to coordinate the independent actions of agents, who aim to optimize their own *personal* objectives, in order to reach allocations that optimize the *social and economic* objectives of welfare and revenue. Designing an auction or market mechanism differs from a pure optimization problem in that the information needed to reach an optimal allocation is not readily available; rather, it is dispersed among the selfish agents and must be elicited via interaction with them. The outcome of such interaction is predicted using the game-theoretic concept of *equilibrium*.

In recent years it has become apparent that realistic design and analysis of an auction or market must take into account not only economic and game-theoretic considerations, but also *informational and computational* ones. Classic mechanism design is often rooted in strong, implicit assumptions regarding the information or the computational resources available to the agents and/or mechanism. In this sense, *the classic theory is not robust* – it does not necessarily hold when these assumptions are violated.

The lack of robustness becomes increasingly problematic as market design becomes more and more ubiquitous: For the past two decades, market design principles have been governing the allocation of electricity, natural resources, medical residencies and seats in public schools, among others. The Internet has further expanded the reach of market design, as increasingly more economic and social activities take place by interaction with online mechanisms. The 2012 Nobel prize in economics was awarded for the development of market design research and practice, and its influence on resource allocation is likely to continue growing in coming years. Robustness is a prerequisite for the theory to stay relevant to this wealth of practical applications.

Applying informational and computational considerations to market design raises new research questions: Which traditional assumptions no longer necessarily hold, and how can we design new allocation mechanisms that do not rely on them? Which designs still work despite violated assumptions, and how can we explain this theoretically? Can the informational/computational perspective help explain anomalies in the classic theory or practice of market design? In short, *what would a theory of market design look like when subjected to realistic informational and computational constraints?* The social significance attached to recent applications of market design underscores the importance of providing answers to these questions.

This dissertation lies in the intersection of computer science, economics and operations research, and explores the above questions by focusing on two domains – *revenue* maximization, and economic efficiency or *welfare* maximization – in the context of discrete resource allocation with money. In the revenue domain, the main goal is to design mechanisms that strike an optimal balance between raising revenue and willingness of buyers to purchase – a task made more challenging *without prior information* regarding the distribution of buyers' willingness to pay. In the welfare domain, the primary goal is to allocate in an economically efficient way despite *complementarities*, a fundamental economic phenomenon that introduces allocation dependencies and thus also computational complexity.

Main Contributions and Organization. Our contributions are organized as four chapters divided into two parts. In Part I, our goal is to study mechanisms for maximizing revenue, without relying on the classic assumption that a seller has prior knowledge about his buyers' willingness to pay, or that competing buyers have knowledge about each other. Instead, we restrict the selling mechanism from building upon distributional information about the buyers' values. In the first chapter of this part (Chapter 3) we rigorously show how *a simple idea of enhancing competition in the market can replace the mechanism's dependence on such prior information*, successfully driving revenue-extracting pricing *even with multiple different resources* for sale.

Often the buyers themselves have partial information of their willingness to pay, due to inherent ambiguity regarding the nature of the resources being allocated (for example, oil drilling rights where the existence of oil is unconfirmed). In the second chapter of this part (Chapter 4), we study how such ambiguity affects the ability to extract revenue from the uncertain buyers. Taken together, the results in this part of the dissertation show that *it is possible to price effectively with hard but realistic knowledge assumptions* – lack of prior information, multiple and complex resources, and uncertainty of buyers regarding their own values.

Part II of this dissertation studies welfare maximization in the presence of economic complementarities. Complementarities occur when the allocation of one resource is economically efficient only if another is simultaneously allocated in a certain way, making the allocation task notoriously difficult: Welfare-maximizing auctions are computationally infeasible to run with complementarities; such auctions have undesirable properties like vulnerability to manipulation by a set of buyers; and competitive markets are not guaranteed to have a welfare-maximizing market equilibrium. *Only the first of these problems has been studied from an algorithmic perspective.*

We frame complementarity issues in general as computational matters. In Chapter 5 we use the greedy algorithmic approach to design the first collusion-proof, computationally-tractable mechanism for approximately maximizing welfare in two-sided trading, despite the presence of complementarities. The mechanism's approximation factor depends on the degree of complementarity. In Chapter 6 we show, for complementarities among goods in a competitive market, that the well-known market failure to support an economically efficient equilibrium is in fact a computational issue. We conclude that *a degree of complementarity might limit, but not necessarily revoke, the welfare guarantees of auctions with desirable properties.* On the other hand, *complementarity rules out welfare-maximizing market equilibria in many competitive markets due to computational considerations.*

The Computational Paradigm. The contributions above emerge from applying computational tools and approaches to market design. The computational paradigm provides an ideal toolbox for the modern market designer, since it specializes in optimization in the face of informational and computational challenges. At the same time, market design applications raise interesting questions within the computational paradigm, for example how to analyze the expected revenue in a probabilistic environment, or what is the power of the greedy approach applied to welfare optimization. We conclude that a unified view of market design – computational, informational, economic and game-theoretical – opens new horizons to explore theoretically, and is crucial for practical purposes, enabling market design solutions to be implemented on unprecedented scale and scope.

The remainder of the introduction is organized as follows: We present our model (Section 1.2), motivating examples and related literature (Section 1.3), the research goals (Section 1.4), and our contributions (Section 1.5). Bibliographical notes appear in Section 1.6.

1.2 Model

The purpose of this section is to give an informal overview of the most important aspects of our model (for a full description see Chapter 2). We begin with a basic resource allocation model in Section 1.2.1. The extension to a Bayesian model presented in Section 1.2.2 is useful for maximizing revenue, and the substitutability assumptions stated in Section 1.2.3 are useful for maximizing welfare. Section 1.2.4 discusses solution concepts – these are important since they determine the possible resource allocation outcomes.

Applications. Below (Section 1.3, Table 1.1) we shall describe four scenarios to which our model applies, and which correspond to Chapters 3 to 6. We mention these briefly now to make the discussion of the model more concrete:

1. A for-profit website like `Expedia.com` matching multiple goods to buyers;
2. A for-profit auction for goods of uncertain nature, like the right to mine for mineral goods or to show an online advertisement to an unknown Internet user;
3. A governmental auction for redistribution of radio spectrum to better fit modern communication needs;
4. A competitive market with "bread and butter"-type complementarities, which are valued more together than they are individually.

1.2.1 Basic Model

We sketch here our standard model for resource allocation. A *resource allocation setting* has two main components: resources to allocate and *players* – buyers and sellers of the resources. A *mechanism* applied to such a setting allocates resources from the sellers to the buyers.

Resources. We focus on resources of the following nature: (i) indivisible, like an online advertisement slot; (ii) can be traded for money, unlike say an organ donation.

From now on we refer to such resources as *goods* or *items*. There can be multiple copies of the same good, which we call *units*.

An *allocation* is a partition of the goods among the players. In the initial allocation, all goods are held by the sellers, and a mechanism creates a new allocation by allocating (part of) these goods to the buyers. Some allocations may be ruled out as infeasible, and then we say that our setting has *feasibility constraints*. The set of goods allocated to a player is often referred to as his *bundle*, and the money charged or paid to him by the mechanism is his *payment*. The payments may be determined by *prices* associated with the individual goods or with their bundles.

Players. Players in our model are characterized by the *utility function* they are maximizing. The utility depends on their *values* for bundles of goods, which in turn depend on their *privately-known information*.

A player's utility function assigns numerical utilities to combinations of bundles and payments. Utilities in our model are *quasi-linear*,[1] i.e., they equal the difference between the player's value for the bundle he is allocated and his payment for it (and if the player was allocated a bundle in the initial allocation, its value is subtracted from the utility). In their interaction with the mechanism, players aim to maximize their utility, or its *expectation* in settings with randomness;[2] we call such utility-maximizing behavior *strategic*.

A player's *valuation* function (sometimes called *cost* function for a seller) assigns numerical values to bundles of goods. Valuations are monotone and normalized, assigning zero value to the empty bundle. Crucially, the values assigned by a valuation function depend on the player's privately-known information, and so are known only to the player himself (and unless stated otherwise, the privately-known information is simply these values). It is assumed that every part of the model besides the privately-known information is publicly-known.

One well-studied class of players is *single-parameter* (as opposed to *multi-parameter*). A player is single-parameter if there is a single value v such that his valuation function

[1] In particular players in our model do not have *budgets*, i.e. they are unlimited in their payments.

[2] Equivalently, players in our model are *risk-neutral*.

assigns to every bundle either 0 or v.[3] The value v is private, but the *winning* bundles to which he assigns the value v are publicly known. A single-parameter setting is a resource allocation setting in which all players are single-parameter, e.g., a setting with a single good for sale.

Mechanisms. Mechanisms govern the interaction among strategic players, and based on this interaction determine an *outcome*. Most mechanisms we consider are *deterministic* and so their outcome is an allocation and payments.[4] The outcome must be achieved in equilibrium, and so the *equilibrium solution concept* of the mechanism, which determines what constitutes an equilibrium, is very significant; possible solution concepts are discussed in Section 1.2.4 below. There are two additional requirements from a mechanism's outcome: (1) feasibility of the allocation; (2) *budget balance* – the total payment charged to the buyers must be at least the total payment paid to the sellers.

There are two central measures of an outcome: social welfare and revenue. The social welfare of an outcome is the players' aggregate value for the allocation, minus their aggregate value for the initial allocation; the revenue is the players' aggregate payment. The two central objectives in the design of mechanisms are to maximize social welfare and revenue. The welfare and revenue performance of a mechanism are measured against the following benchmarks. For every resource allocation setting, an allocation that maximizes welfare is well-defined, and a *welfare-maximizing mechanism* is one whose outcomes reach such an allocation for every setting. The optimal revenue benchmark must rule out mechanisms with unreasonably high revenue – e.g., a mechanism which charges infinitely high payments – and for this the equilibrium solution concept is utilized. For suitable such concepts, the maximum revenue among all mechanisms that adhere to the solution concept will be well-defined for every resource allocation setting, and a *revenue-maximizing mechanism* will be one whose outcomes reach this revenue for every setting.

[3]More general definitions of single-parameter exist; the standard definition we adopt here is sometimes called *binary* single-parameter, since a player can either win or lose, which he values at v or 0, respectively.

[4]The outcome of a randomized mechanism is a distribution over allocations and payments.

We distinguish between two kinds of mechanisms based on their interaction with players – *auctions* and *competitive markets*. In an auction, players strategically report their values for the goods (or other privately-known information) to the mechanism, and prices are set based on these reports.[5] In a competitive market, players strategically react to fixed market prices.[6] In the context of revenue maximization we shall focus exclusively on auctions.

We often focus on settings with multiple buyers and a single seller whose costs are zero – such a seller can be ignored and then the problem boils down to allocating goods among buyers. In particular, all the competitive markets we consider have a single seller with zero costs. Auctions with multiple sellers with non-zero costs are referred to as *double auctions*.

1.2.2 Revenue: Bayesian Model and Interdependence

The basic model above is simple in terms of the information structures it allows. It does not explicitly address the fact that auctions are games of *incomplete information*. In such games, every player has private information known only to himself, about which others may only hold beliefs. For revenue maximization[7] we adopt the more general *Bayesian* model, which explicitly models privately-known information. The Bayesian model assumes that every player's privately-known information is distributed according to an underlying distribution, and we impose the additional standard assumption that the information is distrubuted independently.

The revenue benchmark is then adjusted as follows: Fixing an equilibrium solution concept, a *revenue-maximizing* auction in the Bayesian model has maximum *expected* revenue among all mechanisms that adhere to this solution concept, where the expectation is taken over the privately-known information. In addition, players in the Bayesian model are assumed to be maximizing their *expected* utilities, where the

[5]Such auctions are called *direct*, and focusing on them is without loss of generality due to the *revelation principle* discussed below.

[6]Our competitive market model is also known as an *exchange economy*.

[7]In the welfare maximization settings we consider, incomplete information does not need to be explicitly modeled, since payments can be designed such that (approximately) maximizing welfare is a dominant strategy for all players.

expectation is taken over their beliefs regarding the privately-known information of others. The following assumption is thus central to the design of revenue-maximizing auctions in the Bayesian model:

Assumption 1.2.1 (Common knowledge)**.** *The distribution of every player's privately-known information is known to (i) all players; (ii) the mechanism designer.*

In the standard case where a player's privately-known information is simply his values, Assumption 1.2.1 means that the *value distributions* are public knowledge.

We present two additional standard assumptions in the Bayesian model. First, a natural assumption if the identities or characteristics of players are unknown is the following:

Assumption 1.2.2 (Symmetry)**.** *Players' privately-known information is drawn from the same distributions.*

A second standard assumption is that distributions are not "too heavy-tailed", i.e., their tails are no heavier than that of a power law distribution. Such distributions are called *regular* [Myerson, 1981].

Assumption 1.2.3 (Regularity)**.** *Players' privately-known information is drawn from regular distributions.*

Virtually all commonly-studied distributions are regular, and irregular distributions found in practice are often a simple combination of regular distributions [Sivan and Syrgkanis, 2013].

Interdependence: Correlation and Non-Private Valuations. One natural generalization of the model above, discussed already by Myerson [1979] and elaborated on by Crémer and McLean [1985, 1988], allows the privately-known information of different players to be *correlated*. Milgrom and Weber [1982] study another generalization of the Bayesian model, in which the players' valuations can be *non-private*, i.e., functions not only of their own privately-known information, but also of the privately-known information held by other players [see also Wilson, 1969]. Together these two generalizations form the *interdependent values* model [Milgrom and Weber, 1982]. We refer to the players in this model as *interdependent* (and otherwise as *independent*).

The idea behind interdependence is to capture more sophisticated information structures that naturally arise in practice. For example, consider the above scenario of an auction for mineral rights. Every buyer has privately-known information – typically a noisy signal indicating the presence of minerals acquired by conducting a geological survey. Notice that these privately-known signals are correlated, and that learning other buyers' signals (say by watching them bid in an auction) would cause a buyer to update his value for mining rights, indicating non-private valuations.

The assumption of common knowledge (Assumption 1.2.1) remains central in the interdependent model, and the assumptions of symmetry and regularity (Assumptions 1.2.2 and 1.2.3) generalize to this setting. See Section 4.2.3 below for further details.

1.2.3 Welfare: Substitutes

The basic model in Section 1.2.1 above allows unlimited freedom in specifying which allocations are feasible and what valuations the players have. This generality comes at a cost – as explained in Section 1.3.2 below, it is impossible to design welfare maximizing mechanisms that are applicable in practice for such general resource allocation settings. The two following assumptions are thus central in the design of welfare maximizing mechanisms. For simplicity they are stated for buyers (assuming a single seller), but easily extend to sellers as well.

To state the first assumption we need the following definitions. Given a resource allocation setting, the *coalitional value function* assigns to every subset of buyers the maximum welfare achievable with only these buyers (i.e., by a feasible allocation of the seller's goods among this subset). Recall that such a function f is *submodular* if for every buyer subsets S and T, $f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$.

Assumption 1.2.4 (Buyers are substitutes [Milgrom, 2004]). *The coalitional value function is submodular.*

Assumption 1.2.5 (Goods are substitutes [Kelso and Crawford, 1982])**.** *For every buyer the following holds. Let $\vec{p}, \vec{q}$ be two price vectors for the goods, where $\vec{q} \geq \vec{p}$. Let S be a bundle of goods that maximizes the buyer's utility given prices $\vec{p}$. Then there exists a bundle of goods T that maximizes the buyer's utility given prices $\vec{q}$, such that every good j in S with price $q_j = p_j$ belongs to T.*

We refer to a buyer's valuation as *gross substitutes* if the goods are substitutes for this buyer,[8] and if goods are substitutes we also say that there are no complements.[9] As a concrete example, consider a firm whose valuation for workers is such that an increase in some workers' salaries does not cause it to regret hiring any other workers; in this case we say that the firm has a gross subsitutes valuation for the workers.

Assumptions 1.2.4 and 1.2.5 are closely related: Consider a class of valuations $\mathcal{V}$ containing additive valuations. Then buyers are substitutes in every resource allocation setting with no feasibility constraints and with a valuation profile from $\mathcal{V}$ if and only if $\mathcal{V}$ is a subclass of gross substitutes valuations [Milgrom, 2004, Theorem 8.4].[10] We will therefore sometimes treat Assumptions 1.2.4 and 1.2.5 as a single unified *substitutes assumption.*

Assumption 1.2.4 and Assumption 1.2.5 can also be formulated algorithmically. Very roughly, single-parameter buyers are substitutes if welfare subject to feasibility can be greedily maximized,[11] and goods are substitutes for a buyer if his utility can be greedily maximized – see Section 1.4.2 below for more precise statements.

1.2.4 Solution Concepts

Market design aims to define a mechanism with which players interact, such that the equilibrium allocations and payments guarantee high welfare or revenue. Thus, what counts as an equilibrium solution is crucial — it dictates the possibilities and

[8]The term *gross* substitutes is consistent with the literature, but somewhat of a misnomer since in our model players do not have budgets.

[9]Note that gross substitutes valuations are a strict subclass of *submodular* valuations, which are also referred to as valuations with no hidden complements [Lehmann et al., 2006]. For two goods, these two classes coincide.

[10]It is not known whether this continues to hold under feasibility constraints.

[11]In technical terms, the feasibility set system forms a matroid.

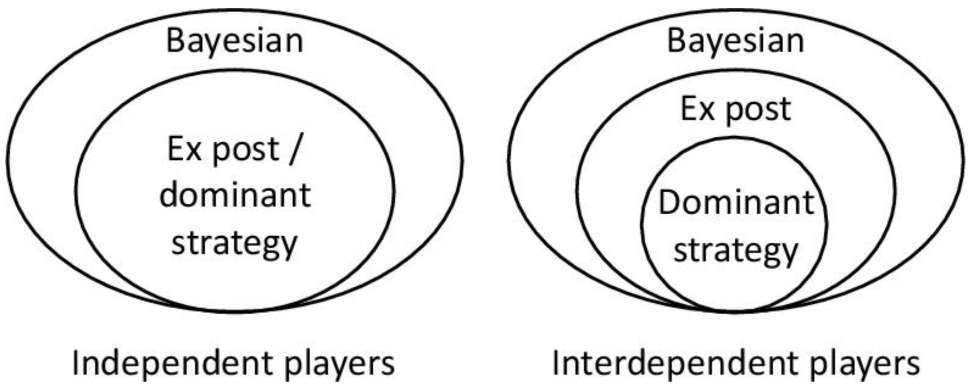

Figure 1.1: Hierarchy of equilibrium, and corresponding truthfulness, concepts.

impossibilities of market design [Chung and Ely, 2006].

Auctions and Truthful Equilibrium. Auctions can be viewed as games of incomplete information in which the players participate, and their game-theoretic equilibria are potential solutions. In the general Bayesian model with interdependence, there are three relevant equilibrium notions – dominant strategy, ex post and Bayes-Nash – and there can be multiple equilibria of each kind. The hierarchy among these equilibrium notions is depicted in Figure 1.1. (Note that Bayes-Nash equilibria are irrelevant in the basic non-Bayesian model, and that dominant strategy and ex post equilibria coincide for independent buyers – see Section 4.3.1 for details.)

To single out one equilibrium as the solution, we focus on direct auctions, and among all relevant equilibria choose the one in which the players *truthfully* report their privately-known information. This is without loss of generality by a fundamental result of Myerson [1979] called the revelation principle. There are thus three main solution concepts:

1. *Dominant strategy truthfulness*: Participating and reporting truthfully forms a dominant strategy equilibrium of the auction. I.e., for any realization of the players' privately-known information and for any possible reports to the mechanism, no player regrets participating truthfully.

2. *Ex post truthfulness*: Participating and reporting truthfully forms a Nash equilibrium in the ex post stage of the auction, where privately-known information

becomes common knowledge. I.e., for any realization of the players' privately-known information, no player regrets participating truthfully as long as the others report truthfully.

3. *Bayesian truthfulness*: Participating and reporting truthfully forms a Bayes-Nash equilibrium in the interim stage of the auction, where each player knows only his own private information. I.e., in expectation over the realization of the other players' privately-known information, no player regrets participating truthfully. Importantly, this solution concept requires the common knowledge assumption (Assumption 1.2.1), since to decide whether to participate and how to report, players must have access to all relevant distributions.

The auctions we consider are truthful (with respect to one of the above notions) and so we do not make a distinction between players' reports and their privately-known information.

Competitive Markets and Market Equilibrium. The basic solution concept that applies to competitive markets is that of a *Walrasian equilibrium*. A Walrasian equilibrium specifies a price per good and an allocation such that (1) for every buyer, his allocated bundle maximizes his utility given the prices; and (2) the market clears, i.e., no positively-priced good is left unallocated (equivalently, the allocation maximizes the revenue given the prices). By the *first welfare theorem*, the equilibrium allocation is welfare-maximizing.

A Walrasian equilibrium is not guaranteed to exist in a competitive market, and so in Section 6.3.3 we discuss generalizations of this market equilibrium notion. Such generalizations entail more general pricing – in full generality, pricing each bundle rather than each good in a buyer-specific way – and so are called *pricing equilibria*. The same two conditions as for Walrasian equilibrium must hold, i.e., utilities and revenue must be maximized by the allocation given the pricing. The first welfare theorem continues to hold as well.

	Scenario 1.3.1 Matching market	Scenario 1.3.2 Inter-dependence	Scenario 1.3.3 Double auction	Scenario 1.3.4 Competitive market
Mechanism	Auction	Auction	Auction	Market
Parameters	Multi	Single	Single	Multi
Objective	Revenue	Revenue	Welfare	Welfare
Bayesian?	Yes	Yes	No	No
Interdependence?	No	Yes	No	No
Many sellers?	No	No	Yes	No
Feasibility	Constrained	Constrained	Constrained	Unconstrained
Substitutes?	Yes	Yes	No	No
Chapter	3	4	5	6
Type of result	Positive	Positive	Positive	Negative

Table 1.1: Representative scenarios studied in this dissertation.

1.3 Motivating Examples and Related Literature

In this section we present four examples of social and economic significance, both to demonstrate our model from Section 1.2 as well as to show scenarios to which an ideal theory of market design should apply. We then give an overview of some of the most salient results in market design from the classic literature and from the more recent algorithmic game theory literature. Comparing the four examples to the literature raises two open research directions, which we shall discuss in Section 1.4.

1.3.1 Motivating Examples

The following four examples correspond to settings studied in Chapters 3 to 6, respectively. See Table 1.1 for their characteristics.

Two Revenue Maximization Scenarios.

Scenario 1.3.1 (Matching)**.** A for-profit travel website has a set of rooms available on a certain date to allocate to a set of buyers. The rooms differ from one another in quality, location etc., and thus are valued differently by the buyers. Since each buyer needs only one room per night, every allocation is in fact a matching (as in the classic work of [Demange et al., 1986]). The buyers' values are highly sensitive to rapidly changing and hard-to-predict parameters such as weather, traveling trends, local events, etc. It is thus assumed to be infeasible to obtain their distributions.

Scenario 1.3.2 (Interdependence)**.** A for-profit auctioneer sells mineral drilling rights, modern artwork, or online advertisement slots. Goods such as these have a high degree of uncertainty attached to their value: the value of drilling rights depends on the presence of minerals [Wilson, 1969]; the value of an abstract painting depends on subtle trends in the art world [Milgrom and Weber, 1982]; and the value of displaying an online advertisement depends on the identity of the online viewer [Abraham et al., 2014]. The buyers each have partial, possible noisy information regarding their uncertain value in the form of their privately-known signals. The buyers are therefore interdependent: their values are correlated with, and influenced by, the signals of other buyers and how much they are willing to pay for the good.

Two Welfare Maximization Scenarios.

Scenario 1.3.3 (Feasibility constraints and double auctions)**.** A scheduling mechanism allocates job slots on a machine to players with one job to run each. The mechanism must ensure that no job is scheduled before its arrival time or after its deadline – such allocations are infeasible. In this scenario, the players are substitutes: intuitively, the marginal coalitional value added by a player decreases as the number of jobs to run on the machine increases [Bikhchandani et al., 2011].

A related scenario involves geographical feasibility constraints imposed in radio spectrum auctions, like the ones conducted by the U.S. Federal Communications Commission (FCC). These are auctions with multiple buyers and possibly multiple sellers (if the range of spectrum for sale has already been allocated in the past). The feasibility constraints prevent two nearby transmission stations from being allocated the same frequencies in order to avoid radio transmission interference. These constraints do not necessarily maintain substitutability among the players: intuitively, a station may not contribute value at all if there is only one additional player who happens to be a neighboring station, but may contribute positively if there are extra players on the market [Milgrom and Segal, 2014].

Scenario 1.3.4 (Competitive markets with complementarities)**.** A toy example is the following: Given prices for grocery items, a buyer is interested in purchasing bread and butter. If the price of bread increases, he no longer wants the bread, and moreover no longer wants the butter even though its price did not change. In the market for foreclosed real estate, many of the buyers are commercial enterprises acquiring properties for the purpose of reselling. As Bikhchandani and Mamer [1997] observe, these buyers often do not view the properties as substitute goods "due to risk aversion, economies of scale and scope, and financial constraints."

1.3.2 Market Design State-of-the-Art

We briefly survey the state-of-the-art in revenue maximization, welfare maximization in auctions, and welfare maximization in competitive markets.

Revenue Maximization. Assume the Bayesian model with regularity. We start with selling a *single good to symmetric buyers*. Consider the dominant strategy truthfulness solution concept. The literature reflects a comprehensive understanding of the optimal revenue, with or without common knowledge (Assumption 1.2.1):

With common knowledge of the value distribution, the revenue-maximizing auction is Myerson's mechanism [Myerson, 1981]. In fact, it maximizes expected revenue even among *Bayesian* truthful auctions. Myerson's mechanism has a simple format: it is precisely the well-known second-price auction [Vickrey, 1961], with an added reserve price that depends on the value distribution. (It is also well-defined and optimal – but more complex – for asymmetric, general single-parameter buyers or without regularity.)

In the same setting (single good and symmetric buyers) but *without common knowledge*, the second-price auction with an extra buyer and no reserve price achieves at least the same revenue in expectation as Myerson's mechanism [Bulow and Klemperer, 1996]. This result generalizes beyond a single good to single-parameter buyers in general [Dughmi et al., 2012]. In the absence of an extra buyer, several mechanisms have been proposed. Some of these use random sampling to obtain an empirical version of Myerson's mechanism, providing asymptotically optimal revenue as the number of buyers goes to infinity [Baliga and Vohra, 2003; Neeman, 2003; Segal, 2003; Goldberg et al., 2006]. Dhangwatnotai et al. [2015] provide the first approximation guarantees that hold for any number of buyers.

When selling a single good to *interdependent buyers*, the linkage principle offers a systematic way of ranking different auction formats according to their expected revenue in equilibrium [Milgrom and Weber, 1982]. The Crémer-McLean mechanism is Bayesian truthful and depends heavily on common knowledge, but manages to extract full welfare as revenue, leaving buyers with zero utilities in expectation [Crémer and McLean, 1985, 1988].

When selling *multiple goods*, Cai et al. [2013] rely on the common knowledge assumption to design a Bayesian truthful[12] extension of Myerson's mechanism from a single good to multiple goods. Alaei et al. [2013] achieve a similar result, from the perspective of marginal revenue maximization.

Welfare Maximization – Auctions. The second-price auction is dominant strategy truthful and maximizes welfare for a single good [Vickrey, 1961]. We shall refer to its generalization to single-parameter settings with feasibility constraints as the Vickrey auction, and to multi-parameter settings as the VCG mechanism [Vickrey, 1961; Clarke, 1971; Groves, 1973]. These mechanisms choose a welfare-maximizing allocation among all feasible allocations, and charge buyers the *externality* imposed on other buyers by their participation (see Section 2.2.2 for details). These are essentially the *unique* welfare-maximizing, dominant strategy truthful auctions (up to a pivot term in the payments) [Green and Laffont, 1977; Holmstrom, 1979].

If the substitutes assumption holds, the Vickrey and VCG mechanisms have nice properties that make them potentially applicable in practice, while if the assumption is violated, in many senses they become impractical [Ausubel and Milgrom,

[12] Participation in their mechanism is a dominant strategy.

2006]. One nice property is that under the substitutes assumption, finding a welfare-maximizing feasible allocation by the Vickrey and VCG mechanisms is computationally tractable [Nisan, 2007].[13]

Without the substitutes assumption, finding a welfare-maximizing feasible allocation becomes computationally hard in general. One of the canonical questions in the field of algorithmic game theory is the following: when do there exist truthful implementations of algorithms for *approximately* maximizing welfare [Nisan and Ronen, 2001; Lehmann et al., 2002]? This question has been extensively studied for the multi-parameter settings known as *combinatorial auctions* [Blumrosen and Nisan, 2007, and subsequent work]. The same question for single-parameter settings is largely solved: there exists an "algorithmically-convenient" characterization of truthful mechanisms as monotone allocation rules with threshold prices [e.g., Krishna, 2010]. See also [Archer and Tardos, 2001].

One single-parameter family of settings for which such a useful characterization does not exist is double auctions, in which there are multiple sellers and buyers of units of a single good. In such settings, even without feasibility constraints, there is no budget-balanced, dominant strategy truthful mechanism that maximizes welfare [Hurwicz, 1972], and this impossibility result continues to hold even if the solution concept is relaxed to Bayesian truthfulness [Myerson and Satterthwaite, 1983]. On the positive side, the McAfee mechanism is budget-balanced and dominant strategy truthful, and realizes all but one of the economically efficient sales [McAfee, 1992]. There is no treatment in the literature of feasibility constraints (and therefore also no treatment of complementarities).

Welfare Maximization – Competitive Markets. In competitive markets where goods are substitutes, a Walrasian equilibrium is guaranteed to exist: there exist prices for the goods that support a welfare-maximizing and market-clearing allocation,

[13] Other nice properties: Outcomes are in the core [Ausubel and Milgrom, 2002; Bikhchandani and Ostroy, 2002], revenues behave monotonically [Ausubel and Milgrom, 2006; Dughmi et al., 2012], and loser collusion is unprofitable [Ausubel and Milgrom, 2006]. In addition, these auctions have ascending implementations that reach the same outcomes (assuming truthfulness, which forms an ex post equilibrium); without the substitutes assumption such ascending auctions do not exist [Vohra, 2011].

in which every player gets a bundle that maximizes his utility given the prices [Kelso and Crawford, 1982]. On the flip side, if there is a single player for which the goods are not substitutes, a Walrasian equilibrium is no longer guaranteed [Milgrom, 2000; Gul and Stacchetti, 1999]. Bikhchandani and Mamer [1997] give a linear programming-based characterization of markets with a Walrasian equilibrium.

Beyond substitutes, a pricing equilibrium is guaranteed to exist with fully general, personalized bundle prices. Such an equilibrium is infeasible due to the prohibitive number of prices required – exponential in the number of goods [Bikhchandani and Ostroy, 2002]. Moreover, the *communication* required to reach a welfare-maximizing allocation is known to be exponential in markets with general complementarities, as well as in several subclasses [Nisan and Segal, 2006]. Starr [1969] studies *quasi-equilibria* in markets with complementarities, and shows that for a fixed number of goods and a sufficiently large number of buyers, there is a sense in which it is possible to get arbitrarily close to a Walrasian equilibrium [see also Arrow and Hahn, 1971, Chapter 7].

1.4 Research Goals

The comparison of the scenarios in the previous section to the state-of-the-art theoretical understanding highlights two research goals, which we now discuss in Sections 1.4.1 and 1.4.2, respectively. In Section 1.4.3 we touch upon the "meta-goal" of designing simple mechanisms.

1.4.1 Goal 1: Revenue without Common Knowledge

Comparing the revenue maximization scenarios (Scenarios 1.3.1 and 1.3.2) to the market design literature above, it becomes apparent that without the common knowledge assumption (Assumption 1.2.1), our picture regarding revenue maximization is far from complete. Indeed, existing auctions for maximizing revenue rely on common knowledge both for multiple goods and for interdependent valuations, and a parallel of the theory developed for a single good without common knowledge does not exist for these important settings. In this section we present strong motivation for relaxing the common knowledge assumption, especially for multiple goods and interdependent valuations, and discuss the design goals for mechanisms in its absence. Our contribution to the study of revenue without common knowledge is summarized in Section 1.5.

Motivation for Relaxing the Common Knowledge Assumption. According to the common knowledge assumption, the distribution of the information privately possessed by a player is known both to the other players and to the mechanism; in a typical case, a player's value for a good is kept private while its distribution is publicly known. The common knowledge assumption has significant benefits for auction design: it is extremely helpful when setting prices for maximizing expected revenue, since it guarantees access to the probability that a buyer will agree to pay any given price; moreover, it enables the solution concept of Bayesian truthfulness, thus expanding the pool of relevant mechanisms and increasing the maximum possible expected revenue.[14] Yet despite its benefits, a long time goal of market design has been to design mechanisms independent of the common knowledge assumption, and the advantages of robustness to distributional information have been recognized universally across disciplines:

In economics, the *Wilson doctrine* [Wilson, 1985, 1987] calls for the development of mechanisms independent of the details of the economic environment, as far as these details cannot truly be expected to be common knowledge among the players and mechanism. Whereas trading rules are considered plausible common knowledge, "one agent's probability assessment about another's preferences or information" is not. The importance of weakening the common knowledge assumption is that only in this way "will the theory approximate reality."

In computer science, worst-case analysis is the dominant approach to assessing algorithm performance, rather than average-case analysis in which the average performance is taken over a known input distribution. This reflects the general-purpose

[14]The interdependent model (already in the correlated private values case) is an extreme example of this, where allowing Bayesian truthfulness guarantees full welfare as revenue!

nature of algorithms, which are expected to work well for a range of applications, as well as the discipline's mistrust in its ability to capture real-world distributions and its tendency to prefer simple designs [Nisan, 2014]. Algorithms are routinely applied in environments where the statistics of the observed data shift very rapidly [Google, 2015]. These considerations apply to mechanisms as well, and give strong motivation for their informational robustness.

In operations research, an early observation of Scarf [1958] (quoted in the epigraph to this chapter) recognizes that value distributions are not necessarily predictable even given access to past data. Bertsimas and Thiele [2014] acknowledge that accurate probabilistic knowledge is rare in practice, leading to estimation errors with unwanted consequences such as stockpiles of unneeded inventory. They conclude that there is a pressing need for an alternative, non-probabilistic theory of decision-making under uncertainty, and in particular an informational-robust theory of pricing.

Relaxing Common Knowledge with Multiple Goods or Interdependence. With multiple goods or interdependent values, there is another motivation to relax the common knowledge assumption – to gain insight about revenue maximization in these settings (see also the discussion of simplicity in Section 1.4.3 below).

Consider a matching market with multiple goods. Even in the relatively simple case where we assume symmetry among the buyers (Assumption 1.2.2), and no correlation among values for different goods, there are still m different distributions from which values are drawn, where m is the number of goods. In such settings, including the matching markets described in Scenario 1.3.1, the revenue-maximizing mechanism for the dominant strategy truthfulness solution concept is not well understood; for the Bayesian truthfulness solution concept it is understood but complex – and has not yet been applied [Cai et al., 2013]. The complication of designing revenue-maximizing mechanisms for such settings is that they need to be tailored to the m distributions. If we relax the common knowledge assumption, and seek mechanisms that do not rely on the distributions, we can hope to achieve transparent designs and gain better understanding of revenue maximization with multiple goods.

The same holds for interdependent values. The Crémer and McLean [1985, 1988] revenue-maximizing mechanism which relies on the common knowledge assumption is arguably not indicative of the real revenue that we can expect to extract: it is implausible to expect buyers to participate even though they may have negative utility, and the lotteries involved in this mechanism do not resemble anything found in practice. When we relax the common knowledge assumption by imposing ex post truthfulness rather than Bayesian truthfulness, the revenue-maximizing mechanism is no longer that of Crémer and McLean [1985, 1988] but rather a generalization of the Myerson mechanism (see Chapter 4), and it can be further simplified by seeking a mechanism that does not depend on distributions but still gets approximately maximal expected revenue.

A Model without Common Knowledge – Ex Post and Prior-Independence. Once we reject the common knowledge assumption, what mechanisms do we wish to design? What is acceptable as a solution concept, and what is the right objective to maximize?

A first observation is that abandoning the common knowledge assumption does not rule out the Bayesian model, since we can still assume values are drawn from distributions, only now these distributions are *unknown*.[15] In the absence of known characteristics that differentiate among the buyers, it is natural to assume that buyers' values are distributed in a symmetric way, and thus Assumption 1.2.2 regarding buyer symmetry still holds. Assumption 1.2.3 regarding regularity is also still very reasonable.

Bayesian truthfulness, however, no longer makes sense as a solution concept: This is because players cannot verify (nor count on the mechanism to verify) that participating truthfully is their best response, in the sense that it maximizes their expected utility where the expectation is taken over the other players' privately-known information. Following the economic literature [Chung and Ely, 2007], we adopt instead the solution concept of dominant strategy truthfulness, or ex post truthfulness in

[15]For an approach that completely diverges from the Bayesian model, and studies conditions under which certain revenue guarantees are possible despite no underlying distributions, see Goldberg et al. [2006].

the interdependent model.[16] This raises the well-known problem that a revenue-maximizing mechanism for these solution concepts does not exist. Since we are still within the Bayesian model, the expected revenue of a mechanism is well-defined, and so we solve the problem by considering the objective to be maximizing the expected revenue.

In summary, we seek dominant strategy truthful mechanisms, which do not rely on knowledge of the buyers' distributions, yet maximize revenue in expectation over these distributions. As stated this goal seems too good to be true, even with the assumptions of symmetry and regularity (Assumptions 1.2.2 and 1.2.3). We are thus willing to accept approximation rather than optimization of the expected revenue. The resulting model for robust auction design is called *prior-independence* [Dhangwatnotai et al., 2015], and it is an interesting mix of the worst-case and average-case methodologies: we aim to (approximately) maximize an average-case measure of expected revenue, but since we do not know the distributions, we aim to achieve this for the worst-case "reasonable" distributions (i.e., those which satisfy regularity). For further details, and alternative models of informational robustness, see Section 2.3.

1.4.2 Goal 2: Welfare without Substitution

Comparing the welfare maximization scenarios in Section 1.3.1 to the market design literature in Section 1.3.2, we see that even though complementarities are common in practice and despite extensive research on combinatorial auctions, substitutability still plays a crucial role in the theory of welfare maximization. The comparison highlights two families of resource allocation settings beyond combinatorial auctions for which welfare must be maximized despite complementarities, neither of which has been thoroughly addressed so far through the computational lens:

1. Single-parameter double auctions, with multiple buyers and sellers.

2. Multi-parameter competitive markets.

[16] Chung and Ely [2007] give a *maxmin* justification for why the solution concept that replaces Bayesian truthfulness is dominant strategy truthfulness: this solution concept maximizes the minimum performance of the mechanism over all possibilities of what players know about other players' valuations. Bergemann and Morris [2005] discuss the foundations of ex post truthfulness.

In addition, there is the open question of designing auctions that share the desirable economic properties of Vickrey and VCG with substitutes, even in the presence of complements [for recent work in this vein see Milgrom and Segal, 2014; Dütting et al., 2014a]. In this section we discuss the motivations for moving away from the substitutability assumption, the connection between substitutes and algorithms, and what we can hope to achieve without substitutability. Our contributions appear in Section 1.5.

Motivation for Relaxing the Substitutes Assumption. The motivation for studying welfare maximization without substitution is two-fold. First, substitutability is a stringent condition: The class of gross substitutes valuations, while far from trivial [Balcan and Harvey, 2011], is only a zero-measure subclass of general valuations [Lehmann et al., 2006], and in many settings of interest fails to hold in practice [Milgrom, 2014]. This is the case for example in the allocation of radio spectrum, where interference constraints do not typically maintain substitutability among the players (see [Milgrom and Segal, 2014] and recall Scenario 1.3.3).

A second motivation stems from the fact that lack of substitutability is closely related to computational complexity. The connection between complementarities and algorithms/complexity has been studied extensively, but almost exclusively in the realm of combinatorial auctions. In this context, there is a mature understanding regarding the computational possibilities and impossibilities of (approximate) welfare maximization for valuation classes with different levels of complementarity. This understanding has been successfully applied to design truthful combinatorial auctions and to derive lower bounds for such designs. What motivates further study is the *vast potential for a similarly fruitful connection beyond combinatorial auctions.*

An Algorithmic View of Substitutes. To better explain the relation between the substitutes assumption and computational complexity, we now outline an algorithmic approach to substitutes.

Our basic model of resource allocation is inherently combinatorial due to the indivisibility of the goods, and this gives rise to computational challenges related to combinatorial optimization.[17] By considering *classes* of valuations and feasibility constraints in place of individual resource allocation settings, we are able to refer to the computational complexity of the underlying combinatorial optimization problem, i.e., the problem of how to allocate goods in order to maximize welfare. Computational intractability of this problem stems from two kinds of complementarities, namely those related to valuations and those arising from feasibility constraints. In this thesis we will treat the two separately[18] by focusing on the following settings:

1. Single-parameter settings with feasibility constraints, which specify subsets of buyers who can win simultaneously and which are assumed to be *downward-closed* – every subset of a feasible winner set is also feasible.

2. Multi-parameter settings, without feasibility constraints.

In particular, we shall focus on single-parameter double auctions with feasibility constraints, and on multi-parameter competitive markets without feasibility constraints.

What makes substitutes computationally tractable? The answer lies in the following algorithmic characterizations based on the *greedy* approach. These characterizations are what enables auctions to maximize welfare while maintaining truthfulness and other desirable properties, and allows prices to guide decentralized markets to equilibrium. The characterizations are based on Algorithms 10 and 11 appearing in Appendix A.1. Algorithm 10 is a simple greedy method for maximizing welfare: in every step, the buyer who adds the most to the welfare without violating feasibility is added to the set of winners, until there is no such buyer who adds a positive amount. Algorithm 11 is a simple greedy method for maximizing utility, similar in form to Algorithm 10. Then:

[17] There are also computational challenges when the goods are divisible: In the absence of convexity, similar computational challenges arise, and it is an interesting open direction to formulate the connection between these two domains.

[18] This separation is not absolute; it is often possible to translate feasibility constraints to valuations and vice versa – for example, additive valuations with a *matching* constraint are equivalent to unit-demand valuations. A direction for further research is to develop a unified treatment of both sources of intractability.

Proposition 1.4.1 (Greedy characterization of substitute feasibility)**.** *Consider a single-parameter setting subject to downward-closed feasibility constraints. The buyers are substitutes for all possible values if and only if Algorithm 10 maximizes welfare for every value profile.*

Proposition 1.4.2 (Greedy characterization of substitute valuation)**.** *A buyer's valuation is gross substitutes if and only if Algorithm 11 maximizes utility for every price vector of the goods.*

The proof of Proposition 1.4.1 appears for completeness in Appendix A.1. For the proof of Proposition 1.4.2 see [Dress and Terhalle, 1995; Paes Leme, 2014, Theorem 3.2].

The following remarks apply to substitute feasibility constraints; parallels for substitute valuations exist [Paes Leme, 2014, see Definition 5.1 and Theorem 5.4].

Remark 1.4.3 (Matroids [Rado, 1957; Gale, 1968; Edmonds, 1971])**.** Buyers are substitutes for all possible values if and only if the set system corresponding to the feasible winners is a *matroid*.[19]

Remark 1.4.4 (Robustness of greedy)**.** If Algorithm 10 maximizes welfare for every value profile, two additional algorithmic properties hold: First, welfare maximization subject to a quota constraint of no more than k winners is achievable by the greedy approach [Schrijver, 2003]. Second, consider the *reversed* greedy approach, which in every step removes from the set of potential winners the buyer who adds the least welfare among those who belong to a minimal infeasible buyer subset;[20] then the reversed greedy approach also maximizes welfare for every value profile.

A Model without Substitutability – Limited Complements. Without the substitutes assumption, how do our market design goals change? In answering this question we use the combinatorial auction literature as a guideline. A starting point

[19]For more on matroids see Section 2.1.1.

[20]A subset of buyers within an infeasible set is *minimal infeasible* if at least one buyer from the subset must be removed to achieve feasibility.

of that literature is the impossibility result for maximizing welfare with general valuations – under standard complexity assumptions, there is no polynomial-time algorithm that can approximate welfare with an approximation factor better than square root of the number of goods. The literature has therefore focused on classes of valuations that are not entirely general, but are still more general than gross substitutes valuations. There are several well-studied intermediate level classes, striking different balances between expressibility and level of complementarity, such as *submodular* or *subadditive* valuations. There are also parameterized examples.[21] Inspired by this literature, our goal will be to gradually expand the envelope of substitutes by studying classes of *limited* complementarities. We now derive specific goals for our two settings of interest – single parameter double auctions with feasibility constraints, and multi-parameter competitive markets.

For double auctions, our objective will be to approximately maximize welfare for an appropriate notion of limited complements, where the optimal welfare benchmark is redefined to reflect the impossibility result of Hurwicz [1972]. Following [McAfee, 1992; Chu, 2009], the optimal welfare is replaced with the optimal number of units sold, and this enables the solution concept of dominant strategy truthfulness (thus avoiding the common knowledge assumption necessary for Bayesian truthfulness). McAfee's mechanism sells one unit less than the optimal number of units, and is dominant strategy truthful and budget-balanced. We aim to achieve a similar result with limited complements, and if possible to design double auction mechanisms with the nice properties of Vickrey with substitutes.

The situation is different for competitive markets due to the nature of the market equilibrium solution concept. Recall that the relevant equilibrium notion with complements is that of a *pricing equilibrium*, which generalizes Walrasian equilibrium and maximizes welfare by the first welfare theorem. The goal of approximately maximizing welfare is therefore irrelevant in this case, and our objective in this domain must be different. Using the fact that a pricing equilibrium with an exponential number of general bundle prices is guaranteed to exist for general valuations [Bikhchandani and Ostroy, 2002], we set as our goal finding an appropriate notion of limited complements

[21] E.g., the *MPH-m* valuation hierarchy [Feige et al., 2014].

that guarantees a pricing equilibrium with less general, more reasonable prices.

1.4.3 Meta-Goal: Simplicity

One of the main goals of practical market and mechanism design has always been simplicity. As Milgrom [2004] notes, "[i]n practice, auction designers place a tremendous value on the simplicity of an auction's design," to the point that "hardly anything matters more" [see also Krishna, 2010; Roth, 2013]. There is also much theoretical merit in simple market design: Hartline [2014] compares, in the context of optimal pricing, a complex optimal solution to a simple 2-approximation. The optimal solution is sensitive to small changes, does not generalize to similar problems, and "does not allow for much intuitive understanding" of what drives effective pricing. The simple solution "gives a crisper picture" in that it provides insight as to the trade-offs driving the revenue guarantee.

Despite its undeniable importance, there is no canonical definition of simplicity [for a recent attempt see Hart and Nisan, 2013]. It is sometimes defined in a circular way – as requiring "easy bidding" for example – or not defined at all in the spirit of "I know it when I see it".[22] Yet there are several considerations that are usually taken into account as part of the simplicity requirement: the mechanics of participating in the mechanism; whether the players feel safe to participate; whether simple strategies (possibly encoded in simple bidding languages [Dütting et al., 2011]) lead to good outcomes of the mechanism; the operational simplicity of running the mechanism; the mechanism's informational robustness as discussed in Section 1.4.1 [Bergemann and Morris, 2005]; and verifiability of the mechanism's outcome by the players [Vulkan et al., 2013].

In this dissertation we adopt a working definition of simplicity captured by the following statements:

1. A mechanism that is successfully applied in practice can be considered simple (at least for the settings it is applied to). For example, the second-price Vickrey

[22] A related issue is the lack of a canonical model that measures the benefits of simplicity; it is known to highly influence participation and running costs, but measuring its affect would require a *bounded rationality* model for the players.

auction (with reserve) is the main mechanism for selling online display advertisement by companies such as `Yahoo! Inc.` [Celis et al., 2014]; and it is indeed considered simple in the literature [e.g., Bulow and Klemperer, 1996].

2. Dominant strategy or ex post truthful auctions are strategically simple, as long as the players' privately-known information can be communicated efficiently.

3. A mechanism composed of well-understood "building blocks" is theoretically (and often also practically) simple. Examples of such building blocks are the Vickrey auction or the greedy algorithm. These building blocks are so well-studied that we consider a mechanism composed of them as transparent and insightful (as well as easy to implement and understand). Note that this simplicity criteria is context-sensitive, i.e., it depends on the standard literature and practice.

4. In a competitive market in which players interact with prices, the prices are *not* considered simple if there are exponentially many in the number of goods, nor if the players must solve a computationally hard problem to verify that they are best responding to the prices.

We aim to design (or identify barriers to designing) mechanisms that are simple according to the above working definition. For example, in accordance with (2), we choose the dominant strategy or ex post truthful solution concepts for our revenue maximization goals in Section 1.4.1.

Related Work. Chawla et al. [2007, 2010a] study posted-price mechanisms, where buyers simply choose from a menu of priced allocations. Hartline and Roughgarden [2009] seek conditions on single-parameter markets such that the simple Vickrey auction with reserve price achieves near-optimal revenue. The extension of the Vickrey auction to multi-item markets, called the VCG mechanism [Vickrey, 1961; Clarke, 1971; Groves, 1973], is arguably no longer as simple as Vickrey [Ausubel and Milgrom, 2006]; yet in several markets of interest (like the matching markets in Scenario 1.3.1), many of the complications of VCG do not occur, i.e., communicating the bids and running the auction are both computationally tractable. Other mechanism formats that have been proposed for their simplicity are running many Vickrey auctions in parallel [Christodoulou et al., 2008, and subsequent work], and a lottery between running Myerson's mechanism for the grand bundle of all goods and between separate runs of Myerson's mechanism for every good [Babaioff et al., 2014a].

1.5 Contributions

In this section we highlight our main contributions towards achieving the goals outlined in Section 1.4: revenue without distributional knowledge, and welfare without substitutability.

1.5.1 Informational Contributions in Revenue Maximization

Our main contribution in the revenue domain is to design simple Vickrey-based mechanisms with approximately-optimal revenue guarantees, for both matching markets (Scenario 1.3.1) and interdependent values (Scenario 1.3.2). Our mechanisms are prior-independent, but their revenue is guaranteed to be approximately as high as the optimal revenue in expectation, for many distributions simultaneously. We give here a brief overview of our results for each scenario and then mention several implications.

Throughout this section we assume that the privately-known information is drawn from underlying distributions (that is, we are in the Bayesian model), and that symmetry and regularity hold (Assumptions 1.2.2 and 1.2.3).

Matching Markets. In Chapter 3 we formalize and establish the following results for matching markets. Consider a matching market with n buyers and m goods. For the solution concept of dominant strategy truthfulness, the optimal expected revenue is a long-time open problem; we show it is approximated by the following versions of the Vickrey auction.

Theorem (Increased demand)**.** *In expectation, the prior-independent Vickrey auction with m more buyers achieves at least the revenue of the optimal prior-dependent mechanism in the matching market.*

If $n < m$ (resp., $n \ll m$), then in expectation the Vickrey auction with n (resp., $2n$) more buyers, and a limit of n on the total number of allocated goods, achieves at least an n/m (resp., constant $\alpha < 1$) fraction of the optimal expected revenue.

This theorem generalizes [Bulow and Klemperer, 1996] beyond single-parameter settings. In the second part of the theorem, we use the notation $n \ll m$ for the case where m is larger than n/α. We see that when the supply m is plentiful compared to the demand n, we can add a number of buyers proportional to the demand, as long as we limit the total number of allocations. On the other hand, if the supply m is smaller than the demand n (i.e., "the competition is fierce" to begin with), we need to add no more than m extra buyers.

The approximation factor in the next theorem also improves as the inherent competition in the market grows:

Theorem (Limited supply)**.** *In expectation, the prior-independent Vickrey auction with no additional buyers, but with a limit of $n/2$ on the total number of allocated goods, achieves at least a constant fraction α' of the revenue of the optimal prior-dependent mechanism in the matching market. If $n \gg m$, it achieves at least an $(n-m)/n$ fraction.*

In the above theorem we use the notation $n \gg m$ for the case where n is larger than $m/(1-\alpha')$.

We summarize our additional contributions for matching markets as follows:

- *Approach and simple implementation*: We explore an *enhanced competition* approach to informational-robust revenue maximization. Instead of relying on prior knowledge regarding buyer valuations in order to set revenue-maximizing prices, we "let the market do the work" in the sense that prices emerge from competition for scarce goods. We study two complementary methods of enhancing competition, both inspired by real marketing practices – increasing demand

and limiting supply. Our mechanisms are a simple implementation of this approach: they first enhance competition using the two methods and then run the standard Vickrey auction.

- *Advantages of our approach*: The resulting mechanisms are computationally tractable and guarantee approximation factors that are either independent of the size of the market (as measured by parameters n and m), or improve as the inherent competition in the market grows. The approach is flexible and works well for other scenarios such as single-parameter settings with feasibility constraints or with buyer asymmetry.

- *Computational analysis*: To prove the first theorem above, we take advantage of the combinatorial structure in matching markets – the stability properties of bipartite matchings play a key role in our analysis. We also use the principle of deferred decision from the analysis of randomized algorithms to overcome probabilistic dependencies in the revenue analysis. To prove the second theorem above, we formalize a connection between increasing demand and limiting supply, by establishing a general reduction from a market with limited supply to one with increased demand, and showing the relation between their revenue guarantees.

Some implications of our results appear below.

Interdependent Values. In Chapter 4 we establish the following for interdependent valuations.

Consider a single-parameter resource allocation setting with feasibility constraints that maintain substitutability. If the values are *not* interdependent, the mechanism with optimal expected revenue is known to be Myerson's mechanism [Myerson, 1981]. Recall that in our symmetric regular setting, with a single good this coincides with the Vickrey auction with a reserve price, where the reserve price depends on the value distribution. In the general (single-parameter but not necessarily single-good) setting, Myerson's mechanism considers the feasible buyer set with the highest values to be the *potential winner set*, and allocates to the potential winners who surpass the reserve price. In both cases winners are charged their *threshold value*, i.e., the lowest value they could have had that would still guarantee winning. It is also known that when there are at least two buyers, using a random buyer's value for setting the reserve price instead of calculating it from the value distribution is approximately revenue-maximizing as well as prior-independent. A mechanism of this form is called the *single-sample mechanism* [Dhangwatnotai et al., 2015].

What we show is that the described Myerson and single-sample mechanisms still work for interdependent values with the appropriate changes (such as using signals rather than values to determine the potential winner set), and under appropriate regularity assumptions. One such assumption is called the *single-crossing condition* and it states the following: Recall that in the interdependent model the privately-known information of every buyer is not necessarily his value but rather some informational signal (in the example above, an indication of the presence of minerals to mine). Roughly, the single-crossing condition requires that a certain transformation[23] of the signals maintains their order (e.g., the highest signal remains highest after the trasnformation). In the Bayesian model with no interdependence, this condition coincides with Assumption 1.2.3 of regularity.

Our main theorems are as follows:

Theorem. *For interdependent values, under the single-crossing condition, Myerson's mechanism achieves the optimal expected revenue among all mechanisms that are ex post truthful.*

(In the interdependent values model ex post truthfulness is the appropriate solution concept rather than dominant strategy truthfulness – see Chapter 4.)

Our prior-independence result holds for a subset of interdependent values with no correlation among the signals – for example, a buyer's value for an abstract painting that depends on his and others' idiosyncratic opinions modeled by independent signals:

Theorem. *For interdependent values with no correlation, under appropriate regularity assumptions, the single sample mechanism achieves approximately the optimal*

[23]The *conditional virtual value* transformation – see Chapter 4.

expected revenue among all mechanisms that are ex post truthful.

Additional contributions are as follows:

- *Computational view*: Interdependent values in general were not studied before from a computational perspective. A subclass – correlated values – had been studied and deemed computationally hard to maximize revenue for in a deterministic manner [Papadimitriou and Pierrakos, 2015] (see also [Dobzinski et al., 2011]), but we show that positive results are possible under regularity assumptions. We also show the first prior-independent mechanism beyond the standard Bayesian model.

- *Concreteness*: The regularity condition of single-crossing is quite opaque; it is important to get a better understanding of it, and in particular to find out when, despite players' partial information, allocating to the player with the highest signal is best in terms of revenue. Based on [Lopomo, 2000; Li, 2013] we show two sets of economically meaningful sufficient conditions, as well as examples of settings that satisfy these conditions.

- *Necessity of regularity*: In the *non*-interdependent case, there is a generalization of the Myerson mechanism based on a method called *ironing*, which maximizes expected revenue without regularity [Myerson, 1981]. Our work explains why this method will not work for interdependent values without regularity.

Implications. The above results regarding revenue without common knowledge have the following implications:

- The prior-independent approach helps achieve a simple understanding of revenue in complex settings, providing a handle on the elusive optimal expected revenue in matching markets with multiple goods, and on what can be considered as plausible optimal expected revenue in markets with interdependence. This understanding is based on the number of competing or sampled buyers needed in order to set revenue-extracting prices. It also gives sellers natural methods for bypassing the lack of prior knowledge (e.g., setting quantities instead of prices), and demonstrates when the use of the welfare-maximizing Vickrey auction is justified for extracting revenue. Our results demonstrate that robustness and simplicity are achievable even in complicated settings, and in fact are advisable when approaching such settings.

- The results for matching markets establish quantitative trade-offs between possible business strategies for extracting revenue, and can be useful for sellers deliberating between sophisticated pricing, advertising to draw more bidders, and creating scarcity on the market by limiting supplies.

- The results for interdependent values are of interest in themselves but also strengthen the approach of relaxing the common knowledge assumption: once this assumption is relaxed for the buyers, the optimal mechanism ceases to be the Crémer and McLean [1985, 1988] mechanism and becomes a Myerson-like mechanism. Imposing robustness as a requirement thus contributes to a comprehensive theory of revenue maximization for interdependent values.

1.5.2 Computational Contributions in Welfare Maximization

Our main contribution in the welfare domain is to use the computational view of substitutes and complements to get positive and negative results for welfare maximization. We first design double auctions with approximately optimal welfare and strong properties, despite the presence of complements arising from feasibility constraints (Scenario 1.3.3). On the other hand, we rule out the guaranteed existence of plausible welfare-maximizing equilibria in many competitive markets with limited complements (Scenario 1.3.4). After an overview of results for each scenario we list several implications.

Double Auctions. Consider a double auction setting with multiple buyers and sellers, each interested in either buying or selling a single unit of a good. The goal of the mechanism is to pair buyers and sellers such that each couple will trade a unit. There are feasibility constraints on which buyers can be paired simultaneously, as well as on which sellers can be paired simultaneously. Such constraints can ensure, for example, that buyers of re-purposed radio spectrum are a mixture of major communication companies and smaller new-comers to the market.

If there were no feasibility constraints, the McAfee [1992] double auction could be applied in order to achieve dominant strategy truthfulness, budget balance, and the realization of all but one of the welfare-optimal trades (this is the best possible [Dütting et al., 2014b]). Our first contribution is to view the McAfee mechanism in a new way – as a combination of two reversed greedy algorithms, one which ranks the buyers by how much they are willing to pay for a unit, and one which ranks the sellers by how much they charge for their unit. Trading pairs are composed by simply pairing players according to their rank. This view of McAfee's mechanism is useful since it illuminates a natural generalization, in which the greedy algorithms implement the feasibility constraints. Another advantage of this view is that it reveals the McAfee mechanism and its generalization to be not only individually truthful, but also *group strategyproof*, meaning it cannot be manipulated by a group of players in a way that increases all their utilities. Group strategyproofness is one of the strong properties possessed by Vickrey with substitutes.

A second contribution of the described approach is that it indicates an appropriate notion of limited complements. Recall the characterization of substitutes by optimality of the greedy algorithm and its variations (Proposition 1.4.1 and Remark 1.4.4), including the reversed version subject to a quota. This version is particularly relevant in the context of the McAfee double auction, since the feasible ranking of sellers effectively sets a quota on the number of trading buyers and vice versa. We use the performance of this version of greedy as a measure of complementarity: For fully general feasibility constraints, it has no performance guarantee; but for constraints inducing limited complements, it is approximately optimal.[24] Two additional motivations to use the success of greedy as a measure of complementarity are that (1) in practice, greedy-based heuristics tend to perform well [Milgrom and Segal, 2014]; and (2) there is good algorithmic understanding of what approximation guarantees

[24] Approximate optimality means here that for every quota, the greedy algorithm ranks a feasible set of buyers (or sellers) with near-optimal total value (or cost).

are achievable by greedy [e.g. Korte and Hausmann, 1978; Borodin et al., 2003].

To informally state the performance guarantee of our McAfee-based design for a class of feasibility constraints, let parameters $\alpha \geq 1$ and $\beta \geq 1$ measure the approximation guarantees for buyers and sellers of reverse greedy subject to a quota ($\alpha = \beta = 1$ for substitutes). Parameters α and β are thus our measure of complementarity. Given a double auction setting, let v_{OPT} and c_{OPT} be the optimal total value, resp. cost, that greedy is trying to approximate, where $v_{\mathrm{OPT}} - c_{\mathrm{OPT}}$ is the optimal welfare. Then:

Theorem 1.5.1. *For every double auction setting as described above, our McAfee-based double auction is group strategyproof, maintains feasibility and budget balance, and realizes trades that achieve welfare of at least* $(1/\alpha) \cdot v_{\mathrm{OPT}} - \beta \cdot c_{\mathrm{OPT}}$ *up to a single trade.*

In Chapter 5 we turn the above theorem into a proper approximation guarantee (see Corollary 5.5.4). We mention two additional contributions:

- *Generalization to a double auction framework*: We generalize the above view of the McAfee mechanism into a general framework for designing modular and greedy-based double auctions, which encompasses the Vickrey auction as well as other mechanisms. Our framework decomposes the design task into designing a greedy ranking algorithm for the buyers, a greedy ranking algorithm for the sellers, and their composition into trading pairs. The modular approach is inspired by the computational method of reducing new problems (double auction design with feasibility constraints) to well-studied ones (greedy algorithms). Notice that the resulting double auctions are simple according to our working definition in Section 1.4.3. This facilitates their analysis and enables us to prove what we call "composition theorems", which relate a desired property of the double auction design to the properties of the modules constructing it. This leads to new results for the classic but timely problem of double auction design.

- *Applications* We instantiate our modular approach for three different feasibility classes based on the following combinatorial structures – matroids, knapsacks and matchings. These classes have different levels of complementarity: For matroids, reverse greedy with a quota is optimal; for matchings, which are intersections of two matroids, it achieves a 2-approximation; etc.

Competitive Markets. Chapter 6 is the only chapter of this dissertation in which we present results that are largely negative. Similarly to the research program for double auctions, for competitive markets we start out by seeking a limited form of complements for which there exists a market equilibrium that maximizes welfare, where in place of classes of feasibility constraints we study classes of valuations. We are unable to find the classes we seek and establish instead impossibility results to explain the difficulty, which turns out to stem directly from the computational nature of complements.

Our main contribution is to show that non-existence of a market equilibrium can be derived from computational hardness/tractability of related computational problems, namely the algorithmic issues that come up in resource allocation via competitive markets. Recall that economically efficient allocation in such markets is based on the existence of a pricing – as simple as item prices or as complex as arbitrary bundle prices – for which all the buyers' utilities as well as the seller's revenue can be maximized at the same time. This raises several algorithmic issues: the buyer's problem of finding (or at least verifying) a utility-maximizing bundle; the seller's problem of finding (or at least verifying) a revenue-maximizing allocation; and the "social planner's problem" of maximizing welfare.

Our first contribution, which demonstrates the general methodology, is for the existence of Walrasian equilibria:[25].

Theorem. *A necessary condition for the guaranteed existence of a Walrasian equilibrium in markets with valuations from class $\mathcal{V}$ is that utility maximization for $\mathcal{V}$ given item prices is at least as hard computationally as welfare maximization for $\mathcal{V}$.*

[25] For the expert reader we remark that such a result is useful even given the results of [Gul and Stacchetti, 1999; Milgrom, 2000], since in recent years there has been much research on Walarasian equilibrium existence in classes that do not include all unit-demand valuations – for details see Chapter 6.

The significance of this theorem is discussed below. Other contributions are as follows.

- *Pricing classes*: We consider general pricing functions, which assign prices to bundles of goods. These pricings cover the whole spectrum between item prices and arbitrary bundle prices. A pricing is very similar mathematically to a valuation, and we can define classes of pricings just as there are classes of valuations. This is what allows us to consider the computational complexity of the revenue maximization problem, which is in fact identical to the welfare maximization problem only with pricings instead of valuations, as well as the complexity of the utility maximization problem (with pricings more general than item prices).

- *Plausible market equilibria*: We identify three properties that make Walrasian equilibrium an economically meaningful, plausible solution concept. A key necessary property is *verifiability*, i.e., the ability of the buyers and seller to verify in a computationally tractable way that they are maximizing their utility or revenue, respectively. Verifiability is closely related to simplicity as discussed in Section 1.4.3. It is not however a sufficient property: there always exists a trivial verifiable market equilibrium, in which the allocation is welfare-maximizing and every buyer is charged his own value, but which does not seem to be a solution that the market is likely to converge to. We therefore identify two other aspects of Walrasian equilibria that make them non-trivial: the prices are anonymous (common to all buyers); the prices are succinct relative to the valuations – indeed the class of item prices is much smaller than the class of gross substitutes, for which Walrasian equilibria are guaranteed to exist.

- *Computational boundary for anonymous or succinct pricing*: We show a computational boundary to the existence of a verifiable pricing equilibria with anonymous prices. Roughly, a necessary condition for its existence in markets with valuations from class $\mathcal{V}$ and prices from class $\mathcal{P}$ is that the welfare maximization problem itself is verifiable, or in coNP. Assuming NP $\not\subseteq$ coNP this means that the welfare maximization problem cannot be NP-hard, ruling out all notions of limited complements that do not maintain computational tractability, including the notion identified in our double auction results. A similar boundary exists for succinct pricing.

- *Another computational barrier with demand queries*: An alternative definition of verifiability allows buyers the ability to answer "demand queries", that is, solve the utility maximization problem even if it is computationally hard. The demand query model is adopted from the combinatorial auctions literature (see Section 6.3.1). Even in this model we show indications that meaningful pricing equlibria are rare at best: their existence may imply a non-standard algorithm for welfare maximization, and such algorithms are currently unknown. There is thus a possible computational explanation for the lack of useful extensions of Walrasian equilibrium – they seem to require the invention of novel welfare-maximization algorithms.

- *Applications*: The literature on market equilibrium existence and non-existence has largely progressed in an ad hoc manner, by studying valuation classes of interest one at a time. Our results give a systematic way of achieving a host of non-existence results, all stemming from the same computational reason.

Implications. Our contributions regarding welfare maximization with complements have the following implications:

- Complements are not always considered a computational issue outside the context of combinatorial auctions, but there is value in studying them through the computational lens in wider settings. The computational view can help us identify limited complements for which positive results exist, and algorithmic barriers to overcome in order to get such positive results.

- Our double auction framework for limited complements not only facilitates the design of new mechanisms but also offers insights into the classic double auction designs of VCG and McAfee, revealing for the latter strong incentive properties.

- Understanding when equilibria are guaranteed to exist is a central theme in economic theory, seemingly unrelated to computation. Our results show that the existence of a Walrasian or pricing equilibrium is inextricably connected to the computational complexity of related optimization problems. This links a purely economic question – existence of equilibrium – to a purely algorithmic one.
- A systematic method for showing that complements cannot be handled well and will inevitably lead to market failure is useful so that alternative resource allocation systems can be deployed, possibly involving legislation, government supervision or taxation.

1.6 Bibliographic Notes

The results presented in this dissertation are based on joint works with Paul Dütting, Tim Roughgarden and Qiqi Yan. Chapter 3 is based on [Roughgarden et al., 2012], Chapter 4 on [Roughgarden and Talgam-Cohen, 2013], Chapter 5 on [Dütting et al., 2014b], and Chapter 6 on [Roughgarden and Talgam-Cohen, 2015].

2

Preliminaries

This chapter expands and formalizes the model described in Section 1.2. Auction preliminaries appear in Sections 2.1 (environments) and 2.2 (mechanisms), and Section 2.3 includes prior independence preliminaries. The model and notation introduced is general, and some of the chapters later define more specialized notation/models (Chapter 4 to capture interdependence, and Chapter 5 to capture multi-seller auctions).

2.1 Auction Environments

An *auction environment* has a set of m items $\{1, \ldots, m\}$, and a set of n bidders $E = \{1, \ldots, n\}$.

2.1.1 Single-Parameter Bayesian Environments

A single-parameter environment is defined by a non-empty collection $\mathcal{I} \subseteq 2^{[n]}$ of feasible bidder sets, each containing bidders who can *win* simultaneously. The sets in $\mathcal{I}$ are sometimes referred to as *feasible allocations*. We assume that every bidder belongs to at least one feasible allocation. Every subset of a feasible allocation is also feasible, i.e., the set system $([n], \mathcal{I}) = (E, \mathcal{I})$ is downward-closed. We assume it is publicly known and represented succinctly by a tractable *feasibility oracle*, which for every set of bidders returns whether this set is feasible or not.

Every bidder i has a private value $v_i \in [0, \infty)$ for winning, which is drawn independently at random from a distribution F_i with a density function f_i positive over a nonzero interval support. The density function is smooth with one exception – a constant amount of probability mass can concentrate on the highest point in the support. The described environment is called single-parameter since the value for winning is fully described by v_i.

Recall that we assume a risk-neutral quasi-linear utility model, in which a bidder's utility for winning is his value minus the payment he is charged, and bidders aim to maximize their expected utilities. We say that single-parameter bidders are *i.i.d.* (or symmetric) if their value distributions are identical. The environment is i.i.d. if the bidders are i.i.d.

Multi-Unit. Recall the distinction between items and units (the latter are different copies of the same item, so bidders have the same value for them). A *multi-unit* (or *k-unit*) environment is a single-parameter environment in which a subset of bidders is a feasible allocation if and only if its size is at most k. This models k units for sale to $n \geq k$ unit-demand bidders who are interested in at most one unit. In Chapter 3 we will sometimes impose an additional supply limit of $\ell \leq k$, restricting feasible allocations to size at most ℓ. For a survey on multi-unit auctions see Nisan [2014].

Matroid. A matroid environment is a single-parameter environment in which the set system $([n], \mathcal{I}) = (E, \mathcal{I})$ of bidders and feasible allocations forms a matroid, and so the bidders are substitutes. Examples of matroid environments include digital goods where $\mathcal{I} = 2^E$, job scheduling markets (Example 3.1.3 below), and k-unit environments as a special case (corresponding to the k-uniform matroid).

In general, a non-empty, downward-closed set system $(E, \mathcal{I})$ is a matroid if the following *exchange property* holds: for every $S, T \in \mathcal{I}$ such that $|S| > |T|$, there is some element $i \in S \setminus T$ such that $T \cup \{i\} \in \mathcal{I}$ (see, e.g., [Oxley, 1992]). The set E of elements (bidders in our context) is called the *ground set* and the sets in $\mathcal{I}$ are called *independent*. All other sets are called *dependent*. A maximal independent set –

that is, an independent set which becomes dependent upon adding any new element from E – is called a *basis*. A minimal dependent set – that is, a dependent set whose proper subsets are all independent – is called a *circuit*. Here are two examples:

1. Vector spaces from linear algebra. Define a matroid by letting the ground set E be the set of vectors, and letting the independent sets in $\mathcal{I}$ be the linearly independent vector subsets. It is not hard to verify that the bases of the matroid coincide with the bases of the vector space, and that the circuits of the matroid coincide with the minimal linearly dependent sets of vectors.

2. Undirected graphs from graph theory. Define a matroid by letting the ground set E be the set of graph edges, and letting the independent sets in $\mathcal{I}$ be all the forests in the graph. It is not hard to verify that the bases of the matroid coincide with the spanning forests of the graph, and that the circuits of the matroid coincide with the simple cycles of the graph.

As discussed in Section 1.4.2 of the Introduction, there is a close relation between matroids and the greedy algorithm. Consider a matroid whose elements have non-negative weights. The greedy algorithm can be used to find a maximum-weight basis – for example, a maximum-weight spanning forest – by starting from the empty set and repeatedly adding a maximum-weight element among the elements whose addition would preserve the independence of the set. Moreover, matroids are precisely the set system for which such a greedy algorithm works for all weights [Edmonds, 1971]. For more on matroids in the context of mechanism design see [Bikhchandani et al., 2011].

2.1.2 Multi-Parameter Matching Environments

In a matching environment there are m different items for sale with one unit available of each item. As a convention we use the index i for bidders and j for items. Feasible allocations are all matchings of items to bidders (each bidder wins at most one item and each item is allocated to at most one bidder). This models unit-demand bidders. In Chapter 3 we will sometimes impose an additional supply limit of $\ell \leq m$, restricting the matchings to size at most ℓ.

Every bidder i has a private value $v_{i,j} \in [0, \infty)$ for winning item j, which is drawn independently at random from a distribution $F_{i,j}$ with a smooth density function $f_{i,j}$ positive over a nonzero interval support. A matching environment is thus multi-parameter.

We say the bidders of a matching environment are i.i.d. (or symmetric) if $F_{i,j}$ does not depend on the identity of bidder i, i.e., each item j has an associated distribution F_j and $F_{i,j} = F_j$. In other words, for every item j the values $\{v_{i,j}\}_{i \in [n]}$ are i.i.d. samples from F_j. Note that different items j, j' have different distributions $F_j, F_{j'}$, as is necessary for most applications (e.g., the travel website example in Section 3.1 below), and that independence of the values is maintained across both bidders and items.

2.2 Optimal Mechanism Design

2.2.1 Mechanisms

Without loss of generality we restrict attention to direct mechanisms, which receive a vector of bids $\mathbf{b}$. In the single-parameter case $\mathbf{b} \in \mathbb{R}^n_{\geq 0}$ where b_i is bidder i's bid for winning, and in the matching case $\mathbf{b} \in \mathbb{R}^{nm}_{\geq 0}$ where $b_{i,j}$ is bidder i's bid for winning item j. The mechanisms we design are deterministic and comprise:

1. An allocation rule $\mathbf{x} = \mathbf{x}(\mathbf{b})$, which maps a bid vector $\mathbf{b}$ to a feasible allocation; in the single-parameter case $\mathbf{x} \in \{0,1\}^n$, where $x_i = x_i(\mathbf{b})$ indicates whether bidder i wins, and in the matching case $\mathbf{x} \in \{0,1\}^{nm}$, where $x_{i,j} = x_{i,j}(\mathbf{b})$ indicates whether bidder i wins item j.

2. A payment rule $\mathbf{p} = \mathbf{p}(\mathbf{b})$, which maps a bid vector $\mathbf{b}$ to a payment vector. The payment vector $\mathbf{p}$ belongs to $\mathbb{R}^n_{\geq 0}$, where $p_i = p_i(\mathbf{b})$ is the payment charged to bidder i.

Fixing a bid vector $\mathbf{b}$, the mechanism's *welfare* in the single-parameter case is $\sum_i x_i v_i$, and in the matching case $\sum_{i,j} x_{i,j} v_{i,j}$. The mechanism's *revenue* is $\sum_i p_i$.

Bidder i's utility in the single-parameter case is $x_i v_i - p_i$, and in the matching case $\sum_j x_{i,j} v_{i,j} - p_i$.

A mechanism is dominant strategy truthful, or simply *truthful*, if for every bidder i and bid profile $\mathbf{b}_{-i}$ of the other bidders, i maximizes his utility by participating and bidding truthfully, i.e., bidding $b_i = v_i$ in the single-parameter case and $b_{i,j} = v_{i,j}$ for all j in the matching case. The other solution concepts of ex post and Bayesian truthfulness are formally defined in Section 4.3.1 below. Due to these solution concepts, we do not distinguish between bids and values, using v_i or $v_{i,j}$ to denote both.

We will mainly be interested in a mechanism's *expected* revenue $\mathbb{E}_{\mathbf{v}}[\sum_i p_i]$, where $\mathbf{p} = \mathbf{p}(\mathbf{v})$ and the expectation is taken over i.i.d. values drawn from the value distributions.

2.2.2 Maximizing Welfare and the Vickrey Auction

The general form of the Vickrey auction is called the VCG mechanism, and it works for any environment whether single-parameter or multi-item. VCG is remarkable in being both truthful and welfare-maximizing for every value profile $\mathbf{v}$. Its allocation rule chooses a feasible allocation that maximizes welfare; its payment rule charges every bidder i a payment equal to i's *externality* – the difference in the maximum welfare *of the other bidders* when i does not participate in the auction and when i does participate in it.

In the context of matching environments, the VCG allocation rule can be implemented as a maximum weighted matching over a complete bipartite graph $G = (V, \mathcal{E})$, where vertices on one side are the bidders, vertices on the other side are the items, and the weight of every edge $(i, j) \in \mathcal{E}$ is $v_{i,j}$ [Bertsekas, 1991]. The payment rule also solves bipartite matching problems to compute the payments: Let $W(\cdot)$ be the weight of a maximum weighted matching; then the payment of bidder i matched to item j is

$$W(V, \mathcal{E} \setminus (i, j)) - [W(V, \mathcal{E}) - v_{i,j}], \tag{2.1}$$

i.e., the weight of a maximum weighted matching after removing edge (i, j) from the graph, minus the weight of a maximum weighted matching in the original graph excluding the contribution of edge (i, j).

For single-parameter k-unit environments, Vickrey's allocation rule finds k bidders with highest values, and for matroid environments it uses a simple greedy algorithm to find a feasible allocation with highest welfare (and similarly for the Vickrey payment rules).

2.2.3 Maximizing Revenue and Myerson's Mechanism

For single-parameter environments, Myerson [1981] characterized the truthful, deterministic mechanism that maximizes expected revenue. In fact, his mechanism maximizes expected revenue over all Bayesian truthful, randomized mechanisms. Let F be a regular distribution with density f (regularity is defined below in Section 2.2.4). Define its *virtual value* function $\phi_F : \mathbb{R}_{\geq 0} \to \mathbb{R}$ to be

$$\phi_F(v) = v - \frac{1 - F(v)}{f(v)}. \tag{2.2}$$

For example, if F is the uniform distribution over $[0, 1]$, then $\phi_F(v) = v - \frac{1-v}{1} = 2v - 1$. Myerson showed the following:

Lemma 2.2.1 (Myerson)**.** *Given a single-parameter environment and a truthful mechanism* $(\mathbf{x}, \mathbf{p})$*, for every bidder* i *and value profile* $\mathbf{v}_{-i}$ *of the other bidders,* $\mathbb{E}_{v_i \sim F_i}[p_i(\mathbf{v})] = \mathbb{E}_{v_i \sim F_i}[x_i(\mathbf{v})\phi_{F_i}(v_i)]$.

Myerson's lemma says that in expectation over bidder i's value, his payment is equal to his virtual value when he is allocated. By summing over all bidders, this lemma implies that in expectation over the value profile, maximizing the revenue is equivalent to maximizing the total virtual value of allocated bidders, a quantity known as the *virtual surplus*. *Myerson's mechanism* maximizes expected revenue by finding the feasible allocation with maximum virtual surplus. For example, in a k-unit environment this will be the set of $\leq k$ bidders with the highest *positive* virtual values.

2.2.4 Regularity

We say that bidders are regular if their values are drawn from regular distributions:

Definition 2.2.2 (Regular Distribution)**.** A distribution F is regular if its virtual value function is monotone non-decreasing.

Roughly, regular distributions are those whose tail is no heavier than that of a power law distribution. The assumption that bidders are regular is standard in optimal mechanism design, analogous to that of downward-sloping marginal revenue in monopoly theory [e.g., Bulow and Klemperer, 1996]. Most commonly-studied distributions are regular, including the uniform, exponential and normal distributions, and distributions with log-concave densities. Without it, no good robust revenue guarantees are possible (see Section 2.3.1 for a negative example).

2.3 Prior-Independent Robustness

2.3.1 Definition

Robustness has been a long-time goal of mechanism and market design. Intuitively, robust mechanisms are mechanisms that "perform well" for a "large range" of economic environments. Their performance is insensitive to the environment's precise details and for this reason robustness is also referred to as *detail-freeness*. To formulate robustness one must specify what it means to perform well and for which range of environments should the performance guarantee hold. There are several alternative formulations in the literature, including robust optimization and others (see related work in Sections 2.3.3 and 3.1.5).

To define the robustness notion of *prior-independence*, we focus for simplicity on the single good case with symmetric buyers. Consider first a particular distribution F from which the buyers' i.i.d. values for the good are drawn. Let OPT_F be the optimal expected revenue that a truthful deterministic mechanism, which has *full knowledge of* F, can achieve in this environment. Let $\alpha \in (0, 1]$ be an approximation factor. A mechanism is *α-optimal* with respect to F if its expected revenue is at least $\alpha \, \mathrm{OPT}_F$.

This is similar to average-case approximation in combinatorial optimization, where an algorithm's approximation guarantee holds for inputs drawn from a known distribution; the difference is that the benchmark OPT_F is with respect to a mechanism instead of an algorithm.

Now let $\mathcal{F}$ be a set of value distributions, called the *priors*. A mechanism is *robustly α-optimal with respect to $\mathcal{F}$* if for every distribution F in this set, the mechanism is α-optimal with respect to F. In this case we also say that it gives an α-approximation to the optimal expected revenue. We have thus defined what it means for a robust mechanism to "perform well": it must achieve expected revenue that either exceeds or approximates the optimal expected revenue *simultaneously for every distribution* in a class of distributions $\mathcal{F}$. An example of the "large range" of distributions $\mathcal{F}$ is all regular distributions.

The above definition of robustness is an interesting mixture of average- and worst-case guarantees. On one hand, performance is measured in expectation over the random input; on the other it is measured in the worst case over all distributions that belong to $\mathcal{F}$. Such robustness is referred to as prior-independence, since there is an underlying assumption that values are sampled from priors, and yet the robust mechanism must be independent of the priors as it must work well for all of them.

Necessity of Regularity. Regularity of the distributions in $\mathcal{F}$ is in fact necessary for designing good prior-independent mechanisms, as demonstrated by Dhangwatnotai et al. [2015]: Fix a value z and a number n of bidders, and define an irregular, long-tailed value distribution F_z such that the probability for z is $1/n^2$ and otherwise the value is zero. Consider a single-item environment with n bidders whose values are drawn from F_z. The optimal auction has expected revenue at least z/n. But any prior-independent truthful auction essentially has to "guess" the value of z, since the probability that the non-winning bids provide information about z is small. Thus its expected revenue cannot be within a constant factor of z/n for every F_z.

2.3.2 Rationales

What are the rationales behind prior-independent robustness? In particular, why measure whether a robust auction is performing well by comparing it to the optimal mechanism with access to the prior distribution? And why choose a large range of distributions with minimal assumptions (rather than incorporate partial information that the seller may have about the distributions in order to narrow it down)?

First, our main results show that for a large range of distributions on which little is assumed, we can get a constant approximation to the ambitious benchmark of OPT_F, even in challenging environments like multi-item markets for which the optimal mechanism remains elusive. The two choices above thus serve to strengthen our results.

An alternative approach to robustness could be, rather than to *approximate* the optimal mechanism for every distribution in $\mathcal{F}$, to design a mechanism that *maximizes* the minimum expected revenue where the minimum is taken over all distributions in the set $\mathcal{F}$. Such an approach would run into the open problem of finding the optimal mechanism for multi-item markets, even in the degenerate case where $\mathcal{F}$ contains only one distribution. Seeking approximation rather than maximization is what enables us to circumvent the open problem. In addition, choosing OPT_F as a benchmark allows the seller to make informed decisions regarding how much to invest in obtaining information on F.

The choice of setting $\mathcal{F}$ to be the class of *all* regular distributions has the advantage of capturing situations in which the seller has no information about the value distribution, as is the case for a new seller or new good on the market and for goods whose distribution is constantly shifting, as well as situations in which the seller's information is prohibitively expensive or highly noisy and thus too risky to rely upon. In other words, the same reason for avoiding dependence on prior distributions in mechanism design – lack of reliable, accessible information – also justifies avoiding dependence on prior information about the distributions. Moreover, assuming no partial information leads to simple and natural mechanisms, thus reinforcing our chosen robustness notion.

2.3.3 Robustness in the Literature

Given its importance, it is not surprising that there is a rich literature on robustness in mechanism design. Our approach contributes to this literature, which has focused on single-parameter environments, by simultaneously achieving robustness and simplicity while being applicable to multi-parameter environments (Chapter 3) and interdependent values (Chapter 4).

Prior-Independence for Single-Parameter Environments. For single-item and other single-parameter environments, Bulow and Klemperer [1996] were the first to study the effect of augmented demand. A simplified proof of their main result was given by Kirkegaard [2006]. Dughmi et al. [2012] generalize this result to environments with matroid-based constraints, and use the generalized version to investigate conditions under which the Vickrey auction inadvertently yields approximately-optimal revenue. Hartline and Roughgarden [2009] develop a similar result for asymmetric buyers who do not share the same value distributions. Fu et al. [2015] study the tightness of the Bulow and Klemperer result. Using techniques related to an early version of Chapter 3 [Roughgarden et al., 2012], Sivan and Syrgkanis [2013] develop a version of Bulow and Klemperer's result for *convex combinations* of distributions satisfying regularity.

A natural approach to prior-independent robustness is to instantiate the optimal mechanism developed by Myerson [1981] with an empirical rather than known distribution, where the samples come from the buyers' bids. This approach is explored by Segal [2003] and Baliga and Vohra [2003], and is asymptotically optimal as the size of the market goes to infinity. It also allows the seller to incorporate into the mechanism prior information about the class of possible value distributions (a "higher level" of prior information – not about the possible values but rather about their possible distributions). Dhangwatnotai et al. [2015] take the sampling approach further by designing simple mechanisms using only a single sample, which are nevertheless robustly α-optimal for small constant $1/\alpha$ parameters.

Other Robustness Notions for Single-Parameter Environments. Value distributions are not used in the design of prior-independent auctions, but they are used in their analysis, namely in the definition of "robustly α-optimal" which is based on comparison to the optimal expected revenue. In prior-*free* auction design, distributions are not even used to evaluate the performance of an auction. This raises the question of which benchmark to use. Neeman [2003] studies the revenue performance of the English auction – an ascending-auction version of the Vickrey auction – comparing it to the benchmark of welfare, which is clearly an upper bound on revenue in auctions in which buyers have no better choice than to participate. A different approach is initiated by Goldberg et al. [2006], who define a notion of "reasonable" auctions to compete against. The relations between the different notions of robustness has also been studied [see, e.g., Hartline and Roughgarden, 2014].

Part I

Informational Challenges in Revenue Maximization

3

Robust Revenue Maximization in Matching Markets

3.1 Chapter Introduction

3.1.1 Overview

Most results in revenue-maximizing mechanism design hinge on "getting the price right" – selling goods to bidders at prices low enough to encourage a sale, but high enough to garner non-trivial revenue. This approach is difficult to implement when the seller has little or no a priori information about bidders valuations, or when the setting is sufficiently complex, such as matching markets with heterogeneous goods. In this chapter we apply a robust approach to designing auctions for revenue. Instead of relying on prior knowledge regarding bidder valuations, we "let the market do the work" and let prices emerge from competition for scarce goods. We analyze the revenue guarantees of one of the simplest imaginable implementations of this idea: first, enhance competition in the market, whether by increasing demand or by limiting supply; second, run a standard second-price (Vickrey) auction. Enhancing competition is a natural way to bypass lack of knowledge – a seller who does not know how to set prices can instead set quantities (of bidders or goods on the market). We prove that simultaneously for many valuation distributions, this achieves expected revenue at least as good as the optimal revenue in the original market or guarantees a constant approximation to it. Our robust and simple approach provides a handle on the elusive optimal revenue in multi-item matching markets, and shows when the use of welfare-maximizing Vickrey auctions is justified even if revenue is a priority. By establishing quantitative trade-offs, our work provides guidelines for a seller in choosing among different revenue-extracting strategies: sophisticated pricing based on market research, advertising to draw additional bidders, and limiting supply to create scarcity on the market.

3.1.2 The Problem

Consider a set of m indivisible goods for sale, and the problem of matching them to n buyers with private values, each of whom wants no more than a single good. This problem has been studied extensively with respect to the goal of maximizing economic efficiency; e.g., this is the topic of the classic paper on "Multi-Item Auctions" of Demange et al. [1986]. Here we focus on an alternative important goal – maximizing the seller's revenue.

To demonstrate our setting, defined formally in Section 2.1.2 above, consider Scenario 1.3.1: A for-profit travel website selling overnight accommodation is faced with the task of assigning m available rooms to n interested buyers. Each buyer needs a single room for the night, and has different private values for different rooms based on their location, size etc. Uncertainty of the seller regarding buyer values is captured by the probabilistic Bayesian model, in which the values for every good $j \in [m]$ are assumed to be independent draws from a distribution F_j (where F_j satisfies regularity). The fact that F_j is common to all buyers makes our model symmetric with respect to buyers (but not with respect to goods). The seller wishes to maximize its expected revenue by designing a deterministic auction in which no buyer can do better than to participate and reveal his true values (i.e., dominant strategy truthful).[1]

When $m = 1$, that is when there is a single good on the market, Myerson [1981]

[1]Recall that by the revelation principle, this requirement is without loss of generality for a seller seeking a dominant strategy implementation.

characterizes the revenue-optimal truthful auction under the assumption that the distribution F_1 from which values for the good are drawn is fully known to the seller. The optimal auction in this case turns out to be the well-known second-price auction [Vickrey, 1961], with an additional reserve price r tailored to the distribution F_1. The resulting auction is very simple: the bidders report their values to the seller, the bidder with the highest bid above r wins, and the winner pays the second-highest bid above r if there is one or r otherwise. Myerson's characterization of optimal mechanisms also applies to markets with multiple copies (units) of the single good, where each bidder seeks at most one copy. More generally, it applies to all single-parameter markets, in which every bidder can either win or lose and has a single private value for winning.[2]

Since Myerson's seminal work there have been efforts to extend it in several directions. A direction that has attracted much attention is to generalize the optimal auction characterization *beyond the* $m = 1$ *case, to multi-parameter markets* [e.g., Vincent and Manelli, 2007]. In particular, there is no known characterization for the matching markets described above, in which there are multiple goods and each buyer seeks at most one good. Another important direction that has become known as "Wilson's doctrine" is to design alternative, *robust* auctions for revenue, in the sense that they do not depend on the seller's full knowledge of the value distributions [Wilson, 1987]. A third direction is inspired by the *simplicity* of Myerson's auction – a second-price auction with reserve – and aims to design similarly simple auctions for revenue in more general settings [e.g., Hartline and Roughgarden, 2009].

In this work we contribute to all three goals above by applying a robust approach to revenue maximization. We develop a framework for designing mechanisms that are robust, simple, and guaranteed to work well for a variety of market environments including matching markets. Our mechanisms are based on the natural idea of enhancing bidder competition[3] for the goods, either by adding competing bidders in the

[2]The Myerson characterization also extends to *asymmetric* environments, where there are multiple distributions $\{F_1^i\}_{i\in[n]}$ and bidder i's value for the good (good 1) is drawn from his distribution F_1^i. However, in this case, the informational burden on the seller is even heavier, as it needs to have full knowledge of the distributions of all buyers, and even getting only approximately close to optimal revenue requires many samples from every distribution [Cole and Roughgarden, 2014].

[3]Note that while we enhance competition among buyers in the market, we do not turn it into

manner of Bulow and Klemperer [1996] or by artificially limiting the supply, and then running a variant of the Vickrey auction. Despite avoiding any reference to the value distributions, the expected revenue achieved by these mechanisms exceeds or approximates the expected revenue of the optimal mechanisms tailored to the distributions in the original environments. Besides leading to good mechanisms, our approach sheds light on trade-offs among possible seller strategies, including how many more buyers are needed, or how many units of a good to produce relative to the market size, in order to replace the need to rigorously learn the preferences of existing buyers.

We now demonstrate our approach via a simple motivating example.

3.1.3 Motivating Example: Multi-Unit Markets

As a simple motivating example we consider symmetric multi-unit markets (Example 3.1.1), and in particular the special case of digital goods (Example 3.1.2), as well as a generalized example with feasibility constraints (Example 3.1.3).

Example 3.1.1 (Multi-unit)**.** There are k identical copies of a single good for sale, and $n \geq k$ bidders who each want at most one unit. The bidders' values for a unit are i.i.d. samples from the value distribution.

For example, the units can be identical rooms at a large hotel. Another example is copies of a digital good such as an e-book, in which case there is no limit on the number of copies that can be made:

Example 3.1.2 (Digital goods)**.** A multi-unit market with $k = n$ units, where n is the number of bidders.

We also consider a generalization in which there are feasibility constraints, i.e., not all sets of bidders can feasibly be allocated units, even if they include less bidders than the number of available units. The next example demonstrates the kind of matroid feasibility constraints we consider:

a "competitive market" in which the buyers become price-takers. The focus of our work is not on large markets and asymptotic results, rather we aim for our results to hold in markets of any size.

Example 3.1.3 (Job Scheduling)**.** A multi-unit market where the units are slots for running jobs on a machine, and a subset of bidders is feasible if each bidder's job can be matched to a slot between its arrival time and deadline.

Technically, a multi-unit market with no constraints can be thought of as finding a matching of bidders to units in a complete bipartite graph, while in the job scheduling example the corresponding bipartite graph has some suitable structure.

Our Approach. We present two approaches to robust revenue maximization in the above examples: augmenting demand and limiting supply. Both approaches are inspired by common business practices – augmenting demand corresponds to advertising the auction and drawing more participants, and limiting supply corresponds to practices like "limited editions", limited runs of artwork, or artificial scarcity (for an example of this phenomenon in the diamond market see [McEachern, 2012]; other examples include scarcity of newly-launched technology products, etc.). On a theoretical level, both approaches rethink the standard definition of an auction environment, in which the demand and supply are considered exogenous, treating these instead as an endogenous part of the mechanism design problem.

Augmenting Demand in Multi-Unit Markets

A well-known result states the following:

Theorem 3.1.4 (Bulow and Klemperer [1996])**.** *When selling a single good to bidders whose values are i.i.d. draws from a distribution satisfying regularity, the expected revenue of the revenue-optimal mechanism with n bidders is at most that of the Vickrey auction with $n + 1$ bidders.*

In other words, when the demand is augmented by adding a single additional bidder competing for the good, the simple Vickrey auction achieves at least the maximum revenue possible with the original demand. This is despite being oblivious to the value distribution, whereas the optimal Myerson mechanism for n bidders depends on this knowledge to set the reserve price.

We remark that the regularity constraint on the distribution is standard; for a bidder whose value is drawn from the distribution, regularity means that the *revenue curve* describing the trade-off between selling to the bidder often at a low price and selling less often at a higher price is concave. This is satisfied by all common distributions (uniform, normal, power-law, etc.), and without it no result along the lines of Theorem 3.1.4 is possible – see Section 2.2.4 for details.

The Bulow and Klemperer theorem generalizes to the multi-unit setting in Example 3.1.1: When there are k units of the good, the expected revenue of the revenue-optimal mechanism with n bidders is at most that of the Vickrey auction with $n + k$ bidders [Bulow and Klemperer, 1996]. It also applies to constrained settings such as the job scheduling example (Example 3.1.3): When the best schedule is able to match $\rho \leq k$ bidders to the k slots without violating an arrival time/deadline, the expected revenue of the revenue-optimal mechanism with n bidders is at most that of the Vickrey auction with $n + \rho$ bidders, where the ρ additional bidders can be scheduled simultaneously [Dughmi et al., 2012].

In general, a *Bulow-Klemperer-type* theorem states that instead of running the optimal mechanism on the original auction environment, we can get as much revenue in expectation by running a variant of the Vickrey auction on an environment suitably augmented with additional bidders. This can be seen as a theoretical justification to treat bidder participation in auctions as a first-order concern when aiming for revenue, perhaps even at the expense of sophisticated pricing.

Limiting Supply in Multi-Unit Markets

The flip side of increasing demand is limiting supply.

Algorithm 1: Supply-limiting mechanism

1. Set a supply limit $\ell = n/2$ equal to half the number of bidders.
2. Run the Vickrey auction subject to supply limit ℓ.

The mechanism in Algorithm 1 is a supply-limiting mechanism for digital goods (Example 3.1.2). In the second step of the algorithm, the *Vickrey auction subject to* *supply limit* ℓ is a simple variation on the standard second-price auction: it assigns copies to the ℓ buyers with highest bids (even though there are enough copies for all buyers), and charges them each the $(\ell + 1)$th highest bid. The resulting mechanism is simple and natural, and does not rely on knowledge of value distributions.

Intuitively, enhancing competition by limiting the supply has a similar effect on revenue as enhancing competition by adding bidders in Bulow-Klemperer-type theorems. The difference between the two approaches is that the former requires augmenting the resources – in this case bidders – available to the auction, while the latter requires the ability to withhold supply (many sellers, e.g. companies like `Apple Inc.`, have the ability to do both). This difference translates into different revenue guarantees: while in Bulow-Klemperer-type theorems the expected revenue of the augmented Vickrey auction usually exceeds that of the optimal mechanism in the original environment, the expected revenue of the supply-limiting mechanism approximates that of the optimal mechanism. In particular, *the supply-limiting mechanism achieves at least half of the optimal revenue in expectation, despite remaining oblivious to the value distribution on which the optimal mechanism depends.*

The performance guarantee for the mechanism in Algorithm 1 can be seen as justification to the following "rule of thumb" for sellers: assuming no production costs and buyers whose values are distributed similarly, produce a number of units equal to a constant fraction of the market size. I.e., when sellers lack the necessary information to set prices, they can set quantities instead, and this works well simultaneously for many value distributions.

The Connection between Augmenting Demand and Limiting Supply

Our technical approach to establishing approximation guarantees of supply-limiting mechanisms utilizes the intuition above, by which limiting supply has a similar effect as increasing demand. This intuition is formulated by basing the proofs of the approximation factors on a reduction among markets, which enables the application of an appropriate Bulow-Klemperer-type theorem.

In particular, the reduction in Algorithm 2 shows that the mechanism in Algorithm 1 guarantees half the optimal revenue in expectation as follows: Starting with the

Algorithm 2: Reduction from limited supply to augmented demand for digital goods

0. **Start with original market with n buyers and n units**

   Denote the optimal expected revenue by OPT.

1. **Restrict to market with $n/2$ buyers and supply limit $n/2$**

   The optimal expected revenue is at least OPT /2, by subadditivity of revenue in the buyers (Lemma 3.5.2 below) and since the supply limit has no effect here.

2. **Augment to get market with n buyers and supply limit $n/2$**

   The expected revenue of the Vickrey auction is at least OPT /2, by the Bulow-Klemperer-type theorem for multi-unit markets applied to the restricted market.

original market, define a new market by dropping half of the bidders and setting a supply limit of $\ell = n/2$. Consider the resulting *restricted market* with half of the original bidders and corresponding supply limit. One can show that if we were to restrict the optimal mechanism to run on this market instead of the original one, its expected revenue would have been at least half of its original expected revenue. Now conceptually add back the $n/2$ removed bidders but without changing the supply to get the *augmented market*, and run the Vickrey auction. It follows from the Bulow-Klemperer theorem for multi-unit markets that the expected revenue is at least as high as the optimal expected revenue for the restricted market. Therefore the supply-limiting mechanism guarantees at least half of the optimal expected revenue in the original market.

3.1.4 Our Contribution

As demonstrated in Section 3.1.3, *our main contribution is in formulating and proving robust revenue guarantees of competition enhancement in auctions, through increased demand or limited supply, for a variety of markets, including types of markets where the optimal mechanism remains unknown (and is presumably very complex).* We show that under minimal regularity assumptions, the simple and robust mechanisms above and their revenue guarantees – Vickrey with additional bidders and Vickrey with a supply limit – generalize to significantly more complex settings. In other words, we identify markets in which such mechanisms are guaranteed to achieve optimal or approximately-optimal expected revenue. We remark that by using these mechanisms for revenue, the seller also guarantees that the welfare is approximately optimal, and in fact one can achieve other trade-offs between revenue and welfare by setting suitable supply limits.

Our technical contribution is in proving novel Bulow-Klemperer-type theorems for different markets, and designing supply-limiting mechanisms whose approximation guarantees follow from the Bulow-Klemperer-type theorems. While Bulow-Klemperer-type theorems have been studied before, they have never been attempted beyond single-parameter buyers, i.e., for multiple different goods. To our knowledge, supply-limiting mechanisms have also not been studied before, nor has the connection between increasing demand and limiting supply been explicitly formulated as in our reductions.

Proving Bulow-Klemperer-type theorems for matching markets is the most technically challenging component of this work. The analysis for multi-parameter settings is challenging due to dependency issues – the competition for item j that drives its price depends on the buyers' values for the other items. We overcome dependency challenges via a technique from the analysis of randomized algorithms called the principle of deferred decision, combined with the combinatorial properties of optimal matchings.

Results: Augmented Demand

We prove the first generalization of Bulow and Klemperer's theorem (Theorem 3.1.4) to multi-parameter markets.

Theorem 3.1.5 (Bulow-Klemperer-Type Theorem for Matching Markets (Informal))**.** *For every matching market with n bidders and m goods, assuming symmetry and regularity, the expected revenue of the Vickrey auction with m additional bidders* *is at least the optimal expected revenue in the original market.*

The formal statement appears in Theorem 3.4.1. We emphasize that the symmetry assumption in this theorem is across bidders, not goods. That is, values of different bidders for the same good are i.i.d. samples from the same distribution, but different goods can have different value distributions. This kind of symmetry makes sense in practical applications, where the seller knows it is selling very different kinds of goods, but sees the bidders – whose identities and characteristics are unknown – as homogeneous [Chung and Ely, 2007].

In addition to Theorem 3.1.5, we prove Bulow-Klemperer-type theorems that achieve better guarantees for matching markets with more supply than demand ($n \leq m$), and that apply to *asymmetric* markets where bidders' values for a good may belong to different distributions (see Section 3.6).

Results: Limited Supply

We design supply-limiting mechanisms for both single-parameter and multi-item markets. The former include digital good markets (recall the mechanism in Algorithm 1), as well as more general multi-unit markets, possibly with constraints or asymmetric bidders (see Section 3.3).

For multi-item matching markets, we first define a notion of setting a limit on supply where the supply is heterogeneous rather than homogeneous. A multi-item auction *subject to supply limit* ℓ means that no more than ℓ goods may be assigned, with no limitation on which ℓ goods these shall be. Intuitively, this lets the market do the work of choosing which part of the supply to limit. This is in line with our robust approach, as a seller with no knowledge of how the values for the different goods are distributed cannot make this decision without risking a big loss in revenue. Notice that the simple supply-limiting mechanism we designed for multi-unit markets (Algorithm 1) is now well-defined for multi-item markets as well, and we can prove the following theorem:

Theorem 3.1.6 (Supply-Limiting Mechanism for Matching Markets (Informal)). *For every matching market with $n \geq 2$ bidders and m goods, the expected revenue of* *the mechanism in Algorithm 1 is at least a constant fraction of the optimal expected revenue.*

Qualitatively, Theorem 3.1.6 is interesting since it shows that a simple robust mechanism can achieve a fraction of the optimal expected revenue that is *independent of the size of the market*, as measured by parameters n and m. Moreover, the constant fractions we achieve are quite good in many cases, e.g., we achieve a fraction of 1/4 when the number of bidders equals the number of goods (Theorem 3.5.1). An interesting open problem is whether this is the best possible by any robust mechanism. The analysis of the approximation guarantees are via a general reduction (the reduction in Algorithm 3, along the lines of the reduction in Algorithm 2), instantiated with appropriate Bulow-Klemperer-type theorems.

Algorithm 3: General reduction from limited supply to augmeneted demand

0. **Start with original market with n buyers**

   Denote the optimal expected revenue by OPT.

1. **Restrict to market with $< n$ buyers and supply limit ℓ**

   The optimal expected revenue is a constant fraction of OPT, by subadditivity.

2. **Augment to get market with n buyers and supply limit ℓ**

   The expected revenue of the Vickrey auction is a constant fraction of OPT, by a suitable Bulow-Klemperer-type theorem.

3.1.5 Related Work

In contemporaneous work with an early version of this paper, Devanur et al. independently consider a similar set of problems as us, but using different mechanisms and analyses (2011). Their mechanisms are arguably more complex and less natural since they are based on carefully-constructed price menus (rather than on enhanced competition). Following our early version, Azar et al. [2014] studied matching markets in which partial information about the value distributions is available to the seller in the form of a limited number of samples.

Closely related to our work, Bandi and Bertsimas [2014] apply the *robust optimization* approach to optimal mechanism design in multi-item markets. Their model differs from our model in several important aspects, including consideration of additive valuations rather than matching markets, and divisible rather than indivisible goods. A main goal of their paper is orthogonal to ours – to study the important issues of budgets and correlated values in mechanism design. They also address auctions without budget constraints but in their setting these reduce to single-item auctions, which is far from the case in our model.

It is also interesting to compare our robustness notion to theirs, where the latter is inspired by the robust optimization paradigm. Bandi and Bertsimas model the seller's knowledge about the values by an uncertainty set, thus accommodating for partial knowledge based on historical bidding data, and then optimally solve the related robust optimization problem. They use simulations to show that their robust optimization approach improves upon the revenue performance of Myerson's mechanism for single items, when the seller's knowledge of the prior distribution is inaccurate. We use a different notion of robustness inspired by approximation algorithms for combinatorial optimization, and our goal for single items is to surpass or approximate the performance of Myerson's mechanism tailored to the accurate distribution (which our mechanism is oblivious to).

Recently there have been significant advances on the problem of prior-*dependent* optimal mechanism characterization for multi-item markets. Cai et al. [2013] give a characterization for optimal mechanisms given access to the prior distributions, and with the relaxed requirement of *Bayesian*, rather than dominant strategy, truthfulness. The relaxed truthfulness notion requires that no buyer can do better *in expectation over the other buyers' valuations* than to participate and bid truthfully in the auction. It thus relies on common knowledge of the prior distributions among the buyers as well as for the seller [*cf.* Chung and Ely, 2007].

Chawla et al. [2010a] give an upper bound on the optimal expected revenue for matching markets, and our techniques utilize one of their reductions. They achieve a prior-*dependent* $1/6.75$-approximation for matching markets with multiple units and asymmetric buyers, and also a $3/32$-approximation for an even more general environment (namely a graphical matroid with unit-demand buyers).

As mentioned above, in the matching markets we consider, many of the complications of VCG do not occur, namely, communicating the bids and running the auction are both computationally tractable, and our competition enhancement methods ensure that the revenue does not collapse. Chawla et al. [2013] analyze the VCG mechanism's performance in a job scheduling context, and some of our techniques are inspired by their analysis.

3.1.6 Organization

Section 3.2 contains several preliminaries. Section 3.3 includes our analysis of competition enhancement for multi-unit markets. Sections 3.4 and 3.5 analyze multi-item matching markets where n is proportional to m and contain our main technical results for increasing demand and limiting supply, respectively. Extensions and generalizations can be found in Section 3.6.

3.2 Preliminaries

Notation. For our supply-limiting mechanisms, we add to the VCG or Vickrey mechanisms a supply limit ℓ and denote them by $\mathrm{VCG}^{\leq\ell}$ and $\mathrm{Vic}^{\leq\ell}$, respectively.

Revenue Benchmark. The revenue benchmark against which we measure the performance of our mechanisms is the following: By optimal expected revenue we mean the maximum expected revenue over all dominant strategy truthful, deterministic mechanisms. Chawla et al. [2010b] show that for matching environments, the expected revenue from the optimal deterministic mechanism is within a constant factor of the expected revenue from the optimal randomized mechanism. Thus our results for deterministic mechanisms apply to randomized mechanisms up to a constant factor.

Technical Tool: Representative Environments. A representative environment is the "single-parameter counterpart" of a matching environment. Consider a matching environment with m items, n symmetric bidders and value distributions $\{F_j\}_{j=1}^m$.

The corresponding representative environment has the same m items, but nm single-parameter bidders – every bidder in the matching environment has m *representatives* in the representative environment. The jth representative of bidder i is only interested in item j and has a value $v_{i,j} \sim F_j$ for winning it. Every allocation in the representative environment can be translated to an allocation in the matching environment – if the jth representative of i wins, then item j is allocated to bidder i in the matching environment – and vice versa. An allocation in the representative environment is *feasible* if the corresponding allocation in the matching environment forms a matching, meaning that only one representative per bidder wins.

Intuitively, the representative environment is more competitive than the matching one, since representatives of the same bidder compete against each other on who will be the winner. Thus the expected revenue achievable in the representative environment should be at least the optimal expected revenue in the matching environment. Chawla et al. [2010a] formalize this intuition by showing that any truthful mechanism M for the matching environment translates to a truthful mechanism M^{rep} for the representative environment, such that the expected revenue of M^{rep} is only higher. Roughly this is by translating the allocation rule of M to an allocation rule in the representative environment as above, and viewing the payment rule of M as a price menu, whose prices are exceeded in the representative environment by charging every representative the minimum value it needs to bid in order to win.

Lemma 3.2.1 (Chawla et al. [2010a]). *The expected revenue of M^{rep} in the single-parameter representative environment is at least the expected revenue of M in the matching environment.*

3.3 Multi-Unit Markets

In this section we formally prove the results presented in Section 3.1.3, demonstrating our general framework. The approach of augmenting demand has been studied for multi-unit auctions, and a slightly generalized version of a result by Bulow and Klemperer [1996] is the following:

Theorem 3.3.1 (Bulow-Klemperer-Type Theorem for Multi-Unit Markets)**.** *For every k-unit environment with i.i.d. regular bidders and supply limit ℓ, the expected revenue of the Vickrey auction with* $\min\{k, \ell\}$ *additional bidders is at least the optimal expected revenue in the original market. In other words, Vickrey with* $\min\{k, \ell\}$ *additional bidders is robustly 1-optimal.*

As for limiting supply, we instantiate our general reduction (Algorithm 3) with the above Bulow-Klemperer-type theorem to prove the following. For simplicity of presentation assume the number of bidders n is even.[4]

Theorem 3.3.2 (Supply-Limiting Mechanism for Multi-Unit Markets)**.** *For every k-unit environment with $n \geq 2$ i.i.d. regular bidders, the expected revenue of the supply-limiting mechanism* $\mathrm{Vic}^{\leq n/2}$ *is at least a* $\max\{\frac{1}{2}, \frac{n-k}{n}\}$*-fraction of the optimal expected revenue. In other words, Vickrey with supply limit $n/2$ is robustly α-optimal for $\alpha = \max\{\frac{1}{2}, \frac{n-k}{n}\}$.*

In the above theorem, the supply limit of $n/2$ "kicks in" when the number of units k exceeds $n/2$, and in this case we get a $1/2$-approximation. If the supply k is limited to $n/2$ to begin with, the competition is inherently high and Vickrey with no supply limit provides an $\frac{n-k}{n}$-approximation.

Proof. We instantiate the reduction in Algorithm 3 as follows. To go from the original market to the restricted market, remove $\min\{\frac{n}{2}, k\}$ bidders from the original market, and if $k > \frac{n}{2}$ set a supply limit of $\ell = \frac{n}{2}$.

Analysis: We first claim that the restriction of the original market maintains at least a fraction of $\max\{\frac{1}{2}, \frac{n-k}{n}\}$ of the optimal expected revenue in the original market. This is because, as shown by Dughmi et al. [2012], the expected optimal revenue as a function of the bidder set is *submodular*.[5] Revenue submodularity means decreasing marginal returns to the expected revenue as more bidders are added, so the first $\max\{\frac{n}{2}, n-k\}$ bidders already capture at least a $\max\{\frac{1}{2}, \frac{n-k}{n}\}$-fraction of the optimal

[4]If n is odd, one can first remove a bidder from the environment, losing at most a $1/n$-fraction of the optimal expected revenue.

[5]Recall that a function f from sets of bidders to $\mathbb{R}$ is *submodular* if for every two sets $S \subset T$ and every bidder $i \notin T$ it holds that $f(S \cup \{i\}) - f(S) \geq f(T \cup \{i\}) - f(T)$. Equivalently, the marginal contribution of a bidder to the value of f is decreasing.

expected revenue. Limiting the supply to $\frac{n}{2}$ when $k > \frac{n}{2}$ has no effect since in this case the number of bidders remaining in the restricted environment is $\frac{n}{2}$.

We can now apply the Bulow-Klemperer-type theorem for multi-unit markets (Theorem 3.3.1) to the restricted environment. In the first case, $k > \frac{n}{2}$ and the restricted environment has $\frac{n}{2}$ bidders, k units and supply limit $\ell = \frac{n}{2}$. In the second case, $k \leq \frac{n}{2}$ and the restricted environment has $n - k$ bidders, k units and no supply limit (i.e., $\ell = k$). In both cases, by Theorem 3.3.1 the expected revenue of Vickrey with $\min\{\frac{n}{2}, k\}$ additional bidders is at least the optimal expected revenue in the restricted environment. So running Vickrey with $\min\{\frac{n}{2}, k\}$ additional bidders on the restricted environment is a $\max\{\frac{1}{2}, \frac{n-k}{n}\}$-approximation to the optimal expected revenue in the original environment. But this is equivalent to running the supply-limiting mechanism $\text{Vic}^{\leq n/2}$ on the original environment, completing the proof. □

The approximation factor in Theorem 3.3.2 is asymptotically tight:

Proposition 3.3.3. *For every $0 < \gamma < 1$, consider the supply-limiting mechanism $\text{Vic}^{\leq \gamma n}$. There exists an n-unit environment with n i.i.d. regular bidders such that the expected revenue of $\text{Vic}^{\leq \gamma n}$ is at most a $(\frac{1}{2} + o(1))$-fraction of the optimal expected revenue.*

Proof. Consider first the case that $1/n \leq \gamma \leq 1/2$, i.e., the supply limit is severe. Let the value distribution F be the uniform distribution over the support $[1, 1+\epsilon]$ for a sufficiently small parameter $\epsilon = \epsilon(n)$. The optimal expected revenue is roughly n, while $\text{Vic}^{\leq \gamma n}$ can extract as revenue at most $\gamma n(1+\epsilon) \leq n/2 + o(1)$.

Now suppose $1/2 < \gamma \leq \frac{n-1}{n}$. For sufficiently large H, let the value distribution be $F(z) = \frac{z}{1+z}$ over the support $[0, H]$ with a point mass of $\frac{1}{1+H}$ at H. The optimal expected revenue is at least the expected revenue achieved by offering a posted price H to every one of the n bidders, which extracts $H(1 - \frac{H}{1+H}) = \frac{H}{1+H} \approx 1$ from every bidder in expectation. In comparison, the expected revenue in $\text{Vic}^{\leq \gamma n}$ comes from the $(\gamma n + 1)$st highest bid. This bid is concentrated around $z = \frac{1-\gamma}{\gamma}$, the value of z such that $F(z) = 1 - \gamma$. So VCG achieves an expected revenue of roughly $\frac{1-\gamma}{\gamma}\gamma n = (1-\gamma)n < \frac{n}{2}$. □

3.4 Matching Markets: Augmenting Demand

In this section we prove a Bulow-Klemperer-type theorem for matching environments – the first generalization of Bulow and Klemperer [1996] to a multi-item market. Recall what we mean by i.i.d. bidders in a matching environment: different items have different distributions, but independence is both across bidders and across items.

Theorem 3.4.1 (Bulow-Klemperer-Type Theorem for Matching Markets)**.** *For every matching environment with i.i.d. regular bidders and m items, the expected revenue of the VCG mechanism with m additional bidders is at least the optimal expected revenue in the original market. In other words, VCG with m additional bidders is robustly optimal.*

Theorem 3.4.1 provides a simple handle on the unknown optimal expected revenue in matching markets. For example, in a market with two goods for sale, the best achievable revenue is at most what VCG can achive with two more bidders. Note that in markets with plentiful supply, i.e. markets in which $m \gg n$, the demand augmentation that is required is substantial. In Section 3.4.5 we present an alternative Bulow-Klemperer-type theorem with weaker requirements for this case.

3.4.1 Overview of the Proof

The proof is divided into two parts. In Section 3.4.2 we identify an upper bound on the optimal expected revenue in the original environment, and a lower bound on the revenue of the VCG mechanism in the augmented environment with m more bidders. These bounds are relatively simple to analyze and are already similar, though not identical, in form. In Section 3.4.3 we carefully relate the two bounds to establish the theorem.

Our proof is based on the following ideas. We first observe there is a simple upper bound on the optimal expected revenue in the matching environment – *the expected revenue from running m Vickrey auctions to sell each of the m goods to m separate sets of $n + 1$ representatives*, who are single-parameter bidders only interested in one particular good (Lemma 3.4.2). Our goal is now to show that VCG with a total of

m additional multi-parameter but unit-demand bidders does just as well in terms of revenue.

Recall that in the VCG mechanism, the winner of a certain good pays the externality he inflicts upon other bidders, which includes in particular the "damages" he causes the losing bidders who are not allocated any good by the mechanism. Thus, the payment for every good j is at least the highest value for j among the losers. In the augmented matching environment to which VCG is applied, it is guaranteed that there will be n losers, since there are m goods and $n + m$ bidders. The expected revenue from running VCG on the augmented environment is thus at least *the expected welfare from running m Vickrey auctions to allocate each of the m goods separately to the n losers* (Lemma 3.4.3). This lower bound is similar to the above upper bound.

The remaining challenge is a dependency issue – by definition, the losers are likely to have lower values for the goods than the $n + 1$ representatives. We use the combinatorial structure of maximum weighted matchings to show that a bidder's values conditional on him losing in the VCG mechanism are, while lower, not likely to be *significantly* so compared to the unconditional case. Thus the losers' damages are enough to cover the expected revenue from the representatives.

On a technical level what we show is that, quite remarkably, the only thing that can be deduced about a bidder's value for an item j from his losing the auction completely is that it is lower than the value of the winner of item j. We establish this by introducing an auxiliary selling mechanism for item j, conceptually and revenue-wise half-way between selling the item separately and selling it as part of the VCG mechanism. The auxiliary mechanism runs a maximum weighted matching algorithm as in VCG, but defers the sale of item j until all other goods have been sold and exactly $n + 1$ bidders remain unallocated. Thus, by construction, these bidders' values for item j are unaffected by the dependency issue described above.

3.4.2 Basic Upper and Lower Bounds

Upper Bound. To upper bound the expected optimal revenue we use the following notation: let $\mathrm{Vic}_j(\eta)$ be the expected revenue from selling item j to η i.i.d. single-parameter representatives with value distribution F_j using the Vickrey auction. Then:

Lemma 3.4.2 (Upper Bound on Optimal Expected Revenue)**.** *For every matching environment with n i.i.d. regular bidders and m items, $\sum_j \mathrm{Vic}_j(n+1)$ is at least the optimal expected revenue in the original market.*

Proof. Let $\{F_j\}_{j=1}^m$ be the regular value distributions of the matching environment. By Lemma 3.2.1, the optimal expected revenue in the matching environment is upper-bounded by the optimal expected revenue in its single-parameter counterpart, the corresponding representative environment. Recall that in the representative environment there are n single-parameter representatives per item j, whose values for j are i.i.d. draws from F_j. The representatives are grouped in n sets of size m corresponding to the original bidders in the matching environment, and feasibility constraints ensure that at most one representative from each set wins.

We now relax these feasibility constraints to get a new single-parameter environment in which the optimal expected revenue has only increased. Relaxing feasibility only increases the optimal expected revenue since by Myerson's lemma (Lemma 2.2.1), it is equal in expectation to the optimal virtual surplus, and clearly the optimum subject to the constraints is bounded from above by the optimum when these are relaxed.

The new environment is equivalent in terms of revenue to a collection of m single-item environments, where in the j-th environment item j is auctioned to its n single-parameter representatives (values are i.i.d. draws from the regular distribution F_j). By Bulow and Klemperer's result (Theorem 3.1.4), the optimal expected revenue from the j-th environment is upper-bounded by $\mathrm{Vic}_j(n+1)$. Summing up over all items completes the proof. □

Lower Bound. We now turn to the VCG mechanism applied to the augmented environment, whose revenue is the sum of VCG payments for the items. The next lemma lower-bounds the VCG payment for item j.

Lemma 3.4.3 (Lower Bound on VCG Revenue)**.** *For every matching environment, the VCG payment for item j is at least the value of any unallocated bidder for j.*

Proof. Say bidder i wins item j. The VCG payment for j is equal to the externality that i imposes on the rest of the bidders by winning j. In particular, i prevents an unallocated bidder i' from being allocated j. Algebraically, using the formula for VCG payments in matching environments (Equation 2.1, where recall that $W(\cdot)$ is the weight of a maximum weighted matching in a bipartite graph), the payment is

$$W(V, \mathcal{E} \setminus (i,j)) - [W(V, \mathcal{E}) - v_{i,j}] \geq [W(V, \mathcal{E}) - v_{i,j} + v_{i',j}] - [W(V, \mathcal{E}) - v_{i,j}] = v_{i',j}.$$

The inequality holds since a valid matching in the graph from which edge (i,j) has been removed is the maximum weighted matching in the original graph in which i wins j and i' is unallocated, but with edge (i',j) in place of (i,j). Thus the payment is at least the value of i' for j. □

3.4.3 Relating the Upper and Lower Bounds via Deferred Allocation

The upper and lower bounds above share a similar form. On the one hand, by definition of the Vickrey auction, the upper bound $\text{Vic}_j(n+1)$ on the expected revenue from separately auctioning off item j is equal to the expected *second-highest value for j among $n+1$ bidders with values drawn independently from F_j*. On the other hand, the lower bound on the VCG payment for item j in the augmented environment is equal to the *highest value for j among n unallocated bidders with values drawn independently from F_j*, where we use that in the augmented environment only m out of $n+m$ bidders are allocated. From this it may appear as if we have already shown that the lower bound exceeds the upper bound. However, a dependency issue arises – conditioned on the event that a bidder in the augmented environment is unallocated by VCG, his value for item j is no longer distributed like a random sample from F_j. We address this issue by introducing a *deferred allocation* selling procedure.

Algorithm 4 describes our selling procedure for item j.

Algorithm 4: Selling item j by deferred allocation

Input: A matching environment with $n+m$ bidders and m items, and an item j.

1. Find a maximum matching of all $m-1$ items other than j to the bidders. Let U be the set of $n+1$ bidders who remain unallocated.

2. Run the Vickrey auction to sell item j to bidder set U.

The following two claims show how deferred allocation resolves the dependency issue; namely, how the revenue from selling item j via the deferred allocation procedure bridges between the upper and lower bounds in Lemmas 3.4.2 and 3.4.3. The relation is also depicted in Figures 3.1a to 3.1c.

Claim 3.4.4 (Deferred Allocation and Upper Bound)**.** *The revenue from selling item j by deferred allocation (Algorithm 4) is equal in expectation to* $\mathrm{Vic}_j(n+1)$.

Proof. The revenue from selling item j to bidder set U by the Vickrey auction is the second-highest value of a bidder in U for j. Since we exclude item j in step (1) of the deferred allocation procedure and allocate it only in step (2), the allocation in step (1) does not depend on the bidders' values for j. Therefore, the values of the unallocated bidders in U for item j are independent random samples from F_j. The expected second-highest among $n+1$ values drawn independently from F_j is equal to $\mathrm{Vic}_j(n+1)$. □

To relate the revenue from deferred allocation to the lower bound in Lemma 3.4.3 we need the following stability property.

Lemma 3.4.5 (Stability of Maximum Matching)**.** *Consider an augmented matching environment with $n+m$ bidders and m items. Let bidder set U be as defined in Algorithm 4. If VCG is run on this environment, the set of bidders left unallocated is U with at most one bidder removed.*

Proof. First note that in the matching instances we consider, we may assume there is a unique maximum weighted matching. This holds with probability 1 since the weights are sampled from distributions as described in Section 3.2.[6]

The following is a well-known stability property of maximum weighted matchings [Lovász and Plummer, 2009]: In a complete weighted bipartite graph with $n+m$ nodes on one side and $m-1$ nodes on the other, consider the maximum weighted matching of size $m-1$. Now add a node to the short side of the graph and find the maximum weighted matching of size m. The set of matched nodes on the long side of the graph remains the same up to a single additional node.

The augmented matching environment corresponds to a complete bipartite graph with bidders on one side and items on the other, with the bidders' values for the items as edge weights. Algorithm 4 finds the maximum weighted matching of size $m-1$ in this graph with item j removed. VCG finds the maximum weighted matching of size m in this graph including item j. The lemma follows by applying the above stability property. □

Claim 3.4.6 (Deferred Allocation and Lower Bound)**.** *Given an augmented matching environment with $n+m$ bidders and m items, the revenue from selling item j by deferred allocation (Algorithm 4) is at most the VCG payment for item j.*

Proof. The revenue from selling item j by deferred allocation is the second-highest value of a bidder in U for j. Let i_1, i_2 be the two bidders in U who value item j the most. By definition, these bidders are left unallocated by the deferred allocation procedure, and by Lemma 3.4.5, one of them (say i_1) is also unallocated by the VCG mechanism. Recall that an unallocated bidder's value for item j gives a lower bound on the VCG payment for j (Lemma 3.4.3). So the VCG payment for j is at least $v_{i_1,j}$, which in turn is at least the second-highest value of a bidder in U for item j. □

[6]It is not hard to adapt the proof to the case in which there are multiple maximum weighted matchings, to show that one possible allocation of VCG run on the augmented matching environment leaves unallocated a set of bidders equal to U with at most one bidder removed.

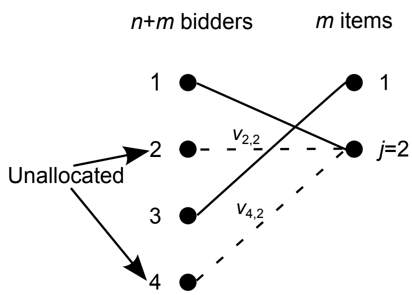

(a) VCG

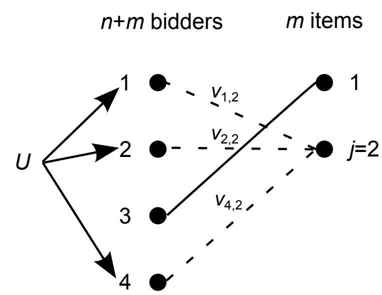

(b) Deferred allocation

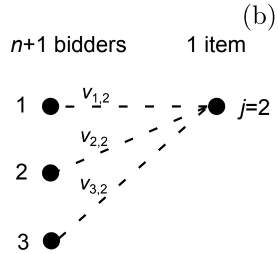

(c) Vickrey

Figure 3.1: Relating bounds by deferred allocation.
Let $n = m = 2$ and item $j = 2$.
(a) **VCG**: Solid edges correspond to the maximum matching. The payment for j is $\geq \max\{v_{2,2}, v_{4,2}\}$ (Lemma 3.4.3), where $v_{2,2}, v_{4,2}$ are *not* i.i.d. samples from F_2 given that bidders 2 and 4 are unallocated.
(b) **Deferred allocation**: Solid edges correspond to the maximum matching excluding j. Bidders unallocated in (a) are a subset of the unallocated set U (Lemma 3.4.5). Since j is sold to U using Vickrey, the payment for j is the 2nd-highest among $v_{1,2}, v_{2,2}, v_{4,2}$, where these are i.i.d. samples from F_2.
(c) **Vickrey**: The payment for j is the 2nd-highest among $v_{1,2}, v_{2,2}, v_{3,2}$, where these are i.i.d. samples from F_2.
Comparing (a) to (b) and (b) to (c) shows that:
- the payment for j in (a) is at least the payment for j in (b) (Claim 3.4.6); and
- in expectation the payment for j in (b) equals the payment for j in (c) (Claim 3.4.4).

3.4.4 Proof of Theorem 3.4.1

Putting everything together, we can now complete the proof of the Bulow-Klemperer-type theorem for matching markets.

Proof. We need to show that for every matching environment with n i.i.d. regular bidders and m items, the expected revenue of the VCG mechanism with m additional bidders is at least the optimal expected revenue in the original environment. By Claim 3.4.6, the VCG payment for item j in the augmented environment is at least the revenue from selling item j by deferred allocation, which by Claim 3.4.4 is equal in expectation to $\mathrm{Vic}_j(n+1)$. Summing up over all items, the total expected VCG revenue in the augmented environment is at least $\sum_j \mathrm{Vic}_j(n+1)$, and by Lemma 3.4.2 this upper-bounds the optimal expected revenue in the original environment. □

3.4.5 The $m \geq n$ Case

In matching markets where items are more plentiful than bidders, the following Bulow-Klemperer-type theorem provides an alternative to Theorem 3.4.1, in which the required demand augmentation is n instead of m bidders.[7] Two additional differences in comparison to previous Bulow-Klemperer-type theorems are that the VCG mechanism is required to be supply-limiting, and the revenue guarantee is an approximation.

Theorem 3.4.7 (Bulow-Klemperer-Type Theorem for Matching with $m \geq n$)**.** *For every matching environment with n i.i.d. regular bidders and $m \geq n$ items, the expected revenue of the* $\mathrm{VCG}^{\leq n}$ *mechanism with n additional bidders is at least an n/m-fraction of the optimal expected revenue in the original market. In other words,* $\mathrm{VCG}^{\leq n}$ *with n additional bidders is robustly n/m-optimal.*

Proof. The proof is similar to that of Theorem 3.4.1, with the following adjustments.

Consider the bounds in Section 3.4.2 above. While the upper bound on the optimal expected revenue in Lemma 3.4.2 holds and is sufficient, the lower bound on VCG payments in Lemma 3.4.3 holds but needs to be strengthened. In the augmented

[7] Similarly, for every $\eta \in [n, m]$ there is a Bulow-Klemperer-type theorem with η additional bidders. This does not improve the guarantees in Section 3.5.

environment, $\text{VCG}^{\leq n}$ allocates n out of the m items to n out of the $2n$ bidders. Therefore the VCG payment for item j not only exceeds the value of any unallocated bidder for j, but also exceeds the value of any unallocated bidder for any unallocated item. We shall refer to the highest such value among unallocated bidders and items as the *global* lower bound on VCG payments, and denote it by G.

Now to relate the bounds as in Section 3.4.3, we use a slightly modified deferred allocation procedure (Algorithm 5), where the difference from the previous procedure (Algorithm 4) is that only n out of the m items are matched, in order to achieve the effect of limited supply.

Algorithm 5: Selling item j by deferred allocation when $m \geq n$

Input: A matching environment with $2n$ bidders and m items, and an item j.

1. Find a maximum matching of $n-1$ items other than j to the bidders. Let U be the set of $n+1$ bidders who remain unallocated.

2. Run the Vickrey auction to sell item j to bidder set U.

Observe that Claims 3.4.4 and 3.4.6 continue to hold when Algorithm 4 is replaced by Algorithm 5. For Theorem 3.4.1 these claims were sufficient to complete the proof, by the following chain of arguments: All items are allocated by VCG in the augmented environment (since it is welfare-maximizing and there are more bidders than items); the VCG payment for item j is at least the revenue from selling j by deferred allocation (by Claim 3.4.6); the deferred allocation revenue is equal in expectation to $\text{Vic}_j(n+1)$ (by Claim 3.4.4); and $\sum_j \text{Vic}_j(n+1)$ is at least the optimal expected revenue in the original environment (by Lemma 3.4.2). For Theorem 3.4.7 we need an additional charging argument – and an approximation factor – since only n out of m items are allocated by $\text{VCG}^{\leq n}$.

For every item $j \in [m]$ there are two cases:

1. If j is allocated, then the VCG payment for j is at least the revenue from selling j by deferred allocation (Claim 3.4.6).

2. If j is not allocated, then the VCG payment for any allocated item j' is at least

the global lower bound G. A straightforward adaptation of the argument in Claim 3.4.6 shows that G is an upper bound on the revenue from selling j by deferred allocation (Algorithm 5).

To complete the proof, we charge the VCG payments for the n allocated items against the aggregate revenue from selling each of the m items by deferred allocation, where the latter is equal in expectation to $\sum_j \mathrm{Vic}_j(n+1)$. This leads to the approximation factor of $\frac{n}{m}$. $\square$

3.5 Matching Markets: Limiting Supply

In this section we present a supply-limiting mechanism for matching environments. For simplicity of presentation assume that the number of bidders n is even.

Theorem 3.5.1 (Supply-Limiting Mechanism for Matching Markets)**.** *For every matching environment with $n \geq 2$ i.i.d. regular bidders and m items, let*

$$\alpha = \max\{\frac{n-m}{n}, \min\{\frac{1}{2}, \frac{n}{4m}\}\}.$$

Then the expected revenue of the supply-limiting mechanism $\mathrm{VCG}^{\leq n/2}$ *is at least an α-fraction of the optimal expected revenue. In other words, VCG with supply limit $n/2$ is robustly α-optimal.*

Intuitively, achieving a good revenue guarantee becomes more difficult as the number of items m grows relatively to the number of bidders n, since the inherent competition among the bidders is split across different items. Accordingly, the fraction α in Theorem 3.5.1 depends on the parameters n and m of the environment as follows:

- If $m \leq \frac{n}{2}$ then $\alpha = \frac{n-m}{n}$, i.e., the approximation gets better as m becomes smaller, and for $m = \frac{n}{2}$ we get $\alpha = \frac{1}{2}$.
- If $m \geq \frac{n}{2}$ then $\alpha = \frac{n}{4m}$, and in particular for $m = n$ we get $\alpha = \frac{1}{4}$.

Also note that when $m \leq n/2$, the supply limit of $\mathrm{VCG}^{\leq n/2}$ has no effect, that is, the revenue guarantee is achieved by simply applying the VCG mechanism. For the case of $m \gg n$, Theorem 3.5.1 does not state a constant approximation. However it still holds in this case that VCG with a supply limit is robustly α'-optimal where $1/\alpha'$ is a constant, albeit with a more involved proof [Yan, 2012]. Theorem 3.5.1 also applies without change to *multi-unit* matching markets, in which there are k_j copies of every item j, and a total of $m = \sum_j k_j$ units overall.

3.5.1 Subadditivity

The following lemma is used to prove Theorem 3.5.1 and may also be of independent interest. It states that in any market environment, including one with asymmetric bidders whose values are drawn independently but not identically, the optimal expected revenue achievable from bidder sets S, T separately is at least the optimal expected revenue achievable from their union. Let $\text{OPT}(\cdot)$ map a bidder set to its optimal expected revenue. Then:

Lemma 3.5.2 (Subadditivity of Optimal Expected Revenue in Bidder Set)**.** *For every auction environment with bidder subsets S and T,* $\text{OPT}(S)+\text{OPT}(T) \geq \text{OPT}(S\cup T)$.

Proof. It is not hard to see that $\text{OPT}(\cdot)$ is monotone, so without loss of generality we can assume that S and T are disjoint. Let M be the optimal mechanism for $S \cup T$. For every value profile $\mathbf{v}_T$ of the bidders in T, we define the mechanism $M_{\mathbf{v}_T}$, which gets bids from the bidders in S and simulates M by using $\mathbf{v}_T$ as the bids of bidders in T. By an averaging argument, there exists a vector $\mathbf{v}_T$ such that mechanism $M_{\mathbf{v}_T}$'s expected revenue is at least the part of the optimal expected revenue of mechanism M that is charged to the bidders in S. On the other hand, the expected revenue of $M_{\mathbf{v}_T}$ is bounded above by $\text{OPT}(S)$. Similarly, the part of the optimal expected revenue that is charged to the bidders in T is bounded above by $\text{OPT}(T)$. This completes the proof. □

A corollary of the subadditivity lemma is that removing bidders from an i.i.d. environment until an α-fraction of the original bidders remains maintains an α-fraction of the optimal expected revenue. By symmetry:

Corollary 3.5.3. *For every auction environment with n i.i.d. bidders and for every integer c that divides n,* $\mathrm{OPT}(n/c) \geq \frac{1}{c}\,\mathrm{OPT}(n)$.

3.5.2 Proof by Reduction

We prove Theorem 3.5.1 by instantiating our general reduction.

Proof of Theorem 3.5.1. Assume first that $m \leq n/2$. We instantiate the reduction in Algorithm 3 as follows: to go from the original market to the restricted market, remove m bidders from the original market. By Corollary 3.5.3, this restriction on the original market maintains at least an $\frac{n-m}{n}$-fraction of the original optimal expected revenue. We can now apply the Bulow-Klemperer-type theorem for matching markets (Theorem 3.4.1) to the restricted market, which has $n-m$ bidders and m items. The expected revenue of VCG with m additional bidders is at least the optimal expected revenue in the restricted market. But this is equivalent to running VCG on the original market, completing the proof for $m \leq n/2$.

Now assume that $m \geq n/2$. We instantiate the reduction in Algorithm 3 as follows: to go from the original market to the restricted market, remove $n/2$ bidders from the original market. By Corollary 3.5.3, this restriction on the original market maintains at least a $\frac{1}{2}$-fraction of the original optimal expected revenue. We can now apply the Bulow-Klemperer-type theorem for matching markets with more items than bidders (Theorem 3.4.7) to the restricted market, which has $n/2$ bidders and $m \geq n/2$ items. The expected revenue of $\mathrm{VCG}^{\leq n/2}$ with $n/2$ additional bidders is at least an $\frac{n}{2m}$-fraction of the optimal expected revenue in the restricted market, and so an $\frac{n}{4m}$-fraction of the original optimal expected revenue. But this is equivalent to running $\mathrm{VCG}^{\leq n/2}$ on the original market, completing the proof. □

3.6 Extensions

In this section we extend our results to markets with a matroid constraint on who can win simultaneously and to asymmetric markets.

3.6.1 Matroid Markets

Recall that a matroid environment is a single-parameter environment in which the set system $([n], \mathcal{I})$ of bidders and feasible allocations forms a matroid (see Section 2.1.1). The *rank* ρ of a matroid is the size of its maximal independent sets or *bases*, and the *packing number* κ of a matroid is its maximum number of disjoint bases. In the job scheduling example presented in Section 3.1.3, if there are four jobs arriving at time 0 of which two must be finished by time 1 and two must be finished by time 2, then the rank is $\rho = 2$ and the packing number is $\kappa = 2$. We will use the fact that an *intersection* of a matroid $([n], \mathcal{I})$ with a u-*uniform* matroid is a new matroid $([n], \mathcal{I}')$, in which a set of bidders S belongs to $\mathcal{I}'$ if and only if $S \in \mathcal{I}$ and $|S| \leq u$.

Dughmi et al. [2012] show a Bulow-Klemperer-type result for matroid environments:

Theorem 3.6.1 (Bulow-Klemperer-Type Theorem for Matroid Markets)**.** *For every matroid environment with i.i.d. regular bidders, the expected revenue of the Vickrey auction with an additional basis of bidders is at least the optimal expected revenue in the original market. In other words, Vickrey with an additional basis of bidders is robustly optimal.*

We use Theorem 3.6.1 with our general reduction to prove the following result, in which the approximation depends on the inherent amount of competition in the market measured not by the number of bidders n but rather by their packing number κ. For simplicity of presentation assume the rank ρ is even. Then:

Theorem 3.6.2 (Supply-Limiting Mechanism for Matroid Markets)**.** *For every matroid environment with $n \geq 2$ i.i.d. regular bidders, rank ρ and packing number κ, let $\ell = \rho/2$ if $\kappa = 1$ and $\ell = \rho$ otherwise. Then the expected revenue of the supply-limiting mechanism* $\text{Vic}^{\leq \ell}$ *is at least a* $\max\{\frac{1}{4}, \frac{\kappa-1}{\kappa}\}$*-fraction of the optimal expected revenue. In other words,* $\text{Vic}^{\leq \ell}$ *is robustly* $\max\{\frac{1}{4}, \frac{\kappa-1}{\kappa}\}$*-optimal.*

Proof. We instantiate the reduction in Algorithm 3 as follows: If $\kappa = 1$, intersect the original matroid with a $\frac{\rho}{2}$-uniform matroid to get a new matroid $([n], \mathcal{I}')$ with rank $\rho' = \frac{\rho}{2}$ and packing number $\kappa' \geq 2$. Otherwise, if $\kappa \geq 2$, simply set the new matroid

$([n], \mathcal{I}')$ to be $([n], \mathcal{I})$. To go from the original market to the restricted market, remove from the original market a basis of bidders of size ρ' according to the matroid $([n], \mathcal{I}')$, and set the matroid of the restricted market to be $([n], \mathcal{I}')$.

Analysis: If $\kappa = 1$, intersecting with the uniform matroid maintains at least a $\frac{1}{2}$-fraction of the original optimal expected revenue. Removing a basis of bidders maintains at least a $\frac{\kappa'-1}{\kappa'}$-fraction. We can now apply the Bulow-Klemperer-type theorem for matroids (Theorem 3.6.1) to the restricted market with $n - \rho'$ bidders and matroid $([n], \mathcal{I}')$. The expected revenue of Vickrey with an additional basis of bidders is at least the optimal expected revenue in the restricted market, and so a $\max\{\frac{1}{4}, \frac{\kappa-1}{\kappa}\}$-fraction of the original optimal expected revenue. But this is equivalent to running $\text{Vic}^{\leq \ell}$ on the original market, completing the proof. □

3.6.2 Asymmetric Bidders: Augmenting Demand

An *attribute-based* environment is a k-unit environment with n bidders, each of whom has a publicly-observable attribute $a = a(i)$ that determines a non-publicly-known value distribution F_a [Dhangwatnotai et al., 2015]. Bidders' values in an attribute-based environment are thus independently but not identically distributed. Attributes enable the incorporation of prior information into our model regarding which bidders are alike, while still avoiding assumptions about the value distributions themselves. In fact, our results can be interpreted as an encouragement to invest in this particular kind of prior information, which entails grouping similar bidders together rather than learning distributions. Examples of attributes are bidding styles such as "bargain-hunter" or "aggressive" on `eBay.com`, or in sponsored search and online advertising, advertiser features such as location. Throughout we assume *non-singular* attribute-based environments, where no bidder's attribute is unique. I.e., for every attribute a, let n_a denote the number of bidders in the environment with attribute a; then $n_a > 0 \implies n_a \geq 2$.

In this section we prove a Bulow-Klemperer-type theorem for attribute-based environments.[8] Let $\text{Vic}^{\leq \ell_a}$ be the Vickrey mechanism with a *local* supply limit ℓ_a for

[8]It is not hard to show that the Bulow-Klemperer-type theorem for asymmetric matroid environments of [Hartline and Roughgarden, 2009, Theorem 4.4] applies to attribute-based environments.

every a, which limits the number of bidders with attribute a who can win simultaneously. The proof of the following theorem uses a *commensuration* argument of Hartline and Roughgarden [2009], and applies the FKG inequality [Alon and Spencer, 2008] to solve dependency issues.

Theorem 3.6.3 (Bulow-Klemperer-Type Theorem for Asymmetric Markets)**.** *For every attribute-based environment with n_a regular bidders per attribute and k units, the expected revenue of the* $\text{Vic}^{\le n_a}$ *auction with* $\min\{n_a, k\}$ *additional bidders per attribute is at least a $\frac{1}{2}$-fraction of the optimal expected revenue in the original market. In other words,* $\text{Vic}^{\le n_a}$ *with* $\min\{n_a, k\}$ *additional bidders per attribute is robustly $\frac{1}{2}$-optimal.*

Proof. Let $W^{\text{OPT}} = W^{\text{OPT}}(\mathbf{v}), W^{\text{Vic}} = W^{\text{Vic}}(\mathbf{v})$ denote the winning bidders chosen by the optimal mechanism in the original environment and by $\text{Vic}^{\le n_a}$ in the augmented environment, respectively, given a value profile $\mathbf{v}$ of both original and augmenting bidders. Hartline and Roughgarden [2009] show that to prove a 1/2-approximation it suffices to establish two commensuration conditions among the two mechanisms:

(C1) $\mathbb{E}_{\mathbf{v}}[\sum_{i \in W^{\text{Vic}} \setminus W^{\text{OPT}}} \phi_i] \ge 0$,

(C2) $\mathbb{E}_{\mathbf{v}}[\sum_{i \in W^{\text{Vic}} \setminus W^{\text{OPT}}} p_i(\mathbf{v})] \ge \mathbb{E}_{\mathbf{v}}[\sum_{i \in W^{\text{OPT}} \setminus W^{\text{Vic}}} \phi_i]$,

where ϕ_i is the virtual value of bidder i.

The proof of (C2) in [Hartline and Roughgarden, 2009, Lemma 4.5] holds in our setting. In contrast, proving (C1) in our setting turns out to be technically challenging due to dependencies among the random bidder sets W^{OPT} and W^{Vic}. We use an auxiliary allocation procedure (Algorithm 6), and rely on the fact that in our setting, $\text{Vic}^{\le n_a}$ applies a simple greedy algorithm: it rejects all but the top n_a bidders per attribute, and allocates the units to the $\le k$ highest remaining bidders.

This theorem requires augmenting the demand with an additional "duplicate" bidder for every original bidder, and adding the constraint that at most one of each such pair wins simultaneously. Our version in Theorem 3.6.3 utilizes the fact that many of the bidders in an attribute-based environment are symmetric – namely all those with the same attribute – in order to avoid the pair constraints, and when k is relatively small requires less bidders to be added to the environment.

For the remainder of the proof, fix an attribute a, and let B_a be the set of $n_a + \min\{n_a, k\}$ bidders with attribute a in the augmented environment. Fix the values of the original bidders in B_a as well as the values of all bidders with different attributes, and let $\mathbf{v}_a$ denote the (random) value profile of the augmenting bidders in B_a. Let $W_a^{\text{OPT}} = B_a \cap W^{\text{OPT}}$ denote the bidders in B_a who win in the optimal mechanism, and let $W_a^{\text{Vic}} = W_a^{\text{Vic}}(\mathbf{v}_a) = B_a \cap W^{\text{Vic}}$ denote the bidders in B_a who win in $\text{Vic}^{\leq n_a}$. We can now define our auxiliary procedure in Algorithm 6. Let $W_a^{\text{Aux}} = W_a^{\text{Aux}}(\mathbf{v}_a)$ denote the bidders in B_a who win in this procedure, and observe $W_a^{\text{OPT}} \subseteq W_a^{\text{Aux}} \subseteq B_a$.

Algorithm 6: Auxiliary allocation procedure for asymmetric bidders

In the augmented environment, given W_a^{OPT}, allocate the k units such that the welfare is maximized subject to the constraint that all bidders in W_a^{OPT} win.

We now establish two claims, the first of which relates the auxiliary procedure to $\text{Vic}^{\leq n_a}$ and the second of which relates it to the optimal mechanism.

Claim 3.6.4. *For every a and every value profile $\mathbf{v}_a$ of the augmenting bidders with attribute a, the bidders in $W_a^{\text{Vic}} \setminus W_a^{\text{Aux}}$ have non-negative virtual values.*

Proof of Claim 3.6.4. Fix $\mathbf{v}_a$ and consider the allocation of $\text{Vic}^{\leq n_a}$ in comparison to that of the auxiliary procedure. $\text{Vic}^{\leq n_a}$ is free to replace bidders in W_a^{OPT}. Since $\text{Vic}^{\leq n_a}$ is greedy, each replacement from B_a will have a higher value than the replaced bidder in W_a^{OPT}, and therefore (using regularity) also a higher virtual value. The proof follows by noticing that all bidders in W_a^{OPT} have non-negative virtual values (Lemma 2.2.1). □

Claim 3.6.5. *For every a, in expectation over the value profile $\mathbf{v}_a$ of the augmenting bidders with attribute a, summing over the highest $h \leq \min\{n_a, k\}$ virtual values of the bidders in $W_a^{\text{Aux}} \setminus W_a^{\text{OPT}}$ results in a non-negative total virtual value.*

Proof of Claim 3.6.5. Given $\mathbf{v}_a$, denote by $\psi_{(1)}, \psi_{(2)}, \ldots$ (where $\psi_{(i)} = \psi_{(i)}(\mathbf{v}_a)$) the virtual values of the bidders in $B_a \setminus W_a^{\text{OPT}}$, sorted in decreasing order of both values and virtual values. Let $\mathbb{1}_{(i)} = \mathbb{1}_{(i)}(\mathbf{v}_a)$ indicate whether the i-th bidder in $B_a \setminus W_a^{\text{OPT}}$ wins in the auxiliary procedure, and let $q_{(i)}$ be the probability that $\mathbb{1}_{(i)} = 1$ over a random choice of $\mathbf{v}_a$. The expected sum of the highest $\min\{n_a, k\}$ virtual values of the bidders in $W_a^{\text{Aux}} \setminus W_a^{\text{OPT}}$ can be written as $\sum_{i \leq \min\{n_a,k\}} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i)} \cdot \mathbb{1}_{(i)}]$. To complete the proof of Claim 3.6.5 it is sufficient to show that this expression is non-negative.

Notice that like $\text{Vic}^{\leq n_a}$, the auxiliary procedure is greedy in nature. Thus if the i-th bidder in $B_a \setminus W_a^{\text{OPT}}$ wins in the auxiliary procedure, and the values of bidders in $B_a \setminus W_a^{\text{OPT}}$ increase, then the i-th highest bidder will still win. Formally, for two value profiles $\mathbf{v}_a \leq \mathbf{v}'_a$ we have $\mathbb{1}_{(i)}(\mathbf{v}_a) \leq \mathbb{1}_{(i)}(\mathbf{v}'_a)$ for every i. The fact that both the virtual value and the probability of winning in the auxiliary procedure are increasing in the values allows us to apply the FKG inequality [Alon and Spencer, 2008], which states that for increasing random variables X, Y, $\mathbb{E}[X \cdot Y] \geq \mathbb{E}[X] \cdot \mathbb{E}[Y]$. Its application yields:

$$\begin{aligned}
\sum_{i \leq \min\{n_a,k\}} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i)} \cdot \mathbb{1}_{(i)}] &\geq \sum_{i \leq \min\{n_a,k\}} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i)}] \cdot \mathbb{E}_{\mathbf{v}_a}[\mathbb{1}_{(i)}] \\
&= \sum_{i \leq \min\{n_a,k\}} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i)}] \cdot q_{(i)} \\
&= \sum_{i \leq \min\{n_a,k\}} \left(\mathbb{E}_{\mathbf{v}_a} \left[\sum_{i'=1}^{i} \psi_{(i')} \right] \cdot \left(q_{(i)} - q_{(i+1)}\right) \right),
\end{aligned}$$

where the first inequality is by FKG, and where we set $q_{(i+1)}$ for $i = \min\{n_a, k\}$ to 0.

It is not hard to see that $q_{(i)}$ is decreasing in i. Therefore it suffices to prove that $\sum_{i'=1}^{i} \mathbb{E}_{\mathbf{v}_a}[\psi_{(i')}] \geq 0$ for every $i \leq \min\{n_a, k\}$. This is the sum of the expected virtual values of the top i bidders in $B_a \setminus W_a^{\text{OPT}}$. Observe that the sum of the expected virtual values of any i augmented bidders in $B_a \setminus W_a^{\text{OPT}}$ equals 0, and there are at least $\min\{n_a, k\}$ such bidders. It follows that this sum for the top i bidders is nonnegative. □

We now use Claims 3.6.4 and 3.6.5 to complete the proof of (C1). Still holding attribute a fixed, rewrite $\mathbb{E}_{\mathbf{v}_a}[\sum_{i \in W_a^{\text{Vic}} \setminus W_a^{\text{OPT}}} \phi_i]$ as

$$\mathbb{E}_{\mathbf{v}_a}\left[\sum_{i \in W_a^{\text{Vic}} \setminus W_a^{\text{Aux}}} \phi_i\right] + \mathbb{E}_{\mathbf{v}_a}\left[\sum_{i \in W_a^{\text{Vic}} \cap (W_a^{\text{Aux}} \setminus W_a^{\text{OPT}})} \phi_i\right].$$

The left-hand side is non-negative by Claim 3.6.4. We consider two cases for the right-hand side, which by greediness of the auxiliary procedure and $\text{Vic}^{\leq n_a}$ are the only possible cases:

1. $(W_a^{\text{Aux}} \setminus W_a^{\text{OPT}}) \subseteq W_a^{\text{Vic}}$: In this case $W_a^{\text{Vic}} \cap (W_a^{\text{Aux}} \setminus W_a^{\text{OPT}}) = W_a^{\text{Aux}} \setminus W_a^{\text{OPT}}$, and so by Claim 3.6.5 the sum of virtual values is non-negative in expectation.

2. $W_a^{\text{Vic}} \subseteq W_a^{\text{Aux}}$: In this case, $W_a^{\text{Vic}} \setminus W_a^{\text{OPT}}$ is a subset of the highest $\min\{n_a, k\}$ bidders in $W_a^{\text{Aux}} \setminus W_a^{\text{OPT}}$, and so by Claim 3.6.5 the sum of virtual values over this subset is non-negative in expectation.

We have shown that $\mathbb{E}_{\mathbf{v}_a}[\sum_{i \in W_a^{\text{Vic}} \setminus W_a^{\text{OPT}}} \phi_i] \geq 0$. Taking expectation and summing over all attributes we get $\mathbb{E}_{\mathbf{v}}[\sum_{i \in W^{\text{Vic}} \setminus W^{\text{OPT}}} \phi_i] \geq 0$, completing the proof of (C1) and Theorem 3.6.3. □

3.6.3 Asymmetric Bidders: Limiting Supply

Consider an attribute-based environment as defined in Section 3.6.2. For simplicity of presentation assume that n_a, the number of bidders with attribute a, is even for every a. Recall that $\text{Vic}^{\leq n_a/2}$ is the Vickrey mechanism with a local supply limit $n_a/2$ for every a, meaning that no more than half the bidders with the same attribute can win simultaneously. We now show that $\text{Vic}^{\leq n_a/2}$ is a good supply-limiting mechanism for attribute-based environments. Note that it is considerably simpler than Myerson's optimal mechanism for asymmetric markets, which requires computing different virtual value functions for different attributes.

Theorem 3.6.6 (Supply-Limiting Mechanism for Asymmetric Markets)**.** *For every attribute-based environment with n_a regular bidders per attribute, the expected revenue of the supply-limiting mechanism $\text{Vic}^{\leq n_a/2}$ is at least a $\frac{1}{4}$-fraction of the optimal expected revenue. In other words, $\text{Vic}^{\leq n_a/2}$ is robustly $\frac{1}{4}$-optimal.*

Proof. We instantiate the reduction in Algorithm 3 as follows: to go from the original market to the restricted market, for every a remove $\min\{\frac{n_a}{2}, k\}$ bidders with attribute a from the original market. By submodularity [Dughmi et al., 2012], this restriction on the original market maintains at least a $\frac{1}{2}$-fraction of the original optimal expected revenue (since we removed at most half of the bidders). We can now apply the Bulow-Klemperer-type theorem for asymmetric markets (Theorem 3.6.3) to the restricted market, which has $\max\{\frac{n_a}{2}, n_a - k\}$ bidders per attribute and k units. The expected revenue of $\text{Vic}^{\leq n_a/2}$ with $\min\{\frac{n_a}{2}, k\}$ additional bidders is at least a $\frac{1}{2}$-fraction of the optimal expected revenue in the restricted market. But this is equivalent to running $\text{Vic}^{\leq n_a/2}$ on the original market, completing the proof. □

4

Robust Revenue Maximization with Interdependence

4.1 Chapter Introduction

4.1.1 Overview

The subject of this chapter is optimal and robust mechanism design in the classic model of interdependent values introduced by Milgrom and Weber [1982]. We extend several fundamental results from the well-studied independent private values model to interdependent settings. For revenue-optimal mechanism design, we give conditions under which Myerson's virtual value-based mechanism remains optimal with interdependent values. One of these conditions is robustness of the truthfulness and individual rationality guarantees, in the sense that they are required to hold ex post. We then consider an even more robust class of mechanisms that are prior-independent, and show that by simply using one of the bidders to set a reserve price, it is possible to extract near-optimal revenue in an interdependent values setting. This shows that a considerable level of robustness is achievable for interdependent values in single-parameter environments. The model of interdependent values is not only of economic importance in itself, but also sheds new light on the inherent trade-off between revenue maximization and robustness in the design of mechanisms.

	Private	Non-private (common)
Independent	Independent private values (IPV)	Non-private values
Correlated	Correlated values	Correlated non-private values

Table 4.1: Valuation models can be divided along two axes. The term *interdependent values* is a generalization of IPV that encompasses all four classes of values described in the table.

4.1.2 Beyond IPV

Economic research on auctions has explored different valuation models over the past decades, which can roughly be divided into independent private[1] values (IPV) versus more general value classes that allow for correlation and non-privacy, as shown in Table 4.1 (see also: [Krishna, 2010, Chapters 2-5; 13-17] versus [Krishna, 2010, Chapters 6-10; 18]). The more nascent research effort in theoretical computer science has focused largely on the more restricted IPV model, recently also venturing into the realm of correlation (see, notably, [Dobzinski et al., 2011; Papadimitriou and Pierrakos, 2015]). A promising research goal is therefore to apply the computer science lens to the study of mechanisms for general values, as captured by the interdependent values model. The results described in this chapter constitute the first steps in this direction, by unifying and generalizing previous results to establish the necessary technical foundations, and by demonstrating natural sufficient conditions under which positive results in the form of mechanisms with very strong guarantees can be achieved.

The importance of the interdependent model in the economic literature stems from the fact that for many high-stake auctions that arise in practice, interdependent values are a more realistic model of bidders' values than IPV. Interdependence captures correlated private values, which arise when bidders' information (signal) about their value for winning the auction is correlated with that of others (lower-left corner of

[1]We differentiate between private values and privately-*known* signals or information. By the former we mean that a bidder's value is idiosyncratic, i.e., his valuation function depends only on his own privately-known information and not on that of others. By the latter we mean that the signal or information is known only to the bidder, creating an information asymmetry among the players. We similarly differentiate between common values and publicly-known signals or information.

Table 4.1). It also captures independent non-private values, where the information held by other bidders directly affects bidders' values; mathematically, values are a function of a bidder's own signal and the independent signals of his competitors (upper-right corner of Table 4.1). Finally it captures combinations of the two (lower-right corner of Table 4.1).

A classic example from the economic literature is the *mineral rights model* [Wilson, 1969]. In auctions for oil drilling, the value of the drilling rights is determined by whether or not there is oil to be found on the drilling site. This value is often common – yet unknown – to all bidders. However, typically every bidder has some private noisy signal regarding the value, achieved by, for example, conducting a geological survey. Not only are these signals positively correlated, but the information gathered by the other bidders would certainly affect a bidder's expected value for winning if he gained access to it.

Note that the IPV model is not rich enough to capture the described informational setting: In the IPV model an attempted approximation is that values all come from a distribution over a high-valued support (oil exists), or a low-valued one (no oil). But to model that the seller does not know which of the two supports is the case will make the bidders correlated in his view. Similarly, we need interdependence to model that bidders do not know their precise value for winning an auction since it depends on others' information. The model of interdependence thus enriches the set of underlying assumptions we are able to make about the informational structure of the auction setting. In this chapter we explicitly treat such informational assumptions and their role in designing mechanisms, as discussed next.

4.1.3 Robustness for Interdependent Values

A second goal of this chapter is the design of *robust* mechanisms. Consider the informational assumptions of a standard Bayesian auction environment, as captured by Assumption 1.2.1. For the purpose of the discussion in this chapter, we restate Assumption 1.2.1 as two separate assumptions:

Informational Assumption 1. The bidders all know the probability distribution of the privately-known information, and make strategic decisions accordingly.

Informational Assumption 2. The seller knows the probability distribution of the privately-known information, and chooses the mechanism (or sets parameters such as reserve prices) accordingly.

We also treat the solution concept of truthfulness as two separate requirements: Incentive compatibility, or *IC*, for truthful reporting, and individual rationality, or *IR*, for participation (see Section 4.3.1 for definitions).

As discussed above (see Sections 1.4 and 2.3), there are many theoretical and practical reasons that motivate relaxing the above assumptions; for example, accurate prior information may be expensive to acquire, or a mechanism may need to be re-used in settings with different distributions. In fact, the issue of robustness to informational assumptions as above is "an old theme in the mechanism design literature" [Bergemann and Morris, 2005]. However, potentially there can be a trade-off between robustness and the objectives of the mechanism, in our case maximizing revenue.

Remarkably, Myerson [1981] showed that in the IPV model there is no trade-off with respect to the first assumption:[2] The optimal mechanism among all mechanisms that make Assumption 1 does not actually use this assumption in any way. More formally, Myerson's optimal mechanism is *ex post* IC and IR, but is optimal among all *Bayesian* IC and *interim* IR mechanisms (see Section 4.3.1 for definitions).

On the other hand, in the context of interdependent values – indeed even correlated private values – the trade-off not only exists but becomes very extreme. A mechanism making Assumption 1 can extract as revenue essentially the full welfare arising from the auction, leaving the bidders with virtually zero utility from participation [Crémer and McLean, 1985, 1988]. Intuitively, dependencies among the bidders "cancel out" their strategic advantage from privately-held information and nullifies their information rents. Without Assumption 1 however, the gap between the optimal expected revenue and full welfare can be arbitrarily large.

[2]Assuming for now that Assumption 2 holds.

While relaxing Assumption 1 leads to a loss of revenue, a possibly surprising property is regained. It turns out that Myerson's result [1981], showing the fundamental connection between the allocation and payment rules needed to induce truthfulness, extends to interdependent values once we restrict attention to mechanisms that are ex post IC and IR. We show that under conditions well-studied in the literature, it is possible to follow the same path as in Myerson's original paper and get an analogue of the Myerson optimal mechanism for interdependent values, where optimality is among all ex post mechanisms. This relates the goal of robustness to that of expanding the theory for interdependent values – the latter is made possible by imposing robustness as a requirement.

We now revisit the second informational assumption, and ask what is achievable for interdependent values when *both* assumptions are relaxed. Even in the IPV model, the Myerson mechanism heavily depends on Assumption 2. The question of designing new mechanisms without this assumption has been studied in the IPV model, leading to the theory of prior-independence. We give the first prior-independence result beyond the IPV model; in particular we show that while interdependence does complicate the problem of relaxing Assumption 2, a considerable level of robustness is still achievable without giving up too much expected revenue.

4.1.4 Our Results

Our work makes two main technical contributions. To describe these and for the remainder of the chapter we use the terminology in Table 4.1.

(1) We develop a general analogue of Myerson's optimal auction theory that applies to many interdependent settings of interest. While Myerson's theory does not hold in general for interdependent values (indeed, there are settings in which the Crémer-McLean mechanism extracts higher revenue than the Myerson mechanism), we show it is partially recovered when we impose ex post rather than Bayesian IC and IR constraints. This synthesizes and generalizes many known results in the economic literature (see related work in Section 4.4).

We first apply standard techniques to characterize ex post IC and IR mechanisms in the interdependent model and to show that their expected revenue equals their expected "conditional virtual surplus". Notably, we use the characterization to identify sufficient conditions under which the simple, "ironless" form of the Myerson mechanism is optimal. Under these conditions, the optimal mechanism simply allocates to the bidders with highest non-negative (conditional) virtual values.

(2) For non-private value settings we analyze a prior-independent auction and show that it is simultaneously near-optimal across a range of possible prior distributions. In particular, we adapt the *single sample* approach of Dhangwatnotai et al. [2015] to interdependent values, and show that with an additional sufficient and necessary MHR assumption, this approach results in an approximately-optimal, prior-independent mechanism.

Our prior-independence result demonstrates that non-trivial research questions can arise even in the simplest interdependent settings. Our Myerson-like characterization suggests that many interesting mechanism design results should be possible, even when bidders have interdependent values.

4.1.5 Organization

In Section 4.2 we present the model and in Section 4.3 the basic ex post solution concept. In Section 4.4 we survey related work. Section 4.5 develops the first result above and the second result appears in Section 4.6.

4.2 Interdependent Values Model

We describe how the interdependent values model is different than the Bayesian model for auctions described in Section 2.1, and re-define some of the notation accordingly.

The bidders in the interdependent values model have possibly correlated, privately-known *signals* $s_1, \dots, s_n$, drawn from a joint distribution F with density f over the support $[0, \omega_i]^n$ (ω_i may be ∞). We adopt the standard assumptions that f is continuous and nowhere zero. Every bidder i has a publicly-known *valuation function* v_i whose arguments are the signals, and his *interdependent value* for winning is $v_i(\vec{s})$.

Interdependent values are also called *information externalities* among the bidders. When the bidders share the same valuation function we say they have a *pure common value*. We impose the following standard assumptions on the valuation function $v_i(\cdot)$:

- Non-negative and normalized ($v_i(\vec{0}) = 0$);
- Twice continuously differentiable;
- Non-decreasing in all variables, strictly increasing in s_i.
- Finite expectation $\mathbb{E}_{\vec{s}}[v_i(\vec{s})] < \infty$.

Encompassed Value Models. The described interdependent values model is very general; it includes several narrower settings of interest (recall Table 4.1):

1. *Correlated values* settings, in which $v_i(\vec{s}) = s_i$ for every i and the signals/values are correlated;
2. *Non-private values* settings, in which the signals are independent (drawn independently from distributions $\{F_i\}$) and the valuation functions $v_i(\vec{s})$ are general;
3. *Independent private values (IPV)* settings, in which both $v_i(\vec{s}) = s_i$ for every i and the signals/values are independent.

Notation. Fix a signal profile s_{-i}. Let $v_{i|s_{-i}}(\cdot)$ denote bidder i's value given s_{-i} as a function of his signal s_i. Since v_i is strictly increasing in s_i, $v_{i|s_{-i}}(\cdot)$ is invertible; denote by $v^{-1}_{i|s_{-i}}(\nu)$ or $v_i^{-1}(\nu \mid s_{-i})$ the signal s_i such that $v_{i|s_{-i}}(s_i) = \nu$. Slightly abusing notation, given s_{-i} denote the derivative of $v_{i|s_{-i}}(\cdot)$ at s_i by

$$\frac{d}{ds_i} v_i(\vec{s}) = \frac{d}{ds_i} v_{i|s_{-i}}(s_i).$$

4.2.1 Motivating Examples

We describe two natural and standard examples of non-private values. In the first example, bidders' values directly depend on the private preferences of the others.

In the second example, bidders' values depend on a hidden stochastic "state of the world", of which others may posses private knowledge. We then give an example of correlated signals.

Example 4.2.1 (Weighted-sum values)**.** Let $\beta \in [0, 1]$. Every bidder's value is a sum of his own signal and a weighted sum of the other signals:

$$v_i(\vec{s}) = s_i + \beta \sum_{j \neq i} s_j.$$

This is a simplified version of Myerson's value with *revision effects* [Myerson, 1981], and when $\beta = 1$, this results in the *wallet game* [Klemperer, 1998]. Weighted-sum values are a plausible value model for a painting sold in auction; a bidder's value for the painting is determined by his own appreciation of it, combined with the painting's "resale value" based on how much others appreciate it.

Example 4.2.2 (Conditionally-independent values, a.k.a mineral rights model)**.** Bidders have a hidden stochastic pure common value v, modeled by a random variable V drawn from a publicly known distribution F_V. An important feature of the mineral rights model is that conditional on the event $V = v$, bidders' signals are independent. Furthermore, each signal is an unbiased estimator of V (its expectation when $V = v$ equals v). The bidders' effective value – their value for all operational purposes – is the pure common value

$$v_i(\vec{s}) = \mathbb{E}_{V \sim F_V}[V \mid \vec{s}].$$

The mineral rights model was developed to capture values in auctions for oil drilling leases [Wilson, 1969]. Such values are determined by the true amount of existing oil, but uncertainty and information asymmetries regarding this amount creates interdependency.

A concrete setting of interest is the one in which F_V is distributed normally with parameters μ_V, σ_V (assume μ_V is far from 0 and σ_V is small), and the signals are $s_i = v + \eta_i$, where $\eta_1, \ldots, \eta_n$ are i.i.d. samples drawn from the normal distribution with parameters $\mu_\eta = 0$ and some small σ_η. In this case the value is a linear combination of the prior and empirical means, similar to Example 4.2.1 above:

$$v_i(\vec{s}) = \frac{n\sigma_V^2}{n\sigma_V^2 + \sigma_\eta^2}(\frac{1}{n}\sum_i s_i) + \frac{\sigma_\eta^2}{n\sigma_V^2 + \sigma_\eta^2}\mu_V.$$

Up to normalization, the coefficient of the empirical mean $\frac{1}{n}\sum_i s_i$ is the prior variance and the coefficient of the prior mean μ_V is the noise variance.

Example 4.2.3 (Signals drawn from a multivariate normal distribution)**.** The signals in Example 4.2.1 above can be arbitrarily correlated. A concrete example of a joint signal distribution is a symmetric multivariate normal distribution. This distribution is "nice" in the sense that its marginals are normal as well, and if all pairwise covariances are nonnegative, then signals drawn from it satisfy a strong form of correlation called *affiliation*. We will make use of these properties below.

4.2.2 Conditional Virtual Values and Regularity

Fix a bidder i and a signal profile s_{-i}. The *conditional marginal density* $f_i(\cdot \mid s_{-i})$ of bidder i's signal given s_{-i} is

$$f_i(s_i \mid s_{-i}) \quad = \quad \frac{f(\vec{s})}{\int_0^{\omega_i} f(t, s_{-i})dt}.$$

We denote the corresponding distribution by $F_i(\cdot \mid s_{-i})$. The *conditional revenue curve* $B_i(\cdot \mid s_{-i})$ of bidder i is

$$B_i(s_i \mid s_{-i}) = v_i(\vec{s}) \int_{s_i}^{\omega_i} f_i(t \mid s_{-i})dt.$$

The conditional revenue curve represents the expected revenue from setting a threshold price $v_i(s_i, s_{-i})$ for bidder i given that the other signals are s_{-i}.[3] We can now

[3]Equation 4.2.2 uses the assumption that v_i is strictly increasing in s_i, so that the probability with which bidder i's value is at least $v_i(\vec{s})$ equals the probability $\int_{s_i}^{\omega_i} f_i(t \mid s_{-i})dt$ that bidder i's signal is at least s_i.

define the *conditional virtual value* of bidder i as

$$\varphi_i(s_i \mid s_{-i}) = -\frac{\frac{d}{ds_i}B_i(s_i \mid s_{-i})}{f_i(s_i \mid s_{-i})} = v_i(\vec{s}) - \frac{1 - F_i(s_i \mid s_{-i})}{f_i(s_i \mid s_{-i})} \cdot \frac{d}{ds_i} v_i(\vec{s}). \tag{4.1}$$

For private values, the conditional virtual value simplifies to a more familiar form (recall Myerson's original definition of a virtual value – see Equation 2.2):

$$\varphi_i(s_i \mid s_{-i}) = s_i - \frac{1 - F_i(s_i \mid s_{-i})}{f_i(s_i \mid s_{-i})}. \tag{4.2}$$

For independent common values it simplifies to the form

$$\varphi_i(s_i \mid s_{-i}) = v_i(\vec{s}) - \frac{1 - F_i(s_i)}{f_i(s_i)} \cdot \frac{d}{ds_i} v_i(\vec{s}). \tag{4.3}$$

Regularity and MHR. We say that $F_i(\cdot \mid s_{-i})$ is *regular* if the conditional virtual value $\varphi_i(\cdot \mid s_{-i})$ is weakly increasing; we say it has monotone hazard rate, or *MHR*, if the inverse hazard rate $(1-F_i(s_i \mid s_{-i}))/f_i(s_i \mid s_{-i})$ is weakly decreasing. Its *monopoly price* is the signal s_i such that the conditional virtual value $\varphi_i(s_i \mid s_{-i})$ equals zero.

Conditional vs. Standard Virtual Values. Equation (4.1) reveals three complications introduced by conditional virtual values in comparison to standard virtual values in the IPV model: first, the value v_i is a function of others' signals s_{-i}; second, the inverse hazard rate is conditional on s_{-i}; and third, there is an extra term dv_i/ds_i. The Myerson-like mechanisms we study below will rank bidders according to their conditional virtual values; while in the IPV model regularity is sufficient for such a mechanism to be IC, these three complications suggest that assumptions beyond regularity will be required. For example, regularity of the signal distribution restricts only the conditional marginals, whereas for a joint distribution more "global" constraints may be necessary.

4.2.3 Auction Settings of Interest

We present several settings of particular interest which are extensively studied in the literature. These settings arise by imposing natural further assumptions on a general single-parameter auction environment, and we will repeatedly refer to them in our results. In addition to these we also refer to matroid settings, as defined in Section 2.1.1.

Settings with Regularity or MHR. Settings in which the signal distribution F is regular (resp., MHR). That is, for every bidder i and signal profile s_{-i}, the conditional marginal distribution $F_i(\cdot \mid s_{-i})$ is regular (resp., MHR). The multivariate normal distribution in Example 4.2.3 is MHR and thus also regular. Regularity arises in the IPV setting as a necessary condition for truthfulness of Myerson's mechanism without an additional ironing procedure.

Settings with Affiliation. Settings which arise by imposing *affiliation* – a form of positive correlation – on the joint signal distribution F with density f. Affiliation was introduced by Milgrom and Weber [1982] and since then has become a standard assumption in the context of correlated and interdependent values, so much so that it is considered "almost synonymous with dependence in auctions" [de Castro, 2010]. It is related to many well-studied mathematical concepts such as association, the FKG inequality and log-supermodularity [Alon and Spencer, 2008].

Intuitively, signals are affiliated when observing a subset of high signals makes it more likely that the remaining signals are also high. Formally, for every pair of signal profiles $\vec{s}, \vec{t} \in [0, \omega_i]^n$ it must hold that

$$f(\vec{s} \vee \vec{t}) f(\vec{s} \wedge \vec{t}) \geq f(\vec{s}) f(\vec{t}), \tag{4.4}$$

where $(\vec{s} \vee \vec{t})$ is the component-wise maximum of $\vec{s}$ and $\vec{t}$, and $(\vec{s} \wedge \vec{t})$ is the component-wise minimum of $\vec{s}$ and $\vec{t}$. Note that the inequality in (4.4) holds with equality for independent signals. It also holds for the multivariate normal distribution in Example 4.2.3, which is affiliated since all pairwise covariances are nonnegative [de Castro and Paarsch, 2010]. An example showing that affiliation is a stronger condition than positive correlation is the bivariate uniform distribution over support $\{(1,2),(2,1),(3,3),(4,4),(5,5),(6,6)\}$; the two variables are positively correlated but not affiliated.

Symmetric Settings. Symmetry involves assumptions on both valuation functions and the signal distribution: (a) For every bidder i it is assumed that $v_i(\vec{s}) = v(s_i, s_{-i})$, where v is common to all bidders and symmetric in its last $n-1$ arguments; (b) The joint density f is assumed to be defined on support $[0,\omega]^n$ and symmetric in all its arguments. In a symmetric setting bidders thus have the same conditional densities, revenue curves and virtual value functions. Their values may be different however, since their own signal plays a distinct role in the valuation function. Notably, Milgrom and Weber [1982] study symmetric settings with affiliation; a concrete example of such a setting is weighted-sum values (Example 4.2.1) with the multivariate normal signals of Example 4.2.3.

Settings with a Single Crossing Condition. Settings in which a *single crossing* condition is imposed on either values or virtual values. Let $x_1(\cdot),\ldots,x_n(\cdot)$ be some functions (such as values or virtual values) of the signals; a single crossing condition captures the idea that bidder i's signal has a greater influence on x_i than on any other bidder's function x_j.[4] Formally, for every i and $j \neq i$, and for every $\vec{s}$,

$$\frac{\partial x_i}{\partial s_i}(\vec{s}) > \frac{\partial x_j}{\partial s_i}(\vec{s}). \tag{4.5}$$

Weaker versions of single crossing may require a non-strict inequality, or that the inequality hold only for $i, j, \vec{s}$ such that $x_i(\vec{s}) = x_j(\vec{s}) = \max_k\{x_k(\vec{s})\}$. Stronger versions may require the left-hand side of Equation 4.5 to be non-negative and the right-hand side non-positive (see, e.g., Lemma 4.5.4), or even that for every $\vec{s}$ such that $s_i > s_j$, $x_i(\vec{s}) > x_j(\vec{s})$ (see, e.g., Lemma 4.5.5). The weighted-sum values in Example 4.2.1 are weakly single crossing.

[4] The term "single crossing" comes from the fact that keeping all other signals fixed and varying only s_i, x_i as a function of s_i is "steeper" than x_j, so the two cross at most once.

4.3 Mechanisms and Solution Concepts

By the revelation principle, we focus without loss of generality on direct mechanisms, in which bidders directly report their signals. An exception is the English auction discussed below (Section 4.3.3). We restrict attention to IC mechanisms and so make no distinction between reported and actual signals.

A *randomized* mechanism M consists of an *allocation rule* $x_i(\cdot)$ and a *payment rule* $p_i(\cdot)$ for every bidder i, where $x_i(\vec{s})$ is the probability over the internal randomness of the mechanism that bidder i wins given the other signals $\vec{s}$, and $p_i(\vec{s})$) is the expected payment of bidder i given s_{-i}. If M is deterministic, $x_i(\vec{s}) \in \{0, 1\}$ and $p_i(\vec{s})$ is the actual payment. We assume risk-neutral bidders with quasi-linear utilities, i.e., bidder i's utility is $x_i(\vec{s})v_i(\vec{s}) - p_i(\vec{s})$.

4.3.1 Solution Concepts

Mechanism design aims to define the rules of a game played by the bidders, such that a solution of the game has desirable properties, in particular high revenue subject to IC and IR. The importance of the solution concept – i.e. what constitutes a solution to the game – is discussed in Section 1.2.4. We now formally define the three major solution concepts described there.[5]

Definition 4.3.1 (Ex post IC and ex post IR mechanism)**.** A mechanism is ex post IC and ex post IR if for every bidder i, true signal s_i, false report $\tilde{s}_i$, and signal profile s_{-i},

$$x_i(\vec{s})v_i(\vec{s}) - p_i(\vec{s}) \geq x_i(\tilde{s}_i, s_{-i})v_i(\vec{s}) - p_i(\tilde{s}_i, s_{-i}); \qquad (4.6)$$

$$x_i(\vec{s})v_i(\vec{s}) - p_i(\vec{s}) \geq 0. \qquad (4.7)$$

(Inequality 4.6 is the ex post IC condition and Inequality 4.7 is the ex post IR condition.) In words: participating and truthtelling is an ex post equilibrium of the corresponding game, that is, it is a Nash equilibrium in the ex post stage of the game where signals are publicly known.

[5] All IC and IR conditions hold in expectation over the internal randomness of the mechanism.

Definition 4.3.2 (Dominant strategy IC mechanism)**.** A mechanism is dominant strategy IC if for every bidder i, true signal profile $\vec{s}$, and reported signal profile $\vec{r}$,

$$x_i(s_i, r_{-i})v_i(\vec{s}) - p_i(s_i, r_{-i}) \geq x_i(\vec{r})v_i(\vec{s}) - p_i(\vec{r}),$$

i.e., truthtelling is a dominant-strategy equilibrium of the corresponding game.

Definition 4.3.3 (Bayesian IC and interim IR mechanism)**.** A mechanism is Bayesian IC and interim IR if for every bidder i, true signal s_i, and false report $\tilde{s}_i$,

$$\begin{aligned} \mathbb{E}_{s_{-i}}[x_i(\vec{s})v_i(\vec{s}) - p_i(\vec{s})] &\geq \mathbb{E}_{s_{-i}}[x_i(\tilde{s}_i, s_{-i})v_i(\vec{s}) - p_i(\tilde{s}_i, s_{-i})]; \\ \mathbb{E}_{s_{-i}}[x_i(\vec{s})v_i(\vec{s}) - p_i(\vec{s})] &\geq 0. \end{aligned}$$

That is: participating and truthtelling is a Bayes-Nash equilibrium of the corresponding game in the interim stage, in which each individual knows his own signal but not the others.

4.3.2 Discussion of the Ex Post Solution Concept

The above definitions show that ex post is a weaker solution concept than dominant strategies (for which truthfulness holds for any reported signal profile), and a stronger one than Bayesian/interim (whose guarantees are in expectation over the true signal profile). We now briefly discuss our focus on this intermediate solution concept. For additional discussion see [Segal, 2003; Milgrom, 2004; Bergemann and Morris, 2005; Chung and Ely, 2006, 2007].

Ex Post vs. Bayesian. The solution concept most widely used in mechanism design theory is Bayes-Nash equilibrium [Chung and Ely, 2006]; in practice, the common first-price and second-price auctions in interdependent settings are Bayesian and not ex post. On the flip side, the Crémer-McLean mechanism has been criticized as impractical for, among other issues, lack of the ex post IR property. This makes the outcome instable in the sense that bidders may regret their participation in hindsight, and attempt to exercise their de facto "veto" power of walking away from the auction

– refusing to collect their winnings and to honor their payments (cf., [Compte and Jehiel, 2009]). The lack of ex post IR also requires the Crémer-McLean mechanism to rely on bidders' knowledge of the joint distribution, without which they cannot determine whether participating is rational.

Focusing on ex post mechanisms prevents these issues. Ex post IC and ex post IR are "no regret" properties – for any realization of the signals, bidders regret neither participating in the auction nor reporting their signals truthfully, even when all signals become publicly known. This makes the mechanism more robust (and thus closer to the computer science worst-case approach). To decide whether to participate and how to report, bidders do not have to know the signal distribution, only the signal support and the valuation functions. This is compatible with Wilson's doctrine of detail-free mechanisms that are robust to detailed knowledge of the distribution [Wilson, 1987]. Among other advantages, robustness saves transaction costs associated with learning about opponents' distributions, and also benefits the seller, who may be wary of using a Bayesian mechanism if unsure how well bidders are informed.

We now mention two caveats to the ex post approach. First, in settings such as the mineral rights model (Example 4.2.2), one can argue that a bidder's knowledge of his own valuation function $v(\vec{s}) = \mathbb{E}_{V \sim F_V}[V \mid \vec{s}]$ depends on his knowledge of the others' distributions – this is necessary for him to derive v from the publicly known distribution F_V. In a model that crucially depends on bidders' knowledge of each other's distributions, and assuming the seller is aware that bidders are well-informed, there is less added robustness in an ex post solution over a Bayesian one. Note however that this issue does not arise in settings such as Example 4.2.1, and also that it arguably indicates the "type space" is simply not rich enough (cf., [Bergemann and Morris, 2005]). A second and related caveat is that it is debatable whether ex post is *necessary* for robustness; this question and more generally the theoretical foundation of robustness in mechanism design is discussed in [Segal, 2003; Bergemann and Morris, 2005; Chung and Ely, 2006, 2007] and references within.

Ex Post vs. Dominant Strategies. For *private* values (whether correlated or independent), dominant strategy IC and ex post IC coincide. For non-private values, however, the concept of dominant strategy IC guarantees an even stronger no regret property than the concept of ex post IC, since it does not depend on the other bidders reporting truthfully. For example, in the weighted-sum values case (Example 4.2.1), if bidder j under-reports his signal s_j, and bidder i somehow knows j's true signal, in an ex post mechanism i may potentially benefit by over-reporting his signal s_i, so that his true value is reflected by the mechanism.

The following example demonstrates that the dominant strategy IC requirement may be too strong for a deterministic mechanism to extract non-trivial revenue. The example involves (by necessity) non-private values, and shows there are cases in which there is an ex post IC and IR mechanism with positive revenue, whereas every dominant strategy IC mechanism has zero revenue.

Example 4.3.4. Two bidders compete for a single item. Their values are $v_1 = s_1 s_2$ and $v_2 = 0$, where $s_1, s_2 \in \{0, 1\}$. Consider a mechanism that allocates the item to bidder 1 and charges $p < 1$ only if both bidders report nonzero signals, and otherwise does not allocate the item. Observe that this mechanism is ex post IC and IR, and has nonzero revenue. Similarly, for a dominant strategy IC mechanism to have nonzero revenue, it must extract from bidder 1 a payment bounded away from 0 for reported signal profile $\vec{r} = (1, 1)$ (if one of the reported signals is 0, by ex post IR the mechanism gets zero revenue). However, if the true signal profile is $\vec{s} = (1, 0)$ but bidder 2 reports $r_2 = 1$, then bidder 1 is better off reporting $r_1 = 0$ untruthfully, in contradiction to dominant strategy IC.

4.3.3 The English Auction

The *English auction* is an ascending price auction that operates in "value space": bidders act upon their postulated values rather then report their signals to the mechanism. Specifically, in order to determine when to (irrevocably) drop out of the ascending auction, bidders must constantly update their conjectured value based on their observations up to the current point in the auction. All other mechanisms we consider are direct revelation mechanisms that work in "signal space", and the English auction provides an indirect implementation for them [Lopomo, 2000; Chung and Ely, 2007].

The most relevant version of the English auction for our purpose is the so-called "Japanese" version (for details see, e.g., [Krishna, 2010]). In this version, the auctioneer gradually raises the price of the item for sale; once a bidder finds the price too high, he indicates (for instance by lowering his hand) that he is no longer participating in the auction. The winner pays the price at which the second-to-last bidder dropped out. The crucial aspect of the English action is that it is *open* – the prices at which bidders drop out are observed by all. The English auction has a unique (cf., [Bikhchandani et al., 2002]) symmetric equilibrium studied by Milgrom and Weber [1982]; thanks to the openness of the auction, this equilibrium has the remarkable property of ensuring no ex post regret.

4.4 Related Work

In this section we discuss related work on revenue guarantees of auctions in single-parameter settings. Section 4.5 of this chapter can be seen as unifying and generalizing many previous results dispersed across the mechanism design literature; we now survey these results as well as present previous work on computational aspects. Note that this section is not a prerequisite to Section 4.5, which is self-contained.

In describing previous results and in Section 4.5 itself we use the following terminology: We refer to results similar to Proposition 4.5.1 as *characterization results*; these state necessary and sufficient conditions on the allocation and/or payment rules such that the resulting mechanism is IC and IR with respect to the desired solution concept. By *virtual surplus results* we mean results similar to Proposition 4.5.2, showing that revenue equals virtual surplus in expectation for an appropriate definition of virtual values. *Optimal mechanism results* state conditions under which the optimal mechanism can be derived from the virtual surplus results.

4.4.1 Independent Values, Bayesian Solution

Myerson's 1981 paper lays the foundations of optimal mechanism design: Myerson considers Bayesian IC and interim IR mechanisms in the IPV model, and establishes characterization, virtual surplus and optimal mechanism results. The optimal mechanism turns out to be deterministic, dominant strategy IC and ex post IR, while achieving optimality among all randomized, Bayesian IC and interim IR counterparts. A regularity condition simplifies the Myerson mechanism but is not required.

Myerson's characterization and virtual surplus (but not optimal mechanism)[6] results apply to non-private (independent) values as well, for an appropriately modified definition of virtual values [Bulow and Klemperer, 1996; Klemperer, 1999]. Additional work on optimal auctions in settings with non-private values includes [Branco, 1996; Ülkü, 2013].

4.4.2 Interdependent Values, Bayesian Solution

Myerson's theory does not directly apply when there's correlation among the bidders. The complicating issue is that in the presence of correlation, the allocation and payment rules for a bidder may depend not only on his reported signal, but also on his true signal through correlation with other bidders' signals. Crémer and McLean [1985, 1988] design an ex post IC but interim IR auction, which extracts full welfare in expectation under a mild "full rank" condition on the correlation; their mechanism is a generalized VCG auction, augmented with carefully designed lotteries. McAfee and Reny [1992] extend this result from discrete to continuous signals.

In a classic paper, building upon early work by Wilson [1969], Milgrom and Weber [1982] lay out a general model of interdependent values, and develop the *linkage principle* in place of the revenue equivalence principle. They apply the linkage principle to rank the common auction formats (first-price, second-price, English and Dutch auctions) according to their expected revenue, when signals are affiliated and bids form a symmetric Bayes-Nash equilibrium (actually an ex post equilibrium for the English auction).

[6] Remarkably, Myerson's model does allow for a natural but restricted form of non-private values.

4.4.3 Interdependent Values, Ex Post Solution

The ex post solution concept has generated much interest in the last decade; we now survey several papers most related to our work.

Correlated Values. Section 4.5 is closely related to the work of Segal [2003] (see also [Chung and Ely, 2007]). Segal studies ex post IC and ex post IR mechanisms for selling multiple units of an item in the correlated values model. He gives a characterization and virtual surplus result based on conditional virtual values as defined in Equation 4.2. Segal also derives an optimal mechanism result under regularity and the assumption that conditional virtual values are single crossing; he notes that for affiliated signals the latter assumption holds. Our results in Section 4.5 can be seen as a generalization of Segal's results beyond multi-unit settings and beyond private values.

Interdependent Values. A characterization result for ex post IC and ex post IR mechanisms in the interdependent values model is found in Chung and Ely [2006] (see also [Lopomo, 2000; Vohra, 2011]). Chung and Ely's result is via an interesting connection among the following characterizations: ex post IC and ex post IR mechanisms for interdependent values, Bayesian IC and interim IR mechanisms for independent values, ex post IC (equivalently, dominant strategy IC) and ex post IR mechanisms for private values, and IC and IR mechanisms for a single bidder (for which all solution concepts converge).

For a single item, Vohra [2011] states virtual surplus and optimal mechanism results, where the former is with respect to conditional virtual values as defined in Equation 4.1, and the latter is under the assumption that conditional virtual values are single-crossing. Vohra notes that for this result to be useful, one must identify restrictions on the distribution and valuation functions that would lead to single crossing. Such restrictions appear in Section 4.5. Recently and concurrently to our work, Csapó and Müller [2013] and Li [2013] also develop virtual surplus results for interdependent values. Csapó and Müller [2013] apply these in the context of supplying a single public good and assuming discrete signals. A more detailed description of the work of Li [2013] appears in Section 4.4.6 below.

Beyond Single-Parameter. Jehiel et al. [2006] show impossibility results for ex post implementation in multi-parameter settings. In particular, in a public decision setting with generic valuations, the only deterministic social choice functions that are ex post implementable are trivial (i.e., constant).

4.4.4 English Auction with Interdependent Values

The importance of the English auction in the context of non-IPV settings has long been recognized. For non-private values, Bulow and Klemperer study symmetric bidders under a strong regularity condition, and show that the English auction's expected revenue (with or without reserve) equals the expected conditional virtual surplus as defined in Equation 4.1 above [Bulow and Klemperer, 1996, Lemmas 1 and 2]. They establish that the English auction with optimally-chosen reserve is optimal among all Bayesian IC and interim IR mechanisms [Bulow and Klemperer, 1996, Theorem 2]. McAfee and Reny [1992] show that while the English auction with reserve is optimal in a symmetric independent setting, minor perturbations of the distribution can introduce correlation and destroy optimality in comparison to other Bayesian mechanisms. They conjecture that the English auction's prevalence in practice has to do with the need to perform well in a variety of circumstances, and call for formalizing a notion of robustness.

The ex post solution concept adopted here is precisely such a robustness notion. The following two results are close to our work, and we re-derive them as corollaries in our framework (see Corollaries 4.5.12 and 4.5.13). For correlated values and under regularity and affiliation assumptions, Chung and Ely [2007] show that the English auction with optimally-chosen reserve is optimal among all ex post IC and ex post IR mechanisms. For interdependent values in the Milgrom and Weber setting, Lopomo [2000] identifies conditions under which the English auction with optimally-chosen reserve is optimal among all mechanisms in a class he terms as "no-regret" (see Definition 4.5.8).

For completeness we describe in more detail the result of Lopomo. He studies mechanisms with a "no-regret" equilibrium in the sense that each bidder has no incentive to revise his decisions after observing his opponents behavior; ex post direct revelation mechanisms are a subclass in which no regret holds after observing all signals. Lopomo shows that payments in a no-regret equilibrium must be determined by the allocation rule and by the bidders' willingness to pay given all information revealed by the others' actions. He then shows that for a fixed allocation rule and assuming affiliation, the expected revenue is maximized by revealing all information to the winning bidder. The next step is to express the expected revenue as conditional virtual surplus, from which the optimal allocation rule can be derived under additional assumptions. Lopomo then shows that in the Milgrom-Weber model, the English auction with reserve implements the optimal allocation rule. Lopomo also demonstrates that the English auction is not optimal among the wider class of interim IR mechanisms with a "losers do not pay" restriction.

4.4.5 Computational Considerations

Oracle vs. Explicit Model. An alternative approach to characterizing ex post IC and ex post IR in order to find the optimal mechanism is designing a computationally-tractable algorithm that computes or approximates such a mechanism. This requires addressing the question of how to represent the joint signal distribution. In the *oracle model*, the distribution is available to the mechanism/algorithm as a black-box, which can be queried with respect to conditional probabilities; upon receiving a signal profile as input, the mechanism/algorithm submits queries and returns an allocation rule and payments. In the *explicit model*, the distribution is explicitly provided as input, upon which the algorithm outputs the mechanism's allocation and payment rules.

Note that a hardness result in the explicit model implies hardness in the oracle model, whereas a positive algorithmic result in the oracle model implies such a result in the explicit model.

Computational Hardness. Papadimitriou and Pierrakos [2015] prove computational hardness in the explicit model – even for correlated values, finding the optimal *deterministic*, ex post IC and ex post IR mechanism when there are at least three bidders is NP-hard. When there are exactly two bidders, the optimal deterministic mechanism can be computed in polynomial time and is optimal among all randomized mechanisms as well [Dobzinski et al., 2011; Papadimitriou and Pierrakos, 2015].

Randomized Mechanisms. Computational hardness does not extend to optimal randomized mechanisms, which can be computed in the explicit model in polynomial time for all single-parameter domains, as well as unit-demand and additive multi-parameter domains [Dobzinski et al., 2011; Papadimitriou and Pierrakos, 2015]. With at least three bidders, randomized mechanisms can strictly outperform deterministic mechanisms in terms of expected revenue, albeit by a small constant factor (see next paragraph). This implies that additional assumptions are needed for a Myerson-like deterministic mechanism to be optimal (see Section 4.5).

Near-Optimal Mechanisms. In the oracle model with correlated values, Ronen [2001] designs the *lookahead auction* – a simple, deterministic, ex post IC and ex post IR mechanism, which guarantees a constant approximation to the optimal expected revenue. Dobzinski et al. [2011] build upon Ronen's work to design, for single-item settings, a deterministic mechanism that achieves a 5/3-approximation, and a randomized one that achieves a $(3/2+\epsilon)$-approximation (for improved bounds see [Chen et al., 2011]).

Informational Hardness for the English Auction. In the oracle model, Ronen and Saberi [2002] show that a deterministic English auction cannot achieve an approximation ratio better than 3/4 with respect to the optimal expected revenue. Due to the oracle setup, this bound is explicit and does not rely on complexity assumptions.

4.4.6 Applications

Myerson's theory has multiple applications in the IPV model. Examples of applications studied in the algorithmic game theory community include simple near-optimal auctions [Neeman, 2003; Hartline and Roughgarden, 2009], prior-independent mechanisms [Segal, 2003; Dhangwatnotai et al., 2015], and prior-free mechanisms [Goldberg et al., 2006].

In Section 4.6, we expand upon the theme of robustness by developing prior-independent mechanisms for interdependent values, using techniques from Dhangwatnotai et al. [2015]. Independently from and orthogonal to our work, Li [2013] shows a simple near-optimal auction for settings with interdependent values. She studies the VCG mechanism with monopoly reserves in matroid settings, where values satisfy a single crossing condition and the valuation distribution satisfies the generalized monotone hazard rate condition. Li shows that in expectation, VCG with monopoly reserves extracts at least $1/e$ of the full surplus.

4.5 Myerson Theory for Interdependent Values

The fundamental results of single-parameter optimal auction theory – Myerson's optimal mechanism and characterization results leading to it – do not carry over to interdependent settings. In this section we show that these results are at least partially recovered with small adaptations once we impose the ex post requirements.

The intuition behind this finding is as follows. The original proofs rely on signal independence so that both the probability x_i of winning and the expected payment p_i depend only on bidder i's reported signal, not on his true one. By switching from Bayesian IC and interim IR to ex post IC and ex post IR, we ensure the guarantees hold for any signal profile s_{-i}. Since we can now fix s_{-i}, rules x_i and p_i once again depend only on bidder i's reported signal, and so the independence assumption is no longer necessary.

We describe the organization of our results using the terminology introduced in Section 4.4. Section 4.5.1 develops characterization and virtual surplus results, and Sections 4.5.2 and 4.5.3 establish optimal mechanism results (Section 4.5.2 addresses correlated values while Section 4.5.3 deals with full interdependence). For completeness, Section 4.5.4 discusses indirect implementation by the English auction. Section 4.5.5 discusses the assumptions of single-crossing and regularity that are used in the optimal mechanism results of 4.5.2 and 4.5.3.

4.5.1 Characterization of Ex Post Mechanisms

We begin by developing the theory as far as we can with no assumptions on the setting, i.e., we refrain from adding any of the constraints in Section 4.2.3. Using an adaptation of the now-standard techniques developed by Myerson [1979], we show in Propositions 4.5.1 and 4.5.2 that characterization and equal-revenue results hold (for an exposition of these techniques see [Nisan, 2007, Theorem 9.39] and [Hartline, 2014]).

Proposition 4.5.1 (Characterization)**.** *For every interdependent values setting, a mechanism is ex post IC and ex post IR if and only if for every i, s_{-i}, the allocation rule x_i is monotone non-decreasing in the signal s_i, and the following payment identity and payment inequality hold:*

$$\begin{aligned} p_i(\vec{s}) &= x_i(\vec{s})v_i(\vec{s}) - \int_{v_i(0,s_{-i})}^{v_i(s_i,s_{-i})} x_i\left(v_i^{-1}(t \mid s_{-i}), s_{-i}\right) dt \\ &\quad - \left(x_i(0,s_{-i})v_i(0,s_{-i}) - p_i(0,s_{-i})\right); \\ p_i(0,s_{-i}) &\le x_i(0,s_{-i})v_i(0,s_{-i}). \end{aligned}$$

The payment identity and inequality imply that the allocation rule for every bidder determines the bidder's payment up to his expected payoff for a zero signal, and that this expected payoff must be non-negative. For private values, with a standard assumption of no positive transfers, the payment constraints simplify to the identity

$$p_i(\vec{s}) = x_i(\vec{s})s_i - \int_0^{s_i} x_i\left(t, s_{-i}\right) dt.$$

Proposition 4.5.2 (Revenue equals virtual surplus in expectation)**.** *For every interdependent values setting, the expected revenue of an ex post IC and ex post IR mechanism equals its expected conditional virtual surplus, up to an additive factor:*

$$\begin{aligned} \mathbb{E}_{\vec{s}}\left[\sum_i p_i(\vec{s})\right] &= \mathbb{E}_{\vec{s}}\left[\sum_i x_i(\vec{s})\varphi_i(s_i \mid s_{-i})\right] \\ &\quad - \sum_i \mathbb{E}_{s_{-i}}\left[\left(x_i(0,s_{-i})v_i(0,s_{-i}) - p_i(0,s_{-i})\right)\right] \end{aligned}$$

Algorithm 7: Myerson mechanism for interdependent values

1. Elicit signal reports $\vec{s}$ from the bidders.
2. Maximize the conditional virtual surplus by allocating to the feasible set S with the highest non-negative conditional virtual value $\sum_{i \in S} \varphi_i(s_i \mid s_{-i})$, breaking ties arbitrarily but consistently.
3. Charge every winner i a payment $p_i(\vec{s}) = v_i(s_i^*, s_{-i})$, where s_i^* is the threshold signal such that given the other signals s_{-i}, if i's signal were below the threshold he would no longer win the auction.

Note that the additive term is just minus the sum of the bidders' expected payoffs for zero signals. For private values this term disappears.

4.5.2 Optimal Mechanism for Correlated Private Values

Proposition 4.5.2 suggests that to optimize expected revenue, the best course of action is to maximize conditional virtual surplus pointwise. However the issue is monotonicity – even in the independent private values model, regularity is necessary for pointwise maximization to form a monotone allocation rule, and in more general models we need more assumptions (see discussion in Section 4.5.5).

In this section we focus on correlated values in matroid settings, with assumptions of regularity and affiliation (recall Section 4.2.3). An example of such a setting – which is symmetric in addition to regular and affiliated – is a single-item setting where bidders' values are drawn from the multivariate normal distribution in Example 4.2.3.

In Algorithm 7 we define a Myerson mechanism for interdependent values. The main result in this section is its optimality for correlated values under the above assumptions.

Theorem 4.5.3 (Myerson mechanism is ex post IC, IR and optimal)**.** *For every matroid setting with correlated values that satisfies regularity and affiliation, the Myerson mechanism is ex post IC, ex post IR, and optimal among all ex post IC and ex post IR mechanisms.*

The following lemma is key to our analysis of the Myerson mechanism's performance.

Lemma 4.5.4 (Single crossing of conditional virtual values)**.** *For every correlated values setting with regular affiliated distribution, raising signal s_i weakly increases bidder i's conditional virtual value, and weakly decreases all other conditional virtual values.*

Proof. First note that by regularity of the signal distribution, raising signal s_i increases bidder i's conditional virtual value $\varphi_i(s_i \mid s_{-i})$. It is left to prove the following claim: For every two bidders i, j, and every signal profile s_{-i}, bidder j's conditional virtual value is weakly decreasing in bidder $i's$ signal s_i.

Let $\tilde{s}_i \geq s_i$, and denote by $\tilde{s}_{-j}, s_{-j}$ the signal profiles excluding j with i's signal set to $\tilde{s}_i, s_i$, respectively. By definition, $\tilde{s}_{-j} \geq s_{-j}$. We show that $\varphi_j(s_j \mid \tilde{s}_{-j}) \leq \varphi_j(s_j \mid s_{-j})$, where

$$\begin{aligned} \varphi_i(s_j \mid \tilde{s}_{-j}) &= s_j - \frac{1 - F_j(s_j \mid \tilde{s}_{-j})}{f_j(s_j \mid \tilde{s}_{-j})}; \\ \varphi_i(s_j \mid s_{-j}) &= s_j - \frac{1 - F_j(s_j \mid s_{-j})}{f_j(s_j \mid s_{-j})}. \end{aligned}$$

By affiliation, $F_j(s_j \mid \tilde{s}_{-j})$ *dominates* $F_j(s_j \mid s_{-j})$ *in terms of hazard rate* [Krishna, 2010, Appendix D], i.e.,

$$\frac{1 - F_j(s_j \mid \tilde{s}_{-j})}{f_j(s_j \mid \tilde{s}_{-j})} \geq \frac{1 - F_j(s_j \mid s_{-j})}{f_j(s_j \mid s_{-j})},$$

and this is sufficient to complete the proof. □

We remark that an even stronger version of single crossing holds if bidders are symmetric.

Lemma 4.5.5 (Order of virtual values matches order of signals)**.** *For every correlated values setting with symmetric bidders and regular affiliated distribution, for every signal profile $\vec{s}$ such that $s_i \geq s_j$, the bidder with higher signal has higher conditional* *virtual value*

$$\varphi_i(s_i \mid s_{-i}) \geq \varphi_j(s_j \mid s_{-j}).$$

Proof. Given a signal profile $\vec{s}$ where $s_i \geq s_j$, we show that $\varphi_i(s_i \mid s_{-i}) \geq \varphi_j(s_j \mid s_{-j})$. Recall

$$\begin{aligned} \varphi_i(s_i \mid s_{-i}) &= s_i - \frac{1 - F_i(s_i \mid s_{-i})}{f_i(s_i \mid s_{-i})}; \\ \varphi_j(s_j \mid s_{-j}) &= s_j - \frac{1 - F_j(s_j \mid s_{-j})}{f_j(s_j \mid s_{-j})}. \end{aligned}$$

The following three inequalities follow from regularity, Lemma 4.5.4 (single crossing conditional virtual values), and symmetry of the bidders and distributions, respectively. These properties allow us to first replace s_i by $s_j \leq s_i$ in bidder i's virtual value, then compare bidder i's virtual value given signal s_j versus signal s_i for bidder j (i.e., replace s_{-i} by s_{-j}), and finally replace F_i, f_i by F_j, f_j, completing the proof:

$$\begin{aligned} s_i - \frac{1 - F_i(s_i \mid s_{-i})}{f_i(s_i \mid s_{-i})} &\geq s_j - \frac{1 - F_i(s_j \mid s_{-i})}{f_i(s_j \mid s_{-i})} \\ &\geq s_j - \frac{1 - F_i(s_j \mid s_{-j})}{f_i(s_j \mid s_{-j})} \\ &= s_j - \frac{1 - F_j(s_j \mid s_{-j})}{f_j(s_j \mid s_{-j})}. \end{aligned}$$

□

Next we establish that the allocation rule of the Myerson mechanism in Algorithm 7 is monotone in matroid settings.

Lemma 4.5.6 (Monotonicity)**.** *For every matroid setting with correlated values that satisfies regularity and affiliation, maximizing conditional virtual surplus is monotone.*

Proof. In a matroid setting, maximizing conditional virtual surplus can be implemented by a greedy algorithm, which considers bidders in non-increasing order of conditional virtual values, and adds them to the winning set if their conditional virtual value is nonnegative and feasibility is maintained. By Lemma 4.5.4, raising signal

s_i can only improve bidder i's ranking in the order of consideration, thus monotonicity holds. □

The example below demonstrates that the condition of a matroid setting in Lemma 4.5.6 is necessary.

Example 4.5.7 (Non-monotonicity beyond matroids)**.** Consider a correlated values setting with three bidders. The signals are drawn from the following regular affiliated distribution: Signal profiles $(1, 1, 1), (1, 2, 1), (2, 1, 1), (2, 2, 1)$ appear with probabilities $0.4, 0.1, 0.1, 0.4$. Assume that the feasible sets of the single-parameter auction environment are bidder sets $\{1, 2\}$ and $\{3\}$.

Now consider signal profile (1,1,1) and assume that bidder 1 raises his reported signal from 1 to 2. The bidders' conditional virtual value profile changes from $(0.75, 0.75, 1)$ to $(2, -3, 1)$. With the original signal reports, the feasible bidder set maximizing non-negative conditional virtual surplus was $\{1, 2\}$, whereas after bidder 1 raises his report it becomes $\{3\}$, contradicting monotonicity.

The problem arises since by raising his own signal, bidder 1 decreased the conditional virtual value of bidder 2, and so 1 and 2 no longer form the feasible set with highest non-negative conditional virtual surplus.

We our now ready to prove the main result regarding optimality of the Myerson mechanism.

Proof of Theorem 4.5.3. By the characterization of ex post mechanisms (Proposition 4.5.1) applied to private values, for every bidder i it is sufficient to show that the allocation rule x_i is monotone in the signal s_i, and that the payment identity $p_i(\vec{s}) = x_i(\vec{s})s_i - \int_0^{s_i} x_i(t, s_{-i})\, dt$ holds. Lemma 4.5.6 establishes monotonicity, and the payment identity holds by the following argument. The Myerson mechanism is deterministic and so either $x_i(\vec{s}) = 0$ or $x_i(\vec{s}) = 1$. In the former case, by monotonicity $x_i(t, s_{-i}) = 0$ for every $t \leq s_i$, so both sides of the identity are equal to zero. In the latter case, since s_i^* is bidder i's threshold signal, the right-hand side is

$$s_i - \int_0^{s_i} x_i(t, s_{-i})\, dt = s_i - (s_i - s_i^*) = s_i^*,$$

and for private values s_i^* is precisely the payment $p_i(\vec{s})$ charged by the Myerson mechanism.

It is left to show optimality. The expected revenue of an ex post IC and ex post IR mechanism is its expected virtual surplus (Proposition 4.5.2), and the Myerson mechanism maximizes virtual surplus for every signal profile. □

4.5.3 Optimal Mechanism for Interdependent Values

The results for correlated values generalize to interdependent values, however this setting is harder and requires further assumptions. Indeed, recall that the general conditional virtual value form in Equation 4.1 includes two extra dependencies on other bidders' signals relative to the form in Equation 4.2 which applies to correlated values. The extra assumptions are needed to establish monotonicity of the allocation rule in Algorithm 7 despite these dependencies.

We adopt the setting studied by Lopomo [2000] in the context of the English auction; namely, the assumptions we impose on our auction setting are that bidders are symmetric and have affiliated signals, and that the following conditions on the valuation function and distribution hold (for definitions see Section 4.2.3).

Definition 4.5.8 (Lopomo assumptions)**.**

1. MHR setting;
2. Bidders with higher signals have higher values, i.e., strong single crossing of values;
3. Bidders with higher signals have lower sensitivity of their value to their signal, i.e., the partial derivative of v_i with respect to s_i is weakly decreasing in s_i, and weakly increasing in s_j for every other j.

For every signal profile $\vec{s}$ such that $s_i \geq s_j$, the Lopomo assumptions imply:

$$\begin{aligned} v_i(\vec{s}) &\geq v_j(\vec{s}); \\ \frac{d}{ds_i} v_i(\vec{s}) &\leq \frac{d}{ds_j} v_j(\vec{s}). \end{aligned} \tag{4.8}$$

An example of a symmetric affiliated setting where the Lopomo assumptions hold is a single-item setting with weighted-sum values (Example 4.2.1), and signals drawn from the multivariate normal distribution of Example 4.2.3. Note that Equation (4.8) holds whenever values are multilinear.[7]

We can now state this section's main result - an analogue of Theorem 4.5.3 for interdependent values, showing that the Myerson mechanism defined in Algorithm 7 is optimal.

Theorem 4.5.9 (Myerson mechanism is ex post IC, IR and optimal)**.** *For every matroid setting with interdependent values that satisfies affiliation, symmetry and the Lopomo assumptions, the Myerson mechanism is ex post IC, ex post IR, and optimal among all ex post IC and ex post IR mechanisms.*

The proof of Theorem 4.5.9, like the proof of its analogue Theorem 4.5.3, boils down to showing monotonicity of the Myerson mechanism. We now turn to establishing monotonicity, using that the order of conditional virtual values coincides with the order of signals. This strong form of single crossing for conditional virtual values (cf., Lemmas 4.5.4 and 4.5.5) is stated in the following lemma, which can be viewed as a generalization of the same result in the IPV model for a symmetric setting that satisfies regularity.

Lemma 4.5.10 (Order of virtual values matches order of signals)**.** *For every symmetric setting with interdependent values that satisfies affiliation and the Lopomo assumptions, for every signal profile $\vec{s}$ such that $s_i \geq s_j$, the bidder with higher signal has higher conditional virtual value*

$$\varphi_i(s_i \mid s_{-i}) \geq \varphi_j(s_j \mid s_{-j}).$$

Proof. Given a signal profile $\vec{s}$ where $s_i \geq s_j$, we show that $\varphi_i(s_i \mid s_{-i}) \geq \varphi_j(s_j \mid s_{-j})$.

[7] Recall that a function is multilinear if it is separately linear in each one of its variables – weighted sums and products are examples.

Recall

$$\begin{aligned}\varphi_i(s_i \mid s_{-i}) &= v(s_i, s_{-i}) - \frac{1 - F_i(s_i \mid s_{-i})}{f_i(s_i \mid s_{-i})} \cdot \frac{d}{ds_i} v(s_i, s_{-i});\\ \varphi_j(s_j \mid s_{-j}) &= v(s_j, s_{-j}) - \frac{1 - F_j(s_j \mid s_{-j})}{f_j(s_j \mid s_{-j})} \cdot \frac{d}{ds_j} v(s_j, s_{-j}).\end{aligned}$$

We now compare the right-hand-side terms of the above equations.

By the second Lopomo assumption, bidders with higher signals have higher values and so $v(s_i, s_{-i}) \geq v(s_j, s_{-j})$. By the third Lopomo assumption, bidders with higher signals have a lower sensitivity of their value to their own signal and so $0 \leq \frac{d}{ds_i} v(s_i, s_{-i}) \leq \frac{d}{ds_j} v(s_j, s_{-j})$ (using the assumption that values are increasing in signals).

It is left to show that $0 \leq \frac{1-F_i(s_i|s_{-i})}{f_i(s_i|s_{-i})} \leq \frac{1-F_j(s_j|s_{-j})}{f_j(s_j|s_{-j})}$. The following three inequalities follow from the first Lopomo assumption of MHR, the affiliation assumption and resulting hazard rate dominance [Krishna, 2010, Appendix D], and the bidders' symmetry and symmetry of distributions, respectively. These properties allow us to first replace s_i by $s_j \leq s_i$ in bidder i's inverse hazard rate, then compare the hazard rates given signal s_j versus signal s_i for bidder j (i.e., replace s_{-i} by s_{-j}), and finally replace F_i, f_i by F_j, f_j to complete the proof:

$$\begin{aligned}\frac{1 - F_i(s_i \mid s_{-i})}{f_i(s_i \mid s_{-i})} &\leq \frac{1 - F_i(s_j \mid s_{-i})}{f_i(s_j \mid s_{-i})}\\ &\leq \frac{1 - F_i(s_j \mid s_{-j})}{f_i(s_j \mid s_{-j})}\\ &= \frac{1 - F_j(s_j \mid s_{-j})}{f_j(s_j \mid s_{-j})}.\end{aligned}$$

□

We remark that in addition to the set of conditions in the statement of Lemma 4.5.10, there are also different (incomparable) conditions that suffice to prove a form of single crossing for conditional virtual values. One example of an alternative set of sufficient conditions excluding bidder symmetry is affiliated signals and the Lopomo assumptions, where we strengthen the first Lopomo assumption to the *generalized MHR* assumption of Li [2013].[8]

Lemma 4.5.10 implies that the Myerson mechanism is monotone.

Lemma 4.5.11 (Monotonicity). *For every symmetric setting with interdependent values that satisfies affiliation, symmetry and the Lopomo assumptions, maximizing conditional virtual surplus is monotone.*

Proof. By Lemma 4.5.10, raising signal s_i only improves bidder i's ranking in the greedy order of consideration by conditional virtual value. This is sufficient for monotonicity by a similar argument to that in the proof of Lemma 4.5.6. □

Monotonicity is sufficient to now prove optimality of the Myerson mechanism.

Proof of Theorem 4.5.9. By the characterization of ex post mechanisms (Proposition 4.5.1) applied to interdependent values, for every bidder i it is sufficient to show that the allocation rule x_i is monotone in the signal s_i, and that the payment identity and payment inequality hold. Lemma 4.5.11 establishes monotonicity. The payment inequality $p_i(0, s_{-i}) \leq x_i(0, s_{-i})v_i(0, s_{-i})$ holds with equality since if $x_i(0, s_{-i}) = 0$ then $p_i(0, s_{-i}) = 0$, and if $x_i(0, s_{-i}) = 1$ then $p_i(0, s_{-i}) = v_i(s_i^*, s_{-i})$ where $s_i^* = 0$. As for the payment identity, by determinism and monotonicity of the Myerson mechanism and assuming $x_i(\vec{s}) = 1$,

$$\begin{aligned} x_i(\vec{s})v_i(\vec{s}) - \int_{v_i(0,s_{-i})}^{v_i(s_i,s_{-i})} x_i\left(v_i^{-1}(t \mid s_{-i}), s_{-i}\right) dt &= v_i(\vec{s}) - (v_i(s_i, s_{-i}) - v_i(s_i^*, s_{-i})) \\ &= v_i(s_i^*, s_{-i}) \\ &= p_i(\vec{s}). \end{aligned}$$

It is left to show optimality. By Proposition 4.5.2, the expected revenue of an ex post IC and ex post IR mechanism is equal to its expected virtual surplus up to an additive term $\mathbb{E}_{s_{-i}}[\sum_i (x_i(0, s_{-i})v_i(0, s_{-i}) - p_i(0, s_{-i}))]$. The Myerson mechanism maximizes the virtual surplus for every signal profile, and sets the non-negative additive term to zero, thus achieving optimality. □

[8] The generalized MHR condition defined and utilized in [Li, 2013] states that the second term in Equation 4.1 is decreasing in s_i for all i and s_{-i}.

4.5.4 Indirect Implementation via the English Auction

For completeness, we conclude this section with results by Lopomo [2000] and Chung and Ely [2007] on implementing the optimal mechanism via the English auction with a carefully chosen reserve, when there's a single item for sale and the bidders are symmetric. The relation between the previous subsections and these results is the same as the relation between Myerson's original mechanism and the following well-known result in the IPV model – the second-price auction with an optimal reserve maximizes the expected revenue from selling a single item when bidders are symmetric and regularity holds. By replacing the second-price auction with the English auction, and setting the reserve price after all bidders but one have dropped out and revealed their information, we get an indirect implementation of the optimal mechanism that works directly in value space.

Corollary 4.5.12 ([Chung and Ely, 2007]). *For every symmetric single-item setting with correlated values that satisfies regularity and affiliation, the English auction with optimal reserve price is optimal among all ex post IC and ex post IR mechanisms.*

Proof. Consider the symmetric ex post equilibrium of the English auction, where assuming δ bidders have dropped out so that their signals are known, every bidder calculates what his value would be if the $n - \delta - 1$ unknown signals were equal to his privately known signal, and drops out if the price reaches this value [Milgrom and Weber, 1982]. The last bidder to remain in the auction is the bidder with highest signal, who by Lemma 4.5.5 also has highest conditional virtual value. Set an optimal reserve price for this bidder. Since the reserve is set after signals of all other bidders are revealed, it guarantees that this bidder wins precisely when his conditional virtual value given the other signals is non-negative. The resulting mechanism thus maximizes conditional virtual surplus for every signal profile, and is equivalent to the optimal Myerson mechanism in Algorithm 7. □

The same proof with Lemma 4.5.5 replaced by Lemma 4.5.10 shows the following.

Corollary 4.5.13 ([Lopomo, 2000]). *For every symmetric single-item setting with*

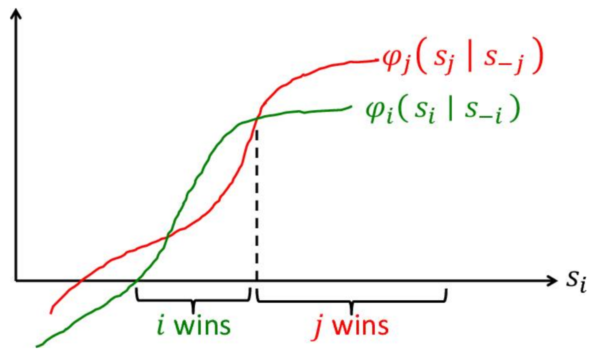

Figure 4.1: Regularity without single crossing. The conditional virtual values of bidders i and j both increase as signal s_i increases, crossing each other more than once. Observe that if i and j are the two bidders with highest conditional virtual values, the virtual-surplus-maximizing allocation rule is not monotone.

interdependent values that satisfies affiliation and the Lopomo assumptions, the English auction with optimal reserve price is optimal among all ex post IC and ex post IR mechanisms.

4.5.5 Discussion of Assumptions

The Myerson mechanism in Algorithm 7 is truthful for interdependent values only if its virtual-surplus-maximizing allocation rule is monotone. Even in the IPV model, a regularity assumption is necessary for the Myerson mechanism to be monotone without an additional *ironing* procedure (we dicuss ironing for interdependent values below). In the interdependent values model, an additional single crossing assumption for conditional virtual values is required. Figure 4.1 and Example 4.5.14 show what could go wrong if this assumption is violated.

Example 4.5.14 (Regularity without single crossing in a wallet game)**.** Consider a symmetric, single-item setting with non-private values. Let the values be weighted-sum values with parameter $\beta = 1$ (Example 4.2.1), and let signals be independently drawn from the regular distribution $G(s) = 1 - s^{-2}$. Plugging into Equation 4.3, the conditional virtual value of bidder i is $\sum_j s_j - (1-G(s_i))/g(s_i) = s_i/2 + \sum_{j \neq i} s_j$. Thus when bidder i's signal increases by Δs_i, his own conditional virtual value increases by $\Delta s_i/2$ while his competitors' conditional virtual values increase by the full difference

Δs_i. Single crossing is thus violated and the Myerson mechanism will not be truthful.[9]

Single crossing as an assumption in itself is quite opaque; above we have identified economically-meaningful conditions on the auction environment that are sufficient for single crossing to hold, both in the special case of correlated values and in the more general case of interdependent values. While the standard assumption of affiliation is sufficient in the correlated values setting, this is no longer the case for full interdependence. However, as we have seen, symmetry together with the Lopomo assumptions are sufficient (and alternative sufficient conditions exist as well). While the known computational hardness results imply that some of these assumptions (or alternative ones) are required for optimality of the deterministic Myerson mechanism, achieving a precise understanding of what is necessary remains an open question.

Ironing. Does the method of ironing developed by Myerson [1981] work for interdependent values?[10] Technically, the ironing method can easily be applied to conditional virtual values, i.e., the Myerson mechanism with ironing is well-defined for interdependent values. Furthermore, it still holds that ironed conditional virtual surplus gives an upper bound on conditional virtual surplus, which is tight for mechanisms that "respect" the ironed intervals (in the sense that the allocation does not change along such an interval).

The crucial difference from the IPV model is that the expected revenue of the Myerson mechanism with ironing can be *strictly lower* than the expected ironed conditional virtual surplus. Thus, even though the Myerson mechanism with ironing truthfully maximizes the latter, it is no longer guaranteed to achieve the maximum expected revenue. This gap arises due to the fact that the Myerson mechanism with ironing does not respect ironed intervals. Indeed, while the increase in a bidder's signal does not change his ironed conditional virtual value within an ironed segment, it may change others' ironed conditional virtual values, thus modifying the allocation.

[9] In fact, in this example the order of conditional virtual values is exactly the order of signals *reversed*; we can show that the optimal allocation rule in this case is to pick a random winner – any other "more sensible" allocation rule will violate monotonicity.

[10] A description of the ironing method is beyond our scope; for an introduction see, e.g., [Hartline, 2014].

Since the Myerson mechanism with ironing is deterministic, it is not surprising that ironing does not allow us to dispose altogether of assumptions on the valuations and/or distributions, as is the case in Myerson's paper (recall the negative results in [Dobzinski et al., 2011; Papadimitriou and Pierrakos, 2015]). It is open whether ironing can help weaken these assumptions.

4.6 Prior-Independence for Non-private Values

In this section we begin to develop a theory of prior-independence for interdependent values. Our main result is, for non-private values and the setting studied above (Section 4.5.3), a prior-independent mechanism that achieves a constant-factor approximation to the optimal expected revenue. An interesting direction for future work is to design good prior-independent mechanisms for general interdependent values.

This section is organized as follows: After presenting the setting and stating the main result, we prove our result for a simple single item setting in which bidders share a pure common value for the item. Section 4.6.5 generalizes the proof to matroid settings with non-private values.

4.6.1 Setting

We study a matroid setting with non-private values where signals are independent, and the Lopomo assumptions (Definition 4.5.8) hold.[11]

Symmetry. As is standard in the prior-independence literature (see, e.g., [Bulow and Klemperer, 1996; Goldberg et al., 2006; Segal, 2003; Dhangwatnotai et al., 2015]), we focus on symmetric environments with $n \geq 2$ bidders, where symmetry here is as defined in Section 4.2.3 (in particular, the feasible region need not be symmetric).

Notation. As above, let F denote the joint distribution of the independent signals. Let G be the distribution from which each of the i.i.d. signals is drawn, and let g be

[11]Note that for digital goods settings, our results hold more generally; i.e., we no longer need all the Lopomo assumptions, as explained in Section 4.6.5.

the corresponding density (G is the marginal distribution of the signals given the joint product distribution F). We denote by $G_{|s_{-i}}(\cdot), g_{|s_{-i}}(\cdot)$ the distribution and density of bidder i's value given the signal profile s_{-i} of the other bidders.

Remark 4.6.1 (Strong MHR guarantee)**.** Observe that by independence, $G_{|s_{-i}}(\cdot)$ is simply the distribution of s_i. It follows that the inverse hazard rate of bidder i's value given signal profile $\vec{s}$ is

$$\frac{1 - G_{|s_{-i}}(v_i(\vec{s}))}{g_{|s_{-i}}(v_i(\vec{s}))} = \frac{1 - G(s_i)}{g(s_i)} \cdot \frac{d}{ds_i} v_i(\vec{s}). \tag{4.9}$$

By the first and third Lopomo assumptions, the inverse hazard rate in Equation 4.9 is weakly decreasing in s_i. We conclude that not only the signal distribution G is MHR, but also, for every bidder i, so is the value distribution $G_{|s_{-i}}$ for every signal profile s_{-i}.

4.6.2 Single Sample Mechanism for Interdependent Values

We describe our prior-independent mechanism for interdependent values in Algorithm 8. It is a natural generalization of the *single sample* mechanism of Dhangwatnotai et al. [2015]. Observe that the mechanism makes no reference to the distribution G.[12]

We are now ready to state this section's main result – that the above prior-independent mechanism is near-optimal. We compare its expected revenue to OPT, the optimal expected revenue achieved by the generalization of Myerson's mechanism to interdependent values (Algorithm 7). In fact, due to the MHR setting, the proof will be able to relate the expected revenue to the expected *welfare*, establishing a stronger property of *effectiveness* as defined by Neeman [2003].

Theorem 4.6.2 (Single sample is near-optimal)**.** *Let $n \geq 2$ and consider a matroid setting with non-private values in which the Lopomo assumptions and symmetry hold. The prior-independent single sample mechanism in Algorithm 8 yields a constant factor approximation to* OPT.

[12]We remark that the mechanism is assumed to know the valuation function v_i. An intriguing open problem is to design a mechanism, perhaps based on an ascending or multi-stage auction, for which this assumption can be dropped.

Algorithm 8: Single sample mechanism for interdependent values

1. Elicit signal reports $\vec{s}$ from the bidders.
2. Choose a *reserve bidder* uniformly at random, denote his signal by s_r.
3. Place the feasible set of non-reserve bidders with highest signals in a "potential winners" set P. Break ties arbitrarily but consistently.
4. Allocate to every bidder $i \in P$ such that $s_i \geq s_r$.
5. Charge every winner i a payment $v_i(\max\{s_r, t_i\}, s_{-i})$, where t_i is the threshold signal below which, given the signals of the other non-reserve bidders, i would not belong to P.

4.6.3 Useful Properties of MHR Distributions

We motivate our focus on MHR settings by the following example, which shows that unlike independent private values for which regularity suffices, for interdependent values a stronger MHR assumption is necessary to guarantee near-optimality of the single sample mechanism. Note that unlike previous sections, the MHR assumption is for the approximation guarantee (weaker assumptions are sufficient for incentives only). More generally, the example demonstrates how interdependence can pose new technical challenges, arising from the information externalities among bidders. The original analysis of Dhangwatnotai et al. [2015] no longer applies; we present properties of MHR distributions that will be useful in the new analysis.

Example 4.6.3 (Non-MHR setting)**.** Consider a digital goods setting with two bidders, whose i.i.d. signals are drawn from the equal revenue distribution $G(s) = 1 - (1/s)$, truncated to a finite range $[1, H]$, where H is an arbitrarily large constant. The bidders have weighted-sum values with $\beta = 1$, i.e., their pure common value is $s_1 + s_2$ (see Example 4.2.1). The optimal expected revenue in this setting is at least $\mathbb{E}[s_1] + \mathbb{E}[s_2] \approx 2 \ln H$, by charging each bidder the signal of the other. However, the expected revenue of the single sample mechanism is $2\mathbb{E}[\min\{s_1, s_2\}]$, which is $\ll H$.

The gap in Example 4.6.3 between the expectation of the distribution and the expectation of the lower among two random samples is due to the long tail of the non-MHR equal revenue distribution; we now show that for MHR distributions this issue does not arise. The next lemma is tight for the exponential distribution.

Lemma 4.6.4 (Lower among two samples from MHR distribution)**.** *Let s, s' be two i.i.d. samples drawn from an MHR distribution, then*

$$\mathbb{E}[s \mid s \leq s'] \geq \frac{1}{2}\mathbb{E}[s].$$

Proof. Let G with density g be the MHR distribution. Let $h(\cdot)$ be the hazard rate function of G and let $H(\cdot)$ be its cumulative hazard rate, i.e., $h(s) = g(s)/(1 - G(s))$, and $H(s) = \int_0^s h(z)dz$. By definition of the hazard rate function, $1 - G(s) = \exp(-H(s))$. Since G is MHR, $h(\cdot)$ is non-negative and weakly increasing, therefore $H(\cdot)$ is weakly increasing and convex. We can now write

$$\mathbb{E}[s \mid s \leq s'] = \int_0^\infty e^{-2H(s)} \geq \int_0^\infty e^{-H(2s)} \geq \frac{1}{2}\int_0^\infty e^{-H(s)} = \frac{1}{2}E[s],$$

where the first equality is by plugging in the distribution of the lower among two samples into $\mathbb{E}[x] = \int_0^\infty 1 - F(z)dz$, the first inequality is by convexity of H, and the second inequality is via integration by substitution. □

We now state a version of the previous lemma with an added threshold (and a necessarily slightly increased constant).

Lemma 4.6.5 (Lower among two samples with threshold)**.** *Let s, s' be two i.i.d. samples drawn from an MHR distribution. For every threshold $t \geq 0$,*

$$\mathbb{E}[\max\{s,t\} \mid \max\{s,t\} \leq s'] \geq \frac{1}{3}\mathbb{E}[s \mid s \geq t].$$

Proof. Since s is drawn from an MHR distribution, we can write

$$\frac{1}{3}(t + \mathbb{E}[s]) \geq \frac{1}{3}\mathbb{E}[s \mid s \geq t].$$

Assume first that $t \leq \mathbb{E}[s]/2$, then

$$\mathbb{E}[\max\{s,t\} \mid \max\{s,t\} \leq s'] \geq \mathbb{E}[s \mid s \leq s'] \geq \mathbb{E}[s]/2 \geq \frac{1}{3}(t + \mathbb{E}[s]),$$

where the second inequality is by Lemma 4.6.4, and the last one is by assumption. We now turn to the case in which $t > \mathbb{E}[s]/2$, and complete the proof by observing that $\mathbb{E}[\max\{s,t\} \mid \max\{s,t\} \leq s'] \geq t \geq \frac{1}{3}(t + \mathbb{E}[s])$. □

The following simple lemmas are stated without proofs.

Lemma 4.6.6 (Conditional MHR distribution)**.** *Let signal s be randomly drawn from an MHR distribution. For any threshold t, the distribution of s conditional on $s \geq t$ is also MHR.*

Lemma 4.6.7 (Optimal revenue approximates welfare – Hartline et al. [2008], Lemma 4.1)**.** *In a single bidder setting where the bidder's value is drawn from an MHR distribution, the optimal expected revenue is at least a $1/e$ fraction of the expected value.*

Lemma 4.6.8 (Median as reserve for regular distributions [Azar et al., 2013])**.** *In a single bidder setting where the bidder's value is drawn from a regular distribution G, a reserve price equal to the median of G guarantees at least half the optimal expected revenue, and the optimal expected revenue is thus at most the median.*

4.6.4 Proof for Single Item with Common Value

In this section we prove Theorem 4.6.2 for a simple setting with a single item for sale, which is valued the same by all bidders. We first state and prove our main lemma.

Main Lemma. Let s_1, s_2 be i.i.d. signals drawn from an MHR distribution G. Consider a single bidder with value $v(s_1, s_2)$, where v is a symmetric valuation function increasing in its arguments. Fixing signal s_1 (resp., s_2), let the value distribution $G_{|s_1}$ (resp., $G_{|s_2}$) be an MHR distribution. Let $c = 1/8e$ (where e is the base of the natural logarithm).[13]

[13]We do not optimize the constant c.

Lemma 4.6.9 (Reusing sample approximates welfare)**.**

$$\mathbb{E}_{s_1,s_2}[v(s_2, s_2) \mid s_2 \leq s_1] \geq c\mathbb{E}_{s_1,s_2}[v(s_1, s_2)]. \tag{4.10}$$

In words, plugging in the lower among s_1, s_2 *into the valuation function decreases the expected value by a factor of no more than* c.

Proof. Let m be the median of distribution G. We begin with the left-hand side of Equation 4.10 and condition on the event that $s_2 \geq m$, which given that $s_2 \leq s_1$ occurs with probability 1/4. Using that v is non-decreasing,

$$\begin{aligned}
\mathbb{E}_{s_1,s_2}[v(s_2, s_2)|s_2 \leq s_1] &\geq \Pr_{s_1,s_2}[s_2 \geq m \mid s_2 \leq s_1]\mathbb{E}_{s_1,s_2}[v(s_2, s_2) \mid m \leq s_2 \leq s_1] \\
&\geq \frac{1}{4}\mathbb{E}_{s_1,s_2}[v(m, s_2) \mid s_2 \leq s_1].
\end{aligned}$$

Now we know that replacing a random MHR sample with the lower among two samples results in a loss of at most 1/2 (Lemma 4.6.4). This can be applied to the distribution $G_{|s_1=m}$ of $v(m, \cdot)$, which is MHR by assumption. We get that

$$\mathbb{E}_{s_1,s_2}[v(m, s_2) \mid s_2 \leq s_1] \geq \frac{1}{2}\mathbb{E}_{s_2}[v(m, s_2)].$$

Now fix s_2. Observe that by the signal independence assumption, $v(m, s_2)$ is the median of the MHR distribution $G_{|s_2}$ of $v(\cdot, s_2)$. Thus

$$v(m, s_2) \geq \frac{1}{e}\mathbb{E}_{s_1}[v(s_1, s_2)],$$

where the inequality follows by combining Lemma 4.6.8, by which the median upper bounds the optimal expected revenue, with Lemma 4.6.7, by which the optimal expected revenue and expected welfare are close. Taking expectation over s_2 completes the proof. □

We are now ready to prove the special case of our main theorem – near-optimality of the single sample mechanism for the simple single-item, common value setting.

Proof of Theorem 4.6.2 for single item with common value. Let $\tilde{v}(\vec{s})$ be the pure common value of the item for the bidders, whose i.i.d. signals $s_1, \dots, s_n$ are drawn from an MHR distribution $\tilde{G}$. In this simple setting, the single sample mechanism in Algorithm 8 reduces to the following mechanism: choose a random reserve bidder; with probability $(n-1)/n$ the bidder with highest signal is not chosen as reserve; he then wins the item and is charged *according to the second highest signal* (whether or not the bidder with second highest signal is chosen as reserve). We claim that in expectation, the revenue achieved by this VCG-like mechanism is a $c(n-1)/n$-fraction of the expected welfare $\mathbb{E}_{\vec{s}}[\tilde{v}(s_1, \dots, s_n)]$, where c is as above.

The proof is by reduction to the single bidder setting of Lemma 4.6.9 (where the single bidder will correspond to the highest bidder). From now on we condition on the highest bidder not being chosen as reserve, incurring a loss of $(n-1)/n$.

Fix the $n-2$ lowest signals, denoted without loss of generality $s_3 \geq \dots \geq s_n$. Let $v(\cdot, \cdot) = \tilde{v}(\cdot, \cdot, s_3, \dots, s_n)$ be the valuation function given the fixed signals. Let G be the distribution $\tilde{G}$ conditioned on exceeding the threshold s_3. By Lemma 4.6.6, G is MHR, and the two highest signals s_1, s_2 can be seen as i.i.d. random draws from G. One of these is the second highest signal, and so we can write the expected revenue of the single sample mechanism as

$$\mathbb{E}_{s_1, s_2 \sim G}[v(s_2, s_2) \mid s_2 \leq s_1]. \tag{4.11}$$

In order to apply the main lemma (Lemma 4.6.9) to lower bound the expected revenue in (4.11), recall that $\tilde{v}$ and hence v are symmetric and increasing. It is left to show that $G_{|s_1}$ and $G_{|s_2}$ are MHR. Without loss of generality consider $G_{|s_2}$. We know from Remark 4.6.1 that given $s_2, \dots, s_n$, the distribution of $\tilde{v}_{|s_{-1}}$ is MHR. If we condition this distribution on $\tilde{v}_{|s_{-1}}$ being at least as high as $\tilde{v}_{|s_{-1}}(s_3)$, we still get an MHR distribution by Lemma 4.6.6. The resulting distribution is precisely the distribution $G_{|s_2}$.

The proof can now be completed by applying Lemma 4.6.9 to get that the expected revenue in (4.11) is at least $c\mathbb{E}_{s_1, s_2 \sim G}[v(s_1, s_2)]$ for any fixed profile $s_3, \dots, s_n$, and finally by taking expectation over $s_3, \dots, s_n$ according to the joint distribution of the

$n-2$ lowest among n draws from $\tilde{G}$. ☐

4.6.5 Proof for General Setting

In this section we prove Theorem 4.6.2 for a general matroid setting in which bidders have symmetric but distinct values. The proof relies on an extension of the main lemma (Lemma 4.6.9) above. Recall that the revenue originates both from the requirement to surpass the reserve signal and from having to surpass a threshold to be included in the potential winner set P. In the previous section, these could be summarized as a single requirement to surpass the second highest signal. In the general setting, the analysis of the expected revenue needs to take both requirements into account, necessitating an extension of the main lemma. Another difference in the general setting is that we need to take into account the contribution to revenue from multiple bidders. For this we use the second Lopomo assumption as well as the matroid assumption in order to relate the revenue contributers to the welfare contributers. This was unnecessary before due to the common value assumption, and is also unnecessary in the digital goods setting when all bidders but the reserve belong to the potential winner set P.

As above, let s_1, s_2 be i.i.d. signals drawn from an MHR distribution G. Consider a single bidder with a symmetric and increasing valuation $v(s_1, s_2)$, and assume the value distributions $G_{|s_1}$ and $G_{|s_2}$ when one of the signals is fixed are MHR. Let $c' = 1/12e$.[14]

Lemma 4.6.10 (Reusing sample with threshold)**.** *For every threshold $t \geq 0$,*

$$\mathbb{E}_{s_1,s_2}[v(\max\{s_2,t\}, s_2) \mid \max\{s_2,t\} \leq s_1] \geq c'\mathbb{E}_{s_1,s_2}[v(s_1,s_2) \mid t \leq s_1].$$

Proof. Let $c' = 1/12e$, and let m be the median of distribution G from which s_1, s_2 are independently drawn. Similarly to the proof of Lemma 4.6.9, we have that

$$\mathbb{E}_{s_1,s_2}[v(\max\{s_2,t\}, s_2) \mid \max\{s_2,t\} \leq s_1] \geq \frac{1}{4}\mathbb{E}_{s_1,s_2}[v(\max\{s_2,t\}, m) \mid \max\{s_2,t\} \leq s_1].$$

[14] As before, the constant is not optimized.

By Lemma 4.6.5 applied to the distribution $G_{|s_2=m}$ of $v(\cdot, m)$, which is MHR by assumption, and using that v is weakly increasing, we get

$$\begin{aligned}\frac{1}{4}\mathbb{E}_{s_1,s_2}[v(\max\{s_2,t\},m) \mid \max\{s_2,t\} \le s_1] &= \\ \frac{1}{4}\mathbb{E}_{s_1,s_2}[\max\{v(s_2,m),v(t,m)\} \mid \max\{v(s_2,m),v(t,m)\} \le v(s_1,m)] &\ge \\ \frac{1}{12}\mathbb{E}_{s_1}[v(s_1,m) \mid s_1 \ge t].\end{aligned}$$

Similarly to the proof of Lemma 4.6.9, we now fix $s_1 \ge t$. Observe that $v(s_1, m)$ is the median of the MHR distribution $G_{|s_1}$ of $v(s_1, \cdot)$. Thus

$$v(s_1,m) \ge \frac{1}{e}\mathbb{E}_{s_2}[v(s_1,s_2)],$$

and taking expectation over s_1 conditional on $s_1 \ge t$ completes the proof. □

The proof of Theorem 4.6.2 for the general setting is by reduction to the single bidder setting and application of Lemma 4.6.10.

Proof of Theorem 4.6.2 for general setting. Recall we wish to analyze the expected revenue of the single sample mechanism in Algorithm 8, for a matroid setting with $n \ge 2$ bidders and symmetry, where the i.i.d. signals $s_1, \ldots, s_n$ are drawn from an MHR distribution G. Denote by $\tilde{v}$ the symmetric valuation function of the bidders. Similarly to the proof in Section 4.6.4, we will reduce this setting to a single bidder setting to which Lemma 4.6.10 is applicable.

Without loss of generality we name the chosen reserve bidder "bidder 2", and consider an arbitrary non-reserve bidder "bidder 1". We condition on the signals of all bidders other than 1 and 2 and omit them from the notation, i.e., we use the notation $v(s_1, s_2)$ for bidder 1's value. In addition, we denote by t the threshold for bidder 1 to belong in the potential winner set P given the fixed signals. We can now write the expected contribution of bidder 1 to the expected revenue of the single sample mechanism as

$$\mathbb{E}_{s_1,s_2}[v(\max\{s_2,t\},s_2) \mid \max\{s_2,t\} \le s_1]\Pr[\max\{s_2,t\} \le s_1]. \tag{4.12}$$

In what follows we show that the expected contribution in (4.12) is a constant fraction of the expected contribution of bidder 1 to the expected maximum welfare excluding bidder 2. To see how this completes the proof, take expectation over the fixed signals and sum up over non-reserve bidders. The total expected revenue of the single sample mechanism is thus a constant fraction of the welfare excluding the reserve bidder, which is in turn a $(n-1)/n$ fraction of the total welfare.

We begin by writing down the expected contribution of bidder 1 to the expected maximum welfare excluding bidder 2. Crucially, the same threshold t as in the single sample mechanism is the threshold for bidder 1 to be included in the welfare-maximizing set of bidders. This is because the potential winner set P of the single sample mechanism is the feasible set of non-reserve bidders with highest signals, and by single-crossing of values (second Lopomo assumption) and the matroid setting, this is also the welfare-maximizing feasible set. The expected welfare contribution is thus

$$\mathbb{E}_{s_1,s_2}[v(s_1,s_2) \mid s_1 \geq t] \Pr[s_1 \geq t]. \tag{4.13}$$

It remains to compare (4.12) to (4.13). Since we know that v is symmetric and increasing and that $v(s_1,\cdot), v(\cdot,s_2)$ are distributed according to an MHR distribution (Remark 4.6.1), we can apply Lemma 4.6.10 to get

$$\begin{aligned}
\mathbb{E}_{s_1,s_2}[v(\max\{s_2,t\},s_2) \mid \max\{s_2,t\} \leq s_1] \Pr[\max\{s_2,t\} \leq s_1] &\geq \\
c'\mathbb{E}_{s_1,s_2}[v(s_1,s_2) \mid s_1 \geq t] \Pr[s_1 \geq t] \Pr[s_1 \geq s_2 \mid s_1 \geq t] &\geq \\
c'\mathbb{E}_{s_1,s_2}[v(s_1,s_2) \mid s_1 \geq t] \Pr[s_1 \geq t] \Pr[s_1 \geq s_2] &= \\
(c'/2)\mathbb{E}_{s_1,s_2}[v(s_1,s_2) \mid s_1 \geq t] \Pr[s_1 \geq t]. &
\end{aligned}$$

□

Part II

Computational Challenges in Welfare Maximization

5

Double Auctions with Complements

5.1 Chapter Introduction

5.1.1 Overview

Designing double auctions is a complex problem, especially when there are feasibility constraints on the sets of buyers and sellers that can trade with one another. In this chapter we develop a modular approach to the design of double auctions, by relating it to the exhaustively-studied problem of designing one-sided mechanisms with a single seller. We consider several desirable properties of a double auction: in addition to feasibility, dominant-strategy incentive compatibility, (approximate) welfare maximization and budget balance, there are the still stronger incentive properties offered by a deferred-acceptance implementation, including group-strategyproofness. For each of these properties, we identify sufficient conditions on two one-sided algorithms – one for ranking the buyers, one for ranking the sellers – and on a method for their composition into trading pairs, which guarantee the desired property of the double auction. These sufficient conditions are often closely related to the performance of the greedy algorithm and the degree of complementarity induced by the feasibility constraints. Our framework also offers new insights into classic double auction designs, such as the VCG and McAfee auctions with unit-demand buyers and unit-supply sellers.

5.1.2 A Modular Approach

Double auctions play an important role in mechanism design theory and practice. They are of theoretical importance because they solve the fundamental problem of how to organize trade between a set of buyers and a set of sellers, when both the buyers and the sellers act strategically. Important practical applications include the New York Stock Exchange (NYSE), where buyers and sellers trade shares, and the upcoming spectrum auction conducted by the US Federal Communication Commission (FCC), which aims at reallocating spectrum licenses from TV broadcasters to mobile communication providers [Milgrom and Segal, 2014].

Designing double auctions can be a complex task, with several competing objectives. These include, but are not limited to: feasibility, dominant-strategy incentive compatibility (DSIC), the still stronger incentive constraints offered by a deferred-acceptance implementation such as weak group-strategyproofness (WGSP) or implementability as a clock auction [Mehta et al., 2009], exact or approximate welfare maximization, and budget balance (BB) (see Section 5.2 for definitions.)

In this work we utilize the fact that the problem of designing *single*-sided mechanisms is well-studied, and develop the following modular approach to designing *double*-sided auctions for complex settings. We split the design task into three algorithmic modules: one for the buyers, one for the sellers and one for their combination. Designing the algorithm for the buyers or the sellers is based on the well-developed theory of designing single-sided mechanisms, in which problems such as feasibility checking have been exhaustively studied – see examples below. The third module incorporates a composition rule that combines buyers and sellers to determine the final allocation. Payments are determined by computing *thresholds* from the three algorithms.[1] This approach can also be described as a "black-box reduction" from designing double

[1] *Threshold payments* are defined in Section 5.2; informally they are based on the threshold bids of the players, which differentiate between acceptance and rejection by the mechanism.

auctions to the problem of designing single-sided mechanisms. The goal of this chapter is to develop the theory that explains when and how such a modular/black-box approach works.

5.1.3 Motivating Examples

Suppose that each buyer i wants to acquire one unit of a good and has a value v_i for it; and each seller j has one unit of this good for sale at a cost of c_j.

The Unconstrained Problem. Assume first there are no restrictions on which buyers and sellers can trade with one another (we refer to this below as the *unconstrained* problem). Is it possible, by composing two single-sided algorithms, to implement the VCG mechanism that maximizes welfare and is DSIC? What about McAfee's [1992] trade reduction mechanism, which accepts all buyer-seller pairs from the welfare-maximizing solution except for the least valuable one, and is DSIC and BB?[2]

The answer to these questions is "yes". We can implement the VCG mechanism using simple *greedy* algorithms that sort the buyers by non-increasing value and the sellers by non-decreasing cost. We iteratively query these algorithms for the next buyer-seller pair and accept it if the buyer has a larger value v_i than the seller's cost c_j, and we apply threshold payments. For McAfee's trade reduction mechanism we use *reverse* versions of these algorithms, that sort the players by non-decreasing value and non-increasing cost. We iteratively query these algorithms for the next buyer-seller pair and reject it if none of the previously inspected buyer-seller pairs had non-negative gain from trade $v_i - c_j \geq 0$, and we again apply threshold payments. See Figure 5.2 for an illustration.

[2]For simplicity we focus on a version of McAfee's trade reduction mechanism in which the least valuable pair is always rejected. *Cf.* the full version which is defined as follows: Sort buyers by non-increasing value $v_1 \geq v_2 \geq \dots$ and sellers by non-decreasing cost $c_1 \leq c_2 \leq \dots$. Let k be the largest index such that $v_k \geq c_k$. Compute $t = (v_{k+1} + c_{k+1})/2$. If $t \in [c_k, v_k]$ let buyers/sellers $1, \dots, k$ trade with each other. Otherwise exclude the buyer-seller pair with the k-th highest value and the k-th lowest cost from trade.

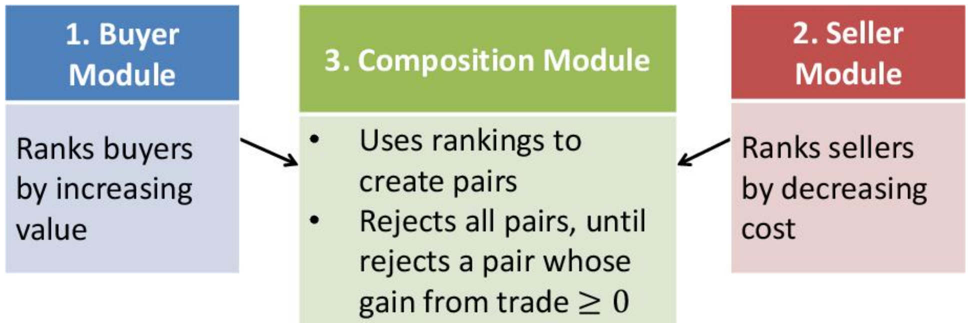

Figure 5.1: A new modular view of McAfee's trade reduction mechanism (assuming for simplicity that the number of buyers equals the number of sellers).

Constraints Maintaining Substitutability. Now, *what happens if we add feasibility constraints on which buyers and sellers can trade*? Such feasibility constraints are important in practical applications, and their potential richness mirrors the richness of real-life economic settings; see, for example, the recent work on the proposed FCC double auctions for spectrum [Milgrom and Segal, 2014; Nguyen and Sandholm, 2014; Leyton-Brown, 2013, 2014].

As a first example, consider the variation of the above problem in which the buyers belong to one of three categories (e.g., they are firms that are either small, medium, or large in market share). To ensure diversity among buyers, the policy maker requires that no more than k_i buyers from each category i shall be accepted. Note that this feasibility constraint maintains substitutability among the buyers. (For additional such *quota* examples see, e.g., [Hatfield and Milgrom, 2005].)

In this example it is still possible to implement the VCG and trade reduction mechanisms by composing two one-sided algorithms. The only change is to the algorithm used for the buyers. In its forward version we would go through the buyers in order of their value and accept the next buyer if and only if we haven't already accepted k_c buyers from that buyer's category c. In its backward version we would go through the buyers in reverse order and reject the next buyer unless there are k_c or fewer buyers from that category left.

Constraints Inducing Complementarity. As a second example, consider the variant of the problem in which sellers have one of two "sizes", s or S, where $S > s$.

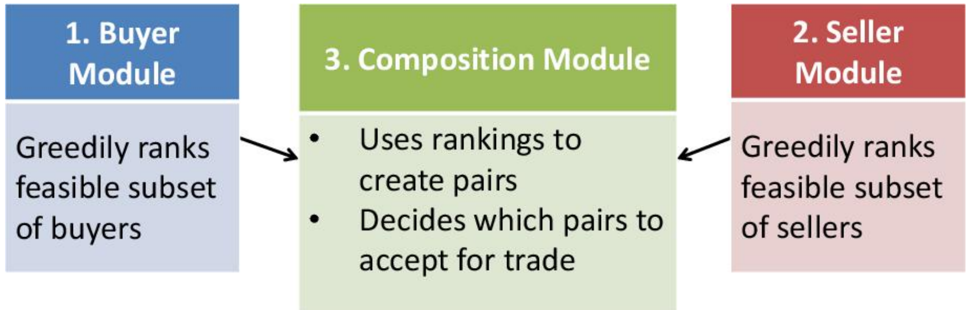

Figure 5.2: A general modular approach that supports feasibility constraints.

For instance, sellers could be firms that pollute the environment to different extents. Suppose there is a cap C on the combined size of the sellers that can be accepted. This feasibility constraint no longer maintains substitutability; rather, it introduces complementarities among the sellers. (For additional such *packing* examples see, e.g., [Aggarwal and Hartline, 2006].)

In this example it is less clear what to do. Even putting aside our goal of modular design, computing the welfare-maximizing solution is an NP-hard packing problem [e.g., Kleinberg and Tardos, 2005], and so specifically the one-sided greedy by cost algorithm is no longer optimal. We thus shift our attention to approximately-maximizing solutions, but it is not clear which one-sided approximation methods – greedy according to cost, greedy according to cost divided by size, non-greedy algorithms, etc. – would offer good approximation guarantees in the double auction context, where the choices of buyers and sellers are entangled. Furthermore, it is not clear if the good properties of the double-sided VCG and McAfee mechanisms, such as DSIC or BB, would continue to hold.

5.1.4 Our Results for Substitutes and Complements

Composition Theorems. We advocate a modular approach to the design of double auctions, which is applicable to complex feasibility constraints on both sides of the market. All of the resulting double auction mechanisms are deterministic.

The modular approach breaks the design task into two subtasks: (a) the design of two one-sided algorithms and (b) the design of a composition rule that pairs buyers and sellers. To identify what we want from the respective subtasks, we prove a number of *compositions theorems* of the following general form:

> If the one-sided algorithms $\mathcal{A}_1$ and $\mathcal{A}_2$ have properties X_1 and X_2 and the composition rule has property Y, then the resulting double auction has property Z.

A main theme of this work is thus to identify sufficient conditions on the two one-sided algorithms and on the method of composition that guarantee a desired property of the double auction.

Strong Incentive Properties. We start with sufficient conditions that ensure that the double auction is DSIC, resp., has the stronger incentive properties shared by *deferred-acceptance* implementations generalizing the Gale-Shapley mechanism (see [Milgrom and Segal, 2014; Gale and Shapley, 1962], and Section 5.4).

Interestingly, *monotonicity* of all involved components (where the standard definition appears in Section 5.2) is not sufficient for DSIC; we also need that the one-sided algorithms return the players in order of their "quality", i.e., their contribution to social welfare.[3] For this reason, *the greedy by quality approach plays an essential role in our designs.* In the above examples, the greedy by value and greedy by cost algorithms follow this approach, while a greedy algorithm based on cost divided by size violates it.

An important consequence of the sufficient conditions we obtain is that trade reduction mechanisms can be implemented within the deferred-acceptance framework, and therefore share the stronger incentive properties of mechanisms within this class. In particular, this approach shows for the first time that McAfee's trade reduction mechanism is WGSP.

BB. We show that *the same conditions* on the one-sided algorithms and the composition, which enable implementation within the deferred-acceptance framework for

[3]A buyer contributes more to social welfare as his value increases; a seller contributes more to social welfare as his cost decreases. In other words, a buyer's quality increases as his ability to extract value from the good increases, and a seller's quality increases as his ability to produce the good at lower cost increases.

properties like WGSP, also lead to BB.

Approximate Welfare. We identify conditions that ensure the double auction obtains a certain fraction of the optimal welfare. These conditions ask that the one-sided algorithms achieve an approximation ratio *uniformly*, that is "at all times" (see Definition 5.5.1). The intuition behind these conditions is that the final number of accepted players is extrinsic, since it depends on the interplay with the other side of the market, and so the algorithm should be close to optimal for any possible number of accepted players rather than just for the final number of accepted players.

We analyze the identified conditions for a number of algorithms, including the greedy by value and greedy by cost algorithms used in the examples above. These algorithms achieve optimal welfare if and only if the players are substitutes, and moreover they achieve this uniformly, and their reversed versions also achieve optimality uniformly. The main question is, then, what happens when feasibility constraints introduce complementarities? I.e., how well can welfare be approximated uniformly in the presence of complementarities, and do the approximation algorithms satisfy the additional conditions needed for strong incentive properties and BB?

5.1.5 Three Applications

To demonstrate the usefulness of our modular approach, we utilize it to design novel double auctions for problems with non-trivial feasibility structure. These serve to illustrate our design framework and are not an exhaustive list of the possible applications. We focus on three types of feasibility constraints, as follows (see Section 5.2.3 for a more formal description):

1. Matroids: The set of feasible subsets of players is downward-closed and satisfies an *exchange* axiom (if two sets S, T are feasible and T is larger than S, then there must be an element $u \in T \setminus S$ such that $S \cup \{u\}$ is feasible). The unconstrained problem discussed above as well as the problem with diversity quota constraints are special cases of this category.

Figure 5.3: Illustration of a matching feasibility constraint. The players are broadcast stations seeking to buy radio spectrum, depicted as edges of a line graph along the U.S. coastline. A player set is infeasible if it contains players that are adjacent in the graph because they interfere with each other's broadcasts. A feasible set thus corresponds to a matching.

2. Knapsacks: Each player has a size and a set of players is feasible if their combined size does not exceed a given threshold. The variation of the unconstrained problem in which sellers have one of two distinct sizes is a special case of this constraint.

3. Matchings: We are given a graph such that each player corresponds to an edge in this graph, and a set of players is feasible if it corresponds to a matching in this graph. A concrete example of this constraint is a setting where the sellers on the market correspond to certain pairs of firms, who can cooperate to produce complementary goods, both of which are required to provide the service being sold on the market. The sellers can thus be thought of as edges of a bipartite graph, where on one side there are firms producing the first complementary good and on the other there are firms producing the other good [*cf.*, Pycia, 2012]. For another example see Figure 5.3.

The first setting is precisely the setting in which greedy by quality is optimal, i.e., substitutability holds: the welfare (the total value of all buyers minus the total cost of all sellers) is a submodular function of the set of buyers or sellers (see Section

1.4.2). The second and third settings can be thought of as different relaxations of the matroid or substitutability constraint (or different limitations of complementarity), in which greedy by quality is not optimal but rather approximately optimal.

Our framework yields novel VCG-style and trade reduction-style mechanisms for all three settings. The former are DSIC, whereas the latter are WGSP, implementable as a clock auction and BB. It also translates approximation guarantees for greedy algorithms into welfare guarantees for these double auctions. These guarantees show that *the welfare degrades gracefully as we move away from substitutability.*

5.1.6 Related Work

The design principle of modularity is embraced in a diverse range of complex design tasks, from mechanical systems through software design to architecture [e.g., Baldwin and Clark, 2000]. Splitting a complex design task into parts or modules, addressing each separately and then combining the modules into a complete system helps make the design and analysis tractable and robust. Economic mechanisms that operate in complex incentive landscapes while balancing multiple objectives are natural candidates for reaping the benefits of modularity. Two predecessors of our work that apply a modular approach to a mechanism design problem are [Mu'alem and Nisan, 2008; Deshmukh et al., 2002], but they consider different settings than ours (one-sided rather than two-sided), or different objectives (profit in the "competitive analysis" framework rather than welfare, strategyproofness and budget balance).

Most prior work on double auctions is motivated by the impossibility results of Hurwicz [1972] and Myerson and Satterthwaite [1983], which state that optimal welfare and BB cannot be achieved simultaneously subject to DSIC or even Bayes-Nash incentive compatibility (BIC). One line of work escapes this impossibility by relaxing the efficiency requirement. This direction can be divided into mechanisms that are BIC and mechanisms that are DSIC. An important example of the former is the buyer's bid double auction [Satterthwaite and Williams, 1989; Rustichini et al., 1994; Satterthwaite and Williams, 2002], which sets a single price to equate supply and demand. More recent work that falls into this category is [Cripps and Swinkels, 2006; Fudenberg et al., 2007]. A prominent example of the latter is McAfee's trade reduction mechanism, which allows all but the least efficient trade. This mechanism has been generalized to more complex settings in [Babaioff and Nisan, 2004; Bredin and Parkes, 2005; Gonen et al., 2007; Babaioff et al., 2009; Chu, 2009]. More recent work that falls into this category is [Kojima and Yamashita, 2013; Blumrosen and Dobzinski, 2014] (where [Kojima and Yamashita, 2013] actually applies *ex post* incentive compatibility, as appropriate for interdependent values). A second line of work that seeks to escape the impossibility results was recently initiated by Carroll [2013], by analyzing the trade-off between incentives and efficiency while insisting on budget balance. Our work is different in that it adds to the double auction design problem the objectives of feasibility and WGSP, and takes an explicitly modular approach to achieve the objectives.

The WGSP property that we highlight was studied in detail in [Juarez, 2013], although a complete characterization of WGSP mechanisms is not known. Deferred-acceptance algorithms on which part of our work is based are proposed in [Milgrom and Segal, 2014], and their performance is analyzed in [Dütting et al., 2014a]. Our work extends the deferred-acceptance framework from one-sided settings to two-sided settings.

The greedy approach has been extensively studied in the context of one-sided mechanism design, for both single- and multi-parameter settings; see, e.g., [Borodin and Lucier, 2010] and references within.

5.1.7 Chapter Organization

Preliminaries appear in Section 5.2, which defines the properties of double auctions that we are interested in; this section can be skipped by the expert reader. Section 5.3 describes our composition framework – one-sided algorithms and different methods of composing them. The next three sections are organized by the desired property of double auctions: Section 5.4 proves the DSIC and WGSP composition theorems; Section 5.5 the welfare composition theorem; and Section 5.6 the BB composition theorem. In Dütting et al. [2014b] we establish a lower bound on the welfare achievable while satisfying the double auction properties of interest.

5.2 Problem Statement

This section defines the double auction settings and the properties of double auction mechanisms that we are interested in. We also define the three settings that will serve as running examples.

5.2.1 Double Auction Settings with Feasibility Constraints

We study single-parameter double auction settings. These are two-sided markets, with n buyers on one side of the market and m sellers on the other (m is not associated with the number of goods in this dissertation chapter). There is a single good for sale. The buyers each want to acquire a single unit of this item, and the sellers each have a single unit to sell. A set of buyers and sellers is *feasible* if the set of buyers is feasible and the set of sellers is feasible, and there are at least as many sellers as there are buyers. The feasible buyer sets are expressed as a set system $(N, \mathcal{I}_N)$, where N is the ground set of all n buyers, and $\mathcal{I}_N \subseteq 2^N$ is a non-empty collection of all the feasible buyer subsets. Similarly, the feasible seller sets are given as a set system $(M, \mathcal{I}_M)$, where M is the ground set of all m sellers, and $\mathcal{I}_M \subseteq 2^M$ is a non-empty collection of all the feasible seller subsets. The set systems are downward-closed, succinctly represented and publicly-known.

Each buyer i has a value $v_i \in [v_i^0, v_i^1]$ where $v_i^1 < \bar{v}$, and each seller j has a cost $c_j \in [c_j^0, c_j^1]$ where $c_j^1 < \bar{c}$, and $\bar{v} = \bar{c}$ are the maximum possible value and cost. For simplicity and without loss of generality, unless stated otherwise we assume that values and costs are non-negative and unique (no two buyers, resp. sellers, have the same value, resp. cost). The *type spaces* $[v_i^0, v_i^1], [c_j^0, c_j^1]$ are publicly known, and the *bid spaces* are equal to the type spaces. A player's *quality* is his value if he is a buyer, and *minus* his cost if he is a seller. We denote by $\vec{v}$ (resp. $\vec{c}$) the value (resp. cost) profile of all buyers (resp. sellers). The players' utilities are quasi-linear, i.e., buyer i's utility from acquiring a unit at price p_i is $v_i - p_i$, and seller j's utility from selling his unit for payment p_j is $p_j - c_j$. The *optimal welfare* is equal to the *gains from trade*, i.e., the maximum difference between the total value and total cost over feasible subsets of players:

$$\mathrm{OPT}(\vec{v}, \vec{c}) = \max_{B \in \mathcal{I}_N, S \in \mathcal{I}_M, |B| \leq |S|} \left\{ \sum_{i \in B} v_i - \sum_{j \in S} c_j \right\}.$$

Note that since we consider downward-closed set systems, the optimal solution will always be attained by buyer set B and seller set S such that $|B| = |S|$.

5.2.2 Double Auction Mechanisms

We study direct and deterministic double auction mechanisms, which consist of an *allocation rule* $x(\cdot, \cdot)$ and a *payment rule* $p(\cdot, \cdot)$. The allocation rule takes a pair of value and cost profiles $\vec{v}, \vec{c}$ as input, and outputs the set of players who are *accepted* for trade, also referred to as *allocated.* A buyer is considered allocated when he buys a unit, and a seller is allocated when he sells his unit. For every buyer i (resp. seller j), $x_i(\vec{v}, \vec{c})$ (resp., $x_j(\vec{v}, \vec{c})$) indicates whether he is allocated or not by the mechanism. The payment rule also takes a pair of value and cost profiles $\vec{v}, \vec{c}$ as input, and computes payments that the mechanism charges the buyers and pays to the sellers. We use $p_i(\vec{v}, \vec{c})$ to denote the payment buyer i is charged, and $p_j(\vec{v}, \vec{c})$ to denote the payment seller j is paid. A buyer who is not accepted is charged 0 and a seller who is not accepted is paid 0.

The *welfare of a mechanism* is the total value of buyers that it accepts minus the total cost of the sellers that it accepts. That is,

$$W(\vec{v}, \vec{c}) = \sum_{i \in N} x_i(\vec{v}, \vec{c}) \cdot v_i - \sum_{j \in M} x_j(\vec{v}, \vec{c}) \cdot c_j.$$

Non-Strategic Properties. We study the following non-strategic properties of double auction mechanisms:

- *Feasibility.* A double auction mechanism is feasible if for every value and cost profiles $\vec{v}, \vec{c}$, the set of accepted buyers and sellers is feasible. Formally, if B is the set of accepted buyers and S is the set of accepted sellers, then $B \in \mathcal{I}_N$, $S \in \mathcal{I}_M$ and $|B| \leq |S|$.

- *Budget balance (BB).* A double auction mechanism is budget balanced if for every value and cost profiles $\vec{v}, \vec{c}$, the difference between the sum of payments charged from the accepted buyers and the sum of payments paid to the accepted sellers is non-negative.

- *Efficiency.* A double auction mechanism is δ-approximately efficient if for every value and cost profiles $\vec{v}, \vec{c}$, its welfare $W(\vec{v}, \vec{c})$ is at least a $(1/\delta)$-fraction of the optimal welfare $\text{OPT}(\vec{v}, \vec{c})$. Clearly, for feasible mechanisms $\delta \geq 1$, and $\delta = 1$ precisely if the mechanism achieves optimal welfare.

Strategic Properties. We also study the following strategic properties of double auction mechanisms:

- *Individual rationality (IR).* A double auction mechanism is IR if for every value and cost profiles $\vec{v}, \vec{c}$, every accepted buyer i is not charged more than his value v_i, and every accepted seller j is paid at least his cost c_j. Non-accepted players are charged/paid zero.

- *Dominant-strategy incentive compatible (DSIC).* A double auction mechanism is DSIC if for every value and cost profiles $\vec{v}, \vec{c}$ and for every i, j, v_i', c_j', it holds that buyer i is (weakly) better off reporting his true value v_i than any other value v_i', and seller j is (weakly) better off reporting his true cost c_j than any other cost c_j'. Formally,

$$x_i(\vec{v}, \vec{c}) \cdot v_i - p_i(\vec{v}, \vec{c}) \geq x_i((v_i', v_{-i}), \vec{c}) \cdot v_i - p_i((v_i', v_{-i}), \vec{c}),$$

  and similarly for seller j.

- *Weak group-strategyproofness (WGSP).* A double auction mechanism is WGSP if for every value and cost profiles $\vec{v}, \vec{c}$, for every set of buyers and sellers $B \cup S$ and every alternative value and cost reports of these players v_B', c_S', there is at least one player in $B \cup S$ who is (weakly) better off when the players report truthfully as when they report v'_B, c'_S. Intuitively, such a player does not have a strict incentive to join the deviating group.[4] This condition is equivalent to DSIC for the special case of singleton sets.

The following characterization of DSIC and IR double auction mechanisms follows from standard arguments. A similarly simple characterization of WGSP and IR double auction mechanisms is not available.[5]

Definition 5.2.1. The allocation rule $x(\cdot,\cdot)$ is *monotone* if for all value and cost profiles $\vec{v}, \vec{c}$, every accepted buyer who raises his value while other values and costs remain fixed is still accepted, and every accepted seller who lowers his cost while other values and costs remain fixed is still accepted.

Definition 5.2.2. For a monotone allocation rule $x(\cdot,\cdot)$, the *threshold payment* of an allocated buyer i given profiles $\vec{v}_{-i}, \vec{c}$, and the threshold payment of an allocated seller j given profiles $\vec{v}, \vec{c}_{-j}$, are, respectively,

$$\sup_{\{v_i | x_i(\vec{v},\vec{c})=0\}} \{v_i\}, \inf_{\{c_j | x_j(\vec{v},\vec{c})=0\}} \{c_j\}.$$

The threshold payments of unallocated buyers and sellers are 0.

Intuitively, the threshold payment of a player is the highest value (resp. lowest cost) he can report without being accepted. We can now state the characterization result:

Proposition 5.2.3. *A double auction mechanism is DSIC and IR if and only if the allocation rule is monotone and the payment rule applies threshold payments.*

[4] A stronger notion of group strategyproofness requires that no group of buyers and sellers can jointly deviate to make some member of the group strictly better off while all other members are no worse off. This stronger notion is violated by all common double auction formats. For example, if a seller's cost sets the price for a buyer, then the seller can claim to have a lower cost to lower the buyer's payment without affecting his own utility.

[5] See [Juarez, 2013] for recent progress towards characterizing WGSP and BB mechanisms in the context of cost sharing mechanisms.

Note that threshold payments are sufficient to guarantee IR, and since all the mechanisms we consider apply threshold payments, we do not discuss individual rationality further. Also, since we focus on DSIC mechanisms, we use the terms true value (resp. cost), reported value (resp. cost) and bid interchangeably.

5.2.3 Running Examples

We now define formally the examples of feasibility constraints mentioned and motivated in Section 5.1.5. We denote by U the ground set of players (either the buyers or sellers in our context), and by $\mathcal{I}$ the collection of feasible subsets.

1. *Matroids*: The feasible sets are such that the set system $(U, \mathcal{I})$ forms a matroid (see Section 2.1.1).

2. *Knapsacks*: The elements of the ground set U have publicly-known *sizes* denoted by $s_1, \ldots, s_{|U|}$, and the family of feasible sets $\mathcal{I}$ includes every subset $S \subseteq U$ such that its total size $\sum_{i \in S} s_i$ is at most the capacity C of the knapsack. We denote the ratio between the size of the largest element and the size of the knapsack by $\lambda \leq 1$. Let $\mu \leq 1$ be a lower bound on the ratio between the size of the smallest element and the size of the largest element; it is assumed that $1/\mu$ is integral.

3. *(Bipartite) Matchings*: The ground set U is the edges of a bipartite graph, and the family of feasible sets $\mathcal{I}$ are the edge subsets that correspond to bipartite matchings.

5.3 Composition Framework

In this section we describe our framework for designing double auctions via composition. We first describe the one-sided algorithms and then the different ways of composing them.

5.3.1 Ranking Algorithms

The one-sided algorithms we use for our compositions are called *ranking algorithms*. A ranking algorithm for buyers (resp. sellers) is a deterministic algorithm that receives as input a value profile $\vec{v}$ (resp. cost profile $\vec{c}$), and returns as output an ordered set of buyers (resp. sellers), which we refer to as a *stream*. Not all buyers (resp. sellers) necessarily appear in the stream, e.g., due to feasibility considerations. The *rank* of a buyer (resp. seller), denoted by $r_i(\vec{v})$ (resp. $r_j(\vec{c})$), is his position in the stream (e.g., 1 if he appears first), or ∞ if he does not appear in the stream.[6] The closer a player's rank is to 1, the smaller or *lower* his rank. Accessing the next player in the stream is called *querying* the ranking algorithm. When querying the kth player, the query *history* is the identities and qualities (values or costs) of the $k-1$ previously-queried players. We say that a history h is a *prefix* of another history h' when the queries recorded in h are the first queries recorded in h'.

Forward and Backward Feasibility. We now distinguish between two natural *directions* of ranking algorithms by their different feasibility guarantees. The rest of this section is stated for buyers but applies to sellers as well.

Definition 5.3.1. A ranking algorithm is *forward-feasible* for a given feasibility set system $(N, \mathcal{I}_N)$ if for every input $\vec{v}$, it returns a stream of buyers $s_1, \ldots, s_n$ such that the following holds: for every $1 \le i \le n$, the set of buyers $\{s_1, \ldots, s_i\}$ is $\in \mathcal{I}_N$.

That is, a ranking algorithm is forward-feasible if for every input, every prefix of the output stream (including the entire stream) is a feasible set according to the feasibility set system.

Definition 5.3.2. A ranking algorithm is *backward-feasible* for a given set system $(N, \mathcal{I}_N)$ if for every input $\vec{v}$, it returns a stream of buyers $s_1, \ldots, s_n$ and a rank $0 \le \ell < n$ such that the following holds: for every $\ell \le i \le n$, the set of buyers $N \setminus \{s_1, \ldots, s_i\}$ is $\in \mathcal{I}_N$ (if $\ell = 0$ then $N \in \mathcal{I}_N$).

[6]This notion of a player's rank by a ranking algorithm is not to be confused with the notion of a matroid's rank; the latter is not referred to in this chapter.

That is, a ranking algorithm is backward-feasible if for every input, there exists a prefix of the output stream such that after discarding it or any larger prefix that contains it, the remaining buyer set is feasible according to the feasibility set system.

The semantic difference between forward-feasible and backward-feasible ranking algorithms is as follows: The former returns a stream of buyers who can be greedily accepted, resulting in a feasible set of accepted buyers; the latter returns a stream of buyers who can be greedily rejected, such that after sufficiently many rejections the remaining buyers can be accepted, again resulting in a feasible set of accepted buyers.

The following example demonstrates the above concepts.

Example 5.3.3 (Greedy Ranking by Value with Knapsack Constraints)**.** Suppose there are 4 buyers $1, 2, 3, 4$ with sizes $4, 3, 2, 1$ and values $3, 7, 1, 8$, and the feasibility constraint is that the total size of accepted buyers cannot exceed 4. A possible forward-feasible ranking algorithm could rank buyers by highest value first. So it would first return buyer 4 and then buyer 2 at which point the total size has reached 4. In other words, the stream of buyers would be $4, 2$ and the ranks would be $\infty, 2, \infty, 1$. A possible backward-feasible ranking algorithm could rank buyers by lowest value first. In this case the output stream would be $3, 1, 2, 4$ with buyer ranks $2, 3, 1, 4$. The rank ℓ of the largest-rank buyer that must be rejected for feasibility would be $\ell = 2$. Note that in both cases the set of feasible players is the same and achieves optimal welfare. It is not difficult to see that with different values the set of feasible players achieved by these algorithms may be suboptimal.

Running Examples. We discuss two algorithmic approaches to designing ranking algorithms for our running examples.

One approach to obtaining a ranking algorithm, say for the buyers, is to start by computing the set of buyers that maximizes total value among all feasible buyer sets. A forward-feasible ranking algorithm could then output these buyers in arbitrary order. Similarly, a backward-feasible ranking algorithm could first output all the buyers outside of this set (followed by all buyers in the set). This approach is computationally tractable for both matroids and bipartite matching; for knapsack it is an NP-hard problem [Garey and Johnson, 1979].

A different class of methods that naturally lead to computationally tractable ranking algorithms are greedy algorithms. *For matroids* there exist forward- and backward-feasible greedy ranking algorithms that compute an optimal solution [Edmonds, 1971]. The standard greedy algorithm *for knapsack* orders players in decreasing order by their value divided by size, and then adds the players to the stream in this order as long as they still fit into the knapsack; taking the maximum between this and simply outputting the most valuable player gives a 2-approximation. An alternative greedy algorithm adds elements to the stream in non-decreasing order of value as long as they fit, and obtains a $((1-\lambda)\mu)^{-1}$-approximation (recall that μ bounds the ratio between the sizes of the smallest and largest elements, and λ denotes the ratio between the sizes of the largest element and the knapsack itself). It is not difficult to see that both these algorithms correspond to forward-feasible ranking algorithms, and can be turned into backward-feasible ranking algorithms by going through the players in reverse order and rejecting them as long as the remaining players do not fit into the knapsack. For matchings, adding edges to the stream by non-increasing value unless one of their endpoints is already matched yields a 2-approximation. It is not possible to turn this process around into a backward-feasible ranking algorithm, instead we propose a different 2-approximation that is a backward-feasible ranking algorithm – see Algorithm 12 in Appendix A.2.2.

5.3.2 Composition of Ranking Algorithms

We distinguish between compositions of two forward-feasible ranking algorithms, and two backward-feasible ones. We call the former *forward composition* and the latter *backward composition*. A crucial ingredient to both is the following definition of a composition rule.

Definition 5.3.4. A *composition rule* for a ranking algorithm for buyers and a ranking algorithm for sellers is a boolean function, which receives as input a buyer-seller pair (i, j) obtained by querying the two ranking algorithms, the pair's value and cost v_i, c_j, and their query histories, and outputs either 1 ("the pair is accepted") or 0 ("the pair is rejected").

Types of Composition Rules. Specific composition rules that we will use in this work are the t-threshold composition rule, the lookback composition rule, and the lookback t-threshold composition rule defined next.

Definition 5.3.5. The *t-threshold* composition rule accepts a buyer-seller pair (i, j) if and only if the pair's gain from trade $v_i - c_j$ is at least t, where t is a non-negative threshold in $\mathbb{R}$.

Definition 5.3.6. A *lookback* composition rule is any rule that decides whether to accept or reject a buyer-seller pair (i, j) without observing their value and cost v_i, c_j, but rather observing only the history of values and costs of previously-queried players.

Definition 5.3.7. The *lookback t-threshold* composition rule is a lookback composition rule that accepts a buyer-seller pair (i, j) if and only if some part of the history contains a previously-queried pair (i', j'), whose gain from trade $v_{i'} - c_{j'}$ is at least t, where t is a non-negative threshold in $\mathbb{R}$.

Forward and Backward Composition. We are now ready to formally define forward and backward composition. Intuitively, the main difference between forward and backward composition is in the information that is available to the composition rules they use.

Definition 5.3.8. The *forward composition* of two forward-feasible ranking algorithms greedily determines an allocation as follows:

1. It queries the output streams of both forward-feasible ranking algorithms. If either ranking algorithm returns no player, then it stops and rejects all remaining players.

2. Otherwise it applies the composition rule to the resulting buyer-seller pair to decide whether or not to accept it based on its value and cost and the history of previous queries.

3. If it accepts the pair, then it continues with Step 1. Otherwise, it stops and rejects all remaining players.

Definition 5.3.9. The *backward composition* of two backward-feasible ranking algorithms greedily determines an allocation as follows:

0. (*Preprocessing.*) It rejects the first $n - \min\{n - \ell_B, m - \ell_S\}$ buyers, and the first $m - \min\{n - \ell_B, m - \ell_S\}$ sellers by querying the output streams (where recall that ℓ_B, ℓ_S are the ranks of the largest-rank buyer and seller, respectively, that must be rejected for feasibility). The remaining players now form a feasible set.

1. It queries the output streams of both backward-feasible ranking algorithms.

2. It applies the composition rule to the resulting buyer-seller pair to decide whether or not to reject it based on its value and cost and the history of previous queries, excluding preprocessing queries from Step 0.[7]

3. If it rejects the pair, then it continues with Step 1. Otherwise, it stops and accepts all remaining players.

Observation 5.3.10. *Forward and backward composition leads to a feasible set of accepted buyers and sellers.*

We present two illustrative examples that demonstrate what our composition framework is able to achieve in simple settings.

Example 5.3.11 (VCG Mechanism via Forward Composition)**.** Consider an unconstrained double auction setting. For such a setting, the trivial forward-feasible ranking algorithm ranks the players from high to low quality (i.e., from high to low value or from low to high cost). Observe that the VCG double auction is precisely a forward composition of the trivial forward-feasible ranking algorithms using the 0-threshold composition rule and applying threshold payments. Indeed, it sorts the players from high to low quality and greedily accepts trading pairs (i, j) while their gain from trade $v_i - c_j$ is positive.

[7] The exclusion of queries carried out during preprocessing is so that the composition rule will only take into account pairs that could have potentially traded. This is necessary to achieve the budget balance property – see Section 5.6.

Example 5.3.12 (McAfee's Trade Reduction Mechanism via Backward Composition)**.** In the same unconstrained double auction setting, the trivial backward-feasible ranking algorithm ranks the players in reverse order, from low to high quality (i.e., from low to high value or from high to low cost). McAfee's trade reduction double auction is precisely a backward composition of the trivial backward-feasible ranking algorithms using the lookback 0-threshold composition rule and applying threshold payments. Indeed, it sorts the players from low to high quality and greedily rejects trading pairs, until it has rejected a pair (i, j) whose gain from trade $v_i - c_j$ is non-negative.

More generally, the forward-feasible greedy ranking algorithms *for matroids* combined with the 0-threshold rule and threshold payments are in fact the VCG mechanism, while the backward-feasible greedy ranking algorithms for matroids combined with the lookback 0-threshold rule and threshold payments are precisely McAfee's mechanism. The next example illustrates this for McAfee's mechanism.

Example 5.3.13 (Trade Reduction on Matroids)**.** Consider a scenario with 4 buyers and values $8, 5, 3, 2$, and 3 sellers with costs $1, 2, 4$. Suppose we can accept at most two sellers for trade (e.g., because they are firms which need a facility to produce the good and there is space for at most two facilities). Note this is a matroid constraint, with a simple matroid called "2-uniform" [Oxley, 1992]. The backward-feasible greedy ranking algorithm for matroids would return buyers in order $4, 3, 2, 1$ with $\ell_B = 0$, and sellers in order $3, 2, 1$ with $\ell_S = 1$. So in the preprocessing step of the backward composition, buyers $4, 3$ and seller 3 would be rejected. The lookback 0-threshold composition rule would proceed by also rejecting buyer-seller pair $(2, 2)$. Now for the first time there exists a pair rejected by the composition rule with positive gain from trade, and the mechanism would stop and accept the remaining buyer-seller pair $(1, 1)$.

Even more generally, the composition framework enables us to generalize the VCG and trade reduction mechanisms to accommodate feasibility constraints beyond matroids, as long as the constraints have appropriate forward- and backward-feasible ranking algorithms. This leads to Definitions 5.4.4 (VCG-style double auction) and 5.4.10 (trade reduction-style double auction) below.

Running Examples. Our composition framework leads to a wealth of double auction mechanisms via forward or backward composition. Specifically, any of the ranking algorithms described in the previous subsection can be combined with the composition rules described above. However, not all combinations will yield DSIC or WGSP mechanisms or succeed in obtaining a good fraction of the optimal welfare. Our goal in the next few sections will be to develop the theory that explains which properties of the ranking algorithms and the composition rule guarantee that the resulting double auction mechanism has these properties.

5.4 Incentives

Recall that a composition theorem relates the properties of ranking algorithms and a composition rule to those of the composed double auction mechanism. This section presents our composition theorems for the double auction properties of DSIC and WGSP. In Section 5.4.1 we focus on the DSIC property. Our DSIC composition theorem applies equally well to both forward and backward compositions; for simplicity of presentation we state it for forward compositions. In Section 5.4.2 we focus on the WGSP property. Our WGSP composition theorem applies only to backward compositions – this is one of the significant differences between the forward and backward approaches. We discuss the implications of both composition theorems for our three running examples.

5.4.1 DSIC Composition Theorem

Preliminaries. To state the DSIC composition theorem we shall need the following definitions, related to the monotonicity notion from Definition 5.2.1.

Definition 5.4.1. A ranking algorithm is *rank monotone* if the rank of a player changes monotonically with his quality as follows:

- In a forward-feasible ranking algorithm for buyers, $v_i < v'_i \implies r_i(\vec{v}) \geq r_i(v'_i, v_{-i})$ for every $\vec{v}, i$.
- In a forward-feasible ranking algorithm for sellers, $c_j < c'_j \implies r_j(\vec{c}) \leq r_j(c'_j, c_{-j}))$ for every $\vec{c}, j$.
- In a backward-feasible ranking algorithm for buyers, $v_i < v'_i \implies r_i(\vec{v}) \leq r_i(v'_i, v_{-i})$ for every $\vec{v}, i$.
- In a backward-feasible ranking algorithm for sellers, $c_j < c'_j \implies r_j(\vec{c}) \geq r_j(c'_j, c_{-j})$ for every $\vec{c}, j$.

Definition 5.4.2. A ranking algorithm is *consistent* if players' ranks are consistent with their qualities as follows:

- In a forward-feasible ranking algorithm for buyers, for every i, i' with ranks $< \infty$, if buyer i's rank is lower than that of buyer i' then $v_i \geq v_{i'}$.
- In a forward-feasible ranking algorithm for sellers, for every j, j' with ranks $< \infty$, if seller j's rank is lower than that of seller j' then $c_j \leq c_{j'}$.
- In a backward-feasible ranking algorithm for buyers, let ℓ_B be the rank of the largest-rank buyer to discard for feasibility; for every i, i' with rank $> \ell_B$, if buyer i's rank is lower than that of buyer i' then $v_i \leq v_{i'}$.
- In a backward-feasible ranking algorithm for sellers, let ℓ_S be the rank of the largest-rank seller to discard for feasibility; for every j, j' with rank $> \ell_S$, if seller j's rank is lower than that of seller j' then $c_j \geq c_{j'}$.

What is the relation between rank-monotonicity (Definition 5.4.1) and consistency (Definition 5.4.2)? Neither implies the other: A forward-feasible ranking algorithm for buyers which outputs only buyer 1 and only if his value is *lower* than a threshold t is consistent but not rank monotone. A forward-feasible ranking algorithm for buyers which always outputs buyer 1 and then buyer 2, even when $v_2 > v_1$, is rank-monotone but not consistent.

Definition 5.4.3 (Monotonicity for Forward/Backward Composition)**.** Consider a composition rule, and let the following be two different inputs to it: two buyer-seller pairs (i, j) and (i', j'), with qualities (v_i, c_j) and $(v_{i'}, c_{j'})$, and histories (h_i, h_j) and $(h_{i'}, h_{j'})$, respectively. Assume that the second input dominates the first in terms of quality, i.e., $v_{i'} \geq v_i$ and $c_{j'} \leq c_j$.

The composition rule is *monotone for forward composition* if for any such two inputs where h' is a prefix of h, if it accepts the pair (i, j) then it accepts the pair (i', j'). The composition rule is *monotone for backward composition* if for any such two inputs where h is a prefix of h', if it accepts the pair (i, j) then it accepts the pair (i', j').

An example of a composition rule that is monotone for forward or backward composition is the t-threshold composition rule, since it ignores the histories and accepts whenever $v_i - c_j \geq t$, which implies that $v_{i'} - c_{j'} \geq t$. The following definition relates the above concepts to the VCG double auction mechanism, generalized to accommodate feasibility constraints.[8]

Definition 5.4.4. A *VCG-style* double auction mechanism is a forward composition of consistent, rank-monotone ranking algorithms using the 0-threshold composition rule and applying threshold payments.

Theorem Statement and Analysis

We are now ready to state our composition theorem; we state here the version for forward composition, but it applies equally well to backward composition by an analogous argument.

Theorem 5.4.5. *A forward composition of consistent, rank monotone ranking algorithms using a composition rule that is monotone for forward composition and applying threshold payments is a DSIC double auction mechanism.*

[8] The welfare of a VCG-style mechanism as defined in Definition 5.4.4 is not necessarily optimal. Since we focus on computationally tractable mechanisms, this is expected – maximizing welfare subject to feasibility constraints is not always computationally tractable (assuming P≠NP).

Proof. We apply the characterization of DSIC double auctions in Proposition 5.2.3 to show that the composition is DSIC. We only need to show that the allocation rule is monotone – this also means that the payments are well-defined. Fix value and cost profiles $\vec{v}, \vec{c}$. We argue that an accepted buyer who raises his value remains accepted; a similar argument shows that an accepted seller who lowers his cost remains accepted, thus completing the proof.

Denote the accepted buyer by i, and the seller with whom i trades by j. Consider the application of the composition rule to the pair (i, j), and denote by (v_i, c_j) and (h_i, h_j) the pair's qualities and histories, respectively. By assumption, the composition rule accepts (i, j). By rank monotonicity of the forward-feasible ranking algorithm for buyers, when i raises his value to $v_i' > v_i$, his rank weakly decreases. Let j' be the seller with whom i is considered for trade by the composition rule after he raises his value and his rank decreases. Consider the application of the composition rule to the pair (i, j'), and denote by $(v_i', c_{j'})$ and $(h_i', h_{j'})$ the pair's qualities and histories, respectively. Then by consistency of the forward-feasible ranking algorithm for sellers, $c_{j'} \leq c_j$. Since the composition rule is monotone for forward composition, and since the history h_i' is a prefix of h_i, the pair (i, j') must be accepted for trade as well by the composition rule. □

The following is an immediate corollary of Theorem 5.4.5.

Corollary 5.4.6. *Let $t \in \mathbb{R}$ be a threshold. Every forward composition of consistent, rank-monotone ranking algorithms using the t-threshold composition rule and applying threshold payments is DSIC. In particular, VCG-style double auctions are DSIC.*

Necessity of the Conditions. The following examples show that monotonicity of the ranking algorithms and the composition rule is necessary for monotonicity of the double auction's allocation rule and DSIC, and demonstrate why consistency of the ranking algorithms is required.

Example 5.4.7 (Necessity of Rank Monotonicity)**.** Let $n = m = 1$. Let the type spaces be $[v_1^0, v_1^1] = [1, v_{\max}]$ and $[c_1^0, c_1^1] = [0, 0]$. Consider a forward composition using the 0-threshold composition rule of a forward-feasible ranking algorithm for buyers that outputs buyer 1 if and only if $v_1 < v_{\max}$, and a forward-feasible ranking algorithm for sellers that outputs seller 1. Observe that the ranking algorithm for buyers is consistent but not rank monotone, and the ranking algorithm for sellers is consistent and rank monotone. Let buyer 1's value be $v_{\max}$ (seller 1's cost is determined by his type space). Then no players are paired and accepted unless buyer 1 shades his bid and reports a value $< v_{\max}$, in which case he is paired with seller 1 and accepted.

Example 5.4.8 (Necessity of Composition Rule Monotonicity)**.** Let $n = m = 1$. Let the type spaces be $[v_1^0, v_1^1] = [1, v_{\max}]$ and $[c_1^0, c_1^1] = [0, 0]$. Consider a forward composition of forward-feasible ranking algorithms for buyers resp. sellers that output buyer 1 resp. seller 1, using a composition rule that accepts a buyer-seller pair (i, j) if and only if the gain from trade $v_i - c_j$ is $\in [1, v_{\max})$. Let buyer 1's value be $v_{\max}$ (seller 1's cost is determined by his type space). Then buyer 1 is paired with seller 1, and the pair is rejected unless buyer 1 shades his bid and reports a value $< v_{\max}$, in which case the pair is accepted.

Example 5.4.9 (Necessity of Ranking Algorithm Consistency)**.** Let $n = m = 2$. Let the type spaces be $[v_1^0, v_1^1] = [4, 4]$, $[v_2^0, v_2^1] = [1, 2]$, $[c_1^0, c_1^1] = [3, 3]$, and $[c_2^0, c_2^1] = [0, 0]$. Consider a forward composition using the 0-threshold composition rule of a forward-feasible ranking algorithm for buyers that outputs buyer 1 and then buyer 2 if $v_2 < 2$ and outputs only buyer 2 otherwise, and a forward-feasible ranking algorithm for sellers that always outputs seller 1 and then seller 2. Observe that the ranking algorithm for buyers is both rank monotone and consistent, and the ranking algorithm for sellers is rank monotone but not consistent. Let buyer 2's value be 2 (buyer 1's value and the sellers' costs are determined by their type spaces). Then buyer 2 is paired with seller 1 and rejected unless he shades his bid and reports a value < 2, in which case he is paired with seller 2 and accepted.

Running Examples. What implications does the DSIC composition theorem have for our running examples? We sketch here how it applies to all three examples. The forward-feasible ranking algorithms that are based on computing the feasible set of buyers (resp. sellers) with maximum value (resp. minimum cost) are rank-monotone and consistent if the buyers (resp. sellers) are returned from highest to lowest value (resp. lowest to highest cost). An analogous argument applies to the backward-feasible ranking algorithms based on this approach. The greedy algorithms discussed before all lead to rank-monotone ranking algorithms, and with the exception of the greedy algorithm which ranks buyers by value divided by size, also to consistent ranking algorithms.

5.4.2 WGSP Composition Theorem

Deferred Acceptance. For our WGSP composition theorem we leverage the framework of *deferred-acceptance algorithms* [Milgrom and Segal, 2014]. For completeness we describe such algorithms in Algorithm 9. Their relation to one-sided auctions is as follows: Applied to a set of players, they output an ordered stream R of rejected players, and thus can form the basis of one-sided auctions. A *deferred-acceptance auction* for sale (procurement) is a one-sided mechanism whose allocation rule runs a maximization (minimization) deferred-acceptance algorithm to get the set R of players to reject. Non-rejected (or "active") players are accepted. By monotonicity of the scoring functions in Algorithm 9, the allocation rule is monotone and so we can set the payment rule to apply threshold payments.

Deferred-acceptance algorithms can form the basis of backward-feasible ranking algorithms, as follows: A *deferred-acceptance ranking* algorithm for buyers (sellers) is a ranking algorithm that first runs a maximization (minimization) version of a deferred-acceptance algorithm to get R. It sets the first part of the output stream of buyers (sellers) to R, and lets the rank of the largest-rank buyer (seller) to reject be $\ell = |R|$. We require that the result is backward-feasible. The second part of the output stream is obtained by sorting the buyers (sellers) not in R from low to high value (high to low cost). Observe that, by construction and by monotonicity of the scoring functions, deferred-acceptance rankings are rank-monotone and consistent.

The following definition relates the above concepts to the trade reduction mechanism of McAfee, generalized to accommodate feasibility constraints:

Algorithm 9: Deferred-acceptance algorithm (maximization and minimization versions)

Input: Set of players P and their bid profile $\vec{b}$; access to a feasibility set system $(P, \mathcal{I}_P)$ (represented in a computationally tractable way);

Set $A = P$; % Set of active players – initially all players
Set $R = ()$; % Ordered stream of rejected players

while $A \neq \emptyset$ **do**
  for every $i \in A$
    Set score of player $i = s_i^A(b_i, b_{-A})$;

  % Where the scoring function is such that:
  % (a) It does not depend on the bids of the set of active players $A \setminus \{i\}$
  % (b) It may depend on the feasibility set system $(P, \mathcal{I}_P)$
  % (c) It is non-negative and weakly increasing in its first argument

  % Maximization (minimization) version:
  if *all scores are* ∞ *(zero)* **then**
    exit;
  end
  Set i^* = player with lowest finite (highest nonzero) score;
  $A = A \setminus \{i^*\}$; % i^* becomes inactive
  Append i^* to R;
end

Definition 5.4.10. A *trade reduction-style* mechanism is a backward composition of deferred-acceptance ranking algorithms using the lookback 0-threshold composition rule and applying threshold payments.

Theorem Statement and Analysis

We are now ready to state and prove our WGSP composition theorem.

Theorem 5.4.11. *A backward composition of deferred-acceptance ranking algorithms using a lookback composition rule and applying threshold payments is a WGSP double auction mechanism.*

Proof. It is sufficient to show that the allocation rule of the backward composition in the theorem statement can be implemented by a deferred-acceptance algorithm applied to the set of all players $N \cup M$: Consider the one-sided deferred-acceptance auction for sale based on this algorithm; by construction its allocation rule is identical to the original allocation rule of the backward composition, and both mechanisms apply threshold payments. Thus the incentives of the players in both mechanisms are identical. The theorem then follows from Corollary 1 of Milgrom and Segal [2014], by which every deferred-acceptance auction is WGSP.[9]

Consider the backward composition of deferred-acceptance ranking algorithms. We begin by transforming the deferred-acceptance ranking algorithm for sellers into a deferred-acceptance ranking algorithm for *pseudo*-buyers, which runs a maximization rather than minimization version of a deferred-acceptance algorithm but has equivalent output (the transformation is straightforward – see [Dütting et al., 2014b] for details).

Our goal now is to use the deferred-acceptance ranking algorithms for buyers and for pseudo-buyers – in particular the scores $s_i^B(v_i, v_{-B})$ for buyers and those for pseudo-buyers and the ranks ℓ_B, ℓ_P of the largest-rank buyer and pseudo-buyer to reject for feasibility – in order to obtain a maximization deferred-acceptance algorithm

[9]Note that while the deferred-acceptance framework of Milgrom and Segal [2014] primarily focuses on finite bid spaces, some of the results including Corollary 1 apply to infinite bid spaces as well, as in our setting (see Footnote 17 of [Milgrom and Segal, 2014]).

for the set of buyers and pseudo-buyers. This algorithm should be equivalent to the allocation rule of the backward composition in the theorem statement. We thus need to define scoring functions such that Algorithm 9 will implement the steps of the backward composition as described in Definition 5.3.9. Utilizing the lookback composition rule of the backward composition, and the fact that scoring functions are allowed to depend on the set of active players, it is not hard to define scoring functions that are weakly increasing in their first argument, implement the preprocessing (Step 0), maintain the balance $|B| = |P|$, and implement the composition rule (Step 2). For the construction of such scoring functions see [Dütting et al., 2014b]. This provides a deferred-acceptance implementation of the backward composition, as required. □

The following is an immediate corollary of Theorem 5.4.11.

Corollary 5.4.12. *Let $t \in \mathbb{R}$ be a threshold. Every backward composition of deferred-acceptance ranking algorithms using the lookback t-threshold composition rule and applying threshold payments is WGSP. In particular, trade reduction-style double auctions are WGSP.*

Two further corollaries apply to backward compositions of deferred-acceptance ranking algorithms using a lookback composition rule, after restricting attention to finite bid spaces: (i) Such double auctions can be implemented as clock auctions (by Proposition 3 in [Milgrom and Segal, 2014]); (ii) For every such double auction, consider a double auction that uses the same allocation rule but charges first-price payments, then it has a complete-information Nash equilibrium in which the allocation and payments are identical to the DSIC outcome of the backward composition double auction with threshold payments (by Proposition 6 in [Milgrom and Segal, 2014]).

None of the strong properties shown here are shared by forward compositions; furthermore, WGSP (as well as the property of budget balance) cannot be achieved by forward compositions without losing a great deal in welfare [Dütting et al., 2014b].

Running Examples. As in the case of the DSIC composition theorem, we ask what implications does the WGSP composition theorem have for our running examples, and sketch how it applies to all three examples.

The greedy algorithm for matroids can be implemented as a deferred-acceptance algorithm. As we saw above, the greedy algorithm for knapsack which sorts buyers by non-increasing value can also be implemented as a deferred-acceptance algorithm. For matchings, neither the optimal nor the greedy by weight algorithm can be implemented via deferred-acceptance; we therefore design a new deferred-acceptance algorithm for matchings (see Appendix A.2.2).

5.5 Welfare

In this section we discuss the welfare guarantees of double auction mechanisms arising from compositions, and the implications for the three running examples.

Throughout this section we will consider fixed feasibility constraints as given by the set systems $\mathcal{I}_N$, $\mathcal{I}_M$ and the requirement that a set of buyers $B \subseteq N$ and a set of sellers $S \subseteq M$ is feasible if $B \in \mathcal{I}_N$, $S \in \mathcal{I}_M$ and $|B| \leq |S|$.

Our goal will be to relate the welfare $W(\vec{v}, \vec{c})$ achieved by a fixed double auction mechanism obtained through composition to the optimal welfare $\mathrm{OPT}(\vec{v}, \vec{c})$ for value and cost profile $(\vec{v}, \vec{c})$. Specifically, we will identify properties related to the composition rule and the ranking algorithms that guarantee that the resulting double auction mechanism achieves optimal or near-optimal welfare on every input.

Recall that since we assume downward-closed feasibility sets, the optimal welfare will be achieved by a set of buyers and a set of sellers of equal cardinality. In our analysis, we will use $s^*(\vec{v}, \vec{c})$ to denote the cardinality of the pair (B^*, S^*) with the maximum number of trades among all pairs $(B, S) \in \mathcal{I}_N \times \mathcal{I}_M$ with $|B| = |S|$ achieving optimal welfare $\mathrm{OPT}(\vec{v}, \vec{c})$.

5.5.1 Composition Theorem: Approximation Parameters

To state our welfare composition theorem we need several parameters that quantify relevant properties of the composition rule and the ranking algorithms, as well as how hard the problem instance is. In particular, the parameters related to the ranking algorithms quantify how limited are the complements introduced by the feasibility constraints. The degree of complementarity affects the welfare guarantees, as anticipated in Section 1.4.2 when laying out our goals.

Composition Rule. The first pair of parameters is related to how "exhaustive" the composition rule is. The first parameter $s(\vec{v}, \vec{c})$ denotes the number of buyer-seller pairs that the double auction mechanism accepts on input $(\vec{v}, \vec{c})$. The second parameter $s'(\vec{v}, \vec{c})$ denotes the optimal number of buyer-seller pairs that a composition mechanism based on the same ranking algorithms but with a potentially different composition rule could have accepted. In other words, $s'(\vec{v}, \vec{c})$ is the number of buyer-seller pairs that the 0-threshold rule would accept. Note that for unconstrained double auction settings it is possible that

$$\forall \vec{v}, \vec{c} : s'(\vec{v}, \vec{c}) = s^*(\vec{v}, \vec{c});$$

and this is also the case for some constrained settings such as matroids.

In our analysis we will focus on cases where $s'(\vec{v}, \vec{c}) \geq 1$ and $s(\vec{v}, \vec{c}) \leq s'(\vec{v}, \vec{c})$. The former means that a forward (or backward) composition based on these ranking algorithms could have accepted at least one buyer-seller pair with non-negative gain from trade. The latter means that the double auction mechanism under consideration does not accept buyer-seller pairs with negative gain from trade.

Ranking Algorithms. The next two parameters $\alpha \geq 1$ and $\beta \geq 1$ quantify how close to the optimal solution the one-sided ranking algorithms are *at any point q of their execution*, in the worst case over all inputs $\vec{v}$ resp. $\vec{c}$. Intuitively, such an "any time" guarantee is necessary as the final number of accepted buyers depends on the interaction of the two ranking algorithms and is therefore extrinsic to the buyer-ranking algorithm (and similarly for the sellers).

Formally, let $card(\mathcal{I}_N) = \max_{B \in \mathcal{I}_N} |B|$ and $card(\mathcal{I}_M) = \max_{S \in \mathcal{I}_M} |S|$ denote the cardinality of the largest feasible buyer resp. seller set. For every $q \in [card(\mathcal{I}_N)]$, denote by $v_{\mathrm{OPT}}(q)$ the value of the feasible solution of at most q buyers that maximizes total value. For every $q \in [card(\mathcal{I}_M)]$, denote by $c_{\mathrm{OPT}}(q)$ the cost of the feasible solution of at least q sellers that minimizes total cost. That is,

$$v_{\mathrm{OPT}}(q) = \max_{B \in \mathcal{I}_N, |B| \leq q} \sum_{i \in B} v_i \quad \text{and} \quad c_{\mathrm{OPT}}(q) = \min_{S \in \mathcal{I}_M, |S| \geq q} \sum_{j \in S} c_j.$$

For a given forward-feasible ranking algorithm, we append 0's to the buyer stream if it has length less than $card(\mathcal{I}_N)$ and we append $\bar{c}$'s to the seller stream if it has length less than $card(\mathcal{I}_M)$. We denote by $v_{\mathrm{ALG}}(q)$ (resp. $c_{\mathrm{ALG}}(q)$) the value (resp. cost) achieved by greedily allocating to the first $q \in [card(\mathcal{I}_N)]$ buyers (resp. $q \in [card(\mathcal{I}_M)]$ sellers) in the modified output stream. For a given backward-feasible ranking algorithm, the definitions are the same except that the last q buyers (resp. sellers) in the feasible part of the output stream are considered.

Definition 5.5.1. A ranking algorithm for buyers is a *uniform α-approximation* if for every value profile $\vec{v}$ and every $q \leq card(\mathcal{I}_N)$,

$$v_{\mathrm{ALG}}(q) \geq \frac{1}{\alpha} \cdot v_{\mathrm{OPT}}(q).$$

A ranking algorithm for sellers is a *uniform β-approximation* if for every cost profile $\vec{c}$ and every $q \leq card(\mathcal{I}_M)$,

$$c_{\mathrm{ALG}}(q) \leq \beta \cdot c_{\mathrm{OPT}}(q).$$

The closer $\alpha \geq 1$ and $\beta \geq 1$ are to 1, the better the ranking algorithms.

Problem Instance. The final parameter $\gamma(\vec{v}, \vec{c})$ measures how difficult the problem instance $(\vec{v}, \vec{c})$ is; and is a standard tool with mixed-sign objective functions [cf. Roughgarden and Sundararajan, 2009]. It measures how close the optimal solution $\mathrm{OPT}(\vec{v}, \vec{c})$ is to zero. Recall the definition of $s^*(\vec{v}, \vec{c})$ from above, then $\mathrm{OPT}(\vec{v}, \vec{c}) = v_{\mathrm{OPT}}(s^*(\vec{v}, \vec{c})) - c_{\mathrm{OPT}}(s^*(\vec{v}, \vec{c}))$. Let $\gamma(\vec{v}, \vec{c}) = v_{\mathrm{OPT}}(s^*(\vec{v}, \vec{c}))/c_{\mathrm{OPT}}(s^*(\vec{v}, \vec{c}))$. Clearly, $\gamma \geq 1$ as $v_{\mathrm{OPT}}(s^*(\vec{v}, \vec{c})) \geq c_{\mathrm{OPT}}(s^*(\vec{v}, \vec{c}))$. For $\gamma(\vec{v}, \vec{c}) = 1$ we have $\mathrm{OPT}(\vec{v}, \vec{c}) = 0$; hence we focus on the case where $\gamma(\vec{v}, \vec{c}) > 1$. Intuitively, the closer $\gamma(\vec{v}, \vec{c})$ is to 1, the closer the optimal welfare is to 0, and the harder it is to achieve a good relative approximation.

5.5.2 Theorem Statement and Analysis

Theorem 5.5.2. *Consider input $(\vec{v}, \vec{c})$. The forward (backward) composition of two consistent, forward-feasible (backward-feasible) ranking algorithms that are uniform α- and β-approximations, using a composition rule that accepts the $s(\vec{v}, \vec{c}) \leq s'(\vec{v}, \vec{c})$ lowest (resp. highest) ranking buyer-seller pairs, achieves welfare at least*

$$\frac{s(\vec{v}, \vec{c})}{s'(\vec{v}, \vec{c})} \cdot \frac{\frac{\gamma(\vec{v}, \vec{c})}{\alpha} - \beta}{\gamma(\vec{v}, \vec{c}) - 1} \cdot \mathrm{OPT}(\vec{v}, \vec{c}).$$

Note that if $\alpha = \beta = 1$, then the second term in the approximation factor vanishes. For general α and β the bound degrades gracefully from this ideal case, in the sense that the dependence on the approximation ratios $\frac{1}{\alpha}$ and β is linear.

Proof. Without loss of generality we may assume $\mathrm{OPT}(\vec{v}, \vec{c}) > 0$. Our goal is to show that

$$v_{ALG}(s(\vec{v}, \vec{c})) - c_{ALG}(s(\vec{v}, \vec{c})) \geq \frac{s(\vec{v}, \vec{c})}{s'(\vec{v}, \vec{c})} \cdot \frac{\frac{\gamma(\vec{v}, \vec{c})}{\alpha} - \beta}{\gamma(\vec{v}, \vec{c}) - 1} \cdot \Big(v_{\mathrm{OPT}}(s^*(\vec{v}, \vec{c})) - c_{\mathrm{OPT}}(s^*(\vec{v}, \vec{c})) \Big).$$

Since the double auction is composed of forward-feasible (backward-feasible) consistent ranking algorithms, we can number the buyers and sellers from the beginning (end) of the respective streams by $1, 2, \ldots$ such that $v_1 \geq v_2 \geq \cdots \geq v_{s'(\vec{v}, \vec{c})}$ and $c_1 \leq c_2 \leq \cdots \leq c_{s'(\vec{v}, \vec{c})}$. Using this notation,

$$\begin{aligned} v_{ALG}(s(\vec{v}, \vec{c})) - c_{ALG}(s(\vec{v}, \vec{c})) &= \sum_{i=1}^{s(\vec{v}, \vec{c})} (v_i - c_i); \\ v_{ALG}(s'(\vec{v}, \vec{c})) - c_{ALG}(s'(\vec{v}, \vec{c})) &= \sum_{i=1}^{s'(\vec{v}, \vec{c})} (v_i - c_i). \end{aligned}$$

Another implication of the fact that the double auction is composed of consistent ranking algorithms is that the gain from trade is non-increasing. That is, $i < j$

implies $v_i - c_i \geq v_j - c_j$. Hence for all s such that $s(\vec{v},\vec{c}) < s \leq s'(\vec{v},\vec{c})$ we have $v_s - c_s \leq \frac{1}{s(\vec{v},\vec{c})} \sum_{i=1}^{s(\vec{v},\vec{c})} (v_i - c_i)$. It follows that

$$\begin{aligned} v_{ALG}(s(\vec{v},\vec{c})) - c_{ALG}(s(\vec{v},\vec{c})) &= \sum_{i=1}^{s'(\vec{v},\vec{c})} (v_i - c_i) - \sum_{i=s(\vec{v},\vec{c})+1}^{s'(\vec{v},\vec{c})} (v_i - c_i) \\ &\geq \sum_{i=1}^{s'(\vec{v},\vec{c})} (v_i - c_i) - \Big(s'(\vec{v},\vec{c}) - s(\vec{v},\vec{c})\Big) \cdot \frac{1}{s(\vec{v},\vec{c})} \sum_{i=1}^{s(\vec{v},\vec{c})} (v_i - c_i), \end{aligned}$$

and this is equal to

$$\Big(v_{ALG}(s'(\vec{v},\vec{c})) - c_{ALG}(s'(\vec{v},\vec{c}))\Big) - \left(\frac{s'(\vec{v},\vec{c})}{s(\vec{v},\vec{c})} - 1\right) \cdot \Big(v_{ALG}(s(\vec{v},\vec{c})) - c_{ALG}(s(\vec{v},\vec{c}))\Big).$$

Rearranging this shows

$$v_{ALG}(s(\vec{v},\vec{c})) - c_{ALG}(s(\vec{v},\vec{c})) \geq \frac{s(\vec{v},\vec{c})}{s'(\vec{v},\vec{c})} \cdot \Big(v_{ALG}(s'(\vec{v},\vec{c})) - c_{ALG}(s'(\vec{v},\vec{c}))\Big). \tag{5.1}$$

Recall that $s^*(\vec{v},\vec{c})$ is defined as the number of trades in a solution that maximizes welfare, while $s'(\vec{v},\vec{c})$ is the number of trades that maximizes welfare for the given ranking algorithms. By the definition of $s'(\vec{v},\vec{c})$ all trades up to and including $s'(\vec{v},\vec{c})$ are beneficial, and then either one of the streams reached its end or the subsequent trades are no longer beneficial. Hence, by the definition of v_{ALG} and c_{ALG},

$$v_{ALG}(s'(\vec{v},\vec{c})) - c_{ALG}(s'(\vec{v},\vec{c})) \geq v_{ALG}(s^*(\vec{v},\vec{c})) - c_{ALG}(s^*(\vec{v},\vec{c})). \tag{5.2}$$

Finally, we use that the ranking algorithms are uniform α- and β-approximations and the definition of $\gamma(\vec{v},\vec{c})$ to deduce that

$$\begin{aligned} v_{ALG}(s^*(\vec{v},\vec{c})) - c_{ALG}(s^*(\vec{v},\vec{c})) &\geq \frac{1}{\alpha} \cdot v_{\mathrm{OPT}}(s^*(\vec{v},\vec{c})) - \beta \cdot c_{\mathrm{OPT}}(s^*(\vec{v},\vec{c})) \\ &= \left(\frac{\gamma(\vec{v},\vec{c})}{\alpha} - \beta\right) \cdot c_{\mathrm{OPT}}(s^*(\vec{v},\vec{c})) \\ &= \frac{\frac{\gamma(\vec{v},\vec{c})}{\alpha} - \beta}{\gamma(\vec{v},\vec{c}) - 1} \cdot \Big(v_{\mathrm{OPT}}(s^*(\vec{v},\vec{c})) - c_{\mathrm{OPT}}(s^*(\vec{v},\vec{c}))\Big). \end{aligned}$$

Combining this with Inequalities (5.1) and (5.2) completes the proof. $\square$

We obtain the following corollaries for VCG-style and trade reduction-style mechanisms with the 0-threshold or the lookback 0-threshold composition rule.

Corollary 5.5.3. *Consider the forward composition of two forward-feasible, consistent ranking algorithms that are uniform* α*- and* β*-approximations. For input* $(\vec{v}, \vec{c})$*, the* 0*-threshold rule accepts the* $s'(\vec{v}, \vec{c})$ *lowest ranking buyer-seller pairs. Hence its welfare is at least*

$$\frac{\frac{\gamma(\vec{v},\vec{c})}{\alpha} - \beta}{\gamma(\vec{v}, \vec{c}) - 1} \cdot \mathrm{OPT}(\vec{v}, \vec{c}).$$

Corollary 5.5.4. *Consider the backward composition of two backward-feasible, consistent ranking algorithms that are uniform* α*- and* β*-approximations. For input* $(\vec{v}, \vec{c})$*, the lookback* 0*-threshold rule accepts the* $s'(\vec{v}, \vec{c}) - 1$ *lowest ranking buyer-seller pairs. Hence its welfare is at least*

$$\left(1 - \frac{1}{s'(\vec{v}, \vec{c})}\right) \cdot \frac{\frac{\gamma(\vec{v},\vec{c})}{\alpha} - \beta}{\gamma(\vec{v}, \vec{c}) - 1} \cdot \mathrm{OPT}(\vec{v}, \vec{c}).$$

When $\alpha = \beta = 1$, these two corollaries specialize to the traditional guarantees of the VCG and trade reduction mechanisms. For general α and β these bounds again degrade gracefully from this ideal case as the dependence on the approximation ratios $\frac{1}{\alpha}$ and β is again linear.

Running Examples. Since all ranking algorithms that we have not yet ruled out are consistent, it is the uniform approximation property that we have to check. The greedy ranking algorithms for matroids are not only optimal, but also uniformly so. This follows from the fact that if we restrict the independent sets of a matroid to sets of size at most k, the matroid structure is preserved [Schrijver, 2003]. Similarly, the ranking algorithms for knapsack that rank by weight achieve a uniform $((1 - \lambda)\mu)^{-1}$ approximation, as we show in Appendix A.2.1. The ranking algorithms for matchings (greedy and the one proposed in Appendix A.2.2) achieve a uniform 2-approximation.

5.6 Budget Balance

This section studies the budget balance properties of compositions, and derives implications for the running examples.

We say that a backward composition *reduces an efficient trade* if there is a buyer-seller pair with non-negative gain from trade that is rejected by the composition rule (Step 2 in Definition 5.3.9).[10] Then:

Theorem 5.6.1. *A backward composition of deferred-acceptance ranking algorithms using a lookback composition rule that reduces at least one efficient trade and applying threshold payments is a BB double auction mechanism.*

Proof. Without loss of generality, denote the buyers in the output stream of the deferred-acceptance ranking algorithm for buyers by $1, 2, \ldots, n$, and the sellers in the output stream of the deferred-acceptance ranking algorithm for sellers by $1, 2, \ldots, m$. Recall that a backward composition is a composition of two backward-feasible ranking algorithms; the streams returned by such algorithms each have a player, denoted by ℓ_B and ℓ_S respectively for buyers and sellers, such that the following holds: for every $i \geq \ell_B$, the set of buyers $N \setminus \{1, \ldots, i\}$ is feasible, and for every $j \geq \ell_S$, the set of sellers $M \setminus \{1, \ldots, j\}$ is feasible.

Let (ℓ'_B, ℓ'_S) be a buyer-seller pair with non-negative gain from trade that is reduced by the composition rule. Since players $1, \ldots, \max\{\ell_B, \ell_S\}$ from both streams are rejected in the preprocessing stage of the backward composition (Definition 5.3.9), it must be the case that both ℓ'_B and ℓ'_S are strictly greater than $\max\{\ell_B, \ell_S\}$, i.e., they appear in their respective streams after the players rejected in the preprocessing

[10] In the proof we use the property of backward composition by which the composition rule does not take into account efficient trades reduced in the preprocessing stage (Step 0 in Definition 5.3.9). The following example shows why this is necessary: Consider a setting with two buyers and two sellers. The value profile is $(8, 4)$ and the cost profile is $(6, 5)$. It is only feasible to accept up to one buyer, and the deferred-acceptance ranking algorithm for buyers scores the first buyer by his value and the second buyer by 3 times his value. The sellers are unconstrained and ranked by the trivial backward-feasible ranking algorithm (which clearly has a deferred-acceptance implementation). Thus the buyer stream is $1, 2$ and the seller stream is $1, 2$, and the preprocessing step rejects buyer 1 and seller 1, a pair with positive gain from trade. If this pair were part of the history, the pair buyer 2 and seller 2 would be accepted. But this pair has negative gain from trade and so the resulting mechanism is either not IR or not BB.

stage. Denote the value of ℓ'_B by v', and the cost of ℓ'_S by c'. Clearly, $v' \geq c'$.

Our goal is to show that every buyer whom the composition accepts pays at least v'. A symmetric argument shows that every seller whom the composition accepts is paid at most c'. Since $v' \geq c'$ this is sufficient to establish the property of budget balance. In fact, due to threshold payments, it is enough to show that every buyer i whom the composition accepts will be rejected if he reports a value lower than v'.

Consider an accepted buyer i. By the greediness of backward composition, which repeatedly rejects until the first buyer-seller pair is accepted, it must be the case that $i > \ell'_B$, i.e., buyer i appears after the reduced buyer ℓ'_B in the buyer ranking. By consistency of the deferred-acceptance ranking algorithm for buyers, i's original report is thus at least v'. What changes if i reports a value lower than v'? By consistency, the buyer ranking must change in this case, and we denote by r the new rank of buyer i. We distinguish between two cases:

- $r \leq \max\{\ell_B, \ell_S\}$. That is, the new rank of i is smaller than the original rank of the largest-ranked buyer to discard for feasibility. We now exploit a property of deferred-acceptance algorithms together with consistency to establish that in the new buyer stream where buyer i appears in rank r, the first $r-1$ buyers have not changed and are still buyers $1, ..., r-1$. The property we use is that an active player's bid does not affect the scores of any other active player in the deferred-acceptance algorithm. Since we know that rejecting buyers $1, \ldots, r-1$ is not enough for feasibility, buyer i is necessarily rejected.

- $r > \max\{\ell_B, \ell_S\}$. As above, the first $r-1$ buyers in the new buyer stream have not changed and are still buyers $1, \ldots, r-1$. Therefore, by consistency and since i reports a value lower than v', it must hold that $r \leq \ell'_B$. We now use the fact that originally the buyer-seller pair (ℓ'_B, ℓ'_S) was reduced. This means that the decision of the lookback composition rule given the history up to and including rank $r-1$ is to reject, and so buyer i is rejected.

This completes the proof. □

When uniform 1-approximate ranking algorithms are available, the trade reduction mechanism is BB; for cases where either $\alpha > 1$ or $\beta > 1$, Theorem 5.6.1 shows that suitable generalizations of this mechanism are BB.

Running Examples. Our BB composition theorem applies to all ranking algorithms that are implementable within the deferred-acceptance framework: the greedy algorithm for matroids, the greedy by weight algorithm for knapsack, and the new matching algorithm that we describe in Appendix A.2.2.

6

Competitive Market Equilibria with Complements

6.1 Chapter Introduction

Computational complexity has already had plenty to say about economic equilibria. Some of the most celebrated results in algorithmic game theory concern equilibrium concepts that have guaranteed existence, like Nash equilibria in finite games, or market equilibria in markets with divisible goods and concave utilities. These results determine whether or not such equilibria can be computed by polynomial-time algorithms under standard complexity assumptions (see, e.g., [Chen et al., 2009b; Daskalakis et al., 2009; Chen et al., 2009a]). Computational complexity is informative also for equilibrium concepts that do not have guaranteed existence, such as pure Nash equilibria in concise games, or market equilibria in markets with concave production functions (see, e.g., [Fischer et al., 2006; Papadimitriou and Wilkens, 2011]). For example, by proving that it is computationally hard to compute whether or not a given game or market admits an equilibrium of a desired type, one provides evidence that there is no "nice characterization" of the instances in which such an equilibrium exists.

The primary theme of this work is that *computational complexity can also be used to study the equilibrium existence question itself*, in that non-existence results can be derived from the computational intractability of related optimization problems, under suitable complexity assumptions.

We explore this theme in the classic setting of market-clearing prices for markets of m indivisible items, where there are n consumers and each consumer i has a valuation $v_i(S)$ for each bundle S of items. There is a large literature on understanding, for various classes $\mathcal{V}$ of allowable valuations, what types of prices – item prices, anonymous bundle prices, etc. – are necessary and sufficient to guarantee the existence of pricing equilibria when all consumers have valuations in $\mathcal{V}$ (see, e.g., [Parkes and Ungar, 2000; Sun and Yang, 2006; Candogan et al., 2015; Ben-Zwi et al., 2013; Sun and Yang, 2014; Candogan et al., 2014; Candogan and Pekec, 2014]). The goal of this work is to formalize a general version of this question, show that the answer is inextricably linked to the computational complexity of well-studied problems (like demand oracles, revenue-maximization, or welfare-maximization), and derive from this connection a number of results for the (non-)existence of pricing equilibria.

6.1.1 Some Highlights

While the main point of this work is its methodology, many of our specific results are also of independent interest. We now informally describe a few of them; see later sections for the relevant formal definitions.

1. Whenever the welfare-maximization problem for a valuation class $\mathcal{V}$ is strictly computationally harder than the demand problem for $\mathcal{V}$ (with item prices), Walrasian equilibria need not exist (Proposition 6.2.4). For example, for budget-additive valuations, assuming P $\neq$ NP, the welfare-maximization problem (which is strongly NP-hard) is strictly harder than the demand problem (which is weakly NP-hard), thus ruling out the possibility of guaranteed existence of Walrasian equilibria (Corollary 6.2.2).

2. Walrasian equilibria, which employ only anonymous item prices, are guaranteed to exist when consumers only want one item (unit-demand). It follows from our work that they need not exist when consumers only want a pair of items,[1] but

[1]This does not follow from [Gul and Stacchetti, 1999] since consumers do not want a single item,

what if we use a richer set of prices, defined on both items and item pairs? It is easy to see that *non-anonymous* prices on item pairs recover the guaranteed existence of pricing equilibria (Observation 6.5.2), but are anonymous prices also sufficient? Under the assumption that NP $\nsubseteq$ coNP, our general results imply a negative answer to this question (Corollary 6.4.2). This conditional non-existence stems from the facts that revenue-maximization with such prices is polynomial-time solvable, the demand problem with such valuations and prices is polynomial-time solvable, and the welfare-maximization problem with these valuations is NP-hard (Proposition 6.4.5).

3. Walrasian equilibria are remarkable in that, despite using only m-dimensional prices (one per item), they are guaranteed to exist for a valuation class with dimension exponential in m (gross substitutes valuations, see Lemma 6.5.1). Despite much research on pricing equilibria for various valuation and pricing classes, no generalizations of Walrasian equilibria with these properties (succinctness and guaranteed existence) have been found to date. Our methodology provides an explanation by identifying an algorithmic barrier to such results: it would require a novel polynomial-time algorithm for the welfare-maximization problem, beyond solving the standard configuration linear programming relaxation of the problem (Section 6.5.3).

6.1.2 Related Work

Our study of equilibrium existence for different classes of valuations is related to the large literature on how the class of valuations shapes computational aspects of combinatorial auctions and markets [Cramton et al., 2006]. We mention here three such aspects. Additional related work relevant to Sections 6.4 to 6.6 appears in these sections.

1. Communication: Nisan and Segal [2006] study communication aspects of welfare-maximization in market settings. Generalizing a result of [Parkes, 2002], they

and so unit-demand valuations are not a subclass.

prove for many classes of valuations (such as submodular valuations) that welfare-maximization requires exponential communication. In particular, it requires communication of a price system that is exponential in the number of items m.

2. Approximation: Algorithmic aspects of approximate welfare-maximization have been extensively studied, especially for the "complement-free" valuation hierarchy [Blumrosen and Nisan, 2007]. Recent research expands upon this by studying valuation classes with "limited" complements and good approximation guarantees [Abraham et al., 2012; Feige et al., 2014].[2]

3. Representation and elicitation: Both succinctness ("compactness") of valuation classes and their learnability have been studied (see, e.g., [Boutilier et al., 2004; Zinkevich et al., 2003]), and for classes of non-succinct valuations, simple sketches have been pursued ([Cohavi and Dobzinski, 2014] and references within).

6.1.3 Chapter Organization

We begin with a discussion in Section 6.2 of the basic market equilibrium notion of Walrasian equilibrium. We then describe our general formalism in Section 6.3. Anonymous pricing is discussed in Section 6.4, compressed pricing is discussed in Section 6.5, and linear pricing is discussed in Section 6.6. Section 6.7 summarizes.

6.2 Walrasian Equilibrium

This section shows that, even for the exhaustively-studied Walrasian equilibrium concept, computational complexity is a useful tool for (conditionally) ruling out existence in many scenarios. While simple to prove, these results are conceptually interesting, and also develop intuition for our subsequent results about other types of pricing equilibria. To keep this section brief we keep formal definitions to a minimum; see

[2]Note that in our work the "baseline" valuation class – beyond which valuations are considered as exhibiting complements – is that of gross substitutes.

Section 6.3 for precise explanations of all terms. We begin with an overview of the results deferring formal statements to Section 6.2.2.

6.2.1 Overview

Recall that a Walrasian equilibrium is an allocation of items to consumers together with item prices such that (1) every consumer is allocated a bundle in his *demand set*; and (2) the market clears, meaning every unallocated item has price 0. For condition (1), recall that the demand set of a consumer given prices is the set of all bundles S that maximize his payoff – his value $v(S)$ for the bundle minus his total payment for it $p(S)$. The *demand problem* for a valuation class $\mathcal{V}$ is to compute, given prices, a bundle of items in the demand set of a consumer with $v \in \mathcal{V}$. For condition (2), note an equivalent condition that is useful later when we discuss more general types of prices: that the allocation maximizes the seller's revenue (given the prices).

A Walrasian equilibrium may or may not exist in a market – it depends on the structure of the consumers' valuations. At first blush, it might seem that the main results of Gul and Stacchetti [1999] and Milgrom [2000] tell us everything we want to know about the existence question: markets with gross substitutes valuations (Definition 6.3.1) always possess Walrasian equilibria, and for every valuation class $\mathcal{V}$ that contains all unit-demand valuations (Definition 6.3.2) and a non-gross-substitutes valuation, there exists a market with valuations in $\mathcal{V}$ and no Walrasian equilibrium. There are, however, natural valuation classes $\mathcal{V}$ that do not include all unit-demand valuations (many examples appear below). Research on such valuation classes, to which the non-existence results of Gul and Stacchetti [1999] and Milgrom [2000] do not apply, has progressed in a relatively ad hoc fashion, relying on explicit constructions to rule out the existence of Walrasian equilibria in various cases (see, e.g., [Candogan et al., 2015; Ben-Zwi et al., 2013; Sun and Yang, 2014; Candogan et al., 2014; Candogan and Pekec, 2014]). Is a more systematic approach possible?

Let the *allocation problem* for a valuation class $\mathcal{V}$ be to maximize welfare given a market with a valuation profile from $\mathcal{V}$. Our main conceptual point in this section is:

Proposition 6.2.1 (Informal)**.** *A necessary condition for the guaranteed existence of a Walrasian equilibrium in markets with valuations from class $\mathcal{V}$ is that the demand problem is at least as hard computationally as the allocation problem for $\mathcal{V}$.*

Proposition 6.2.1 is conceptually interesting because it links a basic economic question (existence of Walrasian equilibria), which is defined without reference to computation, to computational complexity considerations. We show in Section 6.2.2 how it can be used to re-derive many known non-existence results in an arguably more systematic way, rooted in the mature understanding in theoretical computer science of the computational complexity of various problems. We highlight in particular two applications, following directly from the complexity of well-known computational problems under the widely believed P $\neq$ NP assumption – one for the class of budget-additive valuations and one for the class of positive graphical valuations (for definitions of these classes see the proofs of Corollaries 6.2.2 and 6.2.3 in Section 6.2.2, and also, respectively, [Lehmann et al., 2006] and [Conitzer et al., 2005]).

Corollary 6.2.2. *Assuming* P $\neq$ NP*, there exists a market with budget-additive valuations for which there is no Walrasian equilibrium.*

Corollary 6.2.3. *Assuming* P $\neq$ NP*, there exists a market with positive graphical valuations for which there is no Walrasian equilibrium.*

6.2.2 Results

In this section we formalize Proposition 6.2.1 and prove some applications. The proposition relies on a well-known connection between the allocation problem of maximizing welfare and Walrasian equilibria.

Proposition 6.2.4. *Consider a valuation class $\mathcal{V}$. If for every market with valuations from $\mathcal{V}$ there exists a Walrasian equilibrium, then the allocation problem reduces in polynomial time to the demand problem.*

Proof. (Sketch.) Recall that the allocation problem has a canonical linear programming formulation, known as the *configuration LP* (see, e.g., [Blumrosen and Nisan, 2007], Section 11.3.1). A Walrasian equilibrium exists if and only if the configuration LP has an optimal integral solution ([Blumrosen and Nisan, 2007], Theorem 11.13). For completeness, both the configuration LP and the proof of this statement appear in Appendix A.3.1. By assumption of Walrasian equilibrium existence, if we can solve the LP in polynomial time then we can solve the allocation problem in polynomial time. Since the dual of the configuration LP has exponentially many constraints and polynomially many variables, it is solvable using the ellipsoid method by polynomially many applications of a separation oracle [Nisan and Segal, 2006]. The dual constraints correspond to verifying, given item prices, that the seller is maximizing his revenue and that each consumer's allocation is in his demand set. Revenue maximization with item prices is equivalent to market clearance and thus verifiable. We conclude that the separation oracle, and thus the allocation problem, reduces to solving the demand problem. This completes the proof. □

A valuation class has *demand oracle access* if it is assumed that the demand problem can be solved by a computationally efficient oracle. The following is an immediate corollary of Proposition 6.2.4 for such classes.

Corollary 6.2.5. *Consider a valuation class $\mathcal{V}$ with demand oracle access. If for every market with valuations from $\mathcal{V}$ there exists a Walrasian equilibrium, then the allocation problem for $\mathcal{V}$ can be solved using a polynomial amount of computation and demand queries.*

To derive the applications described above in Corollaries 6.2.2 and 6.2.3, our main interest is in the contrapositive of Proposition 6.2.4 – if the allocation problem cannot be reduced to the demand problem (assuming $\mathcal{P} \neq$ NP), then a Walrasian equilibrium does not exist for every market. We use this to prove Corollaries 6.2.2 and 6.2.3.

Proof of Corollary 6.2.2 for budget-additive valuations – sketch. Recall that a budget-additive valuation v assigns values $\{v_j\}_j$ to the items and has a budget b; the value $v(S)$ of a bundle S is the aggregate value of the items capped by the budget, i.e., $v(S) = \min\{\sum_{j\in S} v_j, b\}$. The first step of the proof is to show that the demand problem for a budget additive valuation v given item prices can be solved in *pseudo-polynomial* time. The second step of the proof is to show that the allocation problem is *strongly* NP-hard, meaning NP-hard even for polynomially-bounded budget-additive valuations. We show this by reduction from the strongly NP-hard *bin packing* problem. Therefore, assuming P $\neq$ NP, the allocation problem cannot be reduced to the demand problem, and so by Proposition 6.2.4 a Walrasian equilibrium is not guaranteed. This completes the proof. □

Proof of Corollary 6.2.3 for positive graphical valuations. Recall that a positive graphical valuation v is represented by a graph $G = (M, E)$ with the set of items M as vertices. The vertices and edges of G are weighted by a non-negative weight function $w(\cdot)$. The value of a bundle S is $v(S) = w(G(S))$, where $G(S)$ is the subgraph induced by the vertices in S, and $w(G(S))$ is the total weight of the subgraph's vertices and edges. The fact that the demand problem given item prices is solvable in polynomial time was observed by [Abraham et al., 2012] (Proposition 5.1): The valuation v defined by the positive graph is supermodular, and hence so is the consumer's utility function after subtracting item prices from valuation v; maximizing supermodular functions can be done in polynomial time. On the other hand, the allocation problem is NP-hard by a reduction of [Conitzer et al., 2005] (Theorem 6) from the problem of *exact cover by 3-sets*. This completes the proof. □

We can also apply Proposition 6.2.4 to the following valuation classes to show that, assuming P $\neq$ NP, a Walrasian equilibrium need not exist: graphical valuations with an underlying tree graph and sign consistent weights [Candogan et al., 2015]; XOS valuations with a sub-polynomial number of clauses; and other hard cases of succinct supermodular valuations.

6.3 General Formalism

Walrasian equilibria do not exist for many important valuation classes (see Figure 6.1). A natural idea is to permit prices that are somewhat more complex than anonymous item prices, while retaining as many of the nice properties of Walrasian equilibria as possible. This section introduces the formalism needed to evaluate the prospects of this idea. We define abstract pricing functions, prove that the first and second

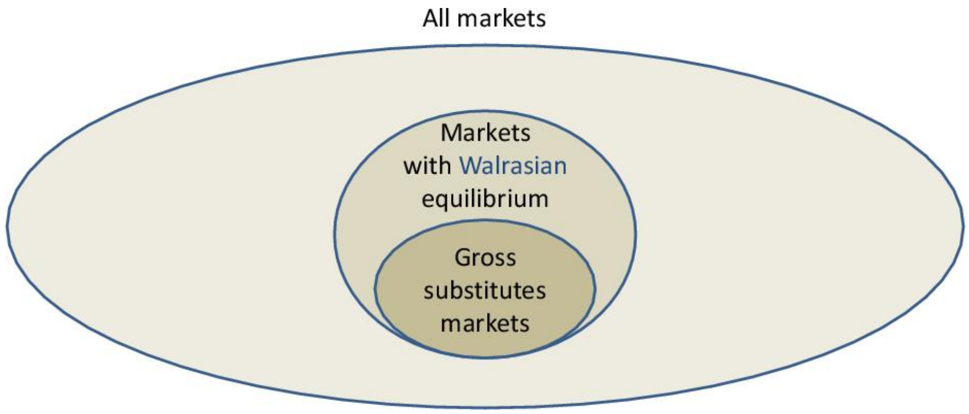

Figure 6.1: Markets with Walrasian equilibrium form a small subclass of all markets.

welfare theorems continue to hold for them, and discuss our requirement of efficient verification. Sections 6.4 and 6.5 build on the concepts of this section to rule out the guaranteed existence of any non-trivial generalization of Walrasian equilibria for many valuation classes.

6.3.1 Valuations and Pricings

Definitions

A *market* is a set M of m items sold by a seller to a set N of n consumers. An *allocation* $\vec{S} = (S_1, \dots, S_n)$ of the items is a partial partition of the item set M among the n consumers (items are allowed to remain unallocated).

Every consumer i has a *valuation* function v_i that maps bundles of items $S \subseteq M$ to their values in $\mathbb{R}$.[3] As is standard, valuation functions are assumed to be normalized ($v_i(\emptyset) = 0$) and monotone ($v_i(S) \leq v_i(T)$ for every $S \subseteq T$), and thus non-negative. A *valuation profile* v is a set of n valuations. The *welfare* of an allocation $\vec{S}$ given a valuation profile v is $\sum_i v_i(S_i)$.

We can describe general pricings in exactly the same way we describe general valuations. Formally, a *pricing profile* p is a set of n *pricing* functions, one function p_i for each consumer, each mapping bundles of items to prices in $\mathbb{R}^+$. An *anonymous*

[3]Throughout, all numerical values are assumed to have a polynomial representation in the parameter m.

pricing profile has the same pricing for every consumer. The *revenue* of an allocation $\vec{S}$ given a pricing p is $\sum_i p_i(S_i)$.

Consider consumer i with valuation v_i and pricing p_i. The *payoff* (also known as *utility*) of this consumer from being allocated bundle S_i is quasi-linear, i.e., is equal to $v_i(S_i) - p_i(S_i)$. His *demand set* is the family of all bundles that maximize his payoff.

We are interested in *classes* $\mathcal{V}$ ($\mathcal{P}$) of valuations (pricings). We say that a profile belongs to a class if all its valuations (pricings) belong to that class.

Examples. Gross substitutes valuations and unit-demand valuations are examples of valuation classes.

Definition 6.3.1. A valuation v is *gross substitutes* if for every two vectors of item prices $\vec{p}, \vec{q}$ such that $\vec{q} \geq \vec{p}$, for every bundle S in the demand set of v given $\vec{p}$, there exists a bundle T in the demand set of v given $\vec{q}$ which contains every item $j \in S$ whose price according to $\vec{q}$ equals its price according to $\vec{p}$.

Definition 6.3.2. A valuation v is *unit demand* if there are item values $\{v_j\}_j$ such that for every bundle S, $v(S) = \max_{j \in S}\{v_j\}$.

Representation

A naïve representation of valuations and pricings is of exponential size and hence computationally uninteresting. One standard way to circumvent this is via *oracle access*. We say that a valuation or pricing class has oracle access of a certain kind if its functions are computed by such oracles, whose representation and queries are considered computationally efficient. The two most common kinds of oracles are as follows.

1. A *value oracle* represents a valuation; it gets a bundle S and returns its value. Similarly, a *price oracle* represents a pricing; it gets a bundle S and returns its price. Unless otherwise noted, we consider only valuation and pricing classes that have such oracle access.

2. A *demand oracle* represents a valuation and is defined with respect to a pricing class $\mathcal{P}$; it gets a pricing from $\mathcal{P}$ (represented by a price oracle), and returns a bundle S in the demand set given this pricing.

Of special interest are valuations (pricings) that belong to *succinct* classes, i.e., whose value (price) oracle has an explicit description polynomial in m that also runs in polynomial time. It is hard to imagine using non-succinct pricings, except in markets with a very small number of items. Similarly, it is natural to assume that actual consumers – subject to bounded rationality constraints – can be modeled faithfully with succinct valuations. Most of the valuation classes we deal with in this chapter are succinct.

6.3.2 The Allocation, Demand and Revenue Problems and their Complexity

The Problems

The *allocation* (*welfare-maximization*) *problem* for a valuation class $\mathcal{V}$ is defined as follows: The optimization version gets as input a market with valuations from $\mathcal{V}$, and outputs a welfare-maximizing allocation. The decision version gets an additional input $w \in \mathbb{R}$, and decides whether or not there exists an allocation with welfare at least w.[4] The *symmetric* allocation problem is a special case of the allocation problem where all consumers on the market have the same valuation.

The *demand* (*utility-maximization*) *problem* for a valuation class $\mathcal{V}$ and pricing class $\mathcal{P}$ is defined as follows: The optimization version gets as input a valuation v from $\mathcal{V}$ and a pricing p from $\mathcal{P}$, and outputs a bundle in the demand set of a consumer with valuation v given pricing p. The decision version gets an additional input $u \in \mathbb{R}$, and decides whether or not the utility of the demand set bundles is at least u.

The *revenue-maximization* problem for a pricing class $\mathcal{P}$ is equivalent to the allocation problem for the valuation class $\mathcal{V} = \mathcal{P}$. The problem is to decide, given a pricing profile from $\mathcal{P}$ and a target revenue $r \in \mathbb{R}$, whether or not there exists an

[4] We will sometimes be informal about distinguishing between optimization and decision versions of a computational problem.

allocation with revenue at least r. The *symmetric* revenue-maximization problem is the special case where all pricings are equal.

Complexity

This chapter focuses on three of the most fundamental computational complexity classes: P, NP and coNP.[5] Consider a succinct valuation class $\mathcal{V}$. The allocation problem (and hence also the revenue-maximization problem when $\mathcal{P}$ is succinct) is in NP since for every market with valuations from $\mathcal{V}$ and a target welfare w, it can be verified in polynomial time with value queries that an allocation's welfare is $\geq w$. Similarly, the demand problem is also in NP. The allocation problem is in coNP if for every such market there is a polynomial-sized *certificate* (also known as *proof*), which verifies in polynomial time with value queries that the optimal welfare is $< w$. We do not know of a natural valuation class for which the allocation problem is known to be in coNP but not known to be in P.

6.3.3 Pricing Equilibrium and Verifiability

Definitions and Basic Properties

Consider a pricing class $\mathcal{P}$. A *pricing equilibrium* for a market is an allocation $(S_1, \ldots, S_n)$ together with a *supporting* pricing profile that belongs to $\mathcal{P}$. We say a pricing profile is supporting if the following two conditions hold:

1. The allocation maximizes the consumers' payoffs given the pricings; in other words, for every consumer i, bundle S_i is in his demand set.

2. The allocation maximizes the seller's revenue given the pricings.

The latter condition can also be stated symmetrically to the former one, by defining the seller's "valuation" v_0 to be zero for every bundle, and then requiring that the

[5]We refer the non-computer-scientist reader to Appendix A.3.2 for short informal definitions, or to [Arora and Barak, 2009] for a detailed exposition.

allocation maximize his payoff.[6] We say that pricing class $\mathcal{P}$ is supporting with respect to valuation class $\mathcal{V}$ if for every market with a valuation profile from $\mathcal{V}$ there is a supporting pricing profile from $\mathcal{P}$.

A Walrasian equilibrium is precisely a pricing equilibrium supported by an anonymous profile from the class of item pricings.

The usefulness of the generalized pricing equilibrium notion stems from the following generalization of the classic welfare theorems (proofs follow those of the standard welfare theorems).

Theorem 6.3.3 (Generalized welfare theorems)**.**

1. *Every pricing equilibrium maximizes welfare.*
2. *For every welfare-maximizing allocation there exists a supporting pricing profile; moreover, every supporting pricing profile supports every welfare-maximizing allocation.*

Verifiability

For a pricing equilibrium concept to be meaningful, it should be efficiently recognizable (in the computational sense) by all parties:[7] given an alleged pricing equilibrium, each consumer should be able to verify efficiently that he is given a utility-maximizing bundle (given his pricing), and the seller should be able to verify efficiently that the proposed allocation maximizes his revenue (given the pricings). For example, with gross substitutes valuations and item prices, Walrasian equilibria can be verified efficiently: the revenue-maximization condition is trivial to check (all unsold items have zero price), and each utility-maximization condition can be checked with a polynomial number of value queries [Bertelsen, 2004].

[6] Algorithmically, however, the seller's problem of maximizing revenue is different from the consumers' problem of maximizing payoff by finding a bundle in demand – the seller must find a way to allocate the items among the consumers such that their total payment according to the pricing profile is maximized.

[7] One could argue that the stronger requirement of polynomial-time computability should also hold. Since our main results are negative, adopting the weaker requirement of efficient verifiability only strengthens our results.

Assume that the pricing class $\mathcal{P}$ is succinct and revenue-maximization is tractable. There are then two levels of verifiability. At the less restrictive level, we say that pricing equilibria for $\mathcal{P}$ and valuation class $\mathcal{V}$ are *verifiable with demand oracle access* if the valuations in $\mathcal{V}$ have demand oracle access with respect to the pricings in $\mathcal{P}$. In this case, all the pricing equilibrium conditions can be verified using a polynomial amount of computation and demand queries. We also say that the *allocation problem* for $\mathcal{V}$ is verifiable with demand oracle access if the welfare-maximizing allocation can be verified using a polynomial amount of computation and demand queries given a succinct certificate. Observe that if a pricing equilibrium for $\mathcal{P}$ and $\mathcal{V}$ is guaranteed to exist, then the allocation problem is verifiable with demand oracle access, since the equilibrium serves as a succinct and verifiable certificate.

A stronger requirement on $\mathcal{V}$ and $\mathcal{P}$ is that pricing equilibria are *verifiable with value oracle access*. This is the case if the demand problem can be solved using a polynomial amount of computation and value queries. As above we can also define allocation problems that are verifiable with value oracle access, and as above it holds that when $\mathcal{P}$ and $\mathcal{V}$ guarantee the existence of a pricing equilibrium then the corresponding allocation problem is verifiable with value oracle access. If in addition $\mathcal{V}$ is succinct, then the pricing equilibria are verifiable in polynomial time, and the allocation problem belongs to coNP.

6.4 Anonymous Pricing

Walrasian equilibria employ item prices that are "simple" in three respects: they are anonymous (common to all consumers); they are succinct; and they make the revenue-maximization problem tractable. This section and the next study, in general, when we can and cannot obtain the first and second property, respectively, while keeping the third property as a constraint.

6.4.1 Related Work

Consider a valuation class and its supporting pricing classes. It is helpful to classify the pricing classes into two categories: In the first category, pricings belong to the same class as the valuations, and in the second category, pricings are untethered from the valuation class and are allowed to belong to a broader or altogether different class. While the second category may be required in order to achieve anonymity, keeping supporting prices simple is important for verifiability of the resulting pricing equilibria.

There are two main examples in the literature of valuation classes supported by anonymous pricing profiles.

1. *Gross substitutes and anonymous item prices*: This canonical example belongs to the first category, since item prices correspond to additive valuations, which belong to the class of gross substitutes.

2. *Supperadditive valuations and anonymous bundle prices*: A valuation v is superadditive if $v(S \cup S') \geq v(S) + v(S')$ for every disjoint bundles $S, S' \subseteq M$ [Parkes and Ungar, 2000; Vohra, 2011; Sun and Yang, 2014]. This example belongs to the second category, since bundle prices correspond to general valuations, which are a strict superclass of superadditive valuations. Even more is known about anonymous bundle prices: Bikhchandani and Ostroy [2002] formulate a linear program that gives an instance by instance characterization of markets for which such supporting prices exist, and Parkes and Ungar [2000] develop an algorithm for finding them. They further show that anonymous bundle prices support certain XOR valuation subclasses.

Candogan [2013] studies anonymous graphical pricings, and establishes the following negative finding: a valuation profile from the class of graphical valuations based on series-parallel graphs is supported by an anonymous graphical pricing profile if and only if it is also supported by anonymous item prices (i.e., a Walrasian equilibrium exists for this market instance).

6.4.2 Overview

In this overview section we present our main findings on anonymous pricing profiles and derive applications; details appear in Section 6.4.3. We focus on succinct pricing classes for which the symmetric revenue-maximization problem is tractable.[8] We establish a similar condition to that of Proposition 6.2.1, where we showed that a necessary condition for the existence of a Walrasian equilibrium is that the demand problem is as hard as the allocation one.

Proposition 6.4.1 (Informal)**.** *Consider a succinct pricing class $\mathcal{P}$ for which the symmetric revenue-maximization problem is tractable. A necessary condition for the guaranteed existence of a pricing equilibrium with an anonymous pricing profile from $\mathcal{P}$ in markets with valuations from class $\mathcal{V}$ is that the demand problem is at least as hard computationally as the allocation problem for these classes.*

We demonstrate three applications of Proposition 6.4.1.

We begin by defining *hypergraph pricings* – a succinct pricing class for which the symmetric revenue-maximization problem is tractable: Recall the class of positive graphical valuations from Corollary 6.2.3. A natural generalization is to consider succinct hypergraphs in place of graphs.[9] By definition, the resulting pricing class is succinct. The symmetric revenue-maximization problem is tractable since revenue is maximized by allocating all items to a single consumer.

As a first application, consider a natural generalization of unit-demand valuations (recall Definition 6.3.2). A valuation v is *pair-demand* if there exist values $\{v_{i,j}\}_{i,j}$ for item pairs such that for every bundle S of size at least 2, $v(S) = \max_{i,j \in S}\{v_{i,j}\}$, and for every singleton S, $v(S) = 0$. Since anonymous item prices support unit-demand valuations, one could hope that anonymous hypergraphical prices support pair-demand valuations. Proposition 6.4.1 refutes this hope.

[8] The previous "success stories" of anonymous pricing profiles beyond item prices – anonymous bundle prices supporting supermodular and other valuation classes – escape our impossibility results in this section by being non-succinct.

[9] Such a pricing is represented by a hypergraph over the set of items M as vertices, with non-negative weights on the vertices and hyperedges. The price of a bundle S is the total weight of the hypergraph induced by it.

Corollary 6.4.2. *Assuming* NP $\nsubseteq$ coNP*, there exists a market with pair-demand valuations for which there is no pricing equilibrium supported by an anonymous profile of hypergraph pricings.*

As a second application, we return to positive graphical valuations, which form a succinct subclass of supermodular valuations and as such may be considered a natural candidate for the existence of a verifiable anonymous pricing equilibrium. However:

Corollary 6.4.3. *Assuming* NP $\nsubseteq$ coNP*, there exists a market with positive graphical valuations for which there is no pricing equilibrium supported by an anonymous profile of hypergraph pricings.*

As a third application, consider a generalization of the Sun and Yang [2006] market: Sun and Yang define a market with two sets of items referred to as "tables" and "chairs", where consumers have unit-demand valuations for both tables and chairs, but consider a pair of (table, chair) to be complementary. They show that a Walrasian equilibrium is guaranteed to exist in their market. An obvious generalization is to allow *three* sets of items with cross-set complementarities (*cf.*, [Teytelboym, 2014]), however:

Corollary 6.4.4. *Assuming* NP $\nsubseteq$ coNP*, there exists a generalized Sun and Yang market for which there is no pricing equilibrium supported by an anonymous profile of hypergraph pricings.*

6.4.3 Results

We first prove a more formal version of Proposition 6.4.1 and then establish the applications described above.

Proposition 6.4.5. *Consider a succinct valuation class $\mathcal{V}$ and a succinct pricing class $\mathcal{P}$ for which the symmetric revenue-maximization problem is tractable. If for every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with an anonymous pricing profile from $\mathcal{P}$, then the allocation problem belongs to* coNP *whenever the demand problem belongs to* coNP.

Proof. Assume the demand problem belongs to coNP. Consider a market with valuations from $\mathcal{V}$. We want to show the existence of a polynomial-sized certificate, which can certify in polynomial time that the optimal welfare is $< w$. By the first welfare theorem (Theorem 6.3.3), every pricing equilibrium maximizes welfare, and by the proposition's assumptions, a polynomial-sized equilibrium with an anonymous pricing profile is guaranteed to exist for $\mathcal{V}$ and $\mathcal{P}$. Consider such a pricing equilibrium with welfare $< w$; we know that symmetric revenue-maximization is tractable and so verifying it reduces to verifying n instances of the demand problem, one per consumer. Since the demand problem is assumed to be in coNP, there are polynomial-sized certificates for verifying these in polynomial time. Together with the pricing equilibrium itself they form the required certificate, completing the proof. □

For valuation classes with demand oracle access, a similar proof shows the following.

Corollary 6.4.6. *Consider a valuation class $\mathcal{V}$ with demand oracle access and a succinct pricing class $\mathcal{P}$ for which the symmetric revenue-maximization problem is tractable. If for every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with an anonymous pricing profile from $\mathcal{P}$, then the allocation problem for $\mathcal{V}$ is verifiable with demand oracle access.*

To derive the applications described above in Corollaries 6.4.2 to 6.4.4, our main interest is in the following immediate corollary of Proposition 6.4.5.

Corollary 6.4.7. *Assuming* NP $\not\subseteq$ coNP*, for classes $\mathcal{V}$ and $\mathcal{P}$ as in Proposition 6.4.5 for which the allocation problem is* NP*-hard and the demand problem is in* P*, there exists a market with valuations in $\mathcal{V}$ for which there is no pricing equilibrium supported by an anonymous pricing profile from $\mathcal{P}$.*

Applications

Proof of Corollary 6.4.2 for pair-demand valuations. The allocation problem for pair-demand valuations is NP-hard, by a reduction from the problem of *three-dimensional matching*. On the other hand, the demand problem for pair-demand valuations with hypergraphical prices is in P, simply by computing and comparing the payoffs from all possible pairs of items. Applying Corollary 6.4.7 completes the proof. □

To prove Corollaries 6.4.3 and 6.4.4, the following definition is useful. A consumer has a *triplet valuation* if there is a triplet of items j, k, l such that his value for a bundle S is as follows: $v_{j,k}$ if $j, k \in S$ and $l \notin S$; $v_{k,l}$ if $k, l \in S$ and $j \notin S$; $v_{j,l}$ if $j, l \in S$ and $k \notin S$; $v_{j,k} + v_{k,l} + v_{j,l}$ if $j, k, l \in S$; and zero otherwise. Since triplet valuations are strict subclasses of positive graphical valuations and of the valuations in the generalized Sun and Yang market, Corollaries 6.4.3 and 6.4.4 follow immediately from the next lemma.

Lemma 6.4.8. *Assuming* NP $\nsubseteq$ coNP*, there exists a market with triplet valuations for which there is no pricing equilibrium supported by an anonymous profile of hypergraph pricings.*

Proof. The allocation problem for triplet valuations is NP-hard by a reduction from the problem of exact cover by 3-sets [Conitzer et al., 2005]. On the other hand, the demand problem for triplet valuations with hypergraphical prices is in $\mathcal{P}$, simply by computing and comparing the payoffs from every pair of items in the triplet and from the entire triplet. Applying Corollary 6.4.7 completes the proof. □

6.5 Compressed Pricing

This section focuses on the low-dimensionality or *compressed* aspect of the prices used in Walrasian equilibria – there is only one price for each of the m items, and yet existence is guaranteed even for the high-dimensional class of gross substitutes valuations. When else can we achieve guaranteed existence of succinct equilibria, even for non-succinct valuations?

6.5.1 Related Work and Background

Compression of valuations is an important theme in mechanism and market design, with a classic trade-off between expressiveness of the valuations and simplicity of the market mechanism.

In the context of market equilibria, the class of gross substitutes stands out as the canonical example for which simple m-dimensional equilibria exist, despite the fact that the dimension of this valuation class is exponential in m.[10] Note that the high-dimensionality rules out the possibility of a compact encoding for gross substitutes (*cf.*, [Hatfield and Milgrom, 2005; Hatfield et al., 2012]):

Lemma 6.5.1. *The class of gross substitutes valuations is not succinct.*

Another research direction is compressed bid spaces for simple auction formats. In this context, the level of compression affects how close equilibria of these auctions can get to a welfare-maximizing allocation. See [Christodoulou et al., 2008; Bhawalkar and Roughgarden, 2011] for item bidding in combinatorial auctions, and [Dütting et al., 2013; Babaioff et al., 2014b] for recent extensions. [Feldman et al., 2015] further extend this line of work by considering posted price mechanisms.

6.5.2 Overview

We begin with a trivial "line in the sand:" when the pricing class is equal to the valuation class – i.e., there is no compression – a (non-anonymous) pricing equilibrium always exists.

Observation 6.5.2. *Consider a valuation class $\mathcal{V}$ and an identical pricing class $\mathcal{P} = \mathcal{V}$. For every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with pricings from $\mathcal{P}$.*

Proof. Given a market with valuations from $\mathcal{V}$, the following is a pricing equilibrium: Pick a welfare-maximizing allocation, and for every consumer i let his pricing p_i be equal to his valuation v_i. The consumers are indifferent among different bundles and so trivially maximize their payoff. The revenue is equal to the welfare and so is maximized by the allocation. □

[10] The gross substitutes class contains all weighted matroid rank functions. Even with constant-size weights and restricted to partition matroids, there are doubly-exponentially many of these [Knuth, 1974].

In fact, every welfare-maximizing allocation can be supported by pricings from the class $\mathcal{P} = \mathcal{V}$. If $\mathcal{V}$ is succinct, this characterizes when finding a verifiable pricing equilibrium is tractable.

Corollary 6.5.3. *Consider a succinct valuation class $\mathcal{V}$. There exists a polynomial time algorithm that finds a verifiable pricing equilibrium for every market with valuations from $\mathcal{V}$ if and only if the allocation problem for $\mathcal{V}$ is in* P.

One example of a tractable allocation problem is for the class of feature-based valuations with a constant number of features (see [Candogan and Pekec, 2014] and Section 6.6).[11]

Given the above observation, the main question in this section is: how much can pricings be compressed in comparison to valuations, without losing the ability to support efficient allocations? I.e., how much can the pricing class be shrunk with respect to the valuation class while still guaranteeing the existence of a pricing equilibrium?

We first present a simple necessary condition on the valuation and pricing classes (Proposition 6.5.4). Like Propositions 6.2.4 and 6.4.5, it is based on computational complexity considerations. We then address, in Section 6.5.3, the question of whether there are classes for which partial compression is possible. We show an example with a supporting pricing class of dimension strictly smaller than the valuation class, yet strictly larger than m item prices. Section 6.6 revisits these questions from an instance-by-instance perspective, and uses linear programming to characterize when succinct "linear" supporting pricings exist.

Proposition 6.5.4. *Consider a succinct valuation class $\mathcal{V}$ and a succinct pricing class $\mathcal{P}$ for which the revenue-maximization problem is tractable. If for every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with pricings from $\mathcal{P}$, then the allocation problem belongs to* coNP *whenever the demand problem belongs to* coNP.

The proof is similar to that of Proposition 6.4.5, and a corollary similar to Corollary 6.4.6 follows for valuation classes with demand oracle access.

[11] Feature-based pricings are of interest in practical settings, where the seller often needs to set prices based on a small set of summarizing features (*cf.*, [Dughmi et al., 2014]). Tractability follows from dynamic programming, utilizing the constant number of different item *types*.

Corollary 6.5.5. *Consider a valuation class $\mathcal{V}$ with demand oracle access and a succinct pricing class $\mathcal{P}$ for which the revenue-maximization problem is tractable. If for every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with pricings from $\mathcal{P}$, then the allocation problem for $\mathcal{V}$ is verifiable with demand oracle access.*

6.5.3 Partial Compression

A Rare Species

Despite the large literature on pricing equilibria for various valuation classes and pricing classes, we are unaware of any previously studied examples of classes $\mathcal{V}$ and $\mathcal{P}$ that meet the following criteria: (i) anonymous item prices are insufficient to support all markets with valuations in $\mathcal{V}$; (ii) the prices in $\mathcal{P}$ are sufficient to support all markets with valuations in $\mathcal{V}$; (iii) $\mathcal{P}$ is succinct and strictly smaller than $\mathcal{V}$ (where $\mathcal{V}$ may or may not be succinct); (iv) pricing equilibria for $\mathcal{V}$ and $\mathcal{P}$ are verifiable with demand oracle access (i.e., the revenue-maximization condition can be efficiently verified). In this sense, there are no known non-trivial generalizations of Walrasian equilibria!

Our complexity-theoretic methodology provides an explanation for the paucity of examples. Suppose the classes $\mathcal{V}$ and $\mathcal{P}$ satisfy (i)–(iv). By (ii)–(iv) and Corollary 6.5.5, the welfare-maximization problem for $\mathcal{V}$ is verifiable with demand oracle access, and this can be thought of as membership in coNP with demand oracle access. Because the problem can also be thought of as belonging to NP with demand oracle access, this suggests (but does not prove, of course) that the welfare-maximization problem can in fact be solved in polynomial time with demand oracle access. On the other hand, by (i), the configuration linear program fails to solve the welfare-maximization problem in polynomial time with demand oracle access (recall the proof of Proposition 6.2.4 and [Blumrosen and Nisan, 2007], Theorem 11.13). Thus, a non-trivial generalization of Walrasian equilibria in the above sense requires a novel polynomial-time algorithm for the welfare-maximization problem (unless such a problem belongs to (coNP $\cap$ NP) $\setminus \mathcal{P}$ with demand oracles)![12]

We remark that the same argument can be repeated for a succinct valuation class $\mathcal{V}$ for which pricing equilibria with $\mathcal{P}$ are verifiable.

An Example

We show there are (somewhat contrived) classes $\mathcal{V}$ and $\mathcal{P}$ that satisfy properties (i)–(iv) above. It is no coincidence that also (v) the welfare-maximization problem for $\mathcal{V}$ can be solved in polynomial time, but not directly by the configuration LP.

Example 6.5.6. We define a family of markets with two consumers, whose valuations belong to a slight variation of the gross substitutes class, which – like the original class – is not succinct. As we will show, the markets in this family are supported by a succinct pricing class but not by item prices.

Define the valuation class $\mathcal{V}$ as follows: There are two special items, without loss of generality items 1 and 2. Every valuation $v \in \mathcal{V}$ is equal to a gross substitutes valuation with the possible deviation that items 1 and 2 are allowed to complement each other, i.e., to be valued positively as a couple but as zero on their own. The following transformation of v must recover a gross substitutes valuation v': unify items 1 and 2 into a single item; for every bundle $S \subseteq M$, let S' be the same bundle after unifying items 1 and 2 and let $v'(S') = v(S)$. M' denotes the reduced item set.

Claim 6.5.7. *There exists a succinct pricing class $\mathcal{P}$ such that for every market with two consumers whose valuations belong to the non-succinct class $\mathcal{V}$ defined above, there are supporting pricings in $\mathcal{P}$. Moreover, there exists such a market for which there is no Walrasian equilibrium.*

Proof. If either both consumers have gross substitutes valuations, or both view items 1 and 2 as complements, item prices suffice by [Gul and Stacchetti, 1999] and [Sun and Yang, 2006], respectively. But item prices cannot stabilize every market in the described family – to see this, consider the classic example of a two-consumer two-item

[12] While there are of course other algorithms (like greedy algorithms) that compute a welfare-maximizing allocation in various special cases, in all known such cases the configuration LP also solves the problem exactly.

market, where the first consumer values the items as a pair at 3 and as individuals at 0, and the second consumer is unit-demand with value 2. We now show that there do exist *succinct* prices which stabilize every market in the described family.

Without loss of generality, assume the first consumer views items 1 and 2 as complements while the second consumer has a gross substitutes valuation. We use anonymous item prices for items $3, \ldots, m$ together with personalized item and bundle prices for the special items 1 and 2. Observe that by definition, both consumers have gross substitutes valuations with respect to M', and therefore there are stabilizing item prices for M'. Let the anonymous item prices of items $3, \ldots, m$ be these prices. Set the personalized prices as follows: Denote by $p_{1,2}$ the price of item $\{1, 2\}$ in M'. For the first consumer, set his prices for items 1 and 2 to zero, and his price for their bundle to $p_{1,2}$. For the second consumer, set his prices for items 1 and 2 as well as his price for their bundle to $p_{1,2}$. We claim that these prices stabilize every welfare maximizing allocation of the items in M: Using the fact that items 1 and 2 are valued at zero by consumer 1, we can assume that in every welfare maximizing allocation, either both are allocated to the first consumer or both are allocated to the second. The allocation thus corresponds to a welfare-maximizing allocation of M', which is stabilized by the prices above.

The final step of the proof is to give a lower bound on the size of $\mathcal{V}$. For this we show a simple injection from gross substitutes valuations over the item set M', to valuations in class $\mathcal{V}$ over the item set M. Given a gross substitutes valuation v' over M', define a valuation v over M by treating item 1 in M as the unified item in M', and setting the marginal value of item 2 to always be 0. The transformation of v by unifying items 1 and 2 returns v', showing that indeed $v \in \mathcal{V}$ and that v is distinct for every distinct v'. We have established that the dimension of $\mathcal{V}$, like the dimension of gross substitutes, is exponential in m, thus completing the proof.

□

6.6 A Note on Linear Pricing

Most of our results thus far have been impossibility results, ruling out the guaranteed existence of pricing equilibria for different classes of valuations and pricings. This section formulates the notion of a succinct linear valuation or pricing, which captures and generalizes many previously studied valuation classes. For this class of pricings, there is an instance-by-instance (rather than class-by-class) characterization of equilibrium existence based on linear programming – see [Roughgarden and Talgam-Cohen, 2015].

Definitions. Consider a set function f on the ground set of items M. A naïve representation is linear in the number of sets and exponential in the number of items m. We propose an alternative representation based on a set $\mathcal{L}$ of *pseudo-items*.

Definition 6.6.1. A *linear representation* is a mapping $L : 2^M \rightarrow 2^{\mathcal{L}}$ from bundles of items in M to bundles of pseudo-items in $\mathcal{L}$.

A linear representation can be seen as a bipartite graph with bundles of items on one side and individual pseudo-items on the other. Each bundle of items is connected to several pseudo-items. Now we can add weights to the pseudo-items. In this way a linear representation associates a value with each bundle of items: the total weight of the pseudo-items the bundle is connected to.

Definition 6.6.2. We say that a set function f on a ground set M has linear representation L if there exist weights $w_\ell \in \mathbb{R}$ for the pseudo-items such that for every bundle of items $S \subseteq M$, $f(S) = \sum_{\ell \in L(S)} w_\ell$.

Observe that every set function f has an exponentially large linear representation as follows: Let $\mathcal{L} = 2^M$, i.e., let there be a pseudo-item ℓ_S for every bundle of items S. For every S let $L(S) = \{\ell_S\}$ and let the corresponding weight be $w_{\ell_S} = f(S)$. Then $\sum_{\ell \in L(S)} w_\ell = f(S)$ as required. However, set functions can have multiple linear representations, and we are most interested in those that are *succinct*.

Definition 6.6.3. A family of linear representations is *succinct* if for every linear representation in the family, the set $\mathcal{L}$ of pseudo-items is of polynomial size in $m = |M|$, and the mapping L can be implemented by a polynomial time algorithm.

Succinct linear functions are functions that have succinct linear representations with polynomially-represented weights $\{w_\ell\}$.

Succinct Linear Valuations in the Literature. Succinct linear valuations capture many previously studied succinct valuation classes, including the following examples:[13]

1. General hypergraphical valuations, and their generalization to additively decomposable valuations [Candogan, 2013, Section 5.5], form a strict subset of succinct linear valuations.
2. Explicit coverage valuations [Dughmi et al., 2011].
3. Feature-based valuations [Candogan and Pekec, 2014].
4. Succinct endowed assignment valuations, proposed as a succinct subclass of gross substitutes by Hatfield and Milgrom [2005].
5. Budget-additive valuations [Lehmann et al., 2006].
6. XOS valuations with polynomially many clauses, and their generalization to maximum over polynomially many positive hypergraph valuations with rank k [Feige et al., 2014].
7. GGS(k, M) valuations with a constant k [Ben-Zwi et al., 2013].
8. Sketches of valuations [Badanidiyuru et al., 2012; Cohavi and Dobzinski, 2014].

Linear Program Characterization. In [Roughgarden and Talgam-Cohen, 2015] we formulate a linear program called LP1 given linear representations $\{L_i\}$, and show that:

[13] An example of a valuation class that is *not* captured is valuations that map bundles to the *product* of their items' values.

Corollary 6.6.4 (Characterization)**.** *A pricing equilibrium whose pricings have linear representations* $\{L_i\}$ *exists if and only if* LP1 *has an integral optimal solution that corresponds to a feasible allocation.*[14]

This is analogous to the configuration LP discussed above (see also Appendix A.3.1). In particular, the pricings correspond to an optimal dual solution.

Implications for Guaranteed Existence. Returning to our theme of necessary conditions for the guaranteed existence of pricing equilibria, we can use Corollary 6.6.4 to extend Proposition 6.2.4 to arbitrary succinct linear pricings.

Corollary 6.6.5. *Consider a valuation class* $\mathcal{V}$ *and a succinct linear pricing class* $\mathcal{P}$ *for which the revenue-maximization problem is tractable. If for every market with valuations from* $\mathcal{V}$ *there exists a pricing equilibrium with pricings from* $\mathcal{P}$*, then the allocation problem reduces in polynomial time to the demand problem.*

The proof follows that of Proposition 6.2.4, using the characterization in Corollary 6.6.4 in place of the classical linear programming characterization of the existence of Walrasian equilibria.

Corollary 6.6.5 differs from Proposition 6.5.4 in two respects: in the hypothesis, the pricing class $\mathcal{P}$ is assumed to be linear in addition to being succinct; and in the conclusion, the computation (rather than just the verification) of a welfare-maximizing allocation reduces to that of the demand problem. Equivalently, the two "coNP" terms of Proposition 6.5.4 are replaced in Corollary 6.6.5 by "P."

6.7 Chapter Conclusion

The well-studied problem of proving or disproving the guaranteed existence of pricing equilibria seems to have nothing to do with computation. As this work demonstrates, however, computational complexity offers numerous insights into the problem, and provides general techniques for proving impossibility results. For example, many

[14]The condition that the solution corresponds to a feasible allocation can alternatively be encoded into the linear program or the pricing.

(conditional) non-existence results for different types of pricing equilibria (Walrasian, anonymous, compressed, etc.) follow easily from the known computational complexity of various optimization problems, obviating the need for ad hoc explicit constructions without equilibria. Similarly, this methodology demystifies the dearth of useful extensions of the Walrasian equilibrium concept, by linking the existence of such extensions to algorithmic progress on the welfare-maximization problem.

7

Conclusion

7.1 Summary

Auctions and markets are economic tools for addressing informational issues in resource allocation, and real-world applications require them to handle computational issues. An informationally and computationally robust theory of auctions and markets is important for making accurate predictions on the performance of existing mechanisms, as well as for designing new ones or alternatively proving that such designs are inherently limited in their capabilities. The toolbox provided by theoretical computer science is highly useful in developing the required robust theory.

In particular, the Bayesian model helps counter the lack of information when maximizing revenue, but in many realistic settings with multiple goods or interdependence among the buyers, the model's common knowledge assumption leads to complicated, implausible or unilluminating mechanisms. A robust approach leads to simplified and realistic designs based on welfare maximization with respect to true or transformed values. The gained insights can serve as specific market design guidelines, and to encourage robustness as a design approach in complicated informational settings.

Similarly, the substitutes assumption helps avoid computational problems, but is quite limiting and in many cases turns out to be overly strong: positive guarantees can be proven beyond substitutes, in terms of both computer science approximations and economic properties like group strategyproofness. This opens the possibility for new applications, such as double auctions with feasibility constraints, and also demonstrates that studying the computational aspects of substitutes and complements beyond combinatorial auctions is worthwhile. On the other hand we show, in the context of competitive market equilibrium, that relaxing the substitutes assumption is not always possible; the methodology we use ties equilibrium existence to the classic P versus NP problem.

7.2 Future Directions

We conclude with a non-exhaustive list of directions for future research. First, our prior-independent designs closely approximate the best possible mechanisms with knowledge of the prior distributions, but are not necessarily the best possible without such knowledge:

Open Problem (Optimal Prior Independence)**.** What is the revenue-maximizing prior-independent auction in different settings? Given x additional bidders? Given y additional samples? Are there standard auction formats that are revenue-maximizing subject to prior-independence?

For preliminary work in this direction – a proof-of-concept result showing a slight improvement over the result of Bulow and Klemperer [1996] – see [Fu et al., 2015]. Cole and Roughgarden [2014] study a related question, asking how many samples of each distribution are needed to achieve a desired approximation parameter. In the domain of contracts rather than auctions, Carroll [2015] shows that a linear format is optimal subject to robustness. Other prior independence questions are whether it is possible to achieve similar results to those in Chapter 3 with respect to the Bayesian truthfulness solution concept rather than the dominant strategy one, and whether our matching market techniques also apply to markets with general gross substitutes valuations and/or positive correlation among the buyers.

We view interdependent values as a new frontier for algorithmic mechanism design, raising the following problem:

Open Problem (Interdependence)**.** Develop a complete algorithmic theory of optimal mechanism design for interdependent values, similar to the existing theory for independent values.

Towards this goal, one direction is the study of approximately revenue-maximizing auctions (with access to prior distributions) as in [Dobzinski et al., 2011]. Chawla et al. [2014] design such auctions under a relaxed single-crossing condition that does not constrain the distributions. Another line of inquiry focuses on optimality rather than approximation, and seeks minimal conditions under which a meaningful description of the revenue-maximizing mechanism can be found.[1] For example, under which conditions does a Myerson-like ironing method maximize the expected revenue? What can be achieved by using randomness to circumvent the computational hardness result for deterministic revenue maximization [Papadimitriou and Pierrakos, 2015]? What can be achieved for natural subclasses of interdependent values such as *conditionally independent* values?[2]

The algorithmic approach coupled with a Myerson-like theory of revenue maximization for interdependent values holds potential for many applications. As one example, good posted-price mechanisms are known for independent values [Chawla et al., 2007], can they be achieved for interdependent values? Interdependent values are also an active research area beyond single-seller auctions, e.g. in the context of double auctions [Satterthwaite et al., 2011], housing [Che et al., 2015] or public goods [Csapó and Müller, 2013], and the algorithmic approach can potentially contribute to such applications as well.

It is also interesting to consider alternative solution concepts to ex post truthfulness, which are either more robust (e.g., rejecting common knowledge of the valuation functions rather than just the signals), or less so (but falling short of the Crémer and McLean [1985, 1988] solution); this is related to different levels of robustness, justified by considerations ranging from risk-averseness to the practical difficulty of collecting payments from losing bidders [Lopomo, 2000].

[1] For that matter, the welfare-maximizing mechanism is also not fully understood.

[2] In the mineral rights example, values are conditionally independent since they are interdependent only through their dependence on the presence or absence of precious minerals.

In Part II of the dissertation we discuss the potential for fruitful research on substitutes and complements beyond combinatorial auctions.

In Chapter 5 we discuss limited complements in the context of feasibility constraints in double auction settings. One open problem is to extend our notion of limited complements to encompass *cross-market* feasibility constraints, which involve both buyers and sellers. Another open problem is whether our notion captures limited complements for which the double auction designs can get *arbitrarily* close to optimal welfare. In other words, our examples do not demonstrate limited complements for which the greedy approach achieves a $(1+\epsilon)$-approximation; is this a shortcoming of our examples or of the limited complements notion, and is it the best possible in this sense? A third (and related) open problem is to extend our results to multi-parameter valuations. This brings us to the following question:

Open Problem (Substitutes)**.** What is an appropriate notion of limited complements in the context of multi-parameter valuations?

A similar question has been studied in combinatorial auctions [Abraham et al., 2012]. Beyond combinatorial auctions, two measures of "appropriateness" are the ability to design auctions for such valuations that share the properties of VCG with substitutes, as well as the existence of meaningful market equilibria. The following related question arises from Chapter 6:

Open Problem (Market Equilibrium)**.** Is there a meaningful market equilibrium or approximate equilibrium notion that is guaranteed to exist for limited complements (not exclusively for very large markets)?

In combinatorial auctions, a hierarchy of valuation classes has been developed and closely studied; central classes in this hierarchy are depicted in Figure 1.1. At the base of the hierarchy is the class of submodular valuations [Lehmann et al., 2006]. Classes of valuations that are "limited in their non-submodularity", or in other words, are close to the class of submodular valuations but above it in the hierarchy, have also been studied [Feige and Izsak, 2013, and subsequent work]. But are limited complements constrained by what is possible for a base class of submodular

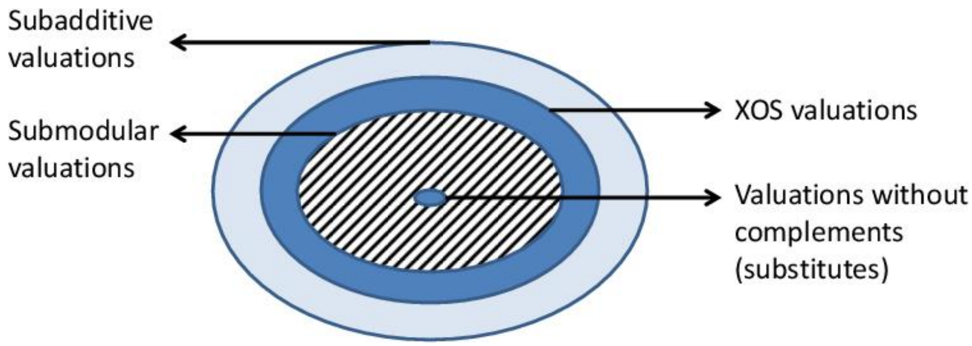

Figure 7.1: Hierarchy of valuation classes. The classes of submodular, XOS and subadditive valuations are well-studied [Blumrosen and Nisan, 2007]. The striped area represents "uncharted territory" of valuations.

valuations? For combinatorial auctions this is not the case [Dughmi et al., 2011]. The above open problems thus call to explore subclasses of submodular valuations, as depicted by the striped area in Figure 7.1.[3]

More broadly, the main question that this dissertation contributes to is:

Open Problem. How can we harness tools from computer science to design better markets?

Economic transactions are driven by computing more than ever before – whether these transactions take place in the realm of e-commerce, government auctions or the sharing economy. These domains and others all rely on the combination of economics, operations and computation to optimize resource allocation in our society. Therefore we believe that this research agenda holds huge theoretical and societal opportunity.

[3]Budget-additive valuations, mentioned in Chapter 6, are an example of a strict subclass of submodular valuations. Note however that they do not encompass gross substitutes.

Appendix A

Supplementary Material

A.1 Appendix for Chapter 1

We include here two missing algorithms and a missing proof from Chapter 1.

Algorithm 10: Greedy welfare maximization subject to feasibility for single-parameter settings.

Input: The n buyers' values for winning
Set $W = \emptyset$;   % Winning buyers
while $|W| < n$ **do**
    F = subset of buyers who can each be added to W while maintaining feasibility of W;
    i = buyer in F that has maximum marginal welfare contribution given W;
    if *marginal welfare contribution of* i *is* ≤ 0 **then**
        exit;
    else
        $W = W \cup \{i\}$;
    end
end

Proof of Proposition 1.4.1. Let N be the buyer set and let $(N, \mathcal{I})$ be the downward-closed feasibility set system corresponding to the feasible winner sets (i.e., $W \in \mathcal{I}$ if and only if W is a feasible winner set). Recall from Remark 1.4.3 that Algorithm 10 maximizes welfare for all values if and only if $(N, \mathcal{I})$ is a matroid [Rado, 1957;

Algorithm 11: Greedy utility maximization (a.k.a. demand oracle) for a buyer.

Input: The m goods' prices
Set $B = \emptyset$; % Bundle
while $|B| < m$ **do**
  $j =$ good that has maximum marginal utility contribution given B;
  if *marginal utility contribution of* j *is* ≤ 0 **then**
    exit;
  else
    $B = B \cup \{j\}$;
  end
end

Gale, 1968; Edmonds, 1971]. If the buyers are substitutes for all values, consider a value of 1 for each buyer. Due to its submodularity, the coalitional value function is a matroid rank function, and so $(N, \mathcal{I})$ is a matroid. For the other direction, if $(N, \mathcal{I})$ is a matroid, then for any values the coalitional value function is equal to the submodular weighted rank function (see, e.g., [Welsh, 2010]). □

A.2 Appendix for Chapter 5

A.2.1 Ranking Algorithm for Knapsacks

We show that the greedy algorithm can be used to achieve a uniform $1/(1-\lambda)\mu$-approximation, where $\lambda \leq 1$ is the ratio of the largest element's size to the knapsack size, and $\mu \leq 1$ is a lower bound on the ratio of the smallest element's size to the largest element's size, such that $1/\mu$ is integral.

Proposition A.2.1. *The ranking algorithms for knapsack based on the greedy by weight one-sided algorithm are a uniform* $\frac{1}{(1-\lambda)\mu}$*-approximation.*

Proof. The proof is by induction on the maximum allowed number k of elements in the knapsack. For every $k \geq 0$ and for every knapsack instance Q, denote by A_k^Q the solution for Q of size at most k found by the greedy algorithm, and by O_k^Q the optimal solution for Q of size at most k. Fix $k > 0$. Induction hypothesis: for every instance Q it holds that $w(A_{k-1}^Q)/(1-\lambda)\mu \geq w(O_{k-1}^Q)$, where $w(A_{k-1}^Q)$ and $w(O_{k-1}^Q)$

denote the total weight of solutions A_{k-1}^Q and O_{k-1}^Q. The hypothesis is easy to verify for $k = 1$.

Consider a knapsack instance Q' in which all elements fit into the knapsack. Without loss of generality we assume that sizes are normalized such that the size of the knapsack is 1, and that the elements are ordered from high to low weight. Thus the first step of the greedy algorithm is to place element 1 into the knapsack. We define a residual instance Q where the allowed number of elements is $k - 1$, and the size of the knapsack decreases by the size of element 1. In addition, the elements whose sizes are larger than the residual knapsack are removed from the residual element set. By the greediness of the algorithm, we know that $w(A_k^{Q'}) = w(A_{k-1}^Q) + w_1$.

Assume that the optimal solution to Q' with up to k elements does not include element 1, by how much is $w(O_{k-1}^Q)$ decreased relative to $w(O_k^{Q'})$ due to placing element 1 in the knapsack? There are two sources of loss. First, element 1 takes up place in the knapsack according to its size, and thus excludes elements from the optimal solution. Second, there may be elements removed from the residual element set since they no longer fit into the knapsack, so that the element sets available to $O_k^{Q'}$ and to O_{k-1}^Q may differ.

We start from the second source of loss. Observe that the minimum size of an element removed from the residual element set is $1 - \lambda$, and the maximum weight of such an element is w_1. There are at most $1/(1 - \lambda)$ such elements in the optimal solution to Q'. If their aggregate size is at least the size of element 1, removing them from the optimal solution also makes room for element 1, and so in this case

$$w(O_{k-1}^Q) \geq w(O_k^{Q'}) - w_1 \frac{1}{1-\lambda} \geq w(O_k^{Q'}) - w_1 \frac{1}{(1-\lambda)\mu}.$$

On the other hand, if the aggregate size of elements to remove from the residual set is less than the size of element 1, then there are at most $\lambda/(1 - \lambda)$ such elements (using that by normalization, the size of element 1 is at most λ). In addition, there is the first source of loss, from excluding a maximum number of $1/\mu$ additional elements from the optimal solution to make room for element 1 (we use here the assumption that $1/\mu$ is integral). Thus, in this case,

$$w(O_{k-1}^{Q}) \geq w(O_k^{Q'}) - w_1 \left(\frac{1}{\mu} + \frac{\lambda}{1-\lambda}\right) \geq w(O_k^{Q'}) - w_1 \frac{1}{(1-\lambda)\mu},$$

where in the second inequality we used that $\lambda \geq \lambda\mu$.

Putting everything together and using the induction assumption we get that

$$w(O_k^{Q'}) \leq w(O_{k-1}^{Q}) + w_1 \frac{1}{(1-\lambda)\mu} \leq w(A_{k-1}^{Q}) \frac{1}{(1-\lambda)\mu} + w_1 \frac{1}{(1-\lambda)\mu} \leq w(A_k^{Q'}) \frac{1}{(1-\lambda)\mu},$$

thus verifying the hypothesis for k and completing the proof. □

A.2.2 Backward-Feasible Ranking Algorithm for Matchings

The design of Algorithm 12 follows that of Drake and Hougardy [2003], with the exception that our algorithm is randomized. It starts with an arbitrary node of the bipartite graph, and grows a path of locally-heaviest edges. When the path cannot be extended any further, it restarts this process at an arbitrary node. In the end it takes all even edges along the paths with probability $1/2$, and all odd edges otherwise, and returns then as a matching.

To turn Algorithm 12 into a backward-feasible ranking algorithm, all that is needed is to rank all edges that do not belong to the output matching first; once the remaining edges form a feasible matching (that is precisely the output matching), continue to rank edges in reverse order of their weight (for the sake of consistency).

Proposition A.2.2. *The backward-feasible ranking algorithm based on the path growing algorithm (Algorithm 12) is consistent, rank monotone, and a uniform 2-approximation.*

Proof. The ranking algorithm is consistent because it outputs the elements of the set M_i that has been picked in order of their weights. It is rank monotone because by increasing its bid a player can only enter and not drop out of either M_1 or M_2.

For the approximation guarantee we assign each edge to some node in the graph in the following way. Whenever a node is removed, all edges that are currently incident to that node are assigned to it. To prove the factor 2, we consider an optimal solution

Algorithm 12: Path growing algorithm for matchings

Input: Graph $G = (V, E)$, weights $w(e) \geq 0$ for all edges $e \in E$
Output: Matching M
Set $M_1 = \emptyset$, $M_2 = \emptyset$, $i = 1$;   % Tentative matchings and an odd/even indicator
while $E \neq \emptyset$ **do**
  Choose $x \in V$ of degree at least 1 arbitrarily;
  while *x has a neighbor* **do**
    Let (x, y) be the heaviest edge incident to x;
    Add (x, y) to M_i;
    Set $i = 3 - i$;   % Switch between 1 (odd) and 2 (even)
    Remove x and incident edges from G;
    Set $x = y$;   % Next node in path
  end
end
Output M_1 with probability 1/2, otherwise output M_2;

of cardinality k. Each of these edges is assigned to a node. If we consider the edges adjacent to these nodes that were added to $M_1 \cup M_2$, then from the fact that we picked the locally heaviest edges we know that their total weight is at least the weight of the optimal edges. The claim now follows from the fact that we pick each of these edges (or a better one) with probability 1/2. □

To implement as a deferred-acceptance algorithm, the randomization can be implemented by tossing a coin at the beginning of the algorithm to choose between M_1 and M_2. The removal of edges during the path growing part of the ranking algorithm can be implemented as described in Dütting et al. [2014a]. The coin toss then determines whether to remove M_1 or M_2. Once the set of active edges becomes feasible, we can continue to score by weight in order to maintain consistency.

A.3 Appendix for Chapter 6

A.3.1 The Configuration LP

Consider a market. Let $x_{i,S}$ be an indicator of whether buyer i receives bundle S; we can write down an integer program for maximizing social welfare with variables

$\{x_{i,S}\}_{i,S}$. Relaxing the variables to be $\in [0,1]$, we get the configuration linear program:

$$\begin{aligned}
&\max \sum_{i,S} x_{i,S} v_i(S) \\
\text{s.t.} \quad & \sum_{i,S|j\in S} x_{i,S} \leq 1 \qquad \forall \text{ item } j && \text{(A.1)}\\
& \sum_{S} x_{i,S} \leq 1 \qquad \forall \text{ buyer } i && \text{(A.2)}\\
& x_{i,S} \geq 0
\end{aligned}$$

Denote the dual variables by p_j (corresponding to Constraint A.1, interpreted as the price of item j), and π_i (corresponding to Constraint A.2, interpreted as the utility of buyer i). The dual of the configuration LP is then:

$$\begin{aligned}
&\min \sum_{j} p_j + \sum_{i} \pi_i \\
\text{s.t.} \quad & \pi_i \geq v_i(S) - \sum_{j\in S} p_j \quad \forall \text{ buyer } i, \text{ bundle } S && \text{(A.3)}\\
& p_j \geq 0, \ \pi_i \geq 0
\end{aligned}$$

For completeness we include the proof of the following fundamental fact (see, e.g., [Blumrosen and Nisan, 2007, Theorem 11.13]):

Proposition A.3.1. *A Walrasian equilibrium exists if and only if the configuration LP has an optimal integral solution.*

Proof. First assume that a Walrasian equilibrium exists. Denote its integral allocation by S^*, and observe that it induces a feasible dual solution $\{p_j^*\}$, $\{\pi_i^*\}$. We claim that the welfare of this allocation exceeds that of any *fractional* allocation $\{x_{i,S}\}$ that is a feasible solution to the configuration LP; and so it induces an optimal integral solution to the configuration LP. Proof of claim: By Constraints A.2 (for $\{x_{i,S}\}$) and A.3 (for S^*, $\{\pi_i^*\}$), we get that for every buyer i,

$$v_i(S_i^*) - \sum_{j\in S_i^*} p_j^* \geq \sum_{S} x_{i,S}(v_i(S) - \sum_{j\in S} p_j^*).$$

Summing up over all buyers we get that

$$\sum_i v_i(S_i^*) - \sum_i \sum_{j\in S_i^*} p_j^* \geq \sum_{i,S} x_{i,S} v_i(S) - \sum_{i,S} x_{i,S} \sum_{j\in S} p_j^*. \tag{A.4}$$

We now invoke Constraint A.1 (for $\{x_{i,S}\}$), by which

$$\sum_i \sum_{j\in S_i^*} p_j^* = \sum_j p_j^* \geq \sum_{i,S} x_{i,S} \sum_{j\in S} p_j^*,$$

(the equality holds since prices of unallocated items in a Walrasian equilibrium are zero). We conclude that in order for Inequality A.4 to hold, it must be that $\sum_i v_i(S_i^*) \geq \sum_{i,S} x_{i,S} v_i(S)$.

In the other direction, if the configuration LP has an optimal integral solution, then there exists an integral allocation and item prices such that every buyer's utility is maximized. To establish the existence of a Walrasian equilibrium it's left to show that the price of every unallocated item is 0. This follows from complementary slackness and Constraint A.1. □

A.3.2 Computational Complexity

For the non-computer-scientist reader, here is a brief review of computational complexity terms, applied to resource allocation problems in Chapter 6. For a detailed exposition see [Arora and Barak, 2009].

An instance of a computational problem is parameterized by its input size, and complexity is measured with respect to this parameter. The class P contains computational problems that can be solved in polynomial time in the input size. The classes NP and coNP both contain P. It is widely believed that P $\neq$ NP and thus both containments are strict. Informally, these classes include "prove/disprove" (decision) problems for which there are short *verifiable* proofs if the problem is in NP, or short verifiable disproofs if it is in coNP. It is also widely believed that NP $\neq$ coNP, i.e., there are problems with short verifiable proofs but not short verifiable disproofs and vice versa. The NP $\neq$ coNP assumption implies P $\neq$ NP but not vice versa.

An important tool in the study of computational complexity is *reductions* among computational problems and the notion of *hardness*. Informally, a problem is hard for a class if it captures the computational complexity of all problems in the class. More formally, all problems in the class can be reduced (transformed) in polynomial time to the hard problem. Thus, a reduction from an NP-hard problem to a problem in P contradicts the P $\neq$ NP assumption, and the existence of short disproofs for an NP-hard problem contradicts the NP $\neq$ coNP assumption.

Bibliography

I. Abraham, M. Babaioff, S. Dughmi, and T. Roughgarden. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Economics and Computation*, pages 3–16, 2012.

I. Abraham, S. Athey, M. Babaioff, and M. Grubb. Peaches, lemons, and cookies: Designing auction markets with dispersed information. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 7–8, 2014. Abstract.

G. Aggarwal and J. D. Hartline. Knapsack auctions. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1083–1092, 2006.

S. Alaei, H. Fu, N. Haghpanah, and J. D. Hartline. The simple economics of approximately optimal auctions. In *Proceedings of the 54th Symposium on Foundations of Computer Science*, pages 628–637, 2013.

N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley, 2008.

A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. 42nd Symposium on Foundations of Computer Science*, pages 482–491, 2001.

S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

K. J. Arrow and F. H. Hahn. *General Competitive Analysis*. Holden Day, 1971.

L. Ausubel and P. R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1):article 1, 2002.

L. M. Ausubel and P. R. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 57–95. MIT Press, Boston, MA, USA, 2006.

P. Azar, C. Daskalakis, S. Micali, and S. M. Weinberg. Optimal and efficient parametric auctions. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 596–604, 2013.

P. D. Azar, R. Kleinberg, and S. M. Weinberg. Prophet inequalities with limited information. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1358–1377, 2014.

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

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

M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Proc. 55th Symposium on Foundations of Computer Science*, pages 21–30, 2014a.

M. Babaioff, B. Lucier, N. Nisan, and R. P. Leme. On the efficiency of the Walrasian mechanism. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 783–800, 2014b.

A. Badanidiyuru, S. Dobzinski, H. Fu, R. Kleinberg, N. Nisan, and T. Roughgarden. Sketching valuation functions. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1025–1035, 2012.

M.-F. Balcan and N. J. A. Harvey. Learning submodular functions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 793–802, 2011.

C. Y. Baldwin and K. B. Clark. *Design Rules: The Power of Modularity*. MIT Press, 2000.

S. Baliga and R. Vohra. Market research and market design. *Advances in Theoretical Economics*, 3(1):article 5, 2003.

C. Bandi and D. Bertsimas. Optimal design for multi-item auctions: A robust optimization approach. *Mathematics of Operations Research*, 39(4):1012–1038, 2014.

O. Ben-Zwi, R. Lavi, and I. Newman. Ascending auctions and Walrasian equilibrium. Working paper, 2013.

D. Bergemann and S. Morris. Robust mechanism design. *Econometrica*, 73(6):1771–1813, 2005.

A. Bertelsen. Substitutes valuations and M-sharp concavity. Master's thesis, Hebrew University of Jerusalem, October 2004.

D. P. Bertsekas. *Linear Network Optimization*. MIT Press, 1991.

D. Bertsimas and A. Thiele. *Robust and Data-Driven Optimization: Modern Decision Making Under Uncertainty*, chapter 5, pages 95–122. INFORMS PubsOnline, 2014. TutORials in Operations Research.

K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–709, 2011.

S. Bikhchandani and J. W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2):385–413, 1997.

S. Bikhchandani and J. M. Ostroy. The package assignment model. *Journal of Economic Theory*, 107(2):377–406, 2002. Following 1998 technical report.

S. Bikhchandani, P. A. Haile, and J. G. Riley. Symmetric separating equilibria in English auctions. *Games and Economic Behavior*, 38:19–27, 2002.

S. Bikhchandani, S. de Vries, J. Schummer, and R. V. Vohra. An ascending Vickrey auction for selling bases of a matroid. *Operations Research*, 59(2):400–413, 2011.

L. Blumrosen and S. Dobzinski. Reallocation mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 617–617, 2014. Abstract.

L. Blumrosen and N. Nisan. Combinatorial auctions. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 11. Cambridge University Press, 2007.

A. Borodin and B. Lucier. Price of anarchy for greedy auctions. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 537–553, 2010.

A. Borodin, M. N. Nielsen, and C. Rackoff. (Incremental) priority algorithms. *Algorithmica*, 37(4):295–326, 2003.

C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

F. Branco. Multiple unit auctions of an indivisible good. *Economic Theory*, 8:77–101, 1996.

J. Bredin and D. C. Parkes. Models for truthful online double auctions. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence*, pages 50–59, 2005.

J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.

Y. Cai, C. Daskalakis, and M. Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *Proc. 53rd Symposium on Foundations of Computer Science*, pages 130–139, 2012.

Y. Cai, C. Daskalakis, and S. M. Weinberg. Understanding incentives: Mechanism design becomes algorithm design. In *Proc. 54th Symposium on Foundations of Computer Science*, pages 618–627, 2013.

O. Candogan. *Dynamic Strategic Interactions: Analysis and Mechanism Design*. PhD thesis, Massachusetts Institute of Technology, June 2013.

O. Candogan and S. Pekec. Efficient iterative auctions for multi-featured items: A network flow approach. In submission, 2014.

O. Candogan, A. Ozdaglar, and P. Parrilo. Iterative auction design for graphical valuations part II: General graphs. In submission, 2014.

O. Candogan, A. Ozdaglar, and P. Parrilo. Iterative auction design for tree valuations. To appear in Operations Research, 2015.

G. Carroll. The efficiency-incentive tradeoff in double auction environments. Working paper, 2013.

G. Carroll. Robustness and linear contracts. *American Economic Review*, 105(2): 536–563, 2015.

L. E. Celis, G. Lewis, M. M. Mobius, and H. Nazerzadeh. Buy-it-now or take-a-chance: Price discrimination through randomized auctions. *Management Science*, 60(12):2927–2948, 2014.

S. Chawla, J. D. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *Proc. 9th ACM Conference on Economics and Computation*, pages 243–251, 2007.

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Sequential posted pricing and multi-parameter mechanism design. In *Proc. 41st Annual ACM Symposium on Theory of Computing*, pages 311–320, 2010a.

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. The power of randomness in Bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2010b.

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Prior-independent mechanisms for scheduling. In *Proc. 44th Annual ACM Symposium on Theory of Computing*, pages 51–60, 2013.

S. Chawla, H. Fu, and A. Karlin. Approximate revenue maximization in interdependent settings. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 277–294, 2014.

Y.-K. Che, J. Kim, and F. Kojima. Efficient assignment with interdependent values. *Journal of Economic Theory*, 158:54–86, 2015.

X. Chen, D. Dai, Y. Du, and S.-H. Teng. Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities. In *Proceedings of the 50th Symposium on Foundations of Computer Science*, pages 273–282, 2009a.

X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009b.

X. Chen, G. Hu, P. Lu, and L. Wang. On the approximation ratio of k-lookahead auction. In *Proceedings of the 7th International Workshop on Internet and Network Economics*, pages 61–71, 2011.

G. Christodoulou, A. Kovacs, and M. Schapira. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, pages 820–832, 2008.

L. Y. Chu. Truthful bundle/multiunit double auction. *Management Science*, 55(7): 1184–1198, 2009.

K.-S. Chung and J. C. Ely. Ex-post incentive compatible mechanism design. Manuscript, 2006.

K.-S. Chung and J. C. Ely. Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476, 2007.

E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

K. Cohavi and S. Dobzinski. Faster and simpler sketches of valuation functions. Working paper, 2014.

R. Cole and T. Roughgarden. The sample complexity of revenue maximization. In *Proc. 45th Annual ACM Symposium on Theory of Computing*, pages 243–252, 2014.

O. Compte and P. Jehiel. Veto constraint in mechanism design: Inefficiency with correlated types. *American Economic Journal: Microeconomics*, 1(1):182–206, 2009.

V. Conitzer, T. Sandholm, and P. Santi. Combinatorial auctions with k-wise dependent valuations. In *Proceedings of the 20th AAAI Conference on Artificial Intelligence*, pages 248–254, 2005.

P. C. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial auctions*. MIT Press, 2006.

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

J. Crémer and R. P. McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.

M. W. Cripps and J. M. Swinkels. Efficiency of large double auctions. *Econometrica*, 74(1):47–92, 2006.

G. Csapó and R. Müller. Optimal mechanism design for the private supply of a public good. *Games and Economic Behavior*, 80:229–242, 2013.

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

L. I. de Castro. Affiliation, equilibrium existence and revenue ranking of auctions. Working paper, 2010.

L. I. de Castro and H. J. Paarsch. Testing affiliation in private-values models of first-price auctions using grid distributions. *Annals of Applied Statistics*, 4(4):2073–2098, 2010.

G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *The Journal of Political Economy*, 94(4):863–872, 1986.

K. Deshmukh, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Truthful and competitive double auctions. In *Proceedings of the 11th European Symposium on Algorithms*, pages 361–373, 2002.

N. R. Devanur, J. D. Hartline, A. R. Karlin, and C. T. Nguyen. Prior-independent multi-parameter mechanism design. In *Proceedings of the 7th International Workshop on Internet and Network Economics*, pages 122–133, 2011.

P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

S. Dobzinski, H. Fu, and R. Kleinberg. Optimal auctions with correlated bidders are easy. In *Proc. 42nd Annual ACM Symposium on Theory of Computing*, pages 129–138, 2011.

D. E. Drake and S. Hougardy. A simple approximation algorithm for the weighted matching problem. *Information Processing Letters*, 85(4):211–213, 2003.

A. Dress and W. Terhalle. Rewarding maps: On greedy optimization of set functions. *Advances in Applied Mathematics*, 16(4):464–483, 1995.

S. Dughmi, T. Roughgarden, and Q. Yan. From convex optimization to randomized mechanisms: Toward optimal combinatorial auctions for submodular bidders. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 149–158, 2011.

S. Dughmi, T. Roughgarden, and M. Sundararajan. Revenue submodularity. *Theory of Computing*, 8(1):95–119, 2012.

S. Dughmi, N. Immorlica, and A. Roth. Constrained signaling in auction design. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1341–1357, 2014.

P. Dütting, F. A. Fischer, and D. C. Parkes. Simplicity-expressiveness tradeoffs in mechanism design. In *Proc. 12th ACM Conference on Economics and Computation*, pages 341–350, 2011.

P. Dütting, M. Henzinger, and M. Starnberger. Valuation compressions in VCG-based combinatorial auctions. In *Proceedings of the 9th International Workshop on Internet and Network Economics*, pages 146–159, 2013.

P. Dütting, V. Gkatzelis, and T. Roughgarden. The performance of deferred-acceptance auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 187–204, 2014a.

P. Dütting, T. Roughgarden, and I. Talgam-Cohen. Modularity and greed in double auctions. In *Proc. 15th ACM Conference on Economics and Computation*, pages 241–258, 2014b.

J. R. Edmonds. Matroids and the greedy algorithm. *Mathematical Programming*, 1: 127–136, 1971.

U. Feige and R. Izsak. Welfare maximization and the supermodular degree. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, pages 247–256, 2013.

U. Feige, M. Feldman, N. Immorlica, R. Izsak, B. Lucier, and V. Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 872–878, 2014.

M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 123–135, 2015.

F. A. Fischer, M. Holzer, and S. Katzenbeisser. The influence of neighbourhood and choice on the complexity of finding pure Nash equilibria. *Information Prrocessing Letters*, 99(6):239–245, 2006.

H. Fu, N. Immorlica, B. Lucier, and P. Strack. Randomization beats second price as a prior-independent auction. In *Proc. 16th ACM Conference on Economics and Computation*, page 323, 2015.

D. Fudenberg, M. Mobius, and A. Szeidl. Existence of equilibrium in large double auctions. *Journal of Economic Theory*, 133(1):550–567, 2007.

D. Gale. Optimal assignments in an ordered set: An application of matroid theory. *J. Combin. Theory*, 4:176–180, 1968.

D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.

M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. 1979.

A. Goldberg, J. D. Hartline, A. R. Karlin, M. E. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

M. Gonen, R. Gonen, and E. Pavlov. Generalized trade reduction mechanisms. In *Proceedings of the 7th ACM Conference on Economics and Computation*, pages 20–29, 2007.

Google. Artificial intelligence and machine learning: What we do. `http://research.google.com/pubs/ArtificialIntelligenceandMachineLearning.html`, accessed May, 2015.

J. Green and J. J. Laffont. Characterization of satisfactory mechanism for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.

S. Hart and N. Nisan. The menu-size complexity of auctions. In *Proc. 14th ACM Conference on Economics and Computation*, pages 565–566, 2013. Extended abstract.

J. D. Hartline. Approximation in economic design. Book manuscript, 2014.

J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 10th ACM Conference on Economics and Computation*, pages 225–234, 2009.

J. D. Hartline and T. Roughgarden. Optimal platform design. Working paper, 2014.

J. D. Hartline, V. S. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th International Conference on World Wide Web*, pages 189–198, 2008.

J. W. Hatfield and P. R. Milgrom. Matching with contracts. *American Economic Review*, 95(4):913–935, 2005.

J. W. Hatfield, N. Immorlica, and S. D. Kominers. Testing substitutability. *Games and Economic Behavior*, 75(2):639–645, 2012.

B. Holmstrom. Groves scheme on restricted domains. *Econometrica*, 47:1137–1144, 1979.

L. Hurwicz. On informationally decentralized systems. In C. B. McGuire and R. Radner, editors, *Decision and Organization*, pages 297–336. University of Minnesota Press, Minneapolis, MN, USA, 1972.

P. Jehiel, M. Meyer-ter-Vehn, B. Moldovanu, and W. Zame. The limits of ex-post implementation. *Econometrica*, 74(3):585–610, 2006.

R. Juarez. Group strategyproof cost sharing: The role of indifferences. *Games and Economic Behavior*, 82:218–239, 2013.

A. Kelso and V. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

R. Kirkegaard. A short proof of the Bulow-Klemperer auctions vs. negotiations result. *Journal of Economic Theory*, 28(2):449–452, 2006.

J. Kleinberg and E. Tardos. *Algorithm Design*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2005.

P. Klemperer. Auctions with almost common values: The "wallet game" and its applications. *European Economic Review*, 42(3):757–769, 1998.

P. Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13:227–286, 1999.

D. E. Knuth. The asymptotic number of geometries. *Journal of Combinatorial Theory (A)*, 16:398–400, 1974.

F. Kojima and T. Yamashita. Double auction with interdependent values: Incentives and efficiency. Working paper, 2013.

B. Korte and D. Hausmann. An analysis of the greedy heuristic for independence systems. *Annals of Discrete Mathematics*, 2:65–74, 1978.

V. Krishna. *Auction Theory*. Academic Press, second edition, 2010.

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

D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577602, 2002.

K. Leyton-Brown. The viability of exact feasibility checking, 2013. Talk at Conference on the Design of the U.S. Incentive Auction, available at `http://www.cs.ubc.ca/~kevinlb/publications.html`.

K. Leyton-Brown. Feasibility checking for spectrum repacking: Methodologies and test results, 2014. Talk at FCC Workshop, available at `http://www.cs.ubc.ca/~kevinlb/publications.html`.

Y. Li. Approximation in mechanism design with interdependent values. In *Proc. 14th ACM Conference on Economics and Computation*, pages 675–676, 2013.

G. Lopomo. Optimality and robustness of the English auction. *Games and Economic Behavior*, 36:219–240, 2000.

L. Lovász and M. D. Plummer. *Matching Theory*. American Mathematical Society, 2009.

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

R. P. McAfee and P. J. Reny. Correlated information and mechanism design. *Econometrica*, 60(2):395–421, 1992.

W. A. McEachern. *Economics: A Contemporary Introduction*. Cengage Learning, 2012.

A. Mehta, T. Roughgarden, and M. Sundararajan. Beyond Moulin mechanisms. *Games and Economic Behavior*, 67(1):125–155, 2009.

P. R. Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(2):245–272, 2000.

P. R. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

P. R. Milgrom. Prices and decentralization without convexity, 2014. Arrow Lecture, Columbia University, available at `http://cgeg.sipa.columbia.edu/events-calendar/7th-kenneth-j-arrow-lecture-prices-and-decentralization-without-convexity`.

P. R. Milgrom and I. Segal. Deferred-acceptance auctions and radio spectrum reallocation. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 185–186, 2014.

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

A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008.

R. B. Myerson. Incentive-compatibility and the bargaining problem. *Econometrica*, 47:61–73, 1979.

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

R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

Z. Neeman. The effectiveness of English auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.

T.-D. Nguyen and T. Sandholm. Optimizing prices in descending clock auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 93–110, 2014.

N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9, pages 209–241. Cambridge University Press, 2007.

N. Nisan. Algorithmic mechanism design: Through the lens of multi-unit auctions. In P. Young and S. Zamir, editors, *Handbook of Game Theory*, chapter 9, pages 477–516. North-Holland, 2014.

N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129:192–224, 2006.

J. G. Oxley. *Matroid Theory*. Oxford, 1992.

R. Paes Leme. Gross substitutability: An algorithmic survey. Working paper, 2014.

C. Papadimitriou and G. Pierrakos. Optimal deterministic auctions with correlated priors. *Games and Economic Behavior*, 92:430–454, 2015.

C. H. Papadimitriou and C. A. Wilkens. Economies with non-convex production and complexity equilibria. In *Proceedings of the 12th ACM Conference on Economics and Computation*, pages 137–146, 2011.

D. C. Parkes. Price-based information certificates for minimal-revelation combinatorial auctions. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 103–122, 2002.

D. C. Parkes and L. H. Ungar. Iterative combinatorial auctions: Theory and practice. In *Proceedings of the 17th AAAI Conference on Artificial Intelligence*, pages 74–81, 2000.

M. Pycia. Stability and preference alignment in matching and coalition formation. *Econometrica*, 80(1):323–362, 2012.

R. Rado. A note on independence functions. *Proc. London Math. Soc.*, 7:300–320, 1957.

A. Ronen. On approximating optimal auctions. In *Proc. 3rd ACM Conference on Economics and Computation*, pages 11–17, 2001.

A. Ronen and A. Saberi. On the hardness of optimal auctions. In *Proc. 43rd Symposium on Foundations of Computer Science*, pages 396–405, 2002.

A. E. Roth. What have we learned from market design? In N. Vulkan, A. E. Roth, and Z. Neeman, editors, *The Handbook of Market Design*, chapter 1. Oxford University Press, 2013.

T. Roughgarden and M. Sundararajan. Quantifying inefficiency in cost-sharing mechanisms. *Journal of the ACM*, 56(4):article 23, 2009.

T. Roughgarden and I. Talgam-Cohen. Optimal and near-optimal mechanism design with interdependent values. In *Proc. 14th ACM Conference on Economics and Computation*, pages 767–784, 2013.

T. Roughgarden and I. Talgam-Cohen. Why prices need algorithms. In *Proc. 16th ACM Conference on Economics and Computation*, pages 19–36, 2015.

T. Roughgarden, I. Talgam-Cohen, and Q. Yan. Supply-limiting mechanisms. In *Proc. 13th ACM Conference on Economics and Computation*, pages 844–861, 2012.

A. Rustichini, M. A. Satterthwaite, and S. R. Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica*, 62(5):1041–63, 1994.

M. A. Satterthwaite and S. R. Williams. The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *Review of Economic Studies*, 56(4):477–98, 1989.

M. A. Satterthwaite and S. R. Williams. The optimality of a simple market mechanism. *Econometrica*, 70(5):1841–1863, 2002.

M. A. Satterthwaite, S. R. Williams, and K. E. Zachariadis. Price discovery. Working paper, 2011.

H. E. Scarf. A min-max solution of an inventory problem. In K. J. Arrow, S. Karlin, and H. E. Scarf, editors, *Studies in the Mathematical Theory of Inventory and Production*, chapter 12, pages 201–209. Stanford University Press, 1958.

A. Schrijver. *Combinatorial Optimziation: Polyhedra and Efficiency*. Springer, 2003.

I. Segal. Optimal pricing mechanisms with unknown demand. *American Economic Review*, 93(3):509–529, 2003.

B. Sivan and V. Syrgkanis. Vickrey auctions for irregular distributions. In *Proceedings of the 9th International Workshop on Internet and Network Economics*, pages 422–435, 2013.

R. M. Starr. Quasi-equilibria in markets with non-convex preferences. *Econometrica*, 37(1):25–38, 1969.

N. Sun and Z. Yang. Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica*, 74(5):1385–1402, 2006.

N. Sun and Z. Yang. An efficient and incentive compatible dynamic auction for multiple complements. *Journal of Political Economy*, 122(2):422–488, 2014.

A. Teytelboym. Gross substitutes and complements: A simple generalization. *Economics Letters*, 123(2):135–138, 2014.

L. Ülkü. Optimal combinatorial mechanism design. *Economic Theory*, 53(2):473–498, 2013.

W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

D. Vincent and A. Manelli. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1): 153–185, 2007.

R. V. Vohra. *Mechanism Design: A Linear Programming Approach*. Econometric Society Monographs, 2011.

N. Vulkan, A. E. Roth, and Z. Neeman. Introduction. In N. Vulkan, A. E. Roth, and Z. Neeman, editors, *The Handbook of Market Design*. Oxford University Press, 2013.

D. J. A. Welsh. *Matroid Theory*. Dover Publications, 2010.

R. B. Wilson. Competitive bidding with disparate information. *Management Science*, 15(7):446–448, 1969.

R. B. Wilson. Incentive efficiency of double auctions. *Econometrica*, 53(5):1101–1115, 1985.

R. B. Wilson. Game-theoretic analyses of trading processes. *Advances in Economic Theory: Fifth World Congress*, pages 33–70, 1987.

Q. Yan. Prior-independence: A new lens for mechanism design, 2012. PhD Thesis, Stanford University.

M. A. Zinkevich, A. Blum, and T. Sandholm. On polynomial time preference elicitation with value queries. In *Proceedings of the 4th ACM Conference on Economics and Computation*, pages 176–185, 2003.